

A Reliable Effective Terascale Linear Learning System

Alekh Agarwal*

ALEKHA@MICROSOFT.COM

*Microsoft Research
641 Avenue of the Americas
New York, NY 10011*

Olivier Chapelle

OLIVIER@CHAPELLE.CC

*Criteo
411 High Street
Palo Alto, CA 94301*

Miroslav Dudík

MDUDIK@MICROSOFT.COM

John Langford

JCL@MICROSOFT.COM

*Microsoft Research
641 Avenue of the Americas
New York, NY 10011*

Editor: Corinna Cortes

Abstract

We present a system and a set of techniques for learning linear predictors with convex losses on terascale data sets, with trillions of features,¹ billions of training examples and millions of parameters in an hour using a cluster of 1000 machines. Individually none of the component techniques are new, but the careful synthesis required to obtain an efficient implementation is. The result is, up to our knowledge, the most scalable and efficient linear learning system reported in the literature.² We describe and thoroughly evaluate the components of the system, showing the importance of the various design choices.

Keywords: distributed machine learning, Hadoop, AllReduce, repeated online averaging, distributed L-BFGS

1. Introduction

Distributed machine learning is a research area that has seen a growing body of literature in recent years. Much work focuses on problems of the form

$$\min_{\mathbf{w} \in \mathbb{R}^d} \sum_{i=1}^n \ell(\mathbf{w}^\top \mathbf{x}_i; y_i) + \lambda R(\mathbf{w}), \quad (1)$$

where \mathbf{x}_i is the feature vector of the i -th example, y_i is the label, \mathbf{w} is the linear predictor, ℓ is a loss function and R is a regularizer. Most distributed methods for optimizing the objective (1) exploit its natural decomposability over examples, partitioning the examples over different nodes in a distributed environment such as a cluster.

*. This work was done while all authors were part of Yahoo! Research.

1. The number of features here refers to the number of non-zero entries in the data matrix.
2. All the empirical evaluation reported in this work was carried out between May-Oct 2011.

Perhaps the simplest strategy when the number of examples n is too large for a given learning algorithm is to reduce the data set size by subsampling. However, this strategy only works if the problem is simple enough or the number of parameters is very small. The setting of interest here is when a large number of examples is really needed to learn a good model. Distributed algorithms are a natural choice for such scenarios.

It might be argued that even for these large problems, it is more desirable to explore multicore solutions developed for single machines with large amounts of fast storage and memory, rather than a fully distributed algorithm which brings additional complexities due to the need for communication over a network. Yet, we claim that there are natural reasons for studying distributed machine learning on a cluster. In many industry-scale applications, the data sets themselves are collected and stored in a decentralized fashion over a cluster, typical examples being logs of user clicks or search queries. When the data storage is distributed, it is much more desirable to also process it in a distributed fashion to avoid the bottleneck of data transfer to a single powerful server. Second, it is often relatively easy to get access to a distributed computing platform such as Amazon EC2, as opposed to procuring a sufficiently powerful server. Finally, the largest problem solvable by a single machine will always be constrained by the rate at which the hardware improves, which has been steadily dwarfed by the rate at which our data sizes have been increasing over the past decade. Overall, we think that there are several very strong reasons to explore the questions of large-scale learning in cluster environments.

Previous literature on cluster learning is broad. Several authors (Mangasarian, 1995; McDonald et al., 2010; Zinkevich et al., 2010) have studied approaches that first solve the learning problem independently on each machine using the portion of the data stored on that machine, and then average the independent local solutions to obtain the global solution. Duchi et al. (2012) propose gossip-style message passing algorithms extending the existing literature on distributed convex optimization (Bertsekas and Tsitsiklis, 1989). Langford et al. (2009) analyze a delayed version of distributed online learning. Dekel et al. (2012) consider mini-batch versions of distributed online algorithms which are extended to delay-based updates in Agarwal and Duchi (2011). A recent article of Boyd et al. (2011) describes an application of the ADMM technique for distributed learning problems. GraphLab (Low et al., 2010) is a parallel computation framework on graphs. The closest to our work are optimization approaches based on centralized algorithms with parallelized gradient computation (Nash and Sofer, 1989; Teo et al., 2007). To our knowledge, all previous versions of algorithms based on parallelized gradient computation rely on MPI implementations.³ Finally, the large-scale learning system Sibyl (currently unpublished, but see the talks Chandra et al., 2010; Canini et al., 2012) implements a distributed boosting approach. It can be used to solve the problems of form (1) at the scales similar to those reported in this paper, but it runs on a proprietary architecture and many implementation details are missing, so doing a fair comparison is currently not possible. We attempt to compare the performance of our algorithm with the published Sibyl performance in Section 3.2.

All of the aforementioned approaches (perhaps with the exception of Sibyl) seem to leave something to be desired empirically when deployed on large clusters. In particular, their *learning throughput*—measured as the input size divided by the wall-clock running

3. See, for example, <http://www.mcs.anl.gov/research/projects/mpi/>.

time of the entire learning algorithm—is smaller than the I/O interface of a single machine for almost all parallel learning algorithms (Bekkerman et al., 2011, Part III, page 8). The I/O interface is an upper bound on the speed of the fastest single-machine algorithm since all single-machine algorithms are limited by the network interface in acquiring data. In contrast, we were able to achieve a learning throughput of 500M features/s, which is about a factor of 5 faster than the 1Gb/s network interface of any one node. This learning throughput was achieved on a cluster of 1000 nodes. Each node accessed its local examples 10 times during the course of the algorithm, so the per-node processing speeds were 5M features/s. We discuss our throughput results in more detail in Section 3.2, and contrast them with results reported for Sibyl.

Two difficulties bedevil easy parallel machine learning:

1. Efficient large-scale parallel learning algorithms must occur on a data-centric computing platform (such as Hadoop) to prevent data transfer overheads. These platforms typically do *not* support the full generality of MPI operations.
2. Existing data-centric platforms often lack efficient mechanisms for state synchronization and force both refactoring and rewriting of existing learning algorithms.

We effectively deal with both of these issues. Our system is compatible with MapReduce clusters such as Hadoop (unlike MPI-based systems) and minimal additional programming effort is required to parallelize existing learning algorithms (unlike MapReduce approaches). In essence, an existing implementation of a learning algorithm need only insert a few strategic library calls to switch from learning on one machine to learning on a thousand machines.

One of the key components in our system is a communication infrastructure that efficiently accumulates and broadcasts values across all nodes of a computation. It is functionally similar to MPI AllReduce (hence we use the name), but it takes advantage of and is compatible with Hadoop so that programs are easily moved to data, automatic restarts on failure provide robustness, and speculative execution speeds up completion. Our optimization algorithm is a hybrid online+batch algorithm with rapid convergence and only small synchronization overhead, which makes it a particularly good fit for the distributed environment.

In Section 2 we describe our approach and our communication infrastructure in more detail. The core of the paper is Section 3 where we conduct many experiments evaluating our design choices and comparing our approach with existing algorithms. In Section 4 we provide some theoretical intuition for our design, and contrast our approach with previous work. We conclude with a discussion in Section 5.

2. Computation and Communication Framework

MapReduce (Dean and Ghemawat, 2008) and its open source implementation Hadoop⁴ have become the overwhelmingly favorite platforms for distributed data processing. However, the abstraction is rather ill-suited for machine learning algorithms as several researchers in the field have observed (Low et al., 2010; Zaharia et al., 2011), because it does not easily allow iterative algorithms, such as typical optimization algorithms used to solve the problem (1).

4. See, for example, <http://hadoop.apache.org/>.




Figure 1: AllReduce operation. Initially, each node holds its own value. Values are passed up the tree and summed, until the global sum is obtained in the root node (reduce phase). The global sum is then passed back down to all other nodes (broadcast phase). At the end, each node contains the global sum.

2.1 Hadoop-compatible AllReduce

AllReduce is a more suitable abstraction for machine learning algorithms. AllReduce is an operation where every node starts with a number and ends up with the sum of the numbers across all the nodes (hence the name). A typical implementation imposes a tree structure on the communicating nodes and proceeds in two phases: numbers are first summed up the tree (the *reduce* phase) and then broadcast down to all the nodes (the *broadcast* phase), see Figure 1 for a graphical illustration. When doing summing or averaging of a long vector, such as the weight vector \mathbf{w} in the optimization (1), the reduce and broadcast operations can be pipelined over the vector entries and hence the latency of going up and down the tree becomes negligible on a typical Hadoop cluster. This is the main optimization we do within the AllReduce architecture. While other (potentially more efficient or simpler) architectures for AllReduce are possible, in our experiments in Section 3 we will see that the time spent in AllReduce operation is negligible compared with the computation time and stalling time while waiting for other nodes. Therefore, we do not attempt to optimize the architecture further.

For problems of the form (1), AllReduce provides straightforward parallelization of gradient-based optimization algorithms such as gradient descent or L-BFGS—gradients are accumulated locally, and the global gradient is obtained by AllReduce. In general, any statistical query algorithm (Kearns, 1993) can be parallelized with AllReduce with only a handful of additional lines of code. This approach also easily implements averaging parameters of online learning algorithms.

An implementation of AllReduce is available in the MPI package. However, it is not easy to run MPI on top of existing Hadoop clusters (Ye et al., 2009). Moreover, MPI implements little fault tolerance, with the bulk of robustness left to the programmer.

To address the reliability issues better, we developed an implementation of AllReduce that is compatible with Hadoop. Our implementation works as follows. We initialize a spanning tree server on the gateway node to the Hadoop cluster. We then launch a map-only (alternatively reduce-only) job where each mapper processes a subset of the data. Each mapper is supplied with the IP address of the gateway node, to which it connects as the first step. Once all the mappers are launched and connected to the spanning tree server, it creates a (nearly balanced) binary tree on these nodes. Each node is given the IP addresses of its parent and child nodes in the tree, allowing it to establish TCP connections with them. All the nodes are now ready to pass messages up and down the tree. The actual communication between the nodes is all implemented directly using C++ sockets and does not rely on any Hadoop services. Implementation of AllReduce using a single tree is clearly less desirable than MapReduce in terms of reliability, because if any individual node fails, the entire computation fails. To deal with this problem, we use a simple trick described below which makes AllReduce reliable enough to use in practice for computations up to 10K node hours.

2.2 Proposed Algorithm

Our main algorithm is a hybrid online+batch approach. Pure online and pure batch learning algorithms have some desirable features, on which we build, and some drawbacks, which we overcome. For instance, an attractive feature of online learning algorithms is that they optimize the objective to a rough precision quite fast, in just a handful of passes over the data. The inherent sequential nature of these algorithms, however, makes them tricky to parallelize and we discuss the drawbacks of some of the attempts at doing so in Section 4. Batch learning algorithms such as Newton and quasi-Newton methods (e.g., L-BFGS), on the other hand, are great at optimizing the objective to a high accuracy, once they are in a *good neighborhood* of the optimal solution. But the algorithms can be quite slow in reaching this good neighborhood. Generalization of these approaches to distributed setups is rather straightforward, only requiring aggregation across nodes after every iteration, as has been noted in previous research (Teo et al., 2007).

We attempt to reap the benefits and avoid the drawbacks of both above approaches through our hybrid method. We start with each node making one online pass over its local data according to adaptive gradient updates (Duchi et al., 2010; McMahan and Streeter, 2010) modified for loss non-linearity (Karampatziakis and Langford, 2011). We notice that each online pass happens completely *asynchronously* without any communication between the nodes, and we can afford to do so since we are only seeking to get into a good neighborhood of the optimal solution rather than recovering it to a high precision at this first stage. AllReduce is used to average these weights non-uniformly according to locally accumulated gradient squares. Concretely, node k maintains a local weight vector \mathbf{w}^k and a diagonal matrix \mathbf{G}^k based on the gradient squares in the adaptive gradient update rule (see Algorithm 1). We compute the following weighted average over all m nodes

$$\bar{\mathbf{w}} = \left(\sum_{k=1}^m \mathbf{G}^k \right)^{-1} \left(\sum_{k=1}^m \mathbf{G}^k \mathbf{w}^k \right). \quad (2)$$

This has the effect of weighting each dimension according to how “confident” each node is in its weight (i.e., more weight is assigned to a given parameter of a given node if that node has seen more examples with the corresponding feature). We note that this averaging can indeed be implemented using AllReduce by two calls to the routine since the matrices \mathbf{G}^k are only diagonal.

This solution $\bar{\mathbf{w}}$ is used to initialize L-BFGS (Nocedal, 1980) with the standard Jacobi preconditioner, with the expectation that the online stage gives us a good *warmstart* for L-BFGS. At each iteration, global gradients are obtained by summing up local gradients via AllReduce, while all the other operations can be done locally at each node. The algorithm benefits from the fast initial reduction of error provided by an online algorithm, and rapid convergence in a good neighborhood guaranteed by quasi-Newton algorithms. We again point out that the number of communication operations is relatively small throughout this process.

In addition to hybrid strategy, we also evaluate repeated online learning with averaging using the adaptive updates. In this setting, each node performs an online pass over its data and then we average the weights according to Equation 2. We average the scaling matrices similarly

$$\bar{\mathbf{G}} = \left(\sum_{k=1}^m \mathbf{G}^k \right)^{-1} \left(\sum_{k=1}^m (\mathbf{G}^k)^2 \right)$$

and use this averaged state to start a new online pass over the data. This strategy is similar to those proposed by McDonald et al. (2010) and Hall et al. (2010) for different online learning algorithms. We will see in the next section that this strategy can be very effective at getting a moderately small test error very fast, but its convergence slows down and it might be too slow at reaching the optimal test error.

All strategies described above share the same processing structure. They carry out several iterations, each of which can be broken into three phases: (1) Pass through the entire local portion of the data set and accumulate the result as a vector of size d (i.e., the same size as the parameter vector). (2) Carry out AllReduce operation on a vector of size d . (3) Do some additional processing and updating of the parameter vector.

The key point to notice is that in typical applications the local data set will be orders of magnitude larger than the parameter vector, hence the communication after each pass is much more compact than transmitting the entire local data set. The second point is that each iteration is naturally a MapReduce operation. The main reason that we expect to benefit by AllReduce is because of the iterative nature of these algorithms and the shared state between iterations.

Our implementation is available as part of the open source project Vowpal Wabbit (Langford et al., 2007) and is summarized in Algorithm 2. It makes use of stochastic gradient descent (Algorithm 1) for the initial pass.

2.3 Speculative Execution

Large clusters of machines are typically busy with many jobs which use the cluster unevenly, resulting in one of a thousand nodes being very slow. To avoid this, Hadoop can speculatively execute a job on identical data, using the first job to finish and killing the other one. In our framework, it can be tricky to handle duplicates once a spanning tree topology

Algorithm 1 Stochastic gradient descent algorithm on a single node using adaptive gradient update (Duchi et al., 2010; McMahan and Streeter, 2010).

Require: Invariance update function s
 (see Karampatziakis and Langford, 2011)

$\mathbf{w} = \mathbf{0}$, $\mathbf{G} = \mathbf{I}$

for all (\mathbf{x}, y) in training set **do**

$\mathbf{g} \leftarrow \nabla_{\mathbf{w}} \ell(\mathbf{w}^\top \mathbf{x}; y)$

$\mathbf{w} \leftarrow \mathbf{w} - s(\mathbf{w}, \mathbf{x}, y) \mathbf{G}^{-1/2} \mathbf{g}$

$G_{jj} \leftarrow G_{jj} + g_j^2$ for all $j = 1, \dots, d$

end for

Algorithm 2 Sketch of the proposed learning architecture

Require: Data split across nodes

for all nodes k **do**

$\mathbf{w}^k =$ result of stochastic gradient descent on the data of node k using Algorithm 1.

 Compute the weighted average $\bar{\mathbf{w}}$ as in (2) using AllReduce.

 Start a preconditioned L-BFGS optimization from $\bar{\mathbf{w}}$.

for $t = 1, \dots, T$ **do**

 Compute \mathbf{g}^k the (local batch) gradient of examples on node k .

 Compute $\mathbf{g} = \sum_{k=1}^m \mathbf{g}^k$ using AllReduce.

 Add the regularization part in the gradient.

 Take an L-BFGS step.

end for

end for

is created for AllReduce. For this reason, we delay the initialization of the spanning tree until each node completes the first pass over the data, building the spanning tree on only the speculative execution survivors. The net effect of this speculative execution trick is perhaps another order of magnitude of scalability and reliability in practice. Indeed, we found the system reliable enough for up to 1000 nodes running failure-free for hundreds of trials (of typical length up to 2 hours). This level of fault tolerance highlights the benefits of a Hadoop-compatible implementation of AllReduce. We will show the substantial gains from speculative execution in mitigating the “slow node” problem in the experiments.

3. Experiments

In this section we will present the empirical evaluation of the system described so far. We begin by describing the datasets used, before evaluating the various properties of our framework both from a systems as well as a machine learning perspective. A more theoretical evaluation of our approach can be found in the next section.

3.1 Data Sets

We evaluated our framework on two datasets. The first dataset is a computational advertising task, with proprietary data from Yahoo! Inc. The second dataset comes from the task of recognizing a human acceptor splice site. Both the datasets are large enough to necessitate

distributed machine learning, and simple approaches such as subsampling do not work to obtain a good model with reasonable predictive accuracy as we describe in detail next.

3.1.1 DISPLAY ADVERTISING

In online advertising, given a user visiting a publisher page, the problem is to select the best advertisement for that user. A key element in this matching problem is the click-through rate (CTR) estimation: what is the probability that a given ad will be clicked on, given some context (user, page visited)? Indeed, in a cost-per-click (CPC) campaign, the advertiser only pays when the ad gets clicked, so even modest improvements in predictive accuracy directly affect revenue.

Training data contains user visits, which either resulted in a click on the ad (positive examples with $y_i = 1$), or did not result in a click (negative examples with $y_i = 0$). We estimate the click probabilities by logistic regression with L_2 regularization. The regularization coefficient is chosen from a small set to obtain the best test performance. The user visit is represented by binary indicator features encoding the user, page, ad, as well as conjunctions of these features. Some of the features include identifiers of the ad, advertiser, publisher and visited page. These features are hashed (Weinberger et al., 2009) and each training example ends up being represented as a sparse binary vector of dimension 2^{24} with around 125 non-zero elements. Let us illustrate the construction of a conjunction feature with an example. Imagine that an ad from `etrade` was placed on `finance.yahoo.com`. Let h be a 24 bit hash of the string “publisher=finance.yahoo.com and advertiser=etrade”. Then the (publisher, advertiser) conjunction is encoded by setting to 1 the h -th entry of the feature vector for that example.

Since the data is unbalanced (low CTR) and because of the large number of examples, we subsample the negative examples resulting in a class ratio of about 2 negatives for 1 positive, and use a large test set drawn from days later than the training set. There are 2.3B examples in the training set. More characteristics of this data set and modeling details can be found in Chapelle et al. (2013).

3.1.2 SPLICE SITE RECOGNITION

The problem consists of recognizing a human acceptor splice site (Sonnenburg and Franc, 2010). We consider this learning task because this is, as far as we know, the largest public data set for which subsampling is not an effective learning strategy. Sonnenburg et al. (2007) introduced the *weighted degree kernel* to learn over DNA sequences. They also proposed an SVM training algorithm for that kernel for which learning over 10M sequences took 24 days. Sonnenburg and Franc (2010) proposed an improved training algorithm, in which the weight vector—in the feature space induced by the kernel—is learned, but the feature vectors are never explicitly computed. This resulted in a faster training: 3 days with 50M sequences.

We solve this problem by L_2 -regularized logistic regression. Again, the regularization coefficient is chosen from a small set to optimize test set performance. We follow the same experimental protocol as in Sonnenburg and Franc (2010): we use the same training and test sets of respectively 50M and 4.6M samples. We also consider the same kernel of degree $d = 20$ and hash size $\gamma = 12$. The feature space induced by this kernel has

	1%	10%	100%
auROC	0.8178	0.8301	0.8344
auPRC	0.4505	0.4753	0.4856
NLL	0.2654	0.2582	0.2554

Table 1: Test performance on the display advertising problem as a function of the subsampling rate, according to three metrics: area under ROC curve (auROC), area under precision/recall curve (auPRC), and negative log likelihood (NLL).

dimensionality 11,725,480. The number of non-zero features per sequence is about 3,300. Unlike Sonnenburg and Franc (2010), we explicitly compute the feature space representation of the examples, yielding about 3TB of data. This explicit representation is a *disadvantage* we impose on our method to simplify implementation.

3.2 Results

We now present a detailed evaluation of our system on both the above datasets. We begin by evaluating the performance of simpler heuristics such as subsampling the data for faster running time, before examining the computational and communication aspects of our system in detail.

3.2.1 EFFECT OF SUBSAMPLING

The easiest way to deal with a very large training set is to reduce it by subsampling as discussed in the introduction. Sometimes similar test errors can be achieved with smaller training sets and there is no need for large-scale learning. For splice site recognition, Table 2 of Sonnenburg and Franc (2010) shows that smaller training sets do hurt the area under the precision/recall curve on the test set.

For display advertising, we subsampled the data at 1% and 10%. The results in Table 1 show that there is a noticeable drop in accuracy after subsampling. Note that even if the drop does not appear large at a first sight, it can cause a substantial loss of revenue. Thus, for both data sets, the entire training data set is needed to achieve optimal performances.

The three metrics reported in Table 1 are area under the ROC curve (auROC), area under the precision/recall curve (auPRC) and negative log-likelihood (NLL). Since auPRC is the most sensitive metric, we report test results using that metric in the rest of the paper. This is also the metric used in Sonnenburg and Franc (2010).

3.2.2 RUNNING TIME

We ran 5 iterations of L-BFGS on the splice site data with 1000 nodes. On each node, we recorded for every iteration the time spent in AllReduce and the computing time—defined as the time not spent in AllReduce. The time spent in AllReduce can further be divided into stall time—waiting for the other nodes to finish their computation—and communication time. The communication time can be estimated by taking the minimum value of the AllReduce times across nodes.

	5%	50%	95%	Max	Comm. time
Without spec. execution	29	34	60	758	26
With spec. execution	29	33	49	63	10

Table 2: Distribution of computing times (in seconds) over 1000 nodes with and without speculative execution. First three columns are quantiles. Times are average per iteration (excluding the first one) for the splice site recognition problem.

Nodes	100	200	500	1000
Comm time / pass	5	12	9	16
Median comp time / pass	167	105	43	34
Max comp time / pass	462	271	172	95
Wall clock time	3677	2120	938	813

Table 3: Computing times to obtain a fixed test error on the splice site recognition data, using different numbers of nodes. The first 3 rows are averages per iteration (excluding the first pass over the data).

The distribution of the computing times is of particular interest because the speed of our algorithm depends on the slowest node. Statistics are shown in Table 2. It appears that computing times are concentrated around the median, but there are a few outliers. Without speculative execution, one single node was about 10 times slower than the other nodes; this has the catastrophic consequence of slowing down the entire process by a factor 10. The table shows that the use of speculative execution successfully mitigates this issue.

We now study the running time as a function of the number of nodes. For the display advertising problem, we varied the number of nodes from 10 to 100 and computed the speed-up factor relative to the run with 10 nodes. In each case, we measured the amount of time needed to get to a fixed test error. Since there can be significant variations from one run to the other—mostly because of the cluster utilization—each run was repeated 10 times. Results are reported in Figure 2. We note that speculative execution was not turned on in this experiment, and we expect better speed-ups with speculative execution. In particular, we expect that the main reason for the departure from the ideal speed-up curve is the “slow node” problem (as opposed to the aspects of the AllReduce communication implementation), which is highlighted also in the next experiment.

Table 3 shows the running times for attaining a fixed test error as a function of the number of nodes on the splice site recognition problem. Unlike Figure 2, these timing experiments have not been repeated and thus there is a relatively large uncertainty in their expected values. It can be seen from Tables 2 and 3 that even with as many as 1000 nodes, communication is not the bottleneck. One of the main challenges instead is the “slow node” issue. This is mitigated to some degree by speculative execution, but as the number of nodes increases, so does the likelihood of hitting slow nodes.




Figure 2: Speed-up for obtaining a fixed test error, on the display advertising problem, relative to the run with 10 nodes, as a function of the number of nodes. The dashed line corresponds to the ideal speed-up, the solid line is the average speed-up over 10 repetitions, and the bars indicate maximum and minimum values.




Figure 3: Effect of initializing the L-BFGS optimization by an average solution from online runs on individual nodes. Suboptimality is the difference between the objective value on the training data and the optimal value obtained by running the algorithm to convergence.

3.2.3 LARGE EXPERIMENT AND COMPARISON WITH SIBYL

We also experimented with an 8 times larger version of the display advertising data (16B examples). Using 1000 nodes and 10 passes over the data, the training took only 70 minutes.⁵ Since each example is described by 125 non-zero features, the average processing speed was

$$16\text{B} \times 10 \times 125 \text{ features}/1000 \text{ nodes}/70 \text{ minutes} = 4.7 \text{ M features}/\text{node}/\text{s} .$$

The overall learning throughput was

$$16\text{B} \times 125 \text{ features}/70 \text{ minutes} = 470 \text{ M features}/\text{s} .$$

We briefly compare this with a result reported for the distributed boosting system Sibyl for a run on 970 cores (Canini et al., 2012, slide 24). The run was done over 129.1B examples, with 54.61 non-zero features per example. The reported processing speed was 2.3M features/core/s (which is a factor of two slower than our achieved processing speed). The reported number of iterations was 10–50, which would lead to the final learning throughput in the range 45–223 M features/s, that is, the result appears to be slower by a factor of 2–10.

3.2.4 ONLINE AND BATCH LEARNING

We now investigate the speed of convergence of three different learning strategies: batch, online and hybrid. We are interested in how fast the algorithms minimize the training objective as well as the test error.

Figure 3 compares how fast the two learning strategies—batch with and without an initial online pass—optimize the training objective. It plots the optimality gap, defined as the difference between the current objective function and the optimal one (i.e., the minimum value of the objective (1)), as a function of the number iterations. From this figure we can see that the initial online pass results in a saving of about 10–15 iterations.

Figure 4 shows the test auPRC on both data sets as a function of the number of iterations for 4 different strategies: only online learning, only L-BFGS learning, and 2 hybrid methods consisting of 1 or 5 passes of online learning followed by L-BFGS optimization. L-BFGS with one online pass appears to be the most effective strategy.

For the splice site recognition problem, an initial online pass and 14 L-BFGS iterations yield an auPRC of 0.581, which is just a bit higher than results of Sonnenburg and Franc (2010). This was achieved in 1960 seconds using 500 machines, resulting in a 68 speed-up factor (132,581 seconds on a single machine reported in Table 2 of Sonnenburg and Franc, 2010). This seems rather poor compared with the ideal 500 speed-up factor, but recall that we used explicit feature representation which creates a significant overhead.

3.3 Comparison with Previous Approaches

In order to better assess the overall system, we next compare its performance with that of some other published baselines. We start by demonstrating the efficacy of AllReduce

5. As mentioned before, there can be substantial variations in timing between different runs; this one was done when the cluster was not too occupied.




Figure 4: Test auPRC for 4 different learning strategies. Note that the online and hybrid curves overlap during the warmstart phase (of either 1 or 5 online passes).

	Full size	10% sample
MapReduce	1690	1322
AllReduce	670	59

Table 4: Average training time per iteration of an internal logistic regression implementation using either MapReduce or AllReduce for gradients aggregation. The data set is the display advertising one and a subset of it.

compared with MapReduce, before comparing to some other distributed machine learning algorithms.

3.3.1 ALLREDUCE VS. MAPREDUCE

The standard way of using MapReduce for iterative machine learning algorithms is the following (Chu et al., 2007): every iteration is a MapReduce job where the mappers compute some local statistics (such as a gradient) and the reducers sum them up. This is ineffective because each iteration has large overheads (job scheduling, data transfer, data parsing, etc.). We have an internal implementation of such a MapReduce algorithm. We updated this code to use AllReduce instead and compared both versions of the code in Table 4. This table confirms that Hadoop MapReduce has substantial overheads since the training time is not much affected by the data set size. The speed-up factor of AllReduce over Hadoop’s MapReduce can become extremely large for smaller data sets, and remains noticeable even for the largest data sets. It is also noteworthy that *all* algorithms described in Chu et al. (2007) can be parallelized with AllReduce, plus further algorithms such as parameter averaging approaches.

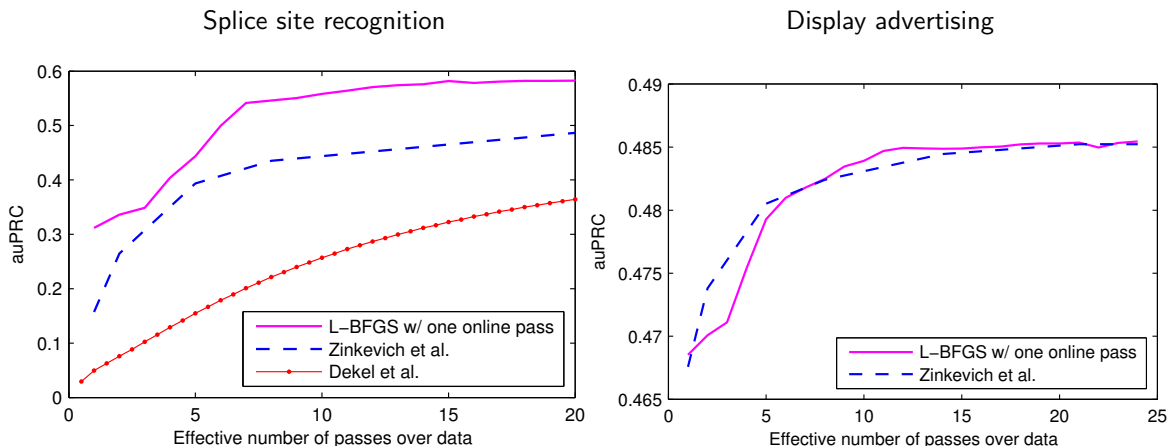


Figure 5: Test auPRC for different learning strategies as a function of the effective number of passes over data. In L-BFGS, it corresponds to iterations of the optimization. In overcomplete SGD with averaging (Zinkevich et al.), it corresponds to the replication coefficient.

3.3.2 OVERCOMPLETE AVERAGE

We implemented oversampled stochastic gradient with final averaging (Zinkevich et al., 2010), and compared its performance to our algorithm. We used stochastic gradient descent with the learning rate in the t -th iteration as

$$\eta_t = \frac{1}{L + \gamma\sqrt{t}} .$$

We tuned γ and L on a small subset of the data set.

In Figure 5, we see that the oversampled SGD is competitive with our approach on the display advertising data, but its convergence is much slower on splice site data.

3.3.3 PARALLEL ONLINE MINI-BATCH

Dekel et al. (2012) propose to perform online convex optimization using stochastic gradients accumulated in small mini-batches across all nodes. We implemented a version of their algorithm using AllReduce. They suggest global minibatch sizes of no more than $b \propto \sqrt{n}$. On m nodes, each node accumulates gradients from b/m examples, then an AllReduce operation is carried out, yielding the mini-batch gradient, and each node performs a stochastic gradient update with the learning rate of the form

$$\eta_t = \frac{1}{L + \gamma\sqrt{t/m}} .$$

We tuned L and γ on a smaller data set. In Figure 5, we report the results on splice site data set, using 500 nodes, and mini-batch size $b = 100k$. Twenty passes over the data thus corresponded to 10k updates. Due to the overwhelming communication overhead

associated with the updates, the overall running time was 40 hours. In contrast, L-BFGS took less than an hour to finish 20 passes while obtaining much superior performance. The difference in the running time between 1h and 40h is solely due to communication. Thus, in this instance, we can conservatively conclude that the communication overhead of 10k mini-batch updates is 39 hours.

We should point out that it is definitely possible that the mini-batched SGD would reach similar accuracy with much smaller mini-batch sizes (for 10k updates theory suggests we should use mini-batch sizes of at most 10k), however, the 39 hour communication overhead would remain. Using larger mini-batches, we do expect that the time to reach 20 passes over data would be smaller (roughly proportional to the number of mini-batch updates), but according to theory (as well as our preliminary experiments on smaller subsets of splice site data), we would have inferior accuracy. Because of the prohibitive running time, we were not able to tune and evaluate this algorithm on display advertising data set.

4. Communication and Computation Complexity

The two key performance characteristics of any distributed algorithm are its communication and computation complexity. The aim of this section is to discuss the complexity of our approach and to compare it with previous solutions. We hope to clarify the reasons underlying our design choices and explain the scalability of our system. We start with a discussion of computational considerations.

4.1 Computational Complexity of the Hybrid Approach

In this section, we explain the convergence properties of the hybrid approach and compare it with other optimization strategies. In order to have a clean discussion, we make some simplifying assumptions. We consider the case of only one online pass at each node. Furthermore, we restrict ourselves to the case of uniform averaging of weights. Similar analysis does extend to the non-uniform weighting scheme that we use, but the details are technical and provide no additional intuition. Before we embark on any details, it should be clear that the hybrid approach is always convergent, owing to the convergence of L-BFGS. All the online learning step initially does is to provide a good warmstart to L-BFGS. This section aims to provide theoretical intuition why the gains of such a warmstart can be substantial in certain problem regimes.

Let $\tilde{\ell}(\mathbf{w}; \mathbf{x}, y) = \ell(\mathbf{w}^\top \mathbf{x}; y) + \lambda R(\mathbf{w})/n$ be the regularized loss function. We analyze a scaled version of the objective (1):

$$f(\mathbf{w}) = \frac{1}{n} \sum_{i=1}^n \ell(\mathbf{w}^\top \mathbf{x}_i; y_i) + \frac{\lambda}{n} R(\mathbf{w}) = \frac{1}{n} \sum_{i=1}^n \tilde{\ell}(\mathbf{w}; \mathbf{x}_i, y_i) .$$

We assume that the cluster is comprised of m nodes, with a total of n data examples distributed uniformly at random across these nodes. Let us denote the local objective function at each node as f^k :

$$f^k(\mathbf{w}) = \frac{m}{n} \sum_{i \in S_k} \tilde{\ell}(\mathbf{w}; \mathbf{x}_i, y_i)$$

where S_k is the set of n/m examples at node k . Note that the global objective $f = (\sum_{k=1}^m f^k)/m$ is the average of the local objectives. We observe that owing to our random data split, we are guaranteed that

$$\mathbb{E}[f^k(\mathbf{w})] = \mathbb{E}[f(\mathbf{w})] = \mathbb{E}_{\mathbf{x},y}[\tilde{\ell}(\mathbf{w}; \mathbf{x}, y)]$$

for each k , where the expectation is taken over the distribution from which our examples are drawn. In order to discuss the convergence properties, we need to make a couple of standard assumptions regarding the functions f^k . First, we assume that the functions f^k are differentiable, with Lipschitz-continuous gradients. We also assume that each f^k is strongly convex, at least in a local neighborhood around the optimum. We note that these assumptions are unavoidable for the convergence of quasi-Newton methods such as L-BFGS.

To understand how many passes over the data are needed for the hybrid approach to minimize f to a precision ϵ , we first analyze the online learning pass at each node. In this pass, we compute a weight vector \mathbf{w}^k by performing n/m steps of stochastic gradient descent or some variant thereof (Duchi et al., 2010; Karampatziakis and Langford, 2011). Since we performed only one pass at each node, the resulting \mathbf{w}^k at each node approximately minimizes $\mathbb{E}[f^k] = \mathbb{E}_{\mathbf{x},y}[\tilde{\ell}]$ to a precision ϵ^k (for the methods we use, we expect $\epsilon^k = \mathcal{O}(\sqrt{m/n})$). Let us now denote the uniform average $\bar{\mathbf{w}} = \sum_{k=1}^m \mathbf{w}^k/m$. For this approach, a direct application of Jensen's inequality yields

$$\begin{aligned} \mathbb{E}_{\mathbf{x},y}[\tilde{\ell}(\bar{\mathbf{w}}; \mathbf{x}, y)] &= \mathbb{E}_{\mathbf{x},y}\left[\tilde{\ell}\left(\frac{\sum_{i=1}^m \mathbf{w}^k}{m}; \mathbf{x}, y\right)\right] \leq \frac{1}{m} \sum_{k=1}^m \mathbb{E}_{\mathbf{x},y}[\tilde{\ell}(\mathbf{w}^k; \mathbf{x}, y)] \\ &\leq \min_{\mathbf{w}} \mathbb{E}_{\mathbf{x},y}[\tilde{\ell}(\mathbf{w}; \mathbf{x}, y)] + \frac{1}{m} \sum_{k=1}^m \epsilon^k = \min_{\mathbf{w}} \mathbb{E}_{\mathbf{x},y}[\tilde{\ell}(\mathbf{w}; \mathbf{x}, y)] + \mathcal{O}\left(\sqrt{\frac{m}{n}}\right). \end{aligned} \quad (3)$$

Furthermore, standard sample complexity arguments (see, e.g., Bartlett and Mendelson, 2002; Devroye et al., 1996) allow us to bound the function value $f(\mathbf{w})$ for an arbitrary \mathbf{w} as

$$\left|f(\mathbf{w}) - \mathbb{E}_{\mathbf{x},y}[\tilde{\ell}(\mathbf{w}; \mathbf{x}, y)]\right| \leq \mathcal{O}\left(\frac{1}{\sqrt{n}}\right). \quad (4)$$

Let \mathbf{w}^* be the minimizer of the training loss function f . Then we can combine the above inequalities as

$$\begin{aligned} f(\bar{\mathbf{w}}) &\leq \mathbb{E}_{\mathbf{x},y}[\tilde{\ell}(\bar{\mathbf{w}}; \mathbf{x}, y)] + \mathcal{O}(1/\sqrt{n}) \\ &\leq \min_{\mathbf{w}} \mathbb{E}_{\mathbf{x},y}[\tilde{\ell}(\mathbf{w}; \mathbf{x}, y)] + \mathcal{O}(\sqrt{m/n}) \\ &\leq \mathbb{E}_{\mathbf{x},y}[\tilde{\ell}(\mathbf{w}^*; \mathbf{x}, y)] + \mathcal{O}(\sqrt{m/n}) \\ &\leq f(\mathbf{w}^*) + \mathcal{O}(\sqrt{m/n}) \end{aligned}$$

where the first inequality follows by (4), the second by (3), and the fourth by (4). For the remainder of the discussion, we denote the overall suboptimality of $\bar{\mathbf{w}}$ relative to \mathbf{w}^* by $\epsilon_0 = \mathcal{O}(\sqrt{m/n})$.

Switching over to the L-BFGS phase, we assume that we are in the linear convergence regime of L-BFGS (Liu and Nocedal, 1989). We denote the contraction factor by κ , so that the number of additional L-BFGS passes over data needed to minimize f to a precision ϵ is at most

$$\kappa \log \frac{\epsilon_0}{\epsilon} .$$

Compared to initializing L-BFGS without any warmstart, our hybrid strategy amounts to overall savings of

$$\kappa \log \frac{1}{\epsilon} - \left(1 + \kappa \log \frac{\epsilon_0}{\epsilon}\right) = \kappa \log \frac{1}{\epsilon_0} - 1 = \mathcal{O}\left(\frac{\kappa}{2} \log \frac{n}{m}\right) - 1$$

passes over data. In typical applications, we expect $n \gg m$ to ensure that computation amortize the cost of communication. As a result, the improvement due to the warmstart can be quite substantial just like we observed in our experiments. Furthermore, this part of our argument is in no way specific to the use of L-BFGS as the batch optimization algorithm. Similar reasoning holds for any reasonable (quasi)-Newton method.

We could also consider the alternative choice of just using parallel online learning without ever switching to a batch optimization method. The theoretical results in this area, however, are relatively harder to compare with the hybrid approach. For general online learning algorithms, previous works study just one local pass of online learning followed by averaging (McDonald et al., 2010), which typically cannot guarantee an error smaller than ϵ_0 in our earlier notation. The repeated averaging approach, discussed and analyze for the specific case of perceptron algorithm in earlier work (McDonald et al., 2010), works well in our experiments on the computational advertising task but does not have easily available convergence rates beyond the special case of separable data and the perceptron algorithm. Nevertheless, one appeal of the hybrid approach is that it is guaranteed to be competitive with such online approaches, by the mere virtue of the first online phase.

Overall, we see that the hybrid approach will generally be competitive with purely online or batch approaches in terms of the computational complexity. As a final point, we discuss two extreme regimes where it can and does offer substantial gains. The first regime is when the data has a significant noise level. In such a scenario, the level ϵ of optimization accuracy desired is typically not too large (intuitive statistical arguments show no reduction in generalization error for $\epsilon \ll 1/n$). Setting $\epsilon = 1/n$ for a clean comparison, we observe that the total number of passes for the hybrid method is at most

$$1 + \frac{\kappa}{2}(\log(m) + \log(n)),$$

as opposed to $\kappa \log(n)$ for just pure batch optimization. When $m \ll n$, this shows that the online warmstart can cut down the number of passes almost by a factor of 2. We do note that in such high noise regimes, pure online approaches can often succeed, as we observed with our advertising data.

The second extreme is when our data is essentially noiseless, so that the desired accuracy ϵ is extremely small. In this case, the relative impact of the online warmstart can be less pronounced (it is certainly strictly better still) over an arbitrary initialization of L-BFGS. However, as we saw on our splice site recognition data, on this extreme, the online learning methods will typically struggle since they are usually quite effective in fitting the data to

Algorithm	Per-node communication cost
Bundle method (Teo et al., 2007)	$O(dT_{\text{bundle}})$
Online with averaging (McDonald et al., 2010; Hall et al., 2010)	$O(dT_{\text{online}})$
Parallel online (Hsu et al., 2011)	$O(ns/m + nT_{\text{online}})$
Overcomplete online with averaging (Zinkevich et al., 2010)	$O(ns + d)$
Distr. minibatch (dense) (Dekel et al., 2012; Agarwal and Duchi, 2011)	$O(dT_{\text{mini}}n/b) = O(dT_{\text{mini}}\sqrt{n})$
Distr. minibatch (sparse) (Dekel et al., 2012; Agarwal and Duchi, 2011)	$O(bsT_{\text{mini}}n/b) = O(nsT_{\text{mini}})$
Hybrid online+batch	$O(dT_{\text{hybrid}})$

Table 5: Communication cost of various learning algorithms. Here n is the number of examples, s is the number of nonzero features per example, d is the number of dimensions, T is the number of times the algorithm examines each example, and b is the minibatch size (in minibatch algorithms).

moderate but not high accuracies (as evident from their $1/\epsilon$ or $1/\epsilon^2$ convergence rates). Overall, we find that even on these two extremes, the hybrid approach is competitive with the better of its two components.

4.2 Communication Cost Comparison with Previous Approaches

In the previous section we discussed the computational complexity of several techniques with an identical communication pattern: communication of the entire weight vector once per pass. In this section we contrast our approach with techniques that use other communication patterns. We focus mainly on communication cost since the computational cost is typically the same as for our algorithm, or the communication dominates the computation.

Since modern network switches are quite good at isolating communicating nodes, the most relevant communication cost is the maximum (over nodes) of the communication cost of a single node.

Several variables (some of them recalled from the previous section) are important:

1. m the number of nodes.
2. n the total number of examples across all nodes.
3. s the number of nonzero features per example.
4. d the parameter dimension.
5. T the number of passes over the examples.

In the large-scale applications that are subject of this paper, we typically have $s \ll d \ll n$, where both d and n are large (see Section 3.1).

The way that data is dispersed across a cluster is relevant in much of this discussion since an algorithm not using the starting format must pay the communication cost of redistributing the data. We assume the data is distributed across the nodes uniformly according to an example partition, as is common.

The per-node communication cost of the hybrid algorithm is $\Theta(dT_{\text{hybrid}})$ where T_{hybrid} is typically about 15 to maximize test accuracy in our experiments. Note that the minimum

possible communication cost is $\Theta(d)$ if we save the model on a single machine. There is no communication involved in getting data to workers based on the data format assumed above. An important point here is that every node has a communication cost functionally smaller than the size of the data set, because there is no dependence on ns .

Similar to our approach, Teo et al. (2007) propose a parallel batch optimization algorithm (specifically, a bundle method) using the MPI implementation of AllReduce. This approach arrives at an accurate solution with $O(dT_{\text{bundle}})$ communication per node. Our approach improves over this in several respects. First, as Figure 4 demonstrates, we obtain a substantial boost thanks to our warmstarting strategy, hence in practice we expect $T_{\text{bundle}} > T_{\text{hybrid}}$. The second distinction is in the AllReduce implementation. Our implementation is well aligned with Hadoop and takes advantage of speculative execution to mitigate the slow node problem. On the other hand, MPI assumes full control over the cluster, which needs to be carefully aligned with Hadoop’s MapReduce scheduling decisions, and by itself, MPI does not provide robustness to slow nodes.

Batch learning can also be implemented using MapReduce on a Hadoop cluster (Chu et al., 2007), for example in the Mahout project.⁶ Elsewhere it has been noted that MapReduce is not well suited for iterative machine learning algorithms (Low et al., 2010; Zaharia et al., 2011). Evidence of this is provided by the Mahout project itself, as their implementation of logistic regression is not parallelized. Indeed, we observe substantial speed-ups from a straightforward substitution of MapReduce by AllReduce on Hadoop. It is also notably easier to program with AllReduce, as code does not require refactoring.

The remaining approaches are based on online convex optimization. McDonald et al. (2010) and Hall et al. (2010) study the approach when each node runs an online learning algorithm on its examples and the results from the individual nodes are averaged. This simple method is empirically rather effective at creating a decent solution. The communication cost is structurally similar to our algorithm $\Theta(dT_{\text{online}})$ when T_{online} passes are done. However, as we saw empirically in Figure 4 and also briefly argued theoretically in Section 4.1, $T_{\text{online}} > T_{\text{hybrid}}$.

Similarly to these, Zinkevich et al. (2010) create an overcomplete partition of the data and carry out separate online optimization on each node followed by global averaging. Our experiments show that this algorithm can have competitive convergence (e.g., on display advertising data), but on more difficult optimization problems it can be much slower than the hybrid algorithm we use here (e.g., on splice site recognition data). This approach also involves deep replication of the data—for example, it may require having 1/4 of all the examples on each of 100 nodes. This is generally undesirable with large data sets. The per-node communication cost is $\Theta(nsT_{\text{rep}}/m + d)$ where T_{rep} is the level of replication and m is the number of nodes. Here, the first term comes from the data transfer required for creating the overcomplete partition whereas the second term comes from parameter averaging. Since T_{rep}/m is often a constant near 1 (0.25 was observed by Zinkevich et al., 2010, and the theory predicts only a constant factor improvement), this implies the communication cost is $\Theta(ns)$, the size of the data set.

Other authors have looked into online mini-batch optimization (Dekel et al., 2012; Agarwal and Duchi, 2011). The key problem here is the communication cost. The per-node

6. For Mahout, see <http://mahout.apache.org/>.

communication cost is $\Theta(T_{\text{mini}}dn/b)$ where b is the minibatch size (number of examples per minibatch summed across all nodes), T_{mini} is the number of passes over the data, n/b is the number of minibatch updates per pass and d is the number of parameters. Theory suggests $b \leq \sqrt{n}$, implying communication costs of $\Theta(T_{\text{mini}}d\sqrt{n})$. While for small minibatch sizes T_{mini} can be quite small (plausibly even smaller than 1), when d is sufficiently large, this communication cost is prohibitively large. This is the reason for the slow performance of mini-batched optimization that we observed in our experiments. Reworking these algorithms with sparse parameter updates, the communication cost per update becomes bs yielding an overall communication cost of $\Theta(T_{\text{mini}}ns)$, which is still several multiples of the data set size. Empirically, it has also been noted that after optimizing learning rate parameters, the optimal minibatch size is often 1 (Hsu et al., 2011).

Another category of algorithms is those which use online learning with a *feature based* partition of examples (Hsu et al., 2011; Dean et al., 2012). The advantage of this class of algorithms is that they can scale to a very large number of parameters, more than can be fit in the memory of a single machine. Several families of algorithms have been tested in Hsu et al. (2011) including delayed updates, minibatch, second-order minibatch, independent learning, and backpropagation. The per-node communication costs differ substantially here. Typical communication costs are $\Theta(ns/m + nT_{\text{online}})$ where the first term is due to shuffling from an example-based format, and the second term is for the run of the actual algorithm. The complexity of our approach is superior to this strategy since $n \gg d$.

5. Discussion

We have shown that a new architecture for parallel learning based on a Hadoop-compatible implementation of AllReduce can yield a combination of accurate prediction and short training time in an easy programming style. The hybrid algorithm we employ allows us to benefit from the rapid initial optimization of online algorithms and the high precision of batch algorithms where the last percent of performance really matters. Our experiments reveal that each component of our system is critical in driving the performance benefits we obtain. Specifically, Table 4 and Figure 3 show the performance gains resulting from our use of AllReduce and the warmstart of the L-BFGS algorithm. The effectiveness of our overall system, as compared to the previous approaches, is confirmed in Figure 5. Two issues we do not discuss in this paper are the overheads of data loading and node scheduling within Hadoop. These issues can indeed affect the performance, but we found that they typically get amortized since they are one-time overheads in the AllReduce approach as opposed to per-iteration overheads in MapReduce. Nonetheless, improvements in the scheduling algorithms can further improve the overall performance of our system.

Our paper carefully considers various design choices affecting the communication and computation speeds of a large-scale linear learning system, drawing from and building upon the available techniques in the literature. The resulting system enables the training of linear predictors on data sets of size unmatched in previous published works. In particular, the results demonstrate the effectiveness of our system compared to other alternatives in the literature. We believe that this provides a very strong and natural baseline which we previously found lacking in the literature on distributed machine learning. The conceptual

simplicity of our framework, and the open-source implementation should further help other researchers in comparing with and building on our system.

References

- A. Agarwal and J. Duchi. Distributed delayed stochastic optimization. In *Advances in Neural Information Processing Systems 24*. 2011.
- P. L. Bartlett and S. Mendelson. Rademacher and Gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3:463–482, 2002.
- R. Bekkerman, M. Bilenko, and J. Langford. A tutorial on scaling up machine learning. Technical report, KDD, 2011. URL http://hunch.net/~large_scale_survey/.
- D. P. Bertsekas and J. N. Tsitsiklis. *Parallel and Distributed Computation: Numerical Methods*. Prentice-Hall, Inc., 1989.
- S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3:1–122, 2011.
- K. Canini, T. Chandra, E. Ie, J. McFadden, K. Goldman, M. Gunter, J. Harmsen, K. LeFevre, D. Lepikhin, T. L. Llinares, I. Mukherjee, F. Pereira, J. Redstone, T. Shaked, and Y. Singer. Sibyl: A system for large scale supervised machine learning. In *MLSS Santa Cruz*, 2012. URL <http://users.soe.ucsc.edu/~niejiazhong/slides/chandra.pdf>. A short presentation.
- T. Chandra, E. Ie, K. Goldman, T. L. Llinares, J. McFadden, F. Pereira, J. Redstone, T. Shaked, and Y. Singer. Sibyl: a system for large scale machine learning. In *LADIS 2010: The 4th ACM SIGOPS/SIGACT Workshop on Large Scale Distributed Systems and Middleware*, 2010. URL <http://www.magicbroom.info/Papers/Ladis10.pdf>. A keynote talk.
- O. Chapelle, E. Manavoglu, and R. Rosales. Simple and scalable response prediction for display advertising. *Transactions on Intelligent Systems and Technology*, 2013. In press.
- C.T. Chu, S.K. Kim, Y.A. Lin, Y.Y. Yu, G. Bradski, A.Y. Ng, and K. Olukotun. Map-reduce for machine learning on multicore. In *Advances in Neural Information Processing Systems 19*, volume 19, page 281, 2007.
- J. Dean and S. Ghemawat. Mapreduce: simplified data processing on large clusters. *Communications of the ACM*, 51:107–113, 2008.
- J. Dean, G. Corrado, R. Monga, K. Chen, M. Devin, Q. Le, M. Mao, M. A. Ranzato, A. Senior, P. Tucker, K. Yang, and A. Ng. Large scale distributed deep networks. In *Advances in Neural Information Processing Systems 25*, pages 1232–1240. 2012.
- O. Dekel, R. Gilad-Bachrach, O. Shamir, and L. Xiao. Optimal distributed online prediction using mini-batches. *Journal of Machine Learning Research*, 13:165–202, 2012.

- L. Devroye, L. Györfi, and G. Lugosi. *A Probabilistic Theory of Pattern Recognition*. Springer, 1996.
- J. Duchi, E. Hazan, and Y. Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12:2121–2159, 2010.
- J.C. Duchi, A. Agarwal, and M.J. Wainwright. Dual averaging for distributed optimization: Convergence analysis and network scaling. *Automatic Control, IEEE Transactions on*, 57(3):592–606, 2012.
- K. Hall, S. Gilpin, and G. Mann. Mapreduce/bigtable for distributed optimization. In *Workshop on Learning on Cores, Clusters, and Clouds*, 2010.
- D. Hsu, N. Karampatziakis, J. Langford, and A. Smola. Parallel online learning. In *Scaling Up Machine Learning*, 2011.
- N. Karampatziakis and J. Langford. Online importance weight aware updates. In *Uncertainty in Artificial Intelligence*, 2011.
- M. Kearns. Efficient noise-tolerant learning from statistical queries. In *Proceedings of the Twenty-Fifth Annual ACM Symposium on the Theory of Computing*, 1993.
- J. Langford, L. Li, and A. Strehl. Vowpal wabbit open source project. Technical report, Yahoo!, 2007.
- J. Langford, A. Smola, and M. Zinkevich. Slow learners are fast. In *Advances in Neural Information Processing Systems 22*, 2009.
- D. C. Liu and J. Nocedal. On the limited memory bfgs method for large scale optimization. *Mathematical Programming*, 45:503–528, 1989.
- Y. Low, J. Gonzalez, A. Kyrola, D. Bickson, C. Guestrin, and J. M. Hellerstein. Graphlab: A new framework for parallel machine learning. In *Uncertainty in Artificial Intelligence*, 2010.
- O. L. Mangasarian. Parallel gradient distribution in unconstrained optimization. *SIAM Journal on Optimization*, 33:1916–1925, 1995.
- R. McDonald, K. Hall, and G. Mann. Distributed training strategies for the structured perceptron. In *North American Chapter of the Association for Computational Linguistics (NAACL)*, 2010.
- H. B. McMahan and M. Streeter. Adaptive bound optimization for online convex optimization. In *Proceedings of the Twenty Third Annual Conference on Computational Learning Theory*, 2010.
- S. G. Nash and A. Sofer. Block truncated-newton methods for parallel optimization. *Mathematical Programming*, 45:529–546, 1989.
- J. Nocedal. Updating quasi-Newton matrices with limited storage. *Math. Comp.*, 35(151):773–782, 1980.

- S. Sonnenburg and V. Franc. COFFIN: a computational framework for linear SVMs. In *International Conference on Machine Learning*, 2010.
- S. Sonnenburg, G. Rätsch, and K. Rieck. Large scale learning with string kernels. In *Large Scale Kernel Machines*, pages 73–103. 2007.
- C. Teo, Q. Le, A. Smola, and SVN Vishwanathan. A scalable modular convex solver for regularized risk minimization. In *ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 2007.
- K. Weinberger, A. Dasgupta, J. Langford, A. Smola, and J. Attenberg. Feature hashing for large scale multitask learning. In *International Conference on Machine Learning*, 2009.
- J. Ye, J.-H. Chow, J. Chen, and Z. Zheng. Stochastic gradient boosted distributed decision trees. In *Proceeding of the 18th ACM Conference on Information and Knowledge Management*, 2009.
- M. Zaharia, M. Chowdhury, T. Das, A. Dave, J. Ma, M. McCauley, M. Franklin, S. Shenker, and I. Stoica. Resilient distributed datasets: A fault-tolerant abstraction for in-memory cluster computing. Technical Report UCB/EECS-2011-82, EECS Department, University of California, Berkeley, 2011.
- M. Zinkevich, M. Weimer, A. Smola, and L. Li. Parallelized stochastic gradient descent. In *Advances in Neural Information Processing Systems 23*. 2010.