

A Comparative Evaluation of Systems for Scalable Linear Algebra-based Analytics

Anthony Thomas

Arun Kumar

University of California, San Diego

{ahthomas, arunkk}@eng.ucsd.edu

ABSTRACT

The growing use of statistical and machine learning (ML) algorithms to analyze large datasets has given rise to new systems to scale such algorithms. But implementing new scalable algorithms in low-level languages is a painful process, especially for enterprise and scientific users. To mitigate this issue, a new breed of systems expose high-level bulk *linear algebra* (LA) primitives that are scalable. By composing such LA primitives, users can write analysis algorithms in a higher-level language, while the system handles scalability issues. But there is little work on a unified comparative evaluation of the scalability, efficiency, and effectiveness of such “scalable LA systems.” We take a major step towards filling this gap. We introduce a suite of LA-specific tests based on our analysis of the data access and communication patterns of LA workloads and their use cases. Using our tests, we perform a comprehensive empirical comparison of a few popular scalable LA systems: MADlib, MLlib, SystemML, ScaLAPACK, SciDB, and TensorFlow using both synthetic data and a large real-world dataset. Our study has revealed several scalability bottlenecks, unusual performance trends, and even bugs in some systems. Our findings have already led to improvements in SystemML, with other systems’ developers also expressing interest. All of our code and data scripts are available for download at <https://adalabucsd.github.io/slab.html>.

1. INTRODUCTION

Supporting large-scale statistical and machine learning (ML) algorithms over structured data has become a mainstream requirement of modern data systems [1, 8]. Consequently, both the database and cloud service industries are scurrying to make it easier to integrate ML algorithms and frameworks such as R with data platforms [2, 9, 13, 14, 16, 32, 50]. While a roster of canned implementations might suffice for popular existing algorithms, many application users, especially in the enterprise and scientific domains often create custom analysis algorithms for their data. Such users typically write their algorithms using *linear algebra* (LA) notation and are accustomed to tools such as R, SAS, and Matlab [11, 17]. LA is a succinct and elegant mathematical language to express matrix transformations for a wide variety of statistical and ML algorithms [32, 37]. However, such LA tools are mostly not scalable to datasets that may not fit in single-node memory, which impedes effective analytics as the volume of datasets continues to grow.

The above situation has led to a new breed of analytics systems: *scalable* LA systems, which allow users to scale

their LA-based algorithms to distributed memory-based or disk-based data without needing to manually handle data distribution and communication. Thus, they offer a measure of *physical data independence* for LA-based analytics. Examples include RIOT [82], EMC’s MADlib [50], Oracle R Enterprise [16], Apache SystemML [32], Microsoft’s Revolution R [14], Spark MLlib [65] and SparkR [18], Apache Mahout Samsara [72], LA on SimSQL [61], and Google’s TensorFlow [27]. Most of these systems expose an LA-based language or API embedded in Python, Scala, or SQL as their front-end. Almost all of them aim to scale LA operations that process the whole dataset *in bulk*, i.e., in one go in a data-parallel manner. A common example of an algorithm that can be expressed in bulk LA is least-squares linear regression (Algorithm 1). This access pattern differs from that of stochastic gradient descent-based training of neural networks, which is not a bulk LA operation but rather repeatedly samples data subsets called mini-batches [66]. Thus, except for TensorFlow, the above systems do not (yet) offer strong support for neural networks, although there are active efforts to extend some of them [70]. Accordingly, the focus of this paper is on systems for bulk LA workloads.

While the recent activity on scalable LA systems has led to many interesting research papers and open-source systems, a fundamental practical question remains largely unanswered: *From a comparative standpoint, how effective and efficient are such systems for scalable LA and LA-based ML?* By effectiveness, we mean how “easy” they are to use in terms of the amount of code needed to express algorithms using their LA abstractions, as well as their setup logistics. By efficiency, we mean how “fast” they run in different settings, especially the increasingly popular distributed memory setting. Different systems have different assumptions and restrictions in their programming and execution models, leading to different runtime and memory usage behaviors. Existing papers and articles on such systems mostly focus on the operating regimes where they perform well, leaving users in the dark on where they may not perform so well. Additionally, most systems were not compared to strong baselines such as pbdR (powered by ScaLAPACK) [69]. This *lack of a unified comparative understanding* of such systems’ behaviors could lead to both user frustration and increased costs when picking a system to use, especially in the cloud.

We take a step towards filling this gap by introducing a suite of performance tests for scalable LA systems and performing a comprehensive empirical comparison of several popular and state-of-the-art LA systems using our tests. Our test suite is inspired by the long line of work on bench-

marks for LA packages in both the single-node in-memory and distributed settings, e.g. benchmarks comparing BLAS, Eigen, Intel MKL, LAPACK, and ScaLAPACK [28, 30, 41, 42, 58, 67]. But unlike such prior work, which focused only on simple LA operations, we want to evaluate the entire hierarchy of task complexity for LA-based analytics. Thus, we delineate three orthogonal axes of interest: *task complexity*, *data scale*, and *computation environment*. For data scale, we vary both the number of examples and sparsity of the data matrix. For the computation environment, we focus on commodity CPU clusters (as of this writing, few of the compared systems offer easy integration with GPUs, FPGAs, or other hardware accelerators in the distributed setting). We vary the number of CPU cores and cluster nodes.

For task complexity, we include 6 common LA operations with differing communication and computation behaviors, as well as 2 simple pipelines of LA operations and 4 LA-based data analysis algorithms, including 2 popular ML algorithms. The tasks at the higher levels of complexity present interesting opportunities for *inter-operation optimizations* (akin to relational query optimization), which are exploited by only some of the scalable LA systems. Moreover, the LA-based algorithms enable us to qualitatively compare how easy or difficult it is to use the LA abstractions provided by such systems to write such algorithms. The LA-based algorithms we include also cover a diverse set of use cases for scalable LA-based statistical and ML analytics: regression, heteroscedasticity-aware error analysis, classification, and feature extraction.

We compare the following systems: MADlib, MLlib, SystemML, ScaLAPACK, SciDB, TensorFlow, R, and NumPy. We picked them mainly due to their popularity but also because they are open source and have active user communities online. Our experimental study has revealed numerous interesting behaviors at large data scales, including unexpected tuning headaches, hidden scalability bottlenecks, unusual relative performance trends, and even bugs in some systems. For instance, we find that tuning memory-related parameters correctly is crucial to get good performance with the Spark-based systems (MLlib and SystemML), but the optimal settings are LA task-specific, making it a non-trivial task for data scientists. RDBMS-based MADlib yielded relatively poor performance for LA operations on dense data and suffers from a scalability bottleneck due to arcane restrictions imposed by the RDBMS, but it is competitive on sparse data. Even more surprisingly, pbdR/ScaLAPACK outperforms all other systems in almost all cases on dense data. Overall, we find that SystemML offers perhaps the best balance between physical data independence and performance at scale among the compared systems.

Based on our empirical study, we identify and summarize the strengths and weaknesses of each scalable LA system to help practitioners decide which one best fits their application needs. We also identify the major gaps in the capabilities of such systems and distill them into a handful of open research questions for those interested in improving such systems or building new ones. In particular, we find that support for large sparse datasets is still relatively poor; this is a pressing issue due to the ubiquity of large-domain categorical features in ML applications. We also find that many systems sometimes exhibit sub-linear multi-node speedups. We suggest that flexible hybrid parallelism models combining data and task parallelism could make such

systems more useful for crucial meta-level ML tasks such as hyper-parameter tuning [55]. Overall, this paper makes the following contributions.

- To the best of our knowledge, this is the first work to create a unified framework for comparative evaluation of popular scalable LA systems. We discuss key design considerations in terms of data access, communication and computation patterns, and use cases, and we create a suite of tests spanning the axes of task complexity, data scale, and computational environment.
- Using our tests, we perform an extensive empirical study comparing MADlib, MLlib, SystemML, pbdR/ScaLAPACK, and SciDB in the distributed memory setting, along with TensorFlow, R, and NumPy on a single node. Apart from synthetic data of various scales, we create two benchmark versions of a large real-world dataset from an online advertising firm, Criteo [5].
- We analyze and distill our results into a concrete set of guidelines for practitioners interested in using such systems. We also identify a handful of major open questions for further research on such systems.
- Our findings have already resulted in bug fixes and feature earmarks for SystemML due to our conversations with its developers. We are also speaking with the developers of some of the other systems to help them improve their systems using our results.

To improve reproducibility and to help others extend our work, all of our code for all tests on all systems compared, including configuration files, as well as all of our scripts for generating synthetic data and pre-processing Criteo data are available for download on our project webpage: <https://adalabucsd.github.io/slab.html>.

Outline. Section 2 provides some background on LA systems and overviews of the systems we compare. Section 3 explains our test suite. Section 4 presents our empirical study with both quantitative results and qualitative discussion. Section 5 presents more analysis and discussion of the implications of our findings for both practitioners and researchers. We discuss other related work in Section 6.

2. BACKGROUND

2.1 LA Systems

A data matrix $\mathbf{X}_{n \times d}$ has n examples and d features. Linear algebra (LA) is elegant and expressive formal language to capture linear transformations of \mathbf{X} (vectors and scalars are special cases). Numerous statistical and ML algorithms can be expressed using LA operations that are “bulk” vectorized computations over \mathbf{X} , including such popular algorithms as ordinary least squares (OLS) linear regression, logistic regression solved with many gradient methods [68], non-negative matrix factorization, k-means clustering, and more [37, 44, 46]. Such algorithms are common in enterprise analytics, as well as social and physical sciences.

Scalable LA systems aim to let users write LA-based algorithms but scale them to larger-than-memory datasets *transparently*, i.e., with some degree of *physical data independence*. Essentially, such systems abstract away lower-level systems issues such as parallel computation, data communication, and fault tolerance (in some cases). In a sense,

scalable LA systems aim to achieve for LA what RDBMSs achieve for relational algebra. Different systems take different routes: some layer LA as an abstraction on top of a RDBMS or Spark, while others implement LA primitives in standalone systems. Our goal is *not* to design new scalable LA systems, but to provide a systematic quantitative and qualitative comparison of state-of-the-art and popular systems on an even footing, using a carefully designed suite of tests at multiple levels of abstraction.

2.2 Overview of Compared Systems

Rationale for Systems Picked. We restrict our study to systems that are open source, well-documented, and have helpful online user communities. This enables us to properly debug unusual system behaviors. Thus, we skip other major open-source systems (e.g., RIOT, Samsara, SimSQL, and RASDAMAN). We also skip industrial tools such as Oracle R Enterprise and Microsoft Revolution R for legal reasons but hope that our work helps spur more industrial standardization in this arena. We study two popular system environments: single-node in-memory and distributed memory (with disk spills allowed) but focus on the latter. For the former, R and Python’s NumPy are the most popular tools [11], with TensorFlow being a new contender. For the latter, the primary scalable LA systems are pbdR, powered by ScaLAPACK [30,69], the array database SciDB [34], the distributed RDBMS-based MADlib [50], and Spark-based MLLib (or SparkML) [65] and SystemML [32]. TensorFlow does not yet offer distributed implementations of bulk LA operations [7]; thus, we skip it for the distributed setting.

R and NumPy. These libraries treat matrices as first-class citizens, offer a rich set of built-in LA operations and algorithms, and are widely used in academia and enterprises [11]. Most LA operations in R and NumPy are thin wrappers around highly optimized LAPACK and BLAS routines [28,39]. R and Python also provide robust visualization/plotting libraries (e.g., ggplot and matplotlib). Both R and Python are interpreted and dynamically typed, which makes optimization of LA scripts challenging.

pbdR (ScaLAPACK). ScaLAPACK extends LAPACK to distributed memory by reimplementing many LA operations and algorithms using a block partitioning scheme for matrices [30]. It follows the “same program multiple data” paradigm in which one logical program is executed by multiple workers (a worker corresponds to a core). A matrix is partitioned in a highly flexible “block-cyclic” fashion (similar to round robin), with each worker process allocated a subset of the blocks. This allocation helps load-balance computation and communication costs regardless of the access patterns of the LA operations. The block size is a user-given parameter. The “Programming with Big Data in R” (pbdR) library provides higher level R interfaces to ScaLAPACK and OpenMPI [69]. The “distributed matrix” package in pbdR overloads several built-in LA operations in R to enable transparent distributed execution. pbdR compiles programs to MPI batch jobs, which are then submitted for execution on an MPI cluster [45].

MADlib. MADlib is a library that implements both LA primitives and popular ML algorithms over PostgreSQL and the parallel RDBMS Greenplum [50]. Matrices are stored as tables and many LA operations are implemented directly in

SQL, enabling MADlib to exploit the RDBMS for memory management and scalability. Some operations, such as eigen or singular value decompositions, require the matrix to fit in single-node memory [62]. Related to MADlib (but not compared here) are RIOT-DB [82], which translates LA scripts written in R into SQL, and SimSQL [61] which implements custom datatypes for block-partitioned matrices and custom user-defined functions for LA operations.

Spark MLLib/SparkML. These are libraries that provide LA primitives and canned implementations of some popular ML algorithms on Spark RDDs. We use MLLib, since SparkML does not yet support distributed matrices. Apart from the *LocalMatrix* datatype for small data, MLLib has three main (physical) datatypes for distributed matrices targeting different data access patterns. The user must select a matrix type explicitly. The operators provided for each type are not consistent, e.g., the type supporting matrix-matrix multiplication does not support decompositions and vice versa. All matrix types support either sparse or dense data, but some operators densify sparse matrices.

SystemML on Spark. Introduced for Hadoop and then ported to Spark, SystemML is perhaps the most mature scalable LA system [32]; we use the recommended Spark version [4]. SystemML offers a “declarative” language called DML with R-like syntax to express LA scripts (there are also APIs in Python and Scala/Java). DML offers full physical data independence, i.e., users do not decide data layout formats or low-level execution details. SystemML stores matrices in a block-partitioned format implemented as a Spark RDD. Inspired by RDBMSs, SystemML has an optimizing compiler that converts a DML script to Spark jobs by applying a suite of logical LA-specific rewrite optimizations and physical execution optimizations. The compiler translates a DML script into a physical execution plan which selects operator implementations based on data and system characteristics (e.g. size, sparsity, RAM). SystemML can transparently switch between local and distributed execution.

TensorFlow. TensorFlow (TF) is a framework for expressing ML algorithms, especially neural networks [27]. It has APIs in Python and C++ for both LA primitives and canned ML implementations; we use the Python API. While TF is primarily meant for easily expressing and training complex neural network architectures using mini-batch stochastic gradient descent (SGD), it can also be used for bulk LA-based algorithms in the single-node setting. But does not provide scalable LA operations. A TF program is modeled as a graph in which nodes represent operations over multidimensional arrays (“tensors”) and edges represent dataflow. TF allows for operators to execute on different cluster nodes and handles communication under the hood. Thus, TF is able to automatically scale mini-batch algorithms to the cluster setting, but it cannot scale bulk LA algorithms, since each *individual* operator must fit on a single node.

SciDB. SciDB is a database system aimed at scientific computing applications [34]. It provides a declarative interface similar to SQL in which users compose algebraic operators that can read, write, and update multidimensional arrays. Internally, SciDB shards a logical array into a set of possibly overlapping chunks that are distributed over instances that possibly reside on different physical nodes. As of this writing, SciDB supports a fairly limited set of LA operations

(matrix-matrix multiplication and singular value decomposition). Furthermore, these operations do not support disk spills, i.e., the data must fit in distributed memory.

3. DESCRIPTION OF TESTS

We delineate our tests along three orthogonal axes; for each axis, we list parameters to vary in order to test both the computational and communication aspects of scalable LA workloads. We now explain each axis in detail.

Axis 1: Task Complexity. At a high-level, we decompose bulk LA workloads into three levels of task complexity. Each level of complexity targets different kinds of implementation choices and optimizations (if any) in scalable LA systems. The first, and lowest, level of tasks consists of basic LA operations over the data matrix. Such operators are the bedrock of LA-based ML; we denote this set of tests **MAT**. The second level consists of “simple” pipelines (compositions) of LA operations, which offer more scope for optimization; we denote this set by **PIPE**. The highest level of task complexity consists of LA-based statistical analysis and ML algorithm scripts; we denote this set by **ALG**. Section 3.1 explains the specific tasks we pick for each level in detail and why. Table 1 summarizes all tasks.

Axis 2: Data Scale Factors. Varying key properties of the data matrix lets us stress test how the systems perform at different data scales. For the sake of simplicity and uniform treatment across all systems, we assume the data matrix X has only of numeric features; categorical features are assumed to be pre-converted using one-hot encoding to obtain sparse 0/1 vectors [15]. There are three key settings: number of rows, number of columns and the fraction of non-zero cells, also called *sparsity*. We note that the number of rows and columns governs the shape of the data matrix X . We focus on varying the number of rows and data sparsity as these parameters govern most practical workloads.

Axis 3: Computational Scale Factors. This axis captures the amount of computational resources available. Once again, for the sake of simplicity and uniform treatment across all systems, we focus on commodity CPU-based compute clusters in the cloud. Thus, we only vary two settings: number of CPU cores and number of worker nodes. While this is a popular setting for many enterprises and domain scientists, we note that GPUs, FPGAs, and custom ASICs (e.g., TPUs [52]) can be used to accelerate LA workloads. However, not all scalable LA systems currently support such hardware. Thus, we leave it to future work to compare the systems in these new environments.

3.1 Task Complexity

While data and computational scale factors are standard for any benchmark comparison of data systems, the task complexity axis is LA-specific. We delineate computational and communication behaviors of scalable LA workloads and capture them using a small set of tests at three levels of complexity, summarized by Table 1. We now elaborate upon our rationale for picking these specific tests.

3.1.1 MAT: Matrix Operations

Scalable implementations of LA operations are the bedrock of scalable LA systems. But there are far too many LA operations for an exhaustive list to be practical. We resolve this issue by observing that LA operations can be grouped into a small set of categories based on commonalities in their data access patterns, memory footprints, and computational and communication behaviors on large data matrices partitioned across workers in a pre-specified way (row-wise, column-wise, or block-partitioned). We define six major categories:

- (1) Unary operations that only need partition-local local reads/writes without communication, e.g., matrix transpose and scalar-matrix multiplication. If the partitioning is row-wise (resp. column-wise), row-wise (resp. column-wise) summation also belongs to this category.
- (2) Unary operations that require communicating a constant-sized state between workers, e.g., norms, trace, and full summation. These are analogous to algebraic aggregates in SQL.
- (3) Unary operations that require communicating a state of data-dependent size between workers, typically as shuffles, e.g., Gramian and outer product. Partial summations that do not belong to the first category belong here.
- (4) Binary operations in which one of the inputs typically fits in single-node memory, e.g., matrix-vector multiplication.
- (5) Binary operations in which both are inputs larger than single-node memory but co-partitioning can avoid communication, e.g., matrix addition and Hadamard product.
- (6) Binary operations in which both inputs may be large and which usually require communicating a state of data-dependent size, e.g., matrix-matrix multiplication, which is one of the most expensive LA operations.

Based on our above analysis, we pick one operation from each category for **MAT**, as listed in Table 1. In choosing between operations, we pick those that arise commonly in LA-based ML algorithms. For instance, Frobenius norm in the second category is often used for normalizing the data, while Gramian in the third category arises in OLS.

3.1.2 PIPE: Pipelines and Decompositions

LA operations can be composed into pipelines, which are typically steps in a more complex algorithm. Such pipelines present opportunities for *inter-operator optimization*. Perhaps the best known example is matrix chain multiplication, which is analogous to join order optimization in RDBMSs [82]. Consider a simple example: $\mathbf{X}_1\mathbf{X}_2\mathbf{X}_3$, wherein \mathbf{X}_1 is $n \times 1$, \mathbf{X}_2 is $1 \times n$, and \mathbf{X}_3 is also $n \times 1$. The naive left-to-right evaluation plan computes an intermediate matrix of size $O(n^2)$, which could exhaust available memory for large n (e.g., 10 million) and waste runtime. But since matrix multiplication is associative, the right-to-left plan will yield the same result, albeit much faster, since the intermediate matrix is only $O(1)$ in size. For longer chains, one can also have bushy plans with different costs, analogous to bushy plans for multi-table joins. The cost of a plan depends on the dimensions of the base matrices. A “smart” LA system should use metadata to determine the lowest cost plan. Thus, we include this simple pipeline as a test workload. Perhaps surprisingly, most popular scalable LA systems do not optimize this pipeline; the user has to manually fix the multiplication order, which could be tedious in general.

Our second test is singular value decomposition (SVD), which has many applications, including principal compo-

Table 1: Tests/settings for Axis 1 (task complexity). Upper case symbols in bold font (e.g., \mathbf{X}) represent matrices; lower case symbols in bold font (e.g., \mathbf{w}) represent vectors. $\mathbf{X}_{i,j}$ is cell (i, j) of \mathbf{X} . ϵ is the residual vector from OLS.

Test Description	Test Semantics
MAT : Matrix Operators \sim Vary : Nodes, Cores, Rows, Sparsity	
Category (1): Matrix Transpose (TRANS)	$\mathbf{X}^T; \mathbf{X}_{i,j}^T = \mathbf{X}_{j,i}$
Category (2): Frobenius Norm (NORM)	$\ \mathbf{X}\ _F = \sqrt{\sum_i \sum_j \mathbf{X}_{i,j}^2 }$
Category (3): Gram Matrix (GRM)	$\mathbf{X}^T \mathbf{X}; (\mathbf{X}^T \mathbf{X})_{i,j} = \sum_k \mathbf{X}_{k,i} \cdot \mathbf{X}_{k,j}$
Category (4): Matrix-Vector Multiplication (MVM)	$\mathbf{X} \mathbf{w}; (\mathbf{X} \mathbf{w})_i = \sum_k \mathbf{X}_{i,k} \cdot \mathbf{w}_k$
Category (5): Matrix Addition (ADD)	$\mathbf{M} + \mathbf{N}; (\mathbf{M} + \mathbf{N})_{i,j} = \mathbf{M}_{i,j} + \mathbf{N}_{i,j}$
Category (6): Matrix Multiplication (GMM)	$\mathbf{M} \mathbf{N}; (\mathbf{M} \mathbf{N})_{i,j} = \sum_k \mathbf{M}_{i,k} \cdot \mathbf{N}_{k,j}$
PIPE : Pipelines and Decompositions \sim Vary : Rows	
Multiplication Chain (MMC)	$\mathbf{X}_1 \mathbf{X}_2 \mathbf{X}_3$
Singular Value Decomposition (SVD)	$\mathbf{U} \Sigma \mathbf{V}^T \leftarrow \mathbf{X}$
ALG : Bulk LA-based ML Algorithms \sim Vary : Nodes, Cores, Rows, Sparsity	
Ordinary Least Squares Regression (OLS)	$(\mathbf{X}^T \mathbf{X})^{-1} (\mathbf{X}^T \mathbf{y})$
Logistic Regression (LR)	See Algorithm 2
Non-negative Matrix Factorization (NMF)	See Algorithm 3
Heteroscedasticity-Robust Standard Errors (HRSE)	$(\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \text{diag}(\epsilon^2) \mathbf{X} (\mathbf{X}^T \mathbf{X})^{-1}$

ment analysis and solving systems of equations. But not all systems support SVD on larger-than-memory data. Thus, we include SVD to serve as a yardstick for the maturity of scalable LA systems.

3.1.3 ALG: Bulk LA-based ML Algorithms

This is the highest level of task complexity: algorithms expressed as bulk LA scripts. For tractability sake, we include only a handful of popular algorithms that satisfy the following desiderata. First, we should cover a variety of use cases, including regression, classification, feature extraction, and statistical analysis. Second, we should cover a spectrum of computation and communication behaviors. Third, we should cover a spectrum of implementation effort. Fourth, at least some of the compared systems should have “native” non-LA-based implementations of some algorithms (say, using SGD) to let us assess the performance costs or gains of using the LA abstractions. Finally, as a converse, we should have an LA-based algorithm that is *not* widely available as a native implementation to exemplify the utility of scalable LA abstractions, since without such abstractions, users might have to write low-level code to scale such algorithms.

Given our above analysis, we include the following tests. First is ordinary least squares (OLS) for linear regression solved using the normal equations [44]. Second is logistic regression (LR) for binary classification solved using gradient descent [68]. Third is non-negative matrix factorization (NMF) for feature extraction [47]. Fourth and final is White’s heteroscedasticity-robust standard error estimator (HRSE) for OLS, which is commonly used in the domain sciences and enterprises for handling heteroscedasticity [80]. Among these four algorithms, HRSE is *not* available as a native implementation in any of the compared systems (except MADlib). We now briefly explain each algorithm and present their respective LA scripts.

Ordinary Least Squares (OLS). OLS is the most popular algorithm for linear regression [11, 44]. Given the data matrix $\mathbf{X}_{n \times d}$ and target $\mathbf{y}_{n \times 1}$ with values in \mathbb{R} , OLS computes the projection of \mathbf{y} onto the space spanned by \mathbf{X} . OLS

is trivial to implement in LA systems, since it has a closed form analytical solution. This algorithm stress tests computation of the Gram matrix $\mathbf{X}^T \mathbf{X}$, a common pattern in ML algorithms. Algorithm 1 presents the LA script for OLS.

Algorithm 1: Ordinary Least Squares (OLS).

Inputs: \mathbf{X}, \mathbf{y}
1 return $(\mathbf{X}^T \mathbf{X})^{-1} (\mathbf{X}^T \mathbf{y})$

Logistic Regression (LR). LR is the most popular ML technique for classification [11]. It is typically solved using iterative optimization algorithms expressed in LA. A simple algorithm is batch gradient descent (BGD); while it is seldom used directly for LR, BGD’s data access and communication behaviors are similar to more popular algorithms, such as limited-memory Broyden-Fletcher-Goldfarb-Shanno (L-BFGS), that often converge in fewer iterations [56, 68]. The convergence properties of optimization algorithms and ML accuracy are *orthogonal* to the LA system used; thus, we focus only on per-iteration runtimes. This algorithm stress tests matrix-vector multiplication. Algorithm 2 presents the LA script for LR with BGD.

Algorithm 2: Logistic Regression (LR) with BGD.

Inputs: $\mathbf{X}, \mathbf{y}, I$: number of iterations, α : step size
1 for $i = 1$ to I **do**
2 | $\mathbf{w} \leftarrow \mathbf{w} - \alpha \cdot \mathbf{X}^T \left(\frac{1}{1 + e^{-\mathbf{x} \mathbf{w}}} - \mathbf{y} \right)$
3 return \mathbf{w}

Non-negative Matrix Factorization (NMF). NMF factorizes $\mathbf{X}_{n \times d}$ into two lower rank matrices such that the reconstruction error is minimized. NMF is common for feature extraction, especially on sparse data (e.g., in text mining). NMF can also use many optimization algorithms;

weighted multiplicative updates is popular [29, 59]. Algorithm 3 presents its LA script. This algorithm stress tests multiple LA operations, including matrix-matrix multiplications and multiplication chains.

Algorithm 3: NMF.

Inputs: \mathbf{X} , I : number of iterations, r : rank
1 for $i = 1$ to I do
2 $\mathbf{W} \leftarrow \mathbf{W} \odot ((\mathbf{X}\mathbf{H}^T)/(\mathbf{W}\mathbf{H}\mathbf{H}^T))$
 $\mathbf{H} \leftarrow \mathbf{H} \odot ((\mathbf{W}^T\mathbf{X})/(\mathbf{W}^T\mathbf{W}\mathbf{H}))$
3 return (\mathbf{W}, \mathbf{H})

Heteroscedasticity-Robust Standard Errors (HRSE).

Standard errors (SE) of OLS coefficients are used for hypothesis tests about the sign and magnitude of the coefficients in many domain sciences and enterprise settings. This often requires addressing heteroscedasticity—a condition that occurs when the variance of the target is correlated with feature values and may lead to incorrect SEs. To mitigate this issue, the procedure of [80], shown in algorithm 4, is often used. While the LA notation is succinct, expressing this computation in lower-level code could be tedious and painful. In spite of the popularity of such procedures, among the systems compared, only MADlib offers a robust OLS variance estimator natively. This affirms the need for scalable LA abstractions to enable data scientists to scale such LA-based procedures easily. This algorithm also stress tests multiple LA optimizations such as avoiding the materialization of the large (but ultra-sparse) matrix $\text{diag}(\epsilon^2)$.

Algorithm 4: HRSE.

Inputs: \mathbf{X} , ϵ : residuals from OLS
1 return $(\mathbf{X}^T\mathbf{X})^{-1}\mathbf{X}^T\text{diag}(\epsilon^2)\mathbf{X}(\mathbf{X}^T\mathbf{X})^{-1}$

4. EXPERIMENTAL COMPARISON

We now perform a comprehensive experimental comparison of the performance and scalability of the LA systems discussed in Section 2 using our suite of tests. Due to space constraints, we only discuss a subset of our results here for each task complexity level and present the other results in our technical report [77]. In particular, we focus primarily on the distributed memory setting here, which means most of our discussion centers on pbdR, MADlib, MLlib, and SystemML. In general, the data for our experiments will fit fully in distributed memory and disk spills should not be necessary. For a handful of tests, intermediate results *may* exceed distributed memory depending on the execution plan selected by the system. All the code for our tests on these systems, as well as our data generation scripts (for synthetic data) and data pre-processing scripts (for Criteo) are available for download at <https://adalabucsd.github.io/slab.html>.

Experimental Environment. All experiments were run on the CloudLab “Clemson” site using the c6320 instance type [71]. Each physical node has the following hardware specifications: two Intel E5-2683 v3 14-core CPUs (2.00 GHz), 256GB ECC memory, two 1 TB 7.2K RPM 3G SATA HDDs, and a dual-port Intel 10Gbe NIC (X520) network

Table 2: Software Packages and Versions

Name	Version	Name	Version
Spark	2.2.0	Hadoop	2.7.3
R	3.4.3	pbdR	0.4-2
Python	2.7.12	TensorFlow	1.4.1
NumPy	1.14.0	OpenBLAS	114fc0bae3a
SystemML	d91d24a9fa	MADlib	1.12
Greenplum	5.1.0	SciDB	18.1

adapter. On top of these nodes, we use OpenStack to create and manage a virtual cluster. Each OpenStack instance corresponds to a physical node and is allocated 24 CPU cores, 200GB RAM and 800GB of disk. The remaining resources are reserved to run OpenStack and cannot be accessed by the benchmarked application. Scripts to create and manage clusters are available on the project GitHub page. Table 2 describes the versions of software packages used.

Methodology. All runs are repeated five times. The first runtime measurement is discarded (for warming caches). The median of the rest is plotted along with error bars depicting the minimum and maximum runtimes. For MLlib, we persist data with MEMORY_AND_DISK_SER storage level to minimize the memory footprint of RDDs. Since Spark does not perform computations until the results are needed, we force the computation of specified operations in MLlib by invoking a simple “count” on the result RDD. Similarly, since SystemML eliminates dead code, we insert a break in its program flow (an “if” clause) and print a small number of matrix cells at random from within this break to force it to compute the outputs. We assume that data has been pre-loaded into the appropriate storage medium for each system (e.g. HDFS or an RDBMS) and do not include data loading times in any of our measurements. But we noticed that when using the GP-load import utility provided by MADlib, importing a 38GB CSV file took about 467 seconds, while uploading the same CSV file to HDFS was a bit faster at about 379 seconds using the HDFS command line utility. In MADlib, we have to write the output matrices to disk as tables. Thus, all runtimes reported for MADlib include this write time, since it is unavoidable. We emphasize that this writing cost is not incurred by the other systems. To simplify installation, SciDB was run from a Docker container. To verify that Docker did not introduce significant overheads, we replicated a distributed benchmark available on the Paradigm4 website and obtained comparable results.

System Configuration Tuning. Almost all compared systems have many tunable configuration parameters that affect performance. Tuning them is non-trivial and requires considerable user time and expertise. We adopted a policy of “reasonable best effort” for tuning based on the best practices obtained from each system’s documentation and online community. We also performed a series of “pre-tests” to tune key parameters, e.g., number of Greenplum segments (for MADlib), number of cores for a Spark Executor (for MLlib), and the LA library to load at runtime (for MLlib and SystemML). There is no universally optimal setting, since it depends on the LA workload run. Thus, for each system, we picked parameters that gave best performance on key LA operations - especially matrix-matrix multiplication. We think this is a reasonable and fair policy, since a typical




Figure 1: Multi-Node Dense Data for MAT Varying Rows




Figure 2: Multi-Node Sparse Data for MAT with Varying Sparsity

user is unlikely to spend much time tuning the system compared to running their actual LA workload. Due to space constraints, we discuss more details of our tuning pre-tests in the technical report [77]. We believe our system tuning efforts are reasonable from the standpoint of *typical data scientist users*. While more extensive tuning of a given system might improve its runtimes slightly, that is orthogonal to our larger goal of a unified comparative understanding. We hope our work spurs more interest among the developers and users of such systems to apply our tests for further empirical analysis.

4.1 Results for MAT

Multi-Node - Dense Data - Vary Rows. We fix the number of cluster nodes at 8, the number of columns (d) in matrices at 100, and vary the number of rows (n) for dense data. The data matrix is stored as a CSV file on disk, and varies in size from about 4 GB to 40 GB. For pbdR we generate random data in memory. For general matrix-multiplication in particular, we fix the number of columns in N at 100 and vary the number of rows, with the dimensions of M being the same as that of N^T . Figure 1 presents the results for four of the **MAT** tests. We see that the runtime for each LA operation increases steadily, but the rate of increase is faster for general matrix-multiplication, which is consistent with their communication and computation behaviors explained in Section 3.1.1.

All systems except MADlib finish in under 5s on norm and matrix-vector multiplication. MADlib’s LA operations are slower than the others systems’, primarily because it has higher per-tuple RDBMS processing overheads and the I/O time of materializing output tables. SystemML is almost always faster than MLLib, MADlib and SciDB, since SystemML mostly pulls computation into its driver and executes operators in local mode. SystemML’s runtimes see a marked rise on ADD at 10 million rows and GMM at 20 million rows; based on the execution logs, we verified that this rise happened because computation was moved to Spark at this data scale. In Spark mode, SystemML ex-

hibits comparable performance to MLLib, which always runs in distributed mode. Thus, SystemML’s ability to transparently switch between local and distributed execution based on its data size estimates is effective. Finally, pbdR exhibits strong performance on all four operators in spite of always executing in distributed mode. This underscores both the overhead imposed by the other scalable LA systems and the importance of including a strong baseline system when evaluating performance.

Multi-Node - Sparse Data - Vary Sparsity. We vary the fraction of nonzero entries in X on the same 8-node cluster. The dimensions are chosen such that X would be about 100 GB, if materialized as a CSV file. For MVM, w is dense. We did not find support for sparse data in pbdR. Figure 2 presents results. MADlib’s performance generally scales linearly with the density of the matrix. This is because MADlib represents sparse matrices using a tuple of values for row index, column index, and value. In this format, most sparse LA operators can be implemented using join-aggregate queries for which the RDBMS has been extensively optimized. This explains MADlib’s stronger showing on these tests relative to the dense matrix operators. MADlib reported an error on MVM because output is returned as an ARRAY rather than a table. At this data scale, the output array exceeds PostgreSQL’s hard size limit of 1 GB for arrays. At the time of this writing, this is a serious scalability bottleneck for MADlib that manifests in several other tests. To avoid this issue, users would need to rewrite algorithms in a less intuitive manner to operate on tables in chunks. We note finally that the gap between MLLib and MADlib can be explained, at least in part, by the fact that MLLib always densifies the right operand in GMM, which leads to extra processing overheads.




Figure 3: Single-Node Dense Data for MAT with Varying Cores




Figure 4: Multi-Node Dense Data for PIPE with Varying Rows

Single-Node - Dense Data - Vary Cores. We fix n at 20 million and d at 100, which is about 16GB. We vary the number of cores on a single node to study multicore speedup behaviors. Figure 3 presents the results for two key LA operations. The table shows runtime on a single core to facilitate comparison of absolute performance. We see that the benefit of additional cores typically plateaus at between 4 and 16 cores depending on the system. Speedup behavior varies considerably by both operator and by system. For SciDB, the strong speedup in the ADD operation can be explained by increased parallelism from adding more SciDB instances. For the in-memory systems, multicore speedups are largely dictated by the underlying LA implementation. For example, R’s implementation of ADD never uses multiple threads while TensorFlow’s does. Such implementation decisions are likely reflective of the different workloads targeted by these systems. R was designed with small to medium size data in mind. In such a context, overhead from setting up threads and copying data may exceed the benefits of parallelism for very simple binary operators like addition. Finally, we note that while TensorFlow’s runtimes are generally fastest, it performs all computations using single precision floating point numbers (`float`) by default, while the other systems all use double precision numbers (`double`).

4.2 Results for PIPE

Back to the 8-node setting, we now run matrix multiplication chain (MMC) and SVD on dense data. Figure 4 (A) presents results for SVD. For SVD we fix $d = 100$ and vary n . We compute only the singular vectors corresponding to the ten largest singular values. SystemML is not shown because it does not provide a distributed SVD. SciDB is not shown because its SVD implementation must compute all three factor matrices independently and so a fair comparison is not possible. MLLib’s strong performance for this operator is due to a novel algorithm (described in [81]), which

is optimized for map-reduce.

Figure 4 (B) presents results for MMC. We omit MADlib and SciDB as the user must specify the multiplication order explicitly. We see that only SystemML is able to determine the optimal multiplication order and avoid creating the large intermediate matrix. Both pbdR and MLLib have orders of magnitude higher runtimes and eventually crash for even modestly large n . We note, however, that neither pbdR nor MLLib claim to be able to detect this optimization; thus, this represents an avenue for improving these systems.

4.3 Results for ALG

Multi-Node - Dense Data - Vary Rows. For the same cluster and data setup as Figure 1, we run the **ALG** tests for dense data. For LR and NMF, the runtimes shown are averaged over three iterations. Figure 5 presents the results. The trends (and explanations) largely follow those of Figure 1, with the relative gap between SystemML and MLLib being larger than for the MAT tests. As with Figure 1, SystemML again performs most computations in the driver program, which explains its high performance, often surpassing pbdR. The jump in runtime for SystemML on HRSE at 10 million rows is again caused by switching to Spark mode. For pbdR, it was necessary to hand-optimize the sparse diagonal matrix multiplication in HRSE to avoid out-of-memory errors. MADlib does not automate composition of LA operators, which forces all intermediate results to be written to disk, leading to high I/O costs and runtimes for iterative algorithms. SciDB does not provide an operator for matrix inversion or solving a linear system. To implement OLS and HRSE on SciDB, we were forced to compute $(\mathbf{X}^T \mathbf{X})^{-1}$ using SVD. Furthermore, to compute the inverse, it is necessary to fully compute $SVD(\mathbf{X}^T \mathbf{X})$ thrice as the SVD operator provided is strangely only able to return a single matrix in the decomposition per call.

Single-Node - Dense Data - Vary Cores. For the same single-node setup and data scales as Figure 3, we run the **ALG** tests for dense data to study the multicore speedup behaviors of all systems. To obtain tractable runtimes for MADlib, we reduce matrix size to 2.5×10^6 . Figure 14 presents the results. Speedup behavior varies substantially by system and algorithm. R is faster than NumPy for LR and HRSE, but not on the others, likely due to the overheads of intermediate data copies for matrix-matrix multiplication [75]. Surprisingly, MLLib is comparable in performance to TensorFlow on OLS and LR. But it is slower on NMF and much slower on HRSE due to inefficiency in the method used to compute the dense-diagonal product in HRSE, which




Figure 5: Multi-Node Dense Data for ALG with varying Rows




Figure 6: Single-Node Dense Data for ALG with varying Cores

we were forced to code by hand. We linked MLLib with a native BLAS library because it yielded substantial performance improvements over the default JVM implementation. For TensorFlow, we used a special “loop” operator provided by the TF API, since it does not require us to copy data between NumPy and TF at each iteration.

4.4 Scalability on Criteo Datasets

While our above results with synthetic data already give interesting insights about the compared systems, we also want to know how these systems perform on a large real-world dataset for ML analytics. In particular, we are also interested in comparing the native, non-LA based implementations offered by some systems against their own LA-based versions. We use a large public dataset from Criteo, whose full size is 1 TB [5]. Each example is an ad display event. The target is binary, indicating if the ad was clicked or not, but we can treat it as numeric too for our purposes. There are 13 numeric features (integer counts) and 26 categorical features (32-bit hash strings) overall. For the sake of tractability in terms of runtimes, we subsample this dataset. Our sample has 200 million examples and is 50 GB in raw form. We pre-process this version to create two benchmark versions for our comparison: *DenseCriteo* and *SparseCriteo*.

DenseCriteo retains the target and 10 numeric features. Its size as a CSV file is 32 GB. We picked this size to ensure all systems are able to work and at least most of them are able to finish within our timeout constraints. We impute missing values with just zeros, since ML accuracy is orthogonal to our focus. *SparseCriteo* retains the target and 4 categorical variables, which are pre-converted with one-hot encoding [15]. The resulting data matrix has 71,000 features with a sparsity of 0.0172%. Since Criteo’s data usage agreement does not authorize data redistribution, we have released all of our scripts for downloading, pre-processing, cleaning, and obtaining the above two benchmark dataset versions on our project webpage to aid reproducibility.

We compare the scalability of the LA-based and native implementations of OLS and LR for MADlib, MLLib, and SystemML on our cluster by varying the number of cluster

nodes from one to eight. We skip GNMf and HRSE because most of these systems do not have native implementations these algorithms. SystemML’s native LR and OLS use a hybrid conjugate gradient-trust region optimization algorithm. MLLib’s native LR uses L-BFGS, while its native OLS uses IRLS. MADlib’s native LR offers a choice between conjugate gradient, iteratively reweighted least squares, and SGD. We use SGD, since it has the lowest per-iteration runtime. MADlib’s native OLS is solved directly as in algorithm 1. Note that these different optimization algorithms have different convergence properties and thus, one will almost surely need different numbers of iterations to reach a similar ML accuracy. But these properties are orthogonal to the scalable LA system used, and thus, orthogonal to our focus. For our experiments, we report the per-iteration runtime of each iterative algorithm, both LA-based and native, averaged from three iterations for a more reliable estimate.

Multi-Node - *DenseCriteo* - Vary Nodes. Figure 17(A) presents the results. We observe three major trends overall. First, across the board, the native implementations of LR and OLS have much lower per-iteration runtimes than the LA-based implementations (except for SystemML on LR). Thus, the scalable LA abstractions, especially on MADlib and MLLib, come at a high performance penalty, which users should be aware of. Second, at the largest setting of 8 nodes, all three systems have comparable performance with their native implementations, while SystemML is the fastest for LA-based implementations. Both of these trends suggest that users might be better off using the native implementations of well-known algorithms such as LR and OLS and resort to the scalable LA abstractions only if absolutely necessary. Third, all systems exhibit near-linear speedups on LR (except LA-based LR in SystemML, which is the fastest), illustrating such speedups are possible at this data scale, but almost none of them exhibit such speed-ups on OLS (except MLLib, which is the slowest). This sub-linearity suggests that communication cost for computing the Gramian in OLS remains a scalability bottleneck. We note that SystemML internally uses a lower number of partitions on the




Figure 7: Multi-Node *Criteo* for LA-based and native implementations of LR and OLS

underlying RDD than does MLLib, which limits the degree of parallelism and explains the lack of speedup after two nodes when executing in Spark mode. Increasing the number of RDD partitions increases parallelism but decreases the size of each individual task, potentially reducing throughput.

Multi-Node - *SparseCriteo* - Vary Nodes. Figure 17(B) presents the results. This plot has much fewer lines than Figure 17 (A) due to two reasons. First, the LA-based implementations of both MADlib nor MLLib, as well the native LR and OLS of MADlib were either incapable of running due to hard-coded dimensionality limits, or they crashed or timed out! For instance, the Gramian computation offered by MLLib’s `IndexedRowMatrix` datatype has a hard limit of 65,535 columns, which is lower than the 70,000 features in *SparseCriteo*. Among the LA-based implementations, only SystemML finished within the timeout but its runtimes were much higher than MLLib’s and SystemML’s respective native implementations, which obfuscated the trends. The large drop in Spark runtime between 1 and 2 nodes is because Spark must use disk cache in the single node configuration, while the multi-node configurations can use memory only.

5. ANALYSIS AND DISCUSSION

We first discuss the implementation effort needed for the ALG tests. We then provide concrete guidance to practitioners on the strengths and weaknesses of all systems and a summary of the key takeaways from our empirical analysis. Finally, we discuss some open research questions.

5.1 Algorithm Implementation Effort

To evaluate the implementation effort required for writing the **ALG**, one should ideally conduct a user study with real-world data scientists to obtain a thorough picture of system usability. In lieu of such an extensive study, we provide a rough quantification by tabulating the number of “tokens” needed to implement each algorithm in each system in Table 3. We exclude “non-essential” tokens such as calls used to induce computation by Spark.

Unsurprisingly, SystemML and pbdR generally had the most parsimonious code, since their languages have a near-math syntax for LA scripts. In contrast, SciDB generally had the most verbose code mainly due to the need to write driver programs in Python and compute matrix inverses using SVD. MADlib’s verbosity stems from its requirement of writing out each matrix operator as a separate PL/SQL call. Finally, MLLib’s verbosity is because it does not yet provide

Table 3: Tokens for Implementation of ALG Tests.

System	LR	OLS	NMF	HRSE
pbdR	42	10	49	17
SystemML	40	10	44	32
MLlib	83	26	40	35
MADlib	96	89	191	103
SciDB	109	103	183	138

a few basic LA operators in its API, which users are forced to write, and because of the need to inject type hints.

5.2 Guidelines for Practitioners

We now summarize each system’s strengths and weaknesses. We organize our discussion around the practical difficulties we faced when implementing the ML algorithms in 3.1.3 and specifically address the following questions. How difficult is system setup and configuration tuning? How robust is the support for large sparse data? To what extent is physical data independence supported? Is there support for ad hoc feature engineering/data pre-processing pipelines that precede ML training? Are fault tolerance and automatic disk spills supported? We elide discussion of the single-node tools (R, NumPy, and TensorFlow), since they are generally well known to practitioners.

pbdR. We found tuning pbdR to be simple. Default parameters yielded high performance in most cases. As of this writing, pbdR does not support sparse data. It also has poor support for data I/O; users will likely need to implement custom data readers/writers. It provides a high level of physical data independence; most conventional R scripts can be ported to pbdR with little effort. But pbdR lacks support for complex feature engineering pipelines that Scikit-learn and Spark support. It does not support fault tolerance or automatic disk spills. It also supports only a small subset of R’s native ML implementations; users will need to manually re-implement many existing algorithms available in single node R. Overall, pbdR is best suited to users who want to rapidly prototype new LA-based analysis algorithms at scale.

MLlib. We found tuning MLLib to be highly complex. We found that MLLib often crashed or exhibited very poor performance with default tuning settings. Furthermore, we found that tuning settings were often workload specific and could only be determined through a tedious process of trial and error. MLLib supports sparse matrix types, but it some-

times densifies data during execution, degrading performance. Physical data independence is limited as users must manually tune partitioning, caching behavior, and select from a set of matrix data structures supporting inconsistent operators. Moreover, some **ALG** tasks required us to write Scala code to type cast data. The Spark ecosystem offers excellent support for data pre-processing and feature engineering pipelines at scale, fault tolerance, and disk spills in many settings. MLib’s native ML implementations were typically faster than LA-based implementations, but we encountered scalability issues with its matrix types, especially for sparse data. Overall, MLib is best suited to users who require robust support for complex data pre-processing but are happy with canned ML algorithm implementations.

SystemML. We found tuning SystemML to be simpler than MLib, since it abstracts away data layout and caching decisions. SystemML has generally good support for sparse data, but the metadata overheads for empty data blocks are non-trivial and caused performance and scalability issues in some cases. SystemML offers APIs to interpolate between Spark and DML. This allows users to leverage the robust support for data pre-processing offered by the Spark ecosystem, while employing DML’s higher level LA abstraction. Since SystemML uses Spark RDDs, it inherits Spark’s fault tolerance and disk spill capabilities. SystemML’s native ML implementations typically had higher runtimes (per iteration) than the LA-based versions in our **ALG** tests, but on datasets with low d , its native ML implementations can typically converge in fewer iterations. Overall, SystemML is best suited to users who want to deploy custom LA-based ML algorithms in a production environment where robust feature pre-processing pipelines and fault-tolerance are essential.

MADlib. We found tuning MADlib on Greenplum to be relatively simple but laborious. There is not much concrete advice on tuning the number of Greenplum segments per node [10]. Determining the optimal number of segments required manually re-initializing and loading the database which was labor-intensive. MADlib has generally good support for sparse data and was competitive on the **MAT** tests with sparse data. However, it suffers from a serious scalability bottleneck due to a hard limit of 1 GB imposed on the maximum size of an array by PostgreSQL. MADlib’s use of an RDBMS gives it a reasonable level of physical data independence, since data partitioning and caching are abstracted away. However, interleaving table creation in SQL and LA computations might be unintuitive for many ML-oriented data scientists. The RDBMS offers excellent support for data pre-processing using SQL, but flexibility for writing feature engineering pipelines is poorer than Spark. Greenplum supports fault tolerance and automates disk spills (except for the array type). MADlib’s native ML implementations were typically significantly faster than LA-based implementations and were competitive with the other systems compared. Overall, MADlib is well suited to users with large RDBMS-resident datasets and who are happy with canned ML algorithm implementations.

SciDB. Similar to MADlib, tuning SciDB is simple but tedious. SciDB requires tuning (at a minimum) the number of DB instances used and the partition size of matrices. While it has support for sparse matrices, it densifies them when

computing SVD and GMM. Surprisingly, GMM and SVD are the only major LA operators provided by SciDB in the distributed setting. This forces users to manually implement other LA operators using these building blocks. SciDB does support automatic disk spills. The community edition of SciDB does not provide any “canned” ML algorithms. Overall, while SciDB is well suited to users that want to manage large multi-dimensional array data, it requires a lot of implementation effort on the part of the user when performing non-trivial LA-based ML algorithms.

5.3 Key Takeaways

We now provide a brief summary of the key insights on system design we gleaned from our experimental analysis and experience using these tools.

(1) Transparently switching between execution in local and distributed mode improves performance and reduces user effort. Recent surveys of data scientists show that a majority work with datasets that could fit on a beefy commodity node’s memory [12, 31]. By providing APIs that work mostly seamlessly in both local and distributed mode, SystemML attained good performance on smaller datasets without requiring us to re-implement code for scaling it to large datasets.

(2) Automatically detecting LA pipeline optimizations such as reordering matrix multiplication chains and multiplication of a dense matrix with a diagonal matrix (as in HRSE) yields performance benefits and reduces user effort. Once again, only SystemML detected such optimization opportunities automatically. The other systems required us to hand-optimize multiplication orders, which often required trudging through their documentation.

(3) Intermediate results should not be needlessly materialized and computations should be pipelined. MADlib’s performance suffered from the need to write all intermediate results to disk only to read them again in the next line of code. This cripples the RDBMS’s ability to perform any LA-specific *inter-operator* optimizations.

(4) Abstractions that purport to support physical data independence must lead to real reductions in user coding effort. SystemML and pbdR both provide strong physical data independence by offering only a single logical “matrix” data type. In contrast, MLib has too many matrix types with inconsistent APIs that yield dramatically different performance depending on the LA operation. This forces users to experiment with different matrix types and to manually implement LA operators not supported on a chosen type.

5.4 Open Research Questions

Extrapolating from our points above, we identify a few key gaps that require more research from the data systems community. We discuss four major groups of issues.

Better Support for Large Sparse Data. None of compared systems offered the coveted combination of strong physical data independence and high performance at scale for large sparse data, which are increasingly common in ML applications. MLib and TensorFlow have relatively poor physical data independence, and while pbdR is a strong baseline for dense data, it lacks support for sparse data. SystemML offers perhaps the best combination so far for large sparse data, but it too has high metadata overhead in some

cases. Overall, more work is needed for fully supporting sparse data for efficient and scalable LA, including optimizing data layout, staging computations and communication, and managing caching for LA scripts. This relates to the challenging problem of predicting sparsity of intermediates.

Auto-tuning Data Layout and System Parameters.

The promise of physical data independence means users should not have to expend much effort for tuning data layout and system parameters. Yet, most systems required us to tweak one or more of such parameters; some just crashed without such tuning. For instance, MLLib required manually tuning partitioning and caching commands, while MADlib required several data layout decisions for interleaving LA scripts with table creation in SQL. Such low-level decisions are likely to be unintuitive for statistical or ML-oriented users. MLLib and TensorFlow also require lower-level programming skills in Python (or Scala for MLLib) to resolve such issues. SystemML offers perhaps the most automation of such decisions, but its dependence on Spark necessitates some memory-related tuning. Overall, more work is still needed to achieve true physical data independence for scalable LA. Extending the lessons of auto-tuning from the RDBMS and MapReduce worlds (e.g., [51, 78]) to scalable LA is another avenue for new research.

Multi-node “COST” and Parallelism Models.

A surprising takeaway relates to the bedrock of database systems: multi-node data parallelism. Conventional wisdom has that using many cheap nodes is likely better than fewer expensive beefy nodes, even for dozens of GBs. But as our Criteo results show, the speedup curves sometimes flatten quickly, at least for some of the compared algorithm implementations. This issue is the multi-node version of the “COST” factor in [63]. It is not clear if this issue will be mitigated at larger scales (TBs or PBs), since that would require even more nodes to get more total memory (at least for SystemML and MLLib). In turn, that could raise communication costs for LA operations, as well as associated metadata and query processing overheads. Furthermore, recent surveys of data science practitioners show that about 70% analyze datasets under only 100 GB [12, 31]. Thus, while faster multi-node data-parallel LA implementations are useful, we think a more pressing research challenge is to support transparent scalability for hybrid *data and task parallelism*. Such support is useful for crucial meta-level ML model selection tasks such as hyper-parameter tuning [31, 53, 55, 74]. A recent Spark-TensorFlow integration does support task-parallel hyperparameter-tuning for neural networks but only for small datasets that can be broadcasted [6]. SystemML also provides framework for multi-node task parallelism, but does not automate hyperparameter tuning [33]. A related research question is determining the optimal cluster size for a given dataset and LA script, similar to how [74] sizes clusters for a few specific ML algorithms.

Including ML Accuracy for Evaluation. The faster performance of MLLib’s and MADlib’s native ML implementations compared to their LA-based versions suggests that including ML accuracy as a criterion could substantially alter the relative performance landscape. But it is non-trivial to standardize comparisons of both runtimes and ML accuracy *simultaneously* because the latter is inextricably tied to hyper-parameter tuning, which is usually a

non-smooth and non-convex meta-level optimization problem [74]. A given system will likely fall on multiple points on the accuracy-runtime *Pareto frontier* for different datasets, ML algorithms, and hyperparameter-tuning choices, which makes apples-to-apples comparisons between systems hard. Moreover, different implementations of the same ML task could have differing hyper-parameters. For instance, how does one standardize how hyperparameter-tuning should be done when comparing, say, MADlib’s SGD-based LR with SystemML’s native LR? Should sub-sampling be allowed to enable more extensive hyper-parameter tuning? We leave such questions to future work but note that there is growing interest in Pareto frontier-based benchmarks, at least for deep learning tasks, although they do not yet account for the costs of hyper-parameter tuning [38].

6. OTHER RELATED WORK

Benchmarks of LA Packages.

LA packages such as BLAS, LAPACK, ScaLAPACK, and Eigen have been extensively benchmarked [28, 30, 41, 42, 58, 67]. For instance, the LINPACK benchmark focuses on the efficiency (measured in MFLOPs per second) of solving a system of linear equations [40, 57]. BLAS has a long history of development, while LAPACK and ScaLAPACK build upon BLAS. Numerous scalable LA tools were subsequently built and benchmarked, including multicore-specific implementations in PLASMA [57] and numerous implementations to exploit GPUs and other hardware accelerators. There is a long line of work by the high-performance computing and supercomputing community on building and benchmarking distributed LA frameworks, primarily for scientific computing applications [43, 57]. Such implementations typically rely on custom compilers and communication frameworks. In contrast, our work focuses on a comprehensive comparative evaluation of recent scalable LA systems built on top of standard data systems (MADlib, MLLib, SystemML) along with TensorFlow on an even footing. We also include R, NumPy, and pbdR/ScaLAPACK as strong baselines. Recent work has also profiled R to understand its memory usage and execution overheads but their goal was to improve single-node R, not comparing scalable LA tools [75].

Comparisons of Analytics Systems.

MLlib, SystemML, MADlib, and TensorFlow were released in their current form only in the last four years. While their reference publications show several results [27, 32, 50, 65, 76], there is no uniformity in the workloads, data scales, and computational environments studied. Our work fills this crucial gap in the literature by systematically comparing them on an even footing. Recent work has also evaluated the performance of Apache Mahout Samsara but only for a few operations [72]. [64] compared SciDB [34], Myria [79], Spark, and TensorFlow for a specific scientific image processing task but not general LA or ML workloads. [35] compare Spark, GraphLab [60], SimSQL [36], and Giraph [3] for Bayesian ML models implemented using lower-level abstractions of such systems; they did not evaluate general LA or LA-based ML workloads. [31] compared MLLib’s native ML implementations with single-node tools such as Vowpal Wabbit [22] to understand the COST factor for distributed ML. One of our findings is similar in spirit but our work is more general, since we cover scalable LA and LA-based data analytics workloads, not just a few specific ML algorithms. We also

include MADlib, SystemML, and the oft-ignored pbdR/ScaLAPACK in our study. Finally, [38] introduce new benchmark criteria for comparing deep learning tools/models on both accuracy and monetary cost for some image and text prediction tasks. However, it does not yet cover bulk LA workloads or structured data analytics. Thus, overall, all these prior benchmarking efforts are largely orthogonal to our work.

Acknowledgments. This work was supported in part by Faculty Research Award gifts from Google and Opera Solutions. We thank the members of UC San Diego’s Database Lab for their feedback on this work. We thank the developers of MADlib, SciDB, and SystemML for helpful conversations about their respective systems.

7. REFERENCES

- [1] 2017 Big Data Analytics Market Survey Summary. <https://www.forbes.com/sites/louiscolombus/2017/12/24/53-of-companies-are-adopting-big-data-analytics/#4b513fce39a1>. Accessed September, 2018.
- [2] Amazon Web Services ML. <https://aws.amazon.com/machine-learning/>. Accessed September, 2018.
- [3] Apache Giraph. <http://giraph.apache.org>. Accessed September, 2018.
- [4] Apache SystemML Webpage. <https://systemml.apache.org/>. Accessed September, 2018.
- [5] Criteo Terrabyte Click Logs. <http://labs.criteo.com/2013/12/download-terabyte-click-logs/>. Accessed September, 2018.
- [6] Deep Learning with Apache Spark and TensorFlow. <https://databricks.com/blog/2016/01/25/deep-learning-with-apache-spark-and-tensorflow.html>. Accessed September, 2018.
- [7] Distributed TensorFlow Webpage. <https://www.tensorflow.org/deploy/distributed>. Accessed September, 2018.
- [8] Gartner Report on Analytics. gartner.com/it/page.jsp?id=1971516. Accessed September, 2018.
- [9] Google Cloud ML Engine. <https://cloud.google.com/ml-engine/>. Accessed September, 2018.
- [10] Greenplum Tuning Guidelines. https://gpdb.docs.pivotal.io/530/admin_guide/intro/arch_overview.html#arch_segments. Accessed September, 2018.
- [11] Kaggle Survey: The State of Data Science and ML. <https://www.kaggle.com/surveys/2017>. Accessed September, 2018.
- [12] KDNuggets Poll of Data Scientists for Largest Dataset Analyzed. <https://www.kdnuggets.com/2016/11/poll-results-largest-dataset-analyzed.html>. Accessed September, 2018.
- [13] Microsoft Azure ML. <https://azure.microsoft.com/en-us/services/machine-learning-studio/>. Accessed September, 2018.
- [14] Microsoft Revolution R. <http://blog.revolutionanalytics.com/2016/01/microsoft-r-open.html>. Accessed September, 2018.
- [15] One-Hot Encoding Example in Scikit-learn. <http://scikit-learn.org/stable/modules/generated/sklearn.preprocessing.OneHotEncoder.html>. Accessed September, 2018.
- [16] Oracle R Enterprise. www.oracle.com/technetwork/database/database-technologies/r/r-enterprise/overview/index.html. Accessed September, 2018.
- [17] Project R Webpage. r-project.org. Accessed September, 2018.
- [18] SparkR Webpage. spark.apache.org/R. Accessed September, 2018.
- [19] TensorFlow Accelerated Linear Algebra (XLA). <https://www.tensorflow.org/performance/xla/>.
- [20] TensorFlow Webpage. <https://www.tensorflow.org/>.
- [21] The Comprehensive R Archive Networks. <https://cran.r-project.org/>.
- [22] Vowpal Wabbit. https://github.com/JohnLangford/vowpal_wabbit/wiki. Accessed September, 2018.
- [23] Install greenplum oss on ubuntu, 2017. <https://greenplum.org/install-greenplum-oss-on-ubuntu/>.
- [24] Installing tensorflow from source, 2017. https://www.tensorflow.org/install/install_sources.
- [25] Spark mllib data types - rdd based api, 2017. <https://spark.apache.org/docs/latest/mllib-data-types.html#data-types-rdd-based-api>.
- [26] Using native blas in systemml, 2017. <https://apache.github.io/systemml/native-backend>.
- [27] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, M. Kudlur, J. Levenberg, R. Monga, S. Moore, D. G. Murray, B. Steiner, P. Tucker, V. Vasudevan, P. Warden, M. Wicke, Y. Yu, and X. Zheng. TensorFlow: A system for large-scale machine learning. In *Proceedings of the 12th USENIX Conference on Operating Systems Design and Implementation*. USENIX Association, 2016.
- [28] E. Anderson, Z. Bai, C. Bischof, S. Blackford, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, and D. Sorensen. *LAPACK Users' Guide*. Society for Industrial and Applied Mathematics, Philadelphia, PA, third edition, 1999.
- [29] M. W. Berry, M. Browne, A. N. Langville, V. P. Pauca, and R. J. Plemmons. Algorithms and applications for approximate nonnegative matrix factorization. *Computational Statistics & Data Analysis*, 52(1):155–173, 2007.
- [30] L. S. Blackford, J. Choi, A. Cleary, E. D’Azevedo, J. Demmel, I. Dhillon, J. Dongarra, S. Hammarling, G. Henry, A. Petitet, K. Stanley, D. Walker, and R. C. Whaley. *ScaLAPACK Users' Guide*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 1997.
- [31] C. Boden, T. Rabl, and V. Markl. Distributed Machine Learning - but at what COST? In *Machine Learning Systems Workshop at the 2017 Conference on Neural Information Processing Systems*, 2017.
- [32] M. Boehm, M. W. Dusenberry, D. Eriksson, A. V. Evfimievski, F. M. Manshadi, N. Pansare, B. Reinwald, F. R. Reiss, P. Sen, A. C. Surve, and S. Tatikonda. SystemML: Declarative machine learning on spark. *PVLDB*, 9(13):1425–1436, 2016.
- [33] M. Boehm, S. Tatikonda, B. Reinwald, P. Sen, Y. Tian, D. R. Burdick, and S. Vaithyanathan. Hybrid parallelization strategies for large-scale machine learning in systemml. *PVLDB*, 7(7):553–564, 2014.
- [34] P. G. Brown. Overview of SciDB: Large scale array storage, processing and analysis. In *Proceedings of the 2010 ACM SIGMOD International Conference on Management of Data*, SIGMOD '10, pages 963–968, New York, NY, USA, 2010. ACM.

- [35] Z. Cai, Z. J. Gao, S. Luo, L. L. Perez, Z. Vagena, and C. Jermaine. A comparison of platforms for implementing and running very large scale machine learning algorithms. In *Proceedings of the 2014 ACM SIGMOD International Conference on Management of Data*, SIGMOD '14, pages 1371–1382, New York, NY, USA, 2014. ACM.
- [36] Z. Cai, Z. Vagena, L. Perez, S. Arumugam, P. J. Haas, and C. Jermaine. Simulation of database-valued markov chains using SimSQL. In *Proceedings of the 2013 ACM SIGMOD International Conference on Management of Data*, SIGMOD '14, pages 637–648, New York, NY, USA, 2013. ACM.
- [37] L. Chen, A. Kumar, J. Naughton, and J. M. Patel. Towards linear algebra over normalized data. *PVLDB*, 10(11):1214–1225, 2017.
- [38] C. Coleman, D. Narayanan, D. Kang, T. Zhao, J. Zhang, L. Nardi, P. Bailis, K. Olukotun, C. Ré, and M. Zaharia. DAWNbench: An end-to-end deep learning benchmark and competition. In *Machine Learning Systems Workshop at the 2017 Conference on Neural Information Processing Systems*, 2017.
- [39] J. Dongarra, J. Du Croz, S. Hammarling, and R. Hanson. An extended set of fortran basic linear algebra subprograms: Model implementation and test programs. Technical report, Argonne National Lab., IL (USA), 1987.
- [40] J. Dongarra, C. Moler, J. Bunch, and G. Stewart. *LINPACK Users' Guide*. Society for Industrial and Applied Mathematics, 1979.
- [41] J. J. Dongarra, J. Du Croz, S. Hammarling, and I. S. Duff. A set of level 3 basic linear algebra subprograms. *ACM Transactions on Mathematical Software*, 16(1):1–17, 1990.
- [42] J. J. Dongarra, J. Du Croz, S. Hammarling, and R. J. Hanson. An extended set of fortran basic linear algebra subprograms. *ACM Transactions on Mathematical Software*, 14(1):1–17, 1988.
- [43] J. J. Dongarra, L. S. Duff, D. C. Sorensen, and H. A. V. Vorst. *Numerical Linear Algebra for High Performance Computers*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 1998.
- [44] J. Friedman, T. Hastie, and R. Tibshirani. *The elements of statistical learning*, volume 1. Springer series in statistics New York, 2001.
- [45] E. Gabriel, G. E. Fagg, G. Bosilca, T. Angskun, J. J. Dongarra, J. M. Squyres, V. Sahay, P. Kambadur, B. Barrett, A. Lumsdaine, R. H. Castain, D. J. Daniel, R. L. Graham, and T. S. Woodall. Open MPI: Goals, concept, and design of a next generation MPI implementation. In *Proceedings, 11th European PVM/MPI Users' Group Meeting*, pages 97–104, 2004.
- [46] A. Ghoting, R. Krishnamurthy, E. P. D. Pednault, B. Reinwald, V. Sindhvani, S. Tatikonda, Y. Tian, and S. Vaithyanathan. SystemML: Declarative machine learning on mapreduce. In *2011 IEEE 27th International Conference on Data Engineering*, pages 231–242, 2011.
- [47] N. Gillis. Introduction to nonnegative matrix factorization. *SIAG/OPT Views and News*, 25(1):7–16, 2017.
- [48] G. Guennebaud, B. Jacob, et al. Eigen v3. <http://eigen.tuxfamily.org>, 2010.
- [49] D. Hall and D. Ramage. *Breeze Documentation*, 2016. <https://github.com/scalanlp/breeze/wiki>.
- [50] J. M. Hellerstein, C. Ré, F. Schoppmann, D. Z. Wang, E. Fratkin, A. Gorajek, K. S. Ng, C. Welton, X. Feng, K. Li, and A. Kumar. The MADlib analytics library: Or MAD skills, the SQL. *PVLDB*, 5(12):1700–1711, 2012.
- [51] H. Herodotou, H. Lim, G. Luo, N. Borisov, L. Dong, F. B. Cetin, and S. Babu. Starfish: A self-tuning system for big data analytics. In *Proceedings of the Fifth Biennial Conference on Innovative Data Systems Research*, pages 261–272, 2011.
- [52] N. P. Jouppi, C. Young, N. Patil, D. Patterson, G. Agrawal, R. Bajwa, S. Bates, S. Bhatia, N. Boden, A. Borchers, R. Boyle, P.-I. Cantin, C. Chao, C. Clark, J. Coriell, M. Daley, M. Dau, J. Dean, B. Gelb, T. V. Ghaemmaghami, R. Gottipati, W. Gulland, R. Hagmann, C. R. Ho, D. Hogberg, J. Hu, R. Hundt, D. Hurt, J. Ibarz, A. Jaffey, A. Jaworski, A. Kaplan, H. Khaitan, D. Killebrew, A. Koch, N. Kumar, S. Lacy, J. Laudon, J. Law, D. Le, C. Leary, Z. Liu, K. Lucke, A. Lundin, G. MacKean, A. Maggiore, M. Mahony, K. Miller, R. Nagarajan, R. Narayanaswami, R. Ni, K. Nix, T. Norrie, M. Omernick, N. Penukonda, A. Phelps, J. Ross, M. Ross, A. Salek, E. Samadiani, C. Severn, G. Sizikov, M. Snellman, J. Souter, D. Steinberg, A. Swing, M. Tan, G. Thorson, B. Tian, H. Toma, E. Tuttle, V. Vasudevan, R. Walter, W. Wang, E. Wilcox, and D. H. Yoon. In-datacenter performance analysis of a tensor processing unit. In *Proceedings of the 44th Annual International Symposium on Computer Architecture*, pages 1–12, New York, NY, USA, 2017. ACM.
- [53] T. Kraska, A. Talwalkar, J. C. Duchi, R. Griffith, M. J. Franklin, and M. I. Jordan. MLbase: A distributed machine-learning system. In *Proceedings of the Sixth Biennial Conference on Innovative Data Systems Research*, 2013.
- [54] A. Kumar, M. Boehm, and J. Yang. Data management in machine learning: Challenges, techniques, and systems. In *Proceedings of the 2017 ACM International Conference on Management of Data*, SIGMOD '17, pages 1717–1722. ACM, 2017.
- [55] A. Kumar, R. McCann, J. Naughton, and J. M. Patel. Model selection management systems: The next frontier of advanced analytics. *ACM SIGMOD Record*, 44(4):17–22, 2016.
- [56] A. Kumar, J. Naughton, and J. M. Patel. Learning generalized linear models over normalized data. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data*, SIGMOD '15, pages 1969–1984, New York, NY, USA, 2015. ACM.
- [57] J. Kurzak, D. A. Bader, and J. Dongarra. *Scientific Computing with Multicore and Accelerators*. CRC Press, Inc., Boca Raton, FL, USA, 2010.
- [58] C. L. Lawson, R. J. Hanson, D. R. Kincaid, and F. T. Krogh. Basic linear algebra subprograms for Fortran usage. *ACM Transactions on Mathematical Software*, 5(3):308–323, 1979.

- [59] D. D. Lee and H. S. Seung. Learning the parts of objects by non-negative matrix factorization. *Nature*, 401:788, 1999.
- [60] Y. Low, J. Gonzalez, A. Kyrola, D. Bickson, C. Guestrin, and J. M. Hellerstein. GraphLab: A new framework for parallel machine learning. In *Proceedings of the Twenty-Sixth Conference on Uncertainty in Artificial Intelligence*, pages 340–349, 2010.
- [61] S. Luo, Z. J. Gao, M. Gubanov, L. L. Perez, and C. Jermaine. Scalable linear algebra on a relational database system. In *2017 IEEE 33rd International Conference on Data Engineering*, pages 523–534, 2017.
- [62] MADLib Development Team. MADLib User Documentation. <http://madlib.incubator.apache.org/docs/latest/index.html>. Accessed September, 2018.
- [63] F. McSherry, M. Isard, and D. G. Murray. Scalability! but at what COST? In *15th Workshop on Hot Topics in Operating Systems*. USENIX Association, 2015.
- [64] P. Mehta, S. Dorkenwald, D. Zhao, T. Kaftan, A. Cheung, M. Balazinska, A. Rokem, A. Connolly, J. Vanderplas, and Y. AlSayyad. Comparative evaluation of big-data systems on scientific image analytics workloads. *PVLDB*, 10(11):1226–1237, 2017.
- [65] X. Meng, J. Bradley, B. Yavuz, E. Sparks, S. Venkataraman, D. Liu, J. Freeman, D. Tsai, M. Amde, S. Owen, et al. MLLib: Machine learning in Apache Spark. *The Journal of Machine Learning Research*, 17(1):1235–1241, 2016.
- [66] T. M. Mitchell. *Machine Learning*. McGraw-Hill, Inc., New York, NY, USA, first edition, 1997.
- [67] MKL Development Team. Intel math kernel library developer reference. <https://software.intel.com/en-us/articles/mkl-reference-manual>. Accessed September, 2018.
- [68] J. Nocedal and S. J. Wright. *Numerical Optimization*. Springer, New York, NY, USA, second edition, 2006.
- [69] G. Ostrouchov, W.-C. Chen, D. Schmidt, and P. Patel. *Programming with Big Data in R*, 2012. <http://r-pbd.org/>, Accessed September, 2018.
- [70] N. Pansare, M. Dusenberry, N. Jindal, M. Boehm, B. Reinwald, and P. Sen. Deep learning with Apache SystemML. In *SysML*, 2018.
- [71] R. Ricci, E. Eide, and The CloudLab Team. Introducing CloudLab: Scientific infrastructure for advancing cloud architectures and applications. *USENIX ;login.*, 39(6), Dec. 2014.
- [72] S. Schelter, A. Palumbo, S. Quinn, S. Marthi, and A. Musselman. Samsara: Declarative machine learning on distributed dataflow systems. In *Machine Learning Systems Workshop at the 2016 Conference on Neural Information Processing Systems*, 2016.
- [73] D. Schmidt, W.-C. Chen, G. Ostrouchov, and P. Patel. *A Quick Guide for the pbdDMAT Package*. R package vignette.
- [74] E. R. Sparks, A. Talwalkar, D. Haas, M. J. Franklin, M. I. Jordan, and T. Kraska. Automating model search for large scale machine learning. In *Proceedings of the Sixth ACM Symposium on Cloud Computing*, SoCC '15, pages 368–380. ACM, 2015.
- [75] S. Sridharan and J. M. Patel. Profiling R on a contemporary processor. *PVLDB*, 8(2):173–184, 2014.
- [76] R. Taft, M. Vartak, N. R. Satish, N. Sundaram, S. Madden, and M. Stonebraker. Genbase: A complex analytics genomics benchmark. In *Proceedings of the 2014 ACM SIGMOD International Conference on Management of Data*, SIGMOD '14, pages 177–188, New York, NY, USA, 2014. ACM.
- [77] A. Thomas and A. Kumar. A comparative evaluation of systems for scalable linear algebra-based analytics - technical report. https://adalabucsd.github.io/papers/TR_2018_SLAB.pdf.
- [78] D. Van Aken, A. Pavlo, G. J. Gordon, and B. Zhang. Automatic database management system tuning through large-scale machine learning. In *Proceedings of the 2017 ACM SIGMOD International Conference on Management of Data*, SIGMOD'17, pages 1009–1024, New York, NY, USA, 2017. ACM.
- [79] J. Wang, T. Baker, M. Balazinska, D. Halperin, B. Haynes, B. Howe, D. Hutchison, S. Jain, R. Maas, P. Mehta, D. Moritz, B. Myers, J. Ortiz, D. Suciu, A. Whitaker, and S. Xu. The myria big data management and analytics system and cloud services. In *CIDR 2017, 8th Biennial Conference on Innovative Data Systems Research*, 2017.
- [80] H. White. Using least squares to approximate unknown regression functions. In *International Economic Review*, pages 149–170. JSTOR, 1980.
- [81] R. B. Zadeh and G. Carlsson. Dimension independent matrix square using mapreduce. *arXiv preprint arXiv:1304.1467*, 2013.
- [82] Y. Zhang, W. Zhang, and J. Yang. I/o-efficient statistical computing with riot. In *2010 IEEE 26th International Conference on Data Engineering*, pages 1157–1160, 2010.

Program 1: OLS in NumPy

```
def reg(X, y):
    return alg.solve(X.T.dot(X), X.T.dot(y))
```

8. APPENDIX

8.1 Additional Background on Systems Compared

Table 4 presents a summary of salient characteristics of each system in the slate we compare. We additionally provide example implementations of OLS in each of these systems. Note that token counts may not agree exactly with the numbers reported in table 3 in the body. The example programs here have been reformatted to display cleanly and exclude some “utility” routines which are needed to make them work in practice. The following sections discuss these examples.


NumPy, R. R and Python’s NumPy stack treat matrices as first class citizens and provide a rich set of built-in LA operations and algorithms. Most LA operations in R and NumPy are just thin wrappers around highly optimized LAPACK and BLAS routines [28, 39]. R and Python also provide robust visualization/plotting libraries (e.g., ggplot and matplotlib) and are Turing-complete. Since R’s syntax is close to math notation, it is especially popular among data scientists with a statistics background and in the domain sciences. In fact, the open source repository CRAN contains numerous R libraries contributed by such researchers and practitioners [21]. NumPy, in contrast, is typically more popular among data scientists with a CS background [11]; it uses a function call-oriented/object oriented syntax. Both R and Python are interpreted and dynamically typed, which makes inter-operator optimization of LA scripts challenging.

Program 1 shows the implementation OLS in NumPy. Compare this script to program 2 which presents the corresponding algorithm in R/pbdR. The `init.grid()` and `finalize()` statements are used by pbdR to initialize and finalize the MPI communicator. The pure R implementation is identical less these statements. These scripts also highlight the differing paradigms between R and NumPy. NumPy favors an object oriented syntax while R’s is imperative.

pbdR (ScaLAPACK) ScaLAPACK extends LAPACK to the distributed memory setting by re-implementing many LA operations and algorithms using a block partitioning scheme to distribute data matrices [30]. It follows the “same program multiple data” paradigm in which a single logical program is executed by multiple workers (each worker corresponds to one core). A matrix is partitioned in a highly flexible “block-cyclic” fashion (similar to round robin), with each worker process allocated a subset of the blocks. This allocation helps load-balance computation and communication costs regardless of the access patterns of the LA operations. The block size is a user-given parameter. Figure 8 gives an example of the block-cycling partitioning scheme from [73] for a 9×9 matrix with a grid of 6 workers and blocks of size 2×2 .

In general, a skewed allocation can cause performance issues. Interprocess communication is handled by the BLACS

Figure 8: Block-cyclic matrix partitioning [73].



Program 2: OLS in pbdR

```
library(pbdDMAT)
init.grid()

reg <- function(X, y) {
  b <- solve(t(X) %*% X, t(X) %*% y)
  return(b)
}

finalize()
```

library with an LA-specific message passing API. Overall, both ScaLAPACK and BLACS are low-level libraries that require knowledge of C, FORTRAN, and the intricacies of parallel computing. Thankfully, the “Programming with Big Data in R” (pbdR) library provides higher level R interfaces to ScaLAPACK and OpenMPI [69]. The “distributed matrix” package in pbdR overloads several built-in LA operations in R to enable transparent distributed execution. However, unlike regular R scripts, which can run interactively in a REPL, pbdR compiles programs into OpenMPI batch jobs that are then submitted for execution [45].

SystemML. Introduced for Hadoop and then ported to Spark, SystemML is perhaps the most mature scalable LA system [32]. We focus on the recommended Spark version [4]. SystemML offers a “declarative” language named DML with R-like syntax to express LA scripts (there are also APIs in Python and Scala). DML offers full physical data independence, i.e., users do not decide data layout formats or low-level execution details. SystemML stores matrices in custom binary formats using Spark RDDs, in particular, as block-partitioned matrices with each block stored in a tuple.

Inspired by RDBMSs, SystemML has an optimizing compiler that converts a DML script to Spark jobs by applying a suite of logical LA-specific rewrite optimizations and physical execution optimizations. The first level of this translation produces a DAG of so-called “high level operations” (HOPs), which represent basic LA operations. Each HOP is associated with one or more physical execution plans called “low-level operations” (LOPs), which are optimized for data and system characteristics. A HOP-DAG is converted to a LOP-DAG based on both rules and cost-based optimizations to minimize runtimes under memory constraints. LOPs are executed as either RDD operations or in-memory computations in the driver program. SystemML also includes more advanced optimizations such as dead code elimination and operation fusion to reduce data access costs.

As can be seen in program 3, SystemML’s syntax is virtually identical to pbdR (and yields comparable performance!). SystemML also allows users to pass metadata about input matrices which can help the optimizer determine better execution plans.

Program 3: OLS in SystemML

```
reg = function(matrix[double] X,  
              matrix[double] y)  
  return (matrix[double] b) {  
    b = solve(t(X) %*% X, t(X) %*% y)  
  }
```

TensorFlow. TensorFlow (TF) is a framework for expressing ML algorithms, especially neural networks [27]. It has APIs in Python and C++ for both LA primitives and canned ML implementations; we use the Python API. While TF is primarily meant for easily expressing and training complex neural network architectures using mini-batch stochastic gradient descent (SGD), it can also be used for bulk LA-based algorithms in the single-node setting. Models in TF are expressed as “computational graphs” in which nodes represent operations over multi-dimensional arrays (“tensors”) and edges represent dataflow. A TF program has two stages. First, the computational graph is specified and placed on available compute devices. Then, a node is “run” (not necessarily a terminal node), with its inputs made available. Separating these stages enables TF to use lazy evaluation to compile the graph and apply some holistic optimizations within Python’s interpreted environment.

TF’s LA API offers a set of high-level routines (“operators”), each with one or more physical implementations (“kernels”) for specific compute devices. TF supports many device backends, including CPUs, GPUs, and Android smart phones. User can assign different operations to different devices; if such assignments are not specified, TF optimizes the placement to reduce data movement and runtimes. TF also detects and removes redundant computations, but overall, its LA-specific optimizations are not yet as extensive as SystemML’s. Also, TF does not yet offer distributed bulk LA operations. Thus, we consider TF a single-node tool. TF is under intense development, with optimizations such as XLA [19] and new extensions being introduced. We refer the interested reader to their webpage for the latest [20].

Program 4 presents OLS in TensorFlow’s low level LA API. While TF bears some superficial similarity to NumPy, its syntax and programming model are quite different. The lines of code between the graph declaration and the `tf.Session()` constructor merely place operations on the computation graph. When `run()` is called on an object from the graph, TensorFlow compiles the graph into an execution plan and runs the actual computation.

MADlib. MADlib is a library that implements both LA primitives and popular ML algorithms over PostgreSQL and the parallel RDBMS Greenplum [50]. Dense matrices are stored as tables with two attributes: an integer row number (the key) and a value attribute that uses the abstract data type `ARRAY`. Sparse matrices are stored as tables with three attributes: row number, column number, and cell value. Thus, MADlib expresses many LA operations directly in SQL and exploits the RDBMS for memory management and scalability. Low-level in-memory LA operations such as inner products exploit Eigen [48]. To write LA scripts, one has to write SQL queries invoking MADlib’s LA routines. Some LA operations such as certain matrix decompositions, however, require the dataset to fit entirely in single-node

Program 4: OLS in TensorFlow’s LA API

```
def reg(Xdata, ydata):  
  G = tf.Graph()  
  with G.as_default():  
    X = tf.placeholder(tf.float32,  
                      shape=Xdata.shape)  
    y = tf.placeholder(tf.float32,  
                      shape=ydata.shape)  
  
    b = tf.matrix_solve(  
        tf.matmul(X, X, transpose_a=True),  
        tf.matmul(X, y, transpose_a=True)  
    )  
  
    init = tf.global_variables_initializer()  
    with tf.Session() as sess:  
        sess.run(init)  
        res = sess.run(b,  
                       feed_dict={X: Xdata, y: ydata})  
  
  return res
```

Program 5: OLS in MADLib

```
DROP TABLE IF EXISTS XT  
DROP TABLE IF EXISTS XTX  
DROP TABLE IF EXISTS XTY  
DROP TABLE IF EXISTS XTX_INV  
DROP TABLE IF EXISTS B  
  
SELECT madlib.matrix_mult(  
  'X', 'trans=True', 'X', NULL, 'XTX');  
SELECT madlib.matrix_mult(  
  'X', 'trans=True', 'y', NULL, 'XTY');  
SELECT madlib.matrix_inverse(  
  'XTX', NULL, 'XTX_INV', NULL);  
SELECT madlib.matrix_mult(  
  'XTX_INV', NULL, 'XTY', NULL, 'B');
```

memory [62]. Related to MADlib is RIOT-DB [82], which avoids SQL as a front-end and opts for the so-called “query generation” approach [54]. A user writes an LA script in R using RIOT-DB’s datatypes, which is then translated and optimized (via lazy evaluation) to produce several (procedural) SQL queries [82]. Also related is the recent SimSQL [61], which relies on custom user-defined datatypes for block-partitioned matrices and custom user-defined functions that implement LA operations.

Program 5 presents an implementation of OLS using MADlib’s matrix API. Note that it is not possible to pipeline operations in MADlib as in the other languages. The MADlib API requires that each intermediate stage of computation be explicitly materialized. We remark that although this code explicitly computes the inverse of the Gram matrix, the number of columns we use is small and so this step remains low cost.

Mllib. Mllib (and the newer SparkML) are libraries that provide some LA primitives and popular ML algorithms over Spark. We focus on Mllib, since it is popular among enterprise users [1], and since SparkML does not yet support distributed matrices. Apart from the *LocalMatrix* datatype for small data, Mllib offers three main (physical) datatypes for distributed matrices targeting different data access pat-

terns. (1) *DistributedRowMatrix* (DRM), an RDD with rows of a logical matrix stored using a *LocalVector* datatype. (2) *CoordinateMatrix* (CM), an RDD with tripes of row number, column number, and data value (like MADlib’s sparse matrix table). (3) *BlockMatrix* (BM), an RDD of matrix blocks stored using *LocalMatrix*.

DRM supports multiplication with a *LocalMatrix* but not fully distributed matrix multiplication. DRM also supports scalable matrix decompositions such as SVD and QR. CM supports no meaningful scalable LA operations except transpose. Users have to cast it to another distributed matrix type. Sparsity is preserved after casting. BM is the only type that supports fully distributed matrix multiplication. The underlying LA operations over local datatypes are implemented using the Breeze library in Scala [49]. MLLib’s datatypes do not yet support many basic LA operations, including scalar-matrix multiplication, norms, and Hadamard product; users have to implement these using RDD operations. Also, the three distributed matrix types are not consistent in the set of LA operations they support, e.g., BM supports transpose but DRM does not, while DRM supports multiplication by a *LocalMatrix* but BM does not.

Program 6 presents an implementation of OLS using the distributed matrix types provided by MLLib. The script highlights two of the distributed types provided by MLLib. We store the input matrix X as a *IndexedRowMatrix* which represents a distributed matrix using an RDD of local vectors. We selected this datatype because it provides an extremely efficient routine for computing the Gram matrix. However, it supports neither transposition nor multiplication with another distributed matrix and so we must cast it to a *BlockMatrix* to transpose and multiply with y . This highlights several important implementation decisions which pbdR and SystemML abstract from the user. It is unclear whether it would be more efficient to simply store everything as a *BlockMatrix* to begin with and eliminate the step of casting X . However, then we forgo the optimized routine to compute a Gram matrix provided by *IndexedRowMatrix* and must make do with the (much) slower general matrix-matrix multiplication method provided by *BlockMatrix*. Additionally, we could have stored y as an *LocalMatrix* and used the “multiply” method of *IndexedRowMatrix* which supports multiplication with a local matrix. However, then our code would break if y exceeds single node memory. Even then, we would have to cast X to a *CoordinateMatrix* and then back to an *IndexedRowMatrix* to perform the operation which may result in shuffle. Without testing it is unclear which approach is better. Finally, we note that it is necessary to select the block sizes used for the *BlockMatrix*. Naively accepting the default of 1024×1024 would be catastrophic as the blocks are created as sparse matrices, but are promoted to dense during multiplication. This would then result in significant wasted space. This lack of strong physical data independence leads to some implementation headaches with MLLib.

SciDB. Program 7 presents an implementation of OLS in SciDB’s “array query language.” There are several noteworthy points about this program. First, it is much more verbose than the other programs. This is because SciDB does not provide an operator for solving a linear system or inverting a matrix directly. This necessitates the use of the singular value decomposition to compute the inverse. Fur-

Program 6: OLS in MLLib

```
def reg(X: IndexedRowMatrix,
       y: IndexedRowMatrix) : Matrix = {
  val XTX = X.computeGramianMatrix()
  val XTY = X.toBlockMatrix(
    1024, X.numCols.toInt).
    transpose.multiply(
      y.toBlockMatrix(1024, 1), 500
    ).toLocalMatrix
  val b = from_breeze(
    to_dense(as_breeze(XTX)) \
    to_dense(as_breeze(XTY)))
  return b
}
```

Program 7: OLS in SciDB

```
store(gemm(X, X, Z, transa:true), XTX);

store(gemm(project(apply(cross_join(
  transpose(gesvd(XTX, 'VT')) as V,
  project(apply(
    gesvd(XTX, 'S'), sigma_inv,
    POW(sigma, -1)), sigma_inv)
  AS SINV, V.i, SINV.i),
  vsinv, v*sigma_inv), vsinv),
  transpose(
    gesvd(XTX, 'U')), Z), XTX_INV);

gemm(XTX_INV, gemm(X, y, Z, transa:true), Z)
```

thermore, because SciDB operators must return only a single table, the decomposition must be performed three times to compute each constituent matrix! We observe as well that the `gemm` operator for matrix-multiplication computes $MN + Q$. The final matrix Q is not optional, and the user must materialize a table of zeros to compute the standard matrix product.

8.2 Additional Detail on Tuning and Configuration

In the following section we describe the process used to tune each system and provide results from a series of “mini-tests” designed to test configuration parameters.

Choice of Linear Algebra Library. Both SystemML and MLLib use a JVM-based linear algebra library by default, but allow users to import a system optimized BLAS library at runtime. Native BLAS may outperform JVM-based implementations depending on the data size and operator and so we first consider the effect of LA library on SystemML and MLLib performance. We compile the popular OpenBLAS library from source following the instructions in [26]. Table 6 shows the effect of using the JVM-based BLAS vs. OpenBLAS for GMM in the single node setting. We find only a small and inconsistent effect of using OpenBLAS with SystemML and so retain the default LA library in all SystemML tests as this makes implementation simpler. Tables 7 and 8 show the effect of using OpenBLAS in the single node and distributed settings for Spark MLLib. We find a strong positive effect of using OpenBLAS in the single node setting, but only small and inconsistent effects in the distributed setting. For GMM this small benefit is likely

Table 4: Key characteristics of systems compared. “Applicable Environments” are the environments the system was primarily designed for: “SM” is single-node in-memory, “SD” is single-node disk-capable, “DM” is distributed memory, and “DD” is distributed disk-capable. By “Partial Declarativity,” we mean that the system optimizes LA scripts only in a limited way or not at all even if it supports alternative physical implementations of LA operations. “Full Declarativity” means the system optimizes LA scripts both logically and physically. Base R does not support sparse matrices but user packages fill this gap. TensorFlow has a sparse matrix library but it has limited support for sparse LA operations.

	R ; NumPy	MADlib	TensorFlow	SystemML	MLlib
Applicable Environment(s)	SM	SD, DD	SM, DM	DM, DD	DM, DD
Interface Language(s)	R ; Python	SQL	Python, C++	DM, Python	Python, Scala
Storage Back-End	In-memory	RDBMS	Flat files	Spark/HDFS	Spark/HDFS
Declarativity	None	Partial	Partial	Full	Partial
Optimization	None	Cost-Based	Cost-Based	Cost-Based	None
Sparse Matrix Support	Yes	Yes	Partial	Yes	Partial
Implementation Language	C	C++/SQL	C++	Java/Scala	Scala
Linear Algebra Library	BLAS	Eigen	Eigen	Apache Commons Math	Breeze (JBLAS)

because the *BlockMatrix* constructor recommended by the documentation construct *sparse* matrix blocks [25]. Sparse matrix computations are not optimized by the native BLAS. Based on these tests we use OpenBLAS for single node MLlib tests but retain the JVM implementation in the distributed setting.

Additional Tuning Consideration for Spark MLlib.

An additional consideration with Spark is the number partitions in RDDs underlying its matrix types. Ad-hoc comparison indicated that between 500 and 1000 partitions yielded good performance. Another important consideration is the amount of RAM and CPUs allocated to each executor. When using the standalone cluster manager, Spark calculates the number of executor instances based on these parameter settings. We tuned these parameters by starting with a single large executor which was given all memory and CPU available and then dividing the memory and CPU allocated to each executor by two until performance stopped improving. We found that using 3 cores and roughly 20GB of RAM per executor led to the best performance. With these settings, Spark allocated eight executors per node. Results from this mini-benchmark are presented in ??.

In yet another tunable setting, MLlib allows users to adjust the number of “mid-dim-splits” used during distributed matrix multiplication. Tuning this parameter can lead to increased parallelism and reduced shuffle. Table 10 compares various settings of this parameter. We find 500 mid-dim-splits to be optimal and so use this number for all distributed matrix multiplications in MLlib.

SciDB. The key tuning considerations for SciDB are the number of instances used per node and the size of the sub-blocks into which arrays are partitioned. Based on our results in figure 3 we found that 24 instances (1 instance per core) was generally optimal in the single node setting. In the eight node cluster, we found that using such a large number of instances led to substantial overheads. The SciDB manager instances must broadcast queries to each worker instance. This can lead to significant overheads on large clusters. We found that 8 instances per cluster was optimal, and replicated a benchmark available on the SciDB website to validate our tuning settings. SciDB requires that matrix

blocks be square for GMM and so we compared block sizes of 100, 500 and 1000 and found 1000 to be generally optimal.

Greenplum. A key tuning consideration for Greenplum is the number of segments used per cluster node. Each segment corresponds to a Postgres database instance which communicates with other segments to perform work. Too few segments may result in low parallelism while too many may lead to excessive communication overhead. Greenplum documentation states that common practice is to use between two and eight segments per host. Table 12 compares the effect of using six, twelve and twenty-four segments per host on a cluster with eight nodes. We find that six segments per node yielded the best performance and so use this setting for all distributed tests. In the single node setting we found that sixteen to twenty four segments yielded good performance and so use twenty four segments in the single node setting. Greenplum also allows users to tune the memory available to the DBMS. We set `gp_vmem_protect_limit=13000` which allows each Postgres instance to use up to 13000MB of RAM. We note that we saw only very little difference between modifying this setting and leaving it at its default. Greenplum additionally allows users to configure memory for individual queries through use of the `SET STATEMENT_MEM = 'X'` command. We tried various values of statement memory but found either little effect or that they resulted in memory errors. We therefore leave `STATEMENT_MEM` at its default value (2000MB).

8.3 Test Environment and Configuration

The following sections discuss relevant installation and configuration details for each software package used. We stress that it is not necessary to manually replicate these steps. Scripts are available which automate configuration of cluster nodes.

Spark. In the single node setting we compile Spark from source using the command `mvn -DskipTests -Pnetlib-lgpl clean package` in order to take advantage of native BLAS accelerations. In the distributed setting we use a precompiled binary downloaded from Apache. We write all Spark code using the Scala API and create fat JAR files using SBT assembly. We run Spark using the standalone cluster manager (as opposed to YARN) and submit JARs us-

ing `spark-submit`. In the single node setting we configure Spark to import OpenBLAS as the LA backend at runtime. Instructions to do this are specific to the OS and LA library. The node configuration script provided in the project repo provides instructions for Ubuntu 16.04.

OpenBLAS. In the single node setting we compile OpenBLAS from source as described in [26] and use OpenMP as the threading implementation. We pull the OpenBLAS repo at revision `114fc0bae3a` and compile using `sudo make USE_OPENMP=1`. We found it was necessary to manually add some symbolic links to coax SystemML and MLLib into using this BLAS.

Greenplum and MADLib. We compile Greenplum from source as described in the cluster configuration script available on the project github page. We note that since the project began, Pivotal has made a PPA for Ubuntu available which allows Greenplum to be easily installed using `apt-get` [23]. Users may wish to try this method first as compiling Greenplum from source is a non-trivial process. Following the advice available on the Greenplum github page we configure several system parameters as described in table 13. We compiled MADLib from source (see `install-madlib.sh`). We note that it is important to use GCC/G++ 4.9 which must be installed separately from the Ubuntu repositories. Using the default GCC 5.x resulted in errors at runtime.

TensorFlow. We compile TensorFlow from source as described in [24]. We do not enable any of the extra packages available during the “configure” stage of installation and do not build with GPU support.

SystemML. We compile SystemML from source using revision `d91d24a9fa`. We built from this version because it contained patches designed to address an issue we encountered using sparse matrices. We compiled using `mvn clean package` and then manually copied the resulting `.jar` file to the `/lib` folder of each SBT assembly directory. SBT will then automatically package the JAR with other code. The precompiled JAR can be obtained by simply downloading the relevant directories from the project github.

pbdr. We install pbdr using R’s built in package manager. R expects to link to a BLAS library at runtime. In the distributed setting we install OpenBLAS using `apt-get` on each node. Because of the “SPMD” programming model used by MPI, each R process must have access to the source file containing code to be executed. To ensure R processes in remote nodes have access to both the test script and the pbdr library source, we NFS share the home directory over the cluster. We use the “OpenMPI” [45] installation available through the Ubuntu repositories.

Misc. Single Node Tools. We install R in the single node setting from the “R-Studio” PPA using `apt-get`. We configured R to link against the version of OpenBLAS compiled from source (in the single node setting only). We installed NumPy using `pip install`. NumPy ships with a built in BLAS implementation which performs quite well and so we did not link against the native library.

8.4 Additional Experimental Results

We here present and discuss additional results not presented in the body.

8.4.1 Additional Results for MAT

Multi-Node Dense Data Vary Rows Figure 9 presents the remaining matrix operators for tests which fix the cluster size at eight nodes and vary the number of rows in input matrices. We point out that Spark MLLib has an optimized Gram matrix computation routine which performs very well - even beating out pbdr which typically led the pack for distributed matrix ops. We make one concluding remark about matrix operator tests for MLLib. For matrix addition, only the *BlockMatrix* type provides a method supporting addition with another distributed matrix. We found this method to be slow in practice. Because of this, we implement a simple “add” function for the indexed row type which joins a pair of row matrices by row id and then maps over the resulting pair RDD, producing a new RDD which is the sum of the row vectors. We present a comparison in table 11. This method substantially out-performed the built in method and so all numbers reported in plots use this approach.

Multi-Node Sparse Data Vary Sparsity Figure 10 presents the remaining matrix operators for tests which fix the cluster size at eight nodes and vary the sparsity of input matrices. We remark that MADLib crashed during computation of Gram matrix for the largest matrix size due to an attempt to materialize a large array. We note that SystemML out-performs MLLib’s optimized Gram matrix computation on sparse data, although this is likely because it is performing computation in the driver.

Multi-Node Dense Data Vary Nodes Figure 11 presents results from tests which fix matrix dimensions at 20 million rows by 100 columns and vary the number of nodes in the cluster from two to eight. We note that speedups from scaling the number of nodes are remarkably small - especially for MADLib. We note as well that MLLib timed out for GMM in the two node setting.

Multi-Node Sparse Data Vary Nodes Figure 12 presents results from tests which fix data sparsity at 1% and vary the number of nodes in the cluster from two to eight. The upper panel plots raw runtime, and the bottom plots speedup curves relative to a single core. MADLib timed out for GRM computation on a single segment. SystemML shows little consistent benefit from scaling the number of nodes. This is because it is pulling data into the driver and performing computation in single node mode. Interestingly, MLLib and MADLib show slightly more benefit from adding cluster nodes in the sparse setting than dense.

Single-Node Dense Data Vary Rows Figure 13 presents results from tests which scale the number of rows in input matrices in the single node setting. Matrix sizes are as in distributed tests. Note that we here introduce new baseline systems - R and NumPy. In this context, R is conventional single node R as opposed to the pbdr flavor used in the distributed setting. In the single node setting we use MLLib’s *local* matrix types. In the local setting we found that 24 segments yielded optimal (or close to optimal) performance for Greenplum and so all numbers reported here use 24 segments. With the exception of MADLib, performance of all systems is fairly consistent. This is unsurprising as all systems are calling out to libraries which have been heavily optimized for the single-node in memory setting. R suffers from some well known “copy overhead” in certain contexts [75] which likely explains its modest gap relative to other

systems on norm and MVM. We note that using a native BLAS was critical to obtaining good performance from ML-Lib in the single node setting. Using the JVM-based implementation bundled with Spark resulted in poor performance relative to other systems.

Single-Node Dense Data Vary Cores Figure 15 presents results from tests which scale the number of CPU cores for selected single node dense matrix operators. The top panel plots raw numbers (including MADLib) while the bottom plots a “speedup curve” relative to time with a single core. As noted previously, we use the unix `taskset` command to pin each process to a specific subset of CPU cores. For Spark-based systems we explicitly restrict the number of driver cores. For Greenplum, we build separate database instances with the stipulated number of segments. As was remarked previously, speedups (at least on the operators considered) are generally sublinear. Interestingly, several systems exhibit different speedup behaviors between the two operators. This likely reflects differing parallelization strategies of the underlying linear algebra library used by each system. For example, Eigen multi-threads only a small subset of the operators parallelized by OpenBLAS.

8.4.2 Additional Results for ALG

Single-Node Dense Data Vary Cores Figure 14 presents complete results from tests which fix the size of input data and vary the number of cores allocated. For all systems but MADLib we fix the size of input data at $10,000,000 \times 100$. For MADLib, we fix the size of data at $2,500,000 \times 100$ to obtain tractable run-times. The top panel plots raw runtimes while the bottom plots “speedup curves” relative to the run-time for a single core.

Multi-Node Sparse Data Vary Sparsity Figure 16 presents results from tests which fix the cluster size at eight nodes and vary data sparsity. We here present results only for ML-Lib and SystemML as MADLib either timed out or crashed when attempting to allocate a large array for both tests. SciDB was not run for this test. We present results only for GNMF and HRSE as LR and REG were examined in tests on *SparseCriteo*. Somewhat surprisingly, we found that SystemML executed substantially faster when forced to run in “spark-only” mode. Using our default driver memory of 32G, SystemML attempted to perform the computation in the driver and crashed. We had to increase driver memory to 80G before SystemML could complete these tests. However, we found that *reducing* driver memory to 10G, which forces SystemML to perform most computation on Spark, resulted in roughly 3x better performance. By default, SystemML will execute computation in the driver if it estimates driver memory is sufficient. The results of this test highlight two important issues - first, even if computation *can* fit in the driver, it may not be optimal to do so, and second, SystemML often does not correctly estimate the memory footprint of computations which necessitates tedious trial and error tuning of JVM heap sizes. We note as well that these tests use a single large spark executor provisioned with all available RAM and CPU. Our default “small-executor” tuning settings sometimes caused out of memory errors and did not substantially improve performance.

Multi-Node DenseCriteo Vary Nodes Figure 17 presents results from tests which compare native and LA-based PCA

implementation on the dense Criteo AdClick dataset. We compute the five strongest principal components. All tests additionally re-project the input data \mathbf{X} under the corresponding principal components. For our LA-based implementation, we compute the principal components using an Eigen decomposition of the data covariance matrix. This test also uses a single large spark executor per node.

Multi-Node SparseCriteo Vary Nodes Table 5 presents results from our LA-based implementations of logistic regression run on the sparse Criteo dataset. We report only logistic regression because this was the only test which could be completed by any system. Both PCA and OLS regression crashed when attempting to materialize the large $\mathbf{X}^T \mathbf{X}$ matrix necessary for both algorithms. The table contains results only for SystemML as we could not obtain results for any other system due to timeouts.

Table 5: Multi-Node *SparseCriteo* for SystemML with Varying Nodes

Nodes	Run Times (Seconds) - LR				
2	13,870.24	7,367.51	6,972.21	7,643.25	7,033.11
4	3,498.50	2,004.59	2,059.89	2,399.09	2,471.37
8	505.08	450.20	432.31	428.14	445.43

Table 6: Effect of LA Library on SystemML Performance (Single Node)

	GMM Run Time (Seconds)				
Open BLAS	10.58	11.93	10.75	11.44	11.17
Commons Math	7.96	8.49	10.14	13.60	14.45

Table 7: Effect of LA Library on MLLib Performance (8 Nodes)

	GMM Run Time (Seconds)				
Open BLAS	3,404.50	1,872.21	2,068.55	1,587.44	1,727.96
Netlib-Java	2,597.51	1,454.14	1,378.21	1,859.64	1,537.38
	TSM Run Time (Seconds)				
Open BLAS	38.71	29.31	25.37	25.27	23.88
Netlib-Java	5.91	2.11	2.07	3.73	3.32
	ADD Run Time (Seconds)				
Open BLAS	77.67	46.97	37.80	29.50	34.36
Netlib-Java	159.37	38.30	46.11	47.27	28.24
	NORM Run Time (Seconds)				
Open BLAS	3.02	4.55	2.34	4.61	3.49
Netlib-Java	5.33	4.55	4.48	4.54	4.47
	MVM Run Time (Seconds)				
Open BLAS	5.45	1.84	2.86	2.40	1.67
Netlib-Java	5.25	3.60	2.57	3.36	2.57

Table 8: Effect of LA Library on MLLib Performance (Single Node)

	GMM Run Time (Seconds)				
Open BLAS	64.84	47.56	7.82	6.68	6.50
Netlib-Java	282.14	280.85	284.02	289.75	276.72
	ADD Run Time (Seconds)				
Open BLAS	223.88	195.60	58.82	85.26	58.76
Netlib-Java	258.83	79.00	64.46	53.96	60.97
	NORM Run Time (Seconds)				
Open BLAS	144.17	312.43	69.13	103.51	65.35
Netlib-Java	86.40	211.62	68.58	124.10	69.49
	MVM Run Time (Seconds)				
Open BLAS	0.72	0.36	0.39	1.45	0.36
Netlib-Java	3.92	3.66	3.75	5.30	3.32

Table 9: Effect of Number of Executor Settings on MLLib Performance

Memory (GB)	Cores	GMM Run Time (Seconds)				
200	24	5,746.80	2,982.14	2,900.22	624.84	243.63
100	12	836.57	290.92	276.03	269.36	245.12
50	6	826.93	342.80	243.16	235.86	255.95
25	3	764.11	230.12	257.69	269.57	257.46

Table 10: Effect of “NumMidDimSplits” on MLLib GMM Performance (8 Nodes)

	GMM Run Time (Seconds)				
1	2,466.73	1,440.77	1,391.03	1,477.93	1,361.21
500	2,308.37	926.07	739.91	630.80	761.83
1000	2,224.66	927.77	1,198.63	1,124.36	918.35




Figure 9: Multi-Node Dense Data for MAT with varying Rows




Figure 10: Multi-Node Sparse Data for MAT with Varying Sparsity




Figure 11: Multi-Node Dense Data for MAT with Varying Nodes

Table 11: Comparison of Matrix Addition Methods

Method	Run Time (Seconds)				
Native	722.58	615.02	413.82	593.22	646.10
Custom	159.37	38.30	46.11	47.27	28.24




Figure 12: Multi-Node Sparse Data for MAT with Varying Nodes




Figure 13: Single-Node Dense Data for MAT with Varying Rows




Figure 14: Single-Node Dense Data for ALG with Varying Cores




Figure 15: Single-Node Dense Data for MAT with Varying Cores




Figure 16: Multi-Node Sparse Data for ALG with Varying Sparsity

Figure 17: Multi-Node *DenseCriteo* for LA-based and Native Implementations of LR




Table 12: Effect of Number of Segments on MADLib Performance

	GMM Run Time (Seconds)				
6	1,234.10	1,097.47	1,153.95	1,165.01	1,208.63
12	--	--	--	--	--
24	2,321.03	2,182.65	2,129.74	2,175.40	2,095.39
	ADD Run Time (Seconds)				
6	251.41	213.41	143.87	137.61	186.11
12	--	--	--	--	--
24	321.93	263.73	282.65	291.18	251.34
	NORM Run Time (Seconds)				
6	40.08	59.03	47.08	51.94	81.89
12	90.17	106.00	86.98	144.68	98.37
24	95.64	94.69	97.86	105.18	94.55
	MVM Run Time (Seconds)				
6	216.97	203.94	227.01	231.20	209.47
12	265.36	247.24	248.48	244.74	253.37
24	236.66	241.99	247.83	244.80	244.49

Table 13: System Parameters Used

Parameter Name	Value
/etc/security/limits.conf	
nofile (soft/hard)	131093
nproc (soft/hard)	131072
/etc/security/limits.conf	
net.ipv6.conf.all.disable_ipv6	1
net.ipv4.tcp_syncookies	0
net.ipv4.conf.default.accept_source_route	0
net.ipv4.tcp_tw_recycle	1
net.ipv4.tcp_max_syn_backlog	4096
net.ipv4.conf.all.arp_filter	1
net.ipv4.ip_local_port_range	1025 65535
net.core.netdev_max_backlog	10000
net.core.rmem_max	2097152
net.core.wmem_max	2097152
vm.overcommit_memory	1
kernel.shmmax	500000000
kernel.shmni	4096
kernel.shmall	4000000000
kernel.sem	250 512000 100 2048
net.ipv6.conf.lo.disable_ipv6	1