

MALT: Distributed Data-Parallelism for existing ML applications

Hao Li*, Asim Kadav, Erik Kruus, Cristian Ungureanu

NEC Labs, Princeton

{hli, asim, kruus, cristian} @nec-labs.com

Abstract

Machine learning methods, such as SVM and neural networks, often improve their accuracy by using models with more parameters trained on large numbers of examples. Building such models on a single machine is often impractical because of the large amount of computation required.

We introduce MALT, a machine learning library that integrates with existing machine learning software and provides data parallel ~~learning-machine learning~~. MALT provides abstractions for fine-grained in-memory updates ~~-, supports sparse and asynchronous dissemination of model updates to all nodes over using~~ one-sided RDMA, ~~uses configurable dataflow for balancing computation limiting data movement costs during incremental computation~~. MALT ~~allows machine learning developers to specify the dataflow and apply communication and communication, and handles failures gracefully. representation~~ ~~optimizations~~. Through its general-purpose API, ~~we show that~~ MALT can be used to provide data-parallelism to existing ~~applications such as ML applications written in C++ and Lua and based on~~ SVM, matrix factorization and neural networks. In our results, we show ~~that~~ MALT provides fault tolerance, network efficiency and ~~parallel~~-speedup to these applications.

1. Introduction

Machine learning (ML) is becoming increasingly popular due to a confluence of factors: an abundance of data produced and captured in digital form [30]; an abundance of compute power and convenient access to it from various devices; advances in the ML field, making it applicable to an ever growing number of situations [28]. The acceptance and success of ML, from natural language processing to image recognition to others, comes from the increasing accu-

racy achieved by ML applications. This accuracy is achieved partly through advances in ML algorithms, but also through using *known algorithms* with *larger models* trained on *larger datasets* [28]. Building these models on a single machine is often impractical because of the large amount of computation required, or may even be impossible for very large models such as those in state-of-the-art image recognition.

Existing data-parallel frameworks such as the map-reduce model have proven to be tremendously useful and popular paradigm for large-scale batch computations. However, existing frameworks are a poor fit for long running machine learning tasks. Machine learning algorithms such as gradient descent are iterative, and make multiple iterations to refine the output before converging to an acceptable value. Machine learning tasks have *all* of the following properties:

- *Fine-grained and Incremental* : Machine learning tasks perform repeated model updates over new input data. Most existing processing frameworks lack abstractions to perform iterative computations over small modifications *efficiently*. This is because in existing map-reduce models, jobs synchronize using the file-system [23] or maintain in-memory copies of intermediate data ~~[?]~~[53]. For computations with large number of iterations and small modifications, techniques such as these are sub-optimal.
- *Asynchronous* : Machine learning tasks ~~running concurrently in a cluster that run in parallel~~ may communicate asynchronously. As an example, models that train in parallel may synchronize model parameters ~~in this fashion~~ ~~asynchronously~~. Enforcing determinism ~~on in~~ the order of parameter updates can cause unnecessary performance overhead.
- *Approximate* : Machine learning tasks ~~can be approximate~~. They may perform computation stochastically and often an approximation of the trained model is ~~good enough~~. ~~Existing solutions sufficient~~. ~~Existing general purpose systems~~ rarely provide abstractions for trading off strong guarantees such as fault tolerance or accuracy for performance (reduced job times). ~~General purpose systems impose strong guarantees such as determinism which may slow down machine learning algorithms which may not need these guarantees~~.
- *Need Rich Developer Environment*: Developing ML applications require a rich set of ML libraries, developer

* Work done as NEC Labs intern. Currently at U. Maryland.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

EuroSys'15, April 21–25, 2015, Bordeaux, France.

Copyright © 2015 ACM 978-1-4503-3238-5/15/04...\$15.00

tools and graphing abilities which is ~~non-existent in the recent~~ often missing in many highly scalable systems. Furthermore, existing ML software such as scikit [2], Torch [31], RAPID [43], R [3] provide an efficient single-system library that performs well over multiple cores. However, most existing distributed learning tasks ~~requires developers to re-write~~ force developers to re-write their existing libraries in a new software stack and expose an unfamiliar environment to the developers.

To address these properties, we propose a system called as MALT (stands for distributed Machine Learning Toolset), that allows ML developers to run their existing ML software in a distributed fashion. MALT provides an efficient shared memory abstraction that runs existing ML software in parallel and allows them to communicate updates periodically. MALT exports a `scatter-gather` API, that allows pushing model parameters or model parameter updates (called as *gradients*) to ~~its parallel replicas~~. The parallel model replicas. These replicas then process the received values by invoking a user-supplied gather function locally. Additionally, the API allows developers to specify the dataflow across these replicas and ~~also~~ specify representation optimizations (such as ~~compression, sparseness etc.~~). Furthermore, sparseness. MALT communication is ~~based over~~ designed using one-sided RDMA writes (no reads for ~~one-half round-trip times~~ faster round-trip times [34]) and provides ~~mechanisms~~ abstractions for asynchronous model training.

Our parallel, peer-to-peer model communication complements the ~~recent~~ master-slave style parameter server approach [22, 36] [21, 22, 36]. In MALT, parallel model replicas send model updates to one-another instead of a central parameter server. This reduces network costs because the machines only communicate model updates ~~instead of complete models. A peer-to-peer model also simplifies fault tolerance. There is no central master component, and back and forth instead of full models. Furthermore,~~ implementing MALT, does not require writing separate master/slave code or dealing with complex ~~recover~~ recovery protocols to deal with master failures.

The contributions of this paper are as follows:

- We describe a general data-parallel machine learning framework that provides simple and flexible API for parallel learning. MALT provides APIs for sending model parameter updates, designing the data flow of the communication, and making this communication synchronous or asynchronous. Furthermore, MALT abstracts RDMA programming, and deals with system issues like recovering from unresponsive or failed nodes.
- MALT provides ML developers data-parallelism in their existing ML software. We demonstrate how MALT can transform existing ML software written in procedural or scripting languages with reasonable developer efforts.

This allows developers to incorporate data-parallelism in their applications in a familiar environment and use rich developer tools provided within their ML software.

- We demonstrate a network efficient, parallel learning implementation of MALT where we trade-off model freshness at replicas with faster model training times, for the same final accuracy. We use MALT to speedup three applications using gradient descent: SVM [10], matrix factorization [46] and neural networks [43] with reasonable developer efforts and speedup over existing methods [16]. We demonstrate that MALT outperforms single machine performance for small workloads and can efficiently train models over large datasets that span multiple machines.

We now present background on machine learning.

2. Distributed Machine Learning

Machine learning algorithms generalize from data. Machine learning algorithms train over data to create a *model* representation that can predict outcomes (regression or classification) for new unseen data. More formally, given a training set $((\vec{x}_1, \vec{y}_1), (\vec{x}_2, \vec{y}_2), \dots, (\vec{x}_n, \vec{y}_n))$, the goal of model training is to determine the distribution function f such that $\vec{y} = f(\vec{x})$. The input \vec{x} may consist of different features and the model consists of parameters \vec{w} , representing the weights of individual features to compute \vec{y} . The goal of model training is to estimate the values of model parameters \vec{w} . During model testing, this model is tested using an unseen set of \vec{x}_i to compare against ground truth (already known \vec{y}_i), to determine the model accuracy. Thus, machine learning algorithms *train* to ~~optimize the minimization of the~~ minimize the *loss*, which represents some function that evaluates the difference between estimated and true values for the test data.

Model training algorithms are iterative, and the algorithm starts with an initial guess of the model parameters and learns incrementally over data, and refines the model every iteration, to converge to a final acceptable value of the model parameters. Model training time can ~~take from hours~~ last from minutes to weeks and is often the most timing consuming aspect of the learning process. Model training time also hurts model refinement process since longer training times limit the number of times the model configuration parameters (called as hyper-parameters) can be tuned through re-execution.

Machine learning algorithms can benefit from a scale-out computing platform support in multiple ways: First, these algorithms train on large amounts of data, that improves model accuracy [28]. Second, they can train large models that have hundreds of billions of parameters or require large computation ~~(such as requiring lots of layers in a neural network)~~ such as very large neural networks for large-scale image classification or genomic applications [15]. Training with more data is done by *data parallelism*, which requires repli-

cating the model over different machines and synchronizing the model parameters after a fixed number of iterations. Training large models requires the model to be split across multiple machines, and is referred to as *model parallelism*.

MALT API limits itself to data parallelism because models that learn over vast amounts of training data are more common than models with 100 billion parameters [7]. Furthermore, a single machine can process models of the order of 10 billion parameters in-memory (about 80 GB for dense representation; a server with 128GB DRAM costs about 3000\$). This is sufficient for most large machine learning models. Second, even though exposing a distributed, replicated, shared array is fairly straight-forward, efficiently providing model parallelism in existing applications is non-trivial. It requires exposing APIs and modifying machine learning algorithms to ensure that the model is split such that the communication costs are limited within each iteration (this is zero for data-parallelism). This is feasible for systems that focus on specific algorithms that are amenable to such splits like convolution networks [15, 22] or systems that write their own algorithms [36]. MALT’s goal is to provide a simple, general purpose API that integrates easily with existing software, with a wide-variety of algorithms with reasonable developer efforts.

With datasets getting larger, there has been a recent focus to investigate online algorithms that can process datasets incrementally such as the *gradient descent* family of algorithms. Gradient descent algorithms compute the gradient of the loss function over the entire set of training examples. This gradient is used to update model parameters to minimize the loss function [9]. Stochastic Gradient Descent (SGD) is a variant of the above algorithm that trains over one single example at time. With each example, the parameter vector is updated until the loss function yields an acceptable (low) value. SGD and its variants are preferred algorithms to train over large data-sets because it can process large training datasets in batches. Furthermore, gradient descent can be used for a wide-range of algorithms such as regression, k-means, SVM, matrix-factorization and neural networks [11, 17].

In data-parallel machine learning, model replicas train over multiple machines. Each replica trains over a subset of data. There are several ways in which the individual model parameters can be synchronized. We describe three such methods. First, models may train independently and synchronize parameters when all parallel models finish training by exhausting their training data. ~~This~~ These methods are commonly used to train over Hadoop where communication costs are prohibitive. In addition, while these models may train quickly ~~due to~~ because of limited communication between replicas, they may require more passes over training data (each pass over training data is called an *epoch*) for acceptable convergence. Furthermore, for non-convex problems, this method may not converge, since the parallel repli-

cas may be trapped in a different local minimas, and averaging these diverging models may return a model with low accuracy.

The second method is the *parameter server* approach [22, 36]. Here, individual models send their updates to a central parameter server (or a group of parameter servers) and receive an updated model from them. A third method is the *peer-to-peer* approach (used in MALT), where parameters from model replicas train in parallel and are mixed every (or every few) iteration [39]. The last two methods achieve good convergence, even when the parameters are synchronized communicated asynchronously [29, 37]. With MALT, we perform asynchronous parameter mixing with multiple parallel instances of model replicas. This design allows developers to write code once, that runs everywhere on parallel replicas (no separate code for parameter server and client). This design also simplifies fault tolerance – a failed replica is removed from the parameter mixing step and its data is redistributed to other replicas. Finally, instead of performing simple gradient descent, MALT can be used to implement averaging of gradients from its peers, which provides speedup in convergence ~~+51~~ for certain workloads [11, 51].

The goal of our work is to provide distributed-machine learning over existing ML systems. MALT exposes an asynchronous parameter mixing API that can be integrated into existing ML applications to provide data-parallel learning. Furthermore, this API is general enough to incorporate different communication and representation choices as desired by the machine learning developer. MALT provides peer-to-peer learning by interleaving gradient (changes to parameters) updates with parameter values to limit network costs. In the next section, we describe MALT design.

3. MALT Architecture

Figure 1 describes MALT architecture. Model replicas train in parallel on different cores (or sets of cores) across different nodes using existing ML libraries. ML libraries use the MALT vector library to create model parameters or gradients (updates to parameters) that need to be synchronized across machines. These vectors communicate over Distributed One-sided Remote Memory or *dstorm*. Furthermore, like other data-parallel frameworks, MALT loads data in model-replicas from a distributed file-system such as NFS or HDFS. Developers use the MALT API to shard data input across replicas and send/receive gradients. Furthermore, developers can also specify the dataflow across replicas and make their algorithms fully asynchronous. We now describe the shared memory design, the MALT API that allows developers access to shared memory, MALT fault tolerance and a flexible network communication design that allows developers to balance communication and computation.

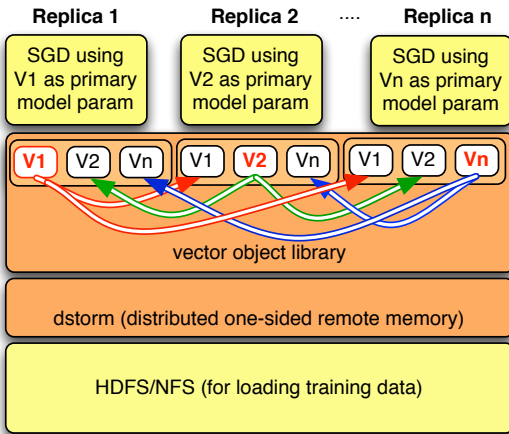


Figure 1. MALT architecture. Existing applications run with modified gradient descent algorithms that receive model update (V) from replicas training on different data. Model vectors are created using a Vector Object Library that allows creation of shared objects. Each replica *scatters* its model update after every (or every few) iteration and *gathers* all received updates before the next iteration.

3.1 Abstractions for Shared Memory with `dstorm`

Machine learning models train in parallel over *sharded* *sharded* data and periodically share model updates after few iterations. The parallel replicas may do so synchronously (referred to as the bulk-synchronous processing [48]). However, this causes the training to proceed at the speed of the slowest machine in that iteration. Relaxing the synchronous requirement speeds up model training but may affect the accuracy of the generated model. Since model weights are approximate, applications developers and researchers pick a point in this trade-off space (accuracy vs speed) depending on their application and system guarantees [20, 46]. Furthermore, this accuracy can be improved by training for multiple epochs or increasing the amount of data at *for training* each model replica.

The original map-reduce design communicates results over GFS/HDFS. However, using disk for communication, results in poor performance especially for machine learning applications which may communicate as often as every iteration. Spark [53] provides immutable objects (RDDs) for an efficient in-memory representation across machines. Spark provides fault tolerance using lineage of RDDs as they are transformed across operations. However, this enforces determinism in the order of operations. As a result, the *immutability-immutability* and determinism makes it less suitable for fine-grained, asynchronous operations [49, 53]. Furthermore, machine learning applications may contain multiple updates to large sparse matrices or *model-updates may propagate may need to propagate model updates* asynchronously across machines *requiring-and need* first-class support for fine-grained and asynchronous operations.

MALT’s design provides efficient mechanisms to transmit model updates. There has been recent trend of wide availability of cheap and fast infiniband hardware and they are being explored for applications beyond HPC environments [24, 40]. RDMA over infiniband allows low latency networking of the order of 1-3 micro-seconds *by* using user-space networking libraries and by re-implementing a portion of the network stack in hardware. Furthermore, the RDMA protocol does not interrupt the remote host CPU while accessing remote memory. RDMA is also available over Ethernet with the newer RDMA over Converged Ethernet (ROCE) NICs that have comparable performance to Infiniband. Infiniband NICs are priced competitively with 10G NICs, costing around 500\$ for *40Gbps-40 Gbps* NICs and 800\$ for *56Gbps-56 Gbps* NICs (as of mid 2014). Finally, writes are faster than reads since they incur *half the-lower* round trip-times (*RTTs*) [34]. MALT uses one-sided RDMA writes to propagate model updates across replicas.

In MALT, every machine can create a shared memory *abstraction-called-as-abstractions called segments via* a `dstorm` object (`dstorm` stands for DiSTributed One-sided Remote Memory). Each `dstorm` *object-segment* is created by supplying the object size and a directed dataflow graph. To facilitate one-sided writes, when a `dstorm` *object-segment* is created, the nodes in dataflow synchronously create a distributed shared memory object. `dstorm` registers a portion of memory on every node with the infiniband interface to facilitate one-sided RDMA operations. When a `dstorm` *object-segment* is transmitted by the sender, it appears at all its receivers (as described by the dataflow), without interrupting any of the receiver’s CPU. We call this operation as *scatter*. Hence, *a dstorm segment* allocates space (a receive queue), for every sender in every machine to facilitate the *scatter* operation. We use per-sender receive queues to avoid invoking the receiver CPU for resolving any write-write conflicts arising from multiple incoming model updates from different senders. *Hence, our design uses extra space with the per-sender receive queues to facilitate lockless model propagation using one-sided RDMA. Both these mechanisms, the one sided RDMA and per-sender receive queues ensure that the scatter operation does not invoke the receive-side CPUs.*

Once the received objects arrive in local per-sender receive queues, *it-they* can be read with a local *gather* operation. The *gather* function uses a user-defined function (UDF), such as an average, to collect the incoming updates. We also use queues on the sender side, allowing senders to perform writes asynchronously. Additionally, the sender-side queues maintain a back-pressure in the network to avoid congestion [47].

The receiver does not know when its per-sender receive queues get filled unless the receiver is actively polling and consuming these items. When the receive queue is full, the default behavior of `dstorm` is to *over-write-over-write*

previously sent items in the queue. We discuss the consistency behavior after we describe the vector abstraction to create shared vectors or tensors (~~multi-dimensional~~ multi-dimensional vectors) over the `dstorm` object.

3.2 Vector Object Library: Programming Dstorm for Machine Learning

We build a vector object library (VOL) over `dstorm` that allows creating vector objects over shared memory. The goal of VOL is to 1) Expose a vector abstraction instead of shared memory abstraction (`dstorm`) and 2) to provide communication and representation optimizations. ML developers can specify gradients or parameters as a VOL vector (or tensor) and specify its representation (sparse or dense). They also specify a dataflow graph describing how the updates should be propagated in the cluster which is used to create the underlying `dstorm` object.

Hence, creating a vector in-turn creates a `dstorm` object segment that allows this vector to be propagated to all machines as described in the dataflow graph. This dataflow describes which machines may send updates to one another (in the simplest case, everyone may send update-their updates to everyone). Hence, an edge in the graph from node A to node nodes B and C implies that when node A pushes a model update, it is received by nodes B and node C. As different machines compute model updates, they may *scatter* this to other remote nodes without acquiring any locks or invoking any operations at the receiver. However, if a machine sends too many updates before the previous ones are consumed, the previous updates are over-written.

VOL inherits `scatter` and `gather` calls from `dstorm` to send the vector to remote machine and gather all the received updates (from local memory). Developers can also specify where to spend the model updates within `scatter` calls. This provides fine-grained access to data-flow to the developers, allowing greater flexibility [41]. Table 1 describes the VOL API. In Section 4, we describe how this API can be used to easily convert serial ML algorithms to data-parallel.

Consistency guarantees: We now describe the consistency guarantees that MALT provides when transmitting model updates to other replicas. With machine learning applications, which are stochastic in nature, model updates maybe be over-written or updated locklessly without affecting overall accuracy of the model output significantly. For example, Hogwild demonstrates that asynchronous, ~~lockless~~ lock-less model updates lead to models that ultimately converge to acceptable accuracy [46]. Hence, MALT need not provide strict consistency guarantees when sending model updates over infiniband (ex. as in key-value stores [24]). However, since MALT is a general-purpose API, it provides mechanisms to enforce consistency.

1. *Torn reads:* When a model replica sends a model update to another model replica, the sender may overwrite the model update while the receiver is reading it in the case

where the replicas operate asynchronously and the receive queue is full. MALT provides an additional `atomic_gather` which reads the shared memory in an atomic fashion.

2. *Bounded staleness:* Model updates carry an `iteration_count` information in the header. When a receiver realizes that a specific model update is arriving too slowly, the receiver may stall its operations until the sender catches up. This design is similar to the bounded-staleness approach explored by recent work [20].

If stricter guarantees are required, the model replicas can train synchronously in *bulk-synchronous* fashion and use the `barrier` construct to do so. The `barrier` construct is a conventional barrier which waits for all model replicas to arrive at a specific point in the training process.

3.3 Fault tolerance

MALT has a straight-forward model for fault tolerance. The training data is present on all machines in a distributed file system. The model replicas train in parallel and perform one-sided writes to all peers in the communication. A fault monitor (similar to a watchdog timer) on every node examines the return values of asynchronous writes on sender-side queues. If the fault monitor observes failed writes, it performs a synchronous health check of the cluster with other monitors on other nodes. A node is considered dead if the node is corrupt (the shared memory or the queue has failed) and the remote fault monitor reports this, or if the node is unreachable by *any* of the other healthy nodes fault monitors. ~~Local~~ Furthermore, to detect failure cases that do not result in a machine/process crash, local fault monitors can detect processor exceptions (such as divide by zero, stack corruption, invalid instructions) and segmentation faults and terminate the local training process.

In case of failure, the working fault monitors create a group of survivor nodes to ensure that all future group operations such as `barrier`, skip the failed nodes. The RDMA interface is registered again (with old memory descriptors) and the queues are re-built. This is to avoid a zombie situation where a dead node may come back and attempt to write to one of the previously registered queues. Finally, the send and receive lists of all model replicas are rebuilt to skip the failed nodes and the training is resumed.

After recovery, if a convergence with an acceptable loss value is not achieved, the training continues on the survivor replicas with more additional training examples until the models converge. This causes a slowdown in the training process proportional to the missing machines apart from a short delay to synchronize and perform recovery (of the order of seconds). MALT only provides fail-stop fault tolerance, i.e. it can only handles failures where a ~~node~~ self-reports-fault monitor detects corruption or is unresponsive because of the MALT process being killed or a machine failure or a network failure. MALT cannot handle Byzantine

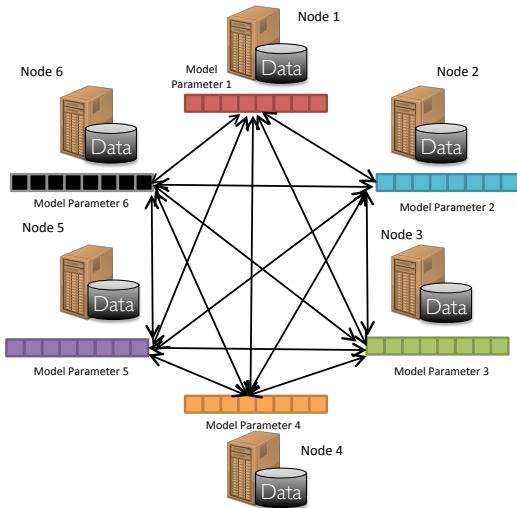


Figure 2. All-reduce exchange of model updates. All arrows indicate bi-directional communication. As number of nodes(N) grow, total number of updates transmitted increases $O(N^2)$.

failures such as when a machine sends corrupt gradients or software corruption of scalar values that cannot be detected by local fault monitors. However, the local fault monitors can be extended to protect against these failures [13].

MALT can afford such simple fault tolerance model because it only provides data parallelism and does not split the model across multiple machines. Furthermore, the model training is stochastic and does not depend on whether the training examples are processed in a specific order, or the training examples are processed more than once, or whether all the training examples have been processed, as long as the model achieves an acceptable accuracy. Furthermore, MALT implements peer-to-peer learning and does not have a central master. As a result, it does not need complex protocols like Paxos [14] to recover from master failures.

3.4 Communication Efficiency in MALT

MALT’s flexible API can model different training configurations such as parameter server [36], mini-batching [19, 38] and peer-peer parameter mixing [39].

When MALT is trained using athe peer-to-peer approach, each machine can send its update to all other machines to ensure that each model receives the most recent updates. We refer to this configuration as MALT_{all-all}. As the number of nodes (N) increases, the gradient communication overhead in MALT_{all-all} increases $O(N^2)$ times, in a naïve all-reduce implementation. Efficient all-reduce primitives such as butterfly [12] or tree style all-reduce [5], reduce the communication cost by propagating model updates in a tree style. However, this increases the latency by a factor of the height of the tree. Furthermore, if the intermediate nodes are affected by stragglers or failures, an efficient all-reduce makes recovery complex.

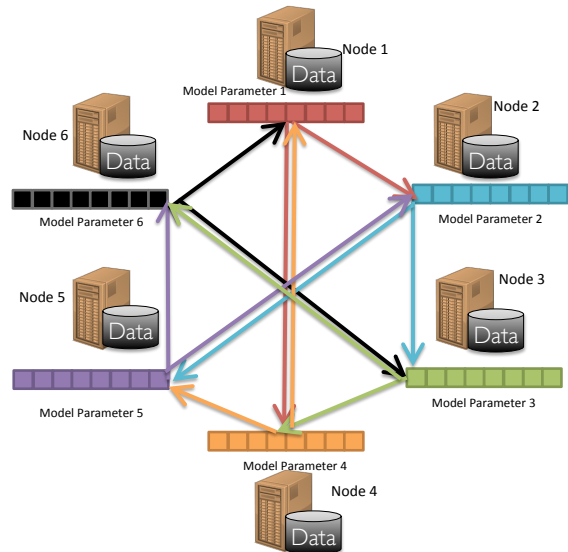


Figure 3. Halton series exchange of model updates ($N = 6$). Each i th machine sends updates to $\log(N)$ (2 for $N = 6$) nodes. (to $N/2 + i$ and $N/4 + i$). As number of nodes N increases, the outbound nodes follows Halton series ($N/2, N/4, 3N/4, N/8, 3N/8..$). All arrows are uni-direction. As number of nodes(N) grow, total number of updates transmitted increases $O(N \log N)$.

In MALT, we propose an efficient mechanism to propagate model updates, what we refer to as *indirect* propagation of model updates. A developer may use the MALT API to send model updates to either all nodes (N) or fewer nodes k , ($1 \leq k \leq N$). MALT API facilitates choosing a value ‘k’ such that a MALT replica (i) disseminate the updates across all the nodes eventually; (ii) optimize specific goals of the system such as freshness, balanced communication/computation ratio in the cluster. By eventually, we mean that all nodes receive a model update from every other node directly or indirectly via an intermediate node. However, when choosing a value ‘k’, less than N, the developer needs to ensure that the communication graph of all nodes is connected.

Hence, instead of performing an all-reduce, MALT limits the reduce operation to fewer nodes. However, naïvely or randomly selecting what nodes to send updates to may either leave out certain nodes from receiving updates from specific nodes (a partitioned graph of nodes) or may propagate updates that may be too stale (a weakly connected node graph). This may adversely affect convergence in parallel learning models. We now describe how MALT can selectively distribute model updates to ensure low communication costs and uniform dissemination of model updates.

MALT API provides a pre-existing data-flow that sends fewer model updates and ensures that all models send/receive model updates in a uniform fashion. To do so, every node picks a node in a uniform fashion to ensure that the updates are distributed across all nodes. For example, if every node

propagates its updates to k nodes ($k < N$), we pick the k node IDs based on a uniform random sequence such as Halton sequence [1] that generates successive points that ~~cover the ID range as uniformly as possible for a given k~~ create a k -node graph with good information dispersal properties. We further propose that each node only send update to $\log(N)$ nodes and maintain a $\log(N)$ sized node list. This node list contains the nodes to send updates to, generated using the Halton series. Hence, if we mark the individual nodes in training cluster as $1 \dots N$, Node 1 sends its updates to $N/2, N/4, 3N/4, N/8, 3N/8, 5N/8, \dots$ and so on (the Halton series with base 2). Hence, in this scheme, the total updates sent in every iteration is only $O(N \log(N))$. We refer to this configuration as $\text{MALT}_{\text{halton}}$. The $\text{MALT}_{\text{halton}}$ scheme ensures that the updates are sent uniformly across the range of nodes. Figures 2 and 3 show the all-to-all and Halton communication schemes. In case of a failure with $\text{MALT}_{\text{halton}}$, the failed node is removed and the send/receive lists are rebuilt.

Using MALT’s network-efficient parallel model training results in faster model training times. This happens because 1) The amount of data transmitted is reduced. 2) The amount of time to compute average of gradients is reduced since the gradient is received from fewer nodes. 3) In a synchronized implementation, this design reduces the number of incoming updates that each node needs to wait for, before going on to the next iteration. Furthermore, our solution reduces the need for high bandwidth interfaces, reducing costs or freeing up the network for other applications. In Section 6, we compare the network costs of indirect updates ($\text{MALT}_{\text{halton}}$) with $\text{MALT}_{\text{all-all}}$ and the parameter server.

Instead of having each node communicate with $\log(N)$ other nodes, developers can program MALT to communicate with higher (or lower) number of nodes. The key idea is to balance communication (sending updates) with computation (computing gradients, applying received gradients). Hence MALT, accepts a dataflow graph as an input while creating vectors for model parameters. However, ~~this requires that developer to ensure that~~ the graph of nodes is-needs to be connected otherwise the model updates-individual model update from a model replica may not get propagated to all nodes, and the model-may-not converge-models may diverge significantly from one another.

4. Programming Interface

The goal of MALT is to provide data-parallelism to any machine learning software or algorithm. Given the MALT library and a list of machines, developers launch multiple replicas of their existing software that perform data-parallel learning. ~~Currently, allows programmers to extend or write programs in C++ and Lua-~~

MALT exposes an API as shown in Table 1. This API can be used to create (or port existing) ML applications for data-parallelism. To do so, the developer creates a param-

eter or gradient object using MALT API. Figure 4 shows a serial SGD algorithm (Algorithm 1) and a parallel SGD written using MALT (Algorithm 2). In the serial algorithm, the training algorithm goes over entire data and for each training sample, it calculates the associated gradient value. It then updates the model parameters, based on this gradient value.

In order to perform this training in a data-parallel fashion, this algorithm can be re-written using MALT API (as shown in algorithm 2). The programmer specifies the representation (sparse vs dense) and the dataflow (ALL – which represents all machines ~~talk-communicate model updates~~ to one-another, HALTON – which represents the network efficient API from previous section or the developer may specify ~~a~~ an arbitrary graph – which represents the dataflow graph). When a job is launched using MALT it runs this code on each machine. Each machine creates a gradient vector object using the MALT API, with the required representation properties (sparse vs dense), and creates communication queues with other machines based on the dataflow specified, and creates receiving queues for incoming gradients.

Each machine trains over a subset of training data and computes the gradient value for each example. ~~For every sample (or after a bunch of samples)~~ After training over each example (or bunch of examples), this gradient value is sent using ~~a-the~~ one-sided RDMA operation. The algorithm then computes an average of the received gradients using the gather function. Instead of an average, one can ~~pass-specify~~ a user-defined function (UDF) to compute the resulting gradient from all incoming gradients. This is useful for algorithms where a simple averaging may not work, such ~~for SVM and as SVM may~~ require an additional re-scaling function apart from averaging parameters-performing an average over the incoming parameters. The training finishes when all machines in the cluster finish training over local examples. The final parameter value W is identical across all machines in the synchronous, all-all case. In other cases, W may differ slightly across machines but is within an acceptable loss value. In such cases, the parameters from any machines may be used as the final model or an additional reduce can be performed over W to obtain final parameter values.

Algorithm 1 Serial SGD

```

1: procedure SERIALSGD
2:   Gradient  $g$ ;
3:   Parameter  $W$ ;
4:
5:   for  $epoch = 1 : \text{maxEpochs}$  do
6:
7:     for  $i = 1 : \text{maxData}$  do
8:        $g = \text{cal\_gradient}(\text{data}[i]);$ 
9:        $W = W + g;$ 
10:
11:   return  $W$ 

```

MALT API call	Purpose of the call
<code>g = createVector (Type)</code>	Creates a globally accessible shared model parameter or gradient (model update) vector. Type signifies sparse or dense.
<code>g.scatter (Dataflow Graph optional)</code>	Send model (or just model updates) to machines as described in graph (default sends to all machines).
<code>g.gather (func)</code>	Apply user-defined function <code>func</code> (like average) over model updates that have arrived (locally) and return a result.
<code>g.barrier ()</code>	Distributed barrier operation to force synchronization.
<code>load_data (f)</code>	Shard and load data from HDFS/NFS from file <code>f</code> .

Table 1. ~~The interface exported by the MALT interface.~~ `g.scatter ()` performs one-sided RDMA writes of gradient `g` to other machines. `g.gather ()`, a local operation, applies average to the received gradients. ~~Optional~~ `g.barrier ()` makes the algorithm synchronous

Algorithm 2 Data-Parallel SGD with MALT

```

1: procedure PARALLELSGD
2:   maltGradient  $g$ (SPARSE, ALL);
3:   Parameter  $W$ ;
4:
5:   for  $epoch = 1 : \max Epochs$  do
6:
7:     for  $i = 1 : \max Data / \text{totalMachines}$  do
8:        $g = \text{cal\_gradient}(data[i])$ ;
9:        $g.scatter(ALL)$ ;
10:       $g.gather(AVG)$ ;
11:       $W = W + g$ ;
12:
13:   return  $W$ 

```

Figure 4: Data-parallel machine learning using MALT. The serial code (in Algorithm 1) is converted to data-parallel using MALT. All machines run the above code (in Algorithm 2). Instead of average, user may specify a function to combine incoming gradients/parameters. Optionally, `g.barrier ()` may be used to run the algorithm in a synchronous fashion.

For more complex algorithms, such as neural networks, which require synchronizing parameters at every layer of neural network, each layer of parameters is represented using a separate `maltGradient` and can have its own dataflow, representation and synchronous/asynchronous behavior.

Finally, it may be difficult to use the `maltGradient` allocation for certain legacy applications that use their own data-structures for parameters or gradients. For such opaque representations, where MALT cannot perform optimizations such as sparseness, developers directly use `dstorm`. `dstorm` provides ~~low-level~~ low-level shared memory access with `scatter` and `gather` operations, allows managing the dataflow and controlling the synchronization. Furthermore, the opaque data-structures need to provide a serialization/de-serialization methods to copy-in/out from `dstorm`. Developers can also implement model-parallelism by carefully sharding their model parameters overs multiple `dstorm`.

~~Serial SGD Gradient g ; Parameter W ; $epoch = 1 : \max Epochs$; $i = 1 : \max Data$ $g = \text{cal_gradient}(data[i])$; $W = W + g$; W~~

~~Data-Parallel SGD with MALT~~ `maltGradient` g (SPARSE, ALL);
~~Parameter W ; $epoch = 1 : \max Epochs$; $i = 1 : \max Data / \text{totalMachines}$~~
 ~~$g = \text{cal_gradient}(data[i])$; $g.scatter(ALL)$; $g.gather(AVG)$;~~
 ~~$W = W + g$; W~~

4.1 Applications

We use the MALT API to make the following algorithms data-parallel. Currently, MALT allows programmers to extend or write programs in C++ and Lua.

4.1.1 Support Vector Machines

We ~~first~~ explore distributed stochastic gradient descent algorithms over ~~simple~~ linear and convex problems ~~like using~~ Support Vector Machines(SVM). We use Leon Bottou’s SVM-SGD [10]. Each machine calculates the partial gradient and sends it to other machines. Each machine averages the received gradients and updates its model weight (w) vector locally.

4.1.2 Matrix Factorization

Matrix factorization involves partitioning a large matrix into its two smaller matrices. This is useful for data composed of two sets of objects, ~~in such a way that observations are observations of couples~~ and their interactions needs to be quantified. As an example, ~~news-personalization~~ movie ratings data contains interactions between users and ~~stories, and the interactions between these entities needs to be quantified.~~ Large-scale movies. By understanding their interactions and calculating the underlying features for every user, one can determine how a user may rate an unseen movie. To scale better, large-scale matrix factorization is not exact, and algorithms approximate the factorizations [26]. SGD gives good performance for matrix factorizations on a single machine [35], and we perform matrix factorization using SGD across multiple machines. We use Hogwild [46] and extend it from a multi-core implementation to a multi-node using MALT. With Hogwild, the gather function is a ~~replace~~ operation that overwrites parameters.

4.1.3 Neural Networks

We train neural networks for text learning. The computation in a neural network occurs over multiple layers forming a network. The training happens in a forward-pass and backward pass. In the forward pass, the input samples are processed at each layer and fed forward into the network, finally returning a predicted result at the end of the network. The difference in the ground truth and this predicted result is used in the *back-propagation* phase to update model weights using gradient descent. Parallel training over neural networks is more difficult than SVM for two reasons. First, a data-parallel neural network requires synchronizing parameters for every layer. Second, finding the model weights for neural networks is a non-convex problem. Hence, just sending the gradients is not sufficient as parallel model replicas maybe ~~training-stuck~~ in different local minimas~~and~~. Hence, gradient synchronization needs to be interleaved with the whole parameters synchronization. We use RAPID [43], and extend its neural-network library with MALT. RAPID is similar in architecture to Torch [31], and provides a C++ libraries with Lua front-end for scripting. MALT exports its calls with Lua bindings and integrates with RAPID.

5. Implementation

MALT is implemented as a library, and is provided as a package to SVM-SGD and RAPID [43], allowing developers to use and extend MALT. dstorm is implemented over GASPI [6], that allows programming shared memory over infiniband. GASPI exposes shared memory segments and supports one-sided RDMA operations. We build dstorm over GASPI, that implements object creation, scatter, gather and other operations. dstorm hides all GASPI memory management from the user and provides APIs for object creation, scatter/gather and dataflow. We choose GASPI over other shared address space implementations such as MPI since unlike MPI, GASPI provides fault tolerance and has superior performance [27].

We ~~also~~ implement the vector object library over dstorm that provides vector abstractions, and provides other APIs for loading data, sparse and dense representations. Overall, MALT library is only 2366 LOC. To integrate with Lua, we have written Lua bindings (in Lua and C++) consisting of 1722 LOC. In Section 6.3, we evaluate the costs of integrating individual applications to MALT.

6. Evaluation

We evaluate MALT along the following criteria:

1. *Speedup with MALT*: What is the speedup provided by using MALT ?
2. *Network Optimizations*: How do the ~~different~~ MALT network optimizations benefit training time?
3. *Developer Effort*: What is the developer effort required to port existing ML applications?

4. *Fault Resilience*: How does MALT behave in the presence of failures?

We use MALT to ~~implement SVM-SGD, matrix factorization and SSI~~ modify SVM-SGD, HogWild (matrix factorization) and RAPID (neural networks). Table 2 lists the application and the datasets used. We perform all experiments on a 8 machine research cluster connected via an infiniband back-plane. We run multiple processes, across these 8 machines, and we refer each process as a *rank* (from the HPC terminology). We run multiple ranks on each machine, especially for models with less than 1M parameters, where a single model replica is unable to saturate the network and CPU. Each machine has an Intel Xeon 8-core, 2.2 GHz Ivy-Bridge processor with support for SSE 4.2/AVX instructions, and 64 GB DDR3 DRAM. Each machine is connected via Mellanox Connect-V3 56 Gbps infiniband card and all machines are connected using a Mellanox managed-switch with copper cables. Our 56 Gbps infiniband network architecture provides a peak throughput of slightly over 40 Gbps after accounting for the bit-encoding overhead for reliable transmission. All machines share storage using a ~~10TB~~ 10 TB NFS partition that we use for loading input data. For all our experiments, we randomize the input data and assign random subsets to each node. All reported times do not account the initial ~~one-time~~ one-time cost for the loading the data-sets in memory. All times are reported in seconds.

We perform our experiments on applications as described in Table 2. We perform data-parallel learning for SVM, matrix factorization and neural networks. We use small and large datasets. The small datasets have well understood convergence behavior and allow us to verify correctness and measure speedup over single machines. The large dataset help us evaluate scalability. To evaluate SVM, we use RCV1, PASCAL ~~and webspam~~ [4] datasets suite (alpha, webspam, DNA) and splice datasets [4]. The compressed dataset sizes are RCV1 – 340MB, ~~PASCAL~~ (700M uncompressed), PASCAL alpha – 1.3 GB ~~and~~ (2 GB uncompressed), webspam – 3.3 GB ~~and~~ (12 GB uncompressed), DNA 2.5 GB (35 GB uncompressed) and splice – 110 GB (250 GB uncompressed). The splice dataset does not fit in a single-machine and requires the entire dataset for an accurate model (sub-sampling of data is unhelpful). For matrix factorization, we use the Netflix dataset (1.5 GB uncompressed). For ~~neural networks~~ neural-networks, we perform click-through rate (CTR) prediction based on the Tencent data released with KDD Cup 2012 challenge [33]. The neural network is a three-layer fully- connected neural network that performs supervised-semantic indexing (SSI) [8]. ~~The training-step requires SSI model trains on 2 GB of processed training data of about 2GB in size.~~ (5.1 GB uncompressed).

6.1 Speedup

Application	Dataset	# of training items	# of testing items	Model	Params
Document classification	RCV1	781K	23K	SVM	47,152
Image classification	PASCAL	250K	250K	SVM	500
DNA	DNA	23M	250K	SVM	800
Genome detection	Splice-site	10M	111K	SVM	11M
Webspam detection	Webspam	250K	100K	SVM	16.6M
Collaborate Collaborative filtering	Netflix	100M	2.8M	MF	14.9M
CTR prediction	KDD12	150M	100K	SSI	12.8M

Table 2. MALT applications and dataset properties.

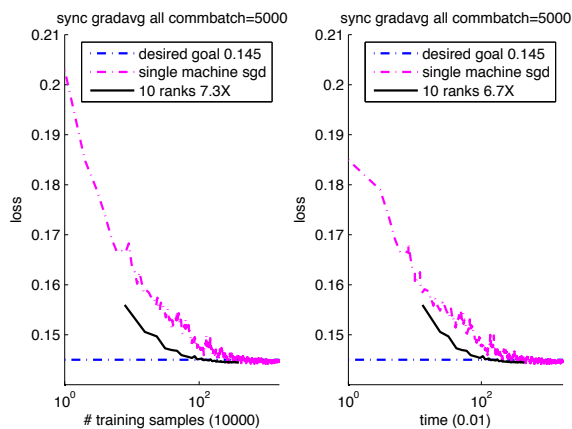


Figure 4. This figure shows convergence for RCV1 workload for MALT_{all-all} with a single machine workload. We find that MALT_{all-all} converges quicker to achieve the desired accuracy.

In this section, we compare the speedup of different datasets over a single machine and existing methods. We also evaluate the time spent across different processing tasks and the benefit of different synchronization methods.

We compare speedup of the systems under test by running them until they reach the same loss value and compare the time and number of samples processed. Figure 4 compares the speedup of MALT_{all-all} with a single machine for RCV1 dataset [10], for a communication batch size or *cb* size of 1000-5000. By *cb* size of 1000-5000, we mean that every model processes 1000-5000 examples from the dataset and then propagates the model updates to all other machines. We find that MALT_{all-all} provides a speedup with good convergence. By 10 ranks, we mean 10 processes, that span our eight machine cluster. The low latency and lock-free update communication framework allows frequent communication and with peers that aids speedup. For RCV1 and other smaller workloads, we find that we do not saturate the network and CPU with a single replica, and run multiple replicas on a single machine.

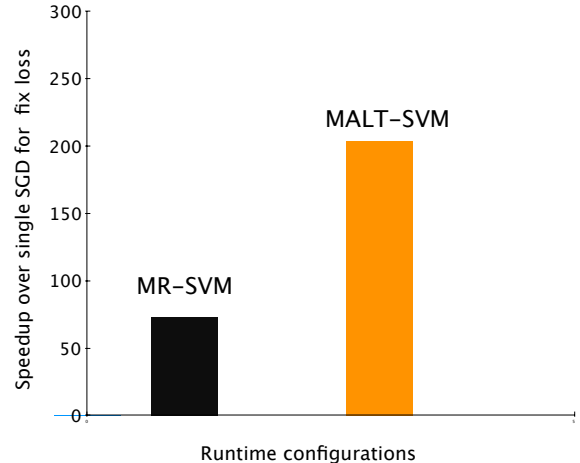


Figure 5. This figure shows convergence with PASCAL workload for MALT_{all-all} SVM with MR-SVM. We achieve super-linear speedup for some workloads because of the averaging effect from parallel replicas [51]. We find that MALT_{all-all} SVM is about 3X faster than map-reduce SVM implemented over infiniband.

We now compare MALT-SVM performance with existing algorithms over-designed for map-reduce SGD algorithm (MR-SVM) [16]. MR-SVM is based on Hadoop and hence communicates gradients after every epoch and needs more iterations to converge. We implement these algorithms over the MALT library and run it over our infiniband cluster. MR-SVM uses one-shot averaging at the end of every epoch to communicate parameters. Figure 5 shows the time to converge for one- workload for MR-SVM (implemented over MALT) and MALT_{all-all}. We find that both workloads achieve super-linear speedup over a single machine SGD using the PASCAL dataset. The underlying library implements averaged SGD that provides this speedup. With the PASCAL alpha dataset, we observe super-linear speedup. This happens because the averaging effect of gradients provides super-linear speedup for certain datasets [51]. In addition, we find that MALT provides 3X speedup over MR-SVM. MALT converges faster since it is designed over low latency communication, and sends gradients more frequently. MR-SVM is based on Hadoop and hence communicates gradients after every epoch and needs more iterations to converge. This result shows that existing algorithms designed for map-reduce may not provide the most optimal speedup for low latency frameworks such as MALT.

Figure 6 shows the time to process examples with neural networks. The neural network is a three-layer fully connected neural network, and is trained for text learning (click prediction). Neural networks need speedup w.r.t time for convergence for ad-click prediction implemented using a fully connected, three layer neural network. This three layer network needs to synchronize parameters at each layer. Furthermore, due to computation in forward and reverse direction, these networks have dense parameters and there is

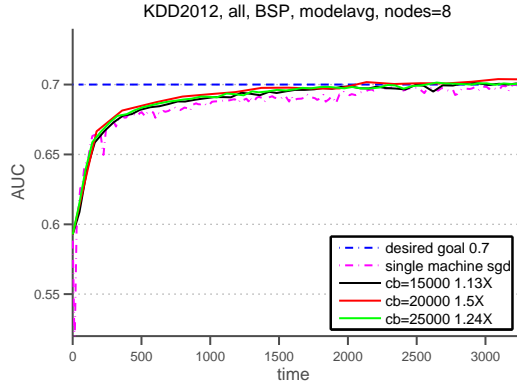


Figure 6. This figure shows the AUC (area under curve) vs time taken to speedup (in seconds) for a three layer neural network for text learning (click prediction). We find that provides speedup over a single node.

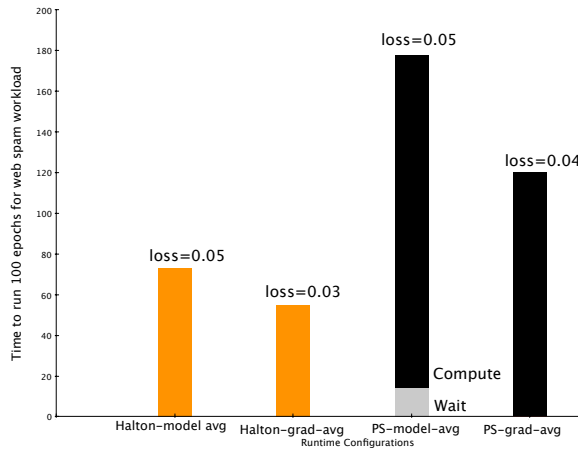


Figure 7. This figure shows the time taken to speedup matrix factorization network compares MALT_{halton} with parameter server (PS) for the Netflix dataset. We train asynchronously over two machines distributed SVM for webspam workload for asynchronous training, and use function with achieved loss values for 20 ranks.

computation in the forward and the reverse direction. Hence, fully-connected neural networks are hard to scale harder to scale than convolution networks [22]. We show the number of examples processed by every node in speedup by using MALT_{all-all} to train over KDD-2012 data on 10 processes 8 processes over single machine. We obtain 1.7 up to 1.5x speedup with 10 processes 8 ranks. The speedup is limited as compared to SVM because 1) SSI is non-convex and requires high-dimensional model communication as opposed to gradient and 2) text processing in a neural network requires very little limited computation and communication costs dominate.

Figure 12 shows the time to process examples on each machine with matrix factorization. This algorithm computes the examples asynchronously across different machines. Each example represents different rows and columns of the original matrix to be factorized. We use as the function that

emulates Hogwild [46]. We show speedup in processing examples for different communication batch sizes.

We now evaluate the time spent by MALT in different steps such as computing the gradient, pushing the gradient values, updating the gradient from received values, etc. Figure 8 shows the time spent for different ML tasks for 20 ranks (when operating synchronously). We find that MALT balances computation with communication and nodes spend most of their time computing gradients and pushing them (as opposed to blocking). In async workloads, MALT configurations do not wait while parameter server clients need to wait for updated models to arrive after sending their gradients as shown in Figure 7. Figure 7 also shows the efficiency of using gradient updates instead of sending whole models over the network. In case of parameter server, even though clients The slaves in the parameter server may send gradients to the server, the server needs to send whole models, and hence but need to receive full model parameters. Furthermore, the parameter server has wait times for model-averaging even in the asynchronous case because the workers need to wait for model parameters to arrive from the server before proceeding to the next iteration. Hence, it suffers from higher network costs for high-dimensional models such as webspam.

Figure 9 shows the speedup with synchronous (BSP), asynchronous (ASP) and bounded-staleness (SSP) models for the splice dataset. This training dataset consists of 10M examples (250 GB) and does not fit in single machine. Our implementation of bounded staleness only merges updates if they are within a specific range of iteration counts of its neighbor. Furthermore, if a specific rank lags too far behind, other ranks stall training for it to catch up. For the splice dataset, we find that SSP converges to the desired value first, followed by ASP and BSP.

To summarize, we find that MALT_{all-all} and MALT_{halton} achieves reasonable speedup over single machine and MR-SVM a single machine despite the additional costs of distributed computing. Furthermore, MALT API is general-purpose and can implement a wide variety of algorithms and applications can process large datasets and models with large parameter values in a reasonable amount of time.

6.2 Network Optimizations

We now evaluate the benefit of our network optimizations. Figure 10 shows the model convergence trends graph for MALT_{all-all} and MALT_{halton} to reach the required loss value for the RCV1 dataset. We find that MALT_{halton} converges more slowly than MALT_{all}, in terms of convergence per iteration. However, the time taken for every iteration overall time to converge is less because: First, parallel nodes (N) only communicate with fewer ($\log(N)$) machines. Second, each node performs model averaging of fewer ($\log(N)$) incoming models. Hence, even though MALT_{halton} may require more iterations than MALT_{all}, the overall time required for every iteration is less, and overall convergence time to

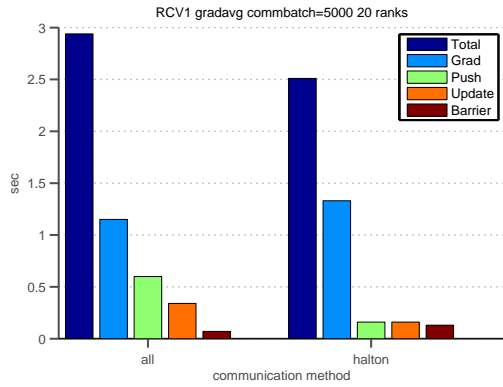


Figure 8. This figure shows the time consumed by different steps in distributed SVM for RCV1 workload for synchronous training.

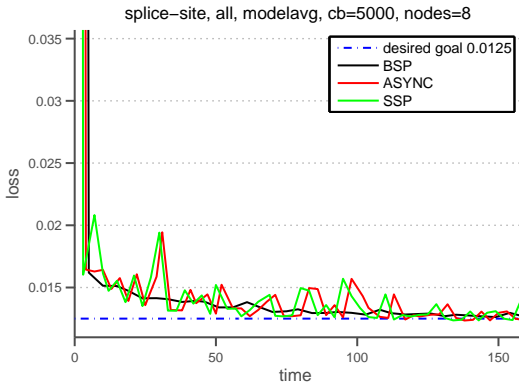


Figure 9. This figure compares with parameter server shows the convergence for bulk-synchronous (PSBSP) for distributed SVM for webspam workload for asynchronous training, with achieved loss values for 20 ranks. PS has wait times for model-averaging even in asynchronous mode because the workers need to wait processing (ASP) and stale synchronous processing (SSP) for parameters to arrive from the server before proceeding to the next iterations splice workload.

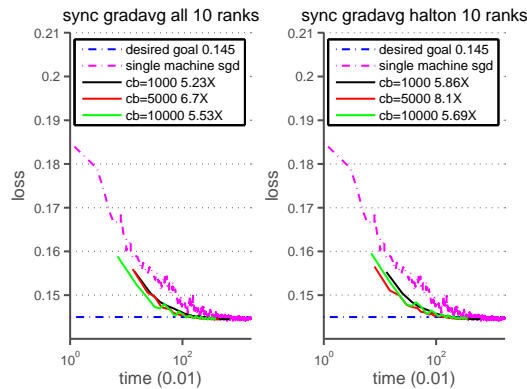


Figure 10. This figure shows convergence (loss vs time in seconds) for RCV1 dataset for MALT-all (left) and MALT-halton(right) for different communication batch sizes. We find that MALT_{halton} converges faster than MALT_{all}.

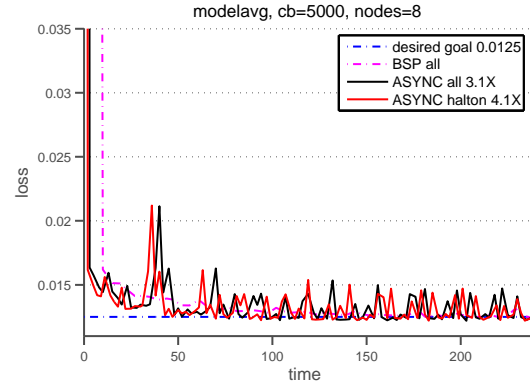


Figure 11. This figure shows convergence (loss vs time in seconds) for splice dataset for MALT-all and MALT-halton. We find that MALT_{halton} converges faster than MALT_{all}.

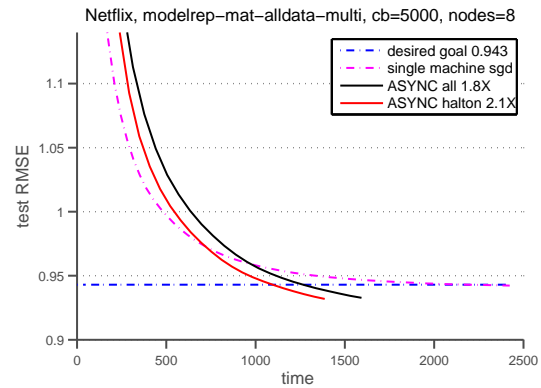


Figure 12. This figure shows the test error vs time in seconds processed for matrix factorization with the Netflix dataset. We train asynchronously over eight machines, and use grad-replace gather function.

reach the desired accuracy is less. Finally, since MALT_{halton} spreads out its updates across all nodes, that aids faster convergence.

Figure 11 shows model convergence for the splice dataset. From the figure, we see that MALT_{halton} converges faster than MALT_{all}-all. Furthermore, we find that that over the course of 32 epochs, with 8 nodes, each node in MALT_{all}-all sends out 305 GB of updates, while MALT_{halton} only sends 131 GB of data. As the number of nodes increase, the logarithmic fan-out of MALT_{halton} should result in lower amounts of data transferred and faster convergence.

Figure 12 shows the examples processed on each machine with matrix factorization. This algorithm computes the examples asynchronously across different machines. Each example represents different rows and columns of the original matrix to be factorized. We use `replace` as the `gather()` function that emulates Hogwild [46]. We show speedup in processing examples for different communication batch sizes. We find that MALT provides speedup over single machine, and using MALT_{halton} provides additional 14% speedup.

MALT_{halton} trades-off freshness of updates at peer replicas with savings in network communication time. For models where the cost of a single iteration is not significant (small model sizes, low averaging costs) model is dense and network communication costs are small compared to update costs, MALT_{all-all} configuration may be used with higher number of nodes than ($\log(N)$). Since, provide similar or better results over MALT is a general purpose API, it allows developers to balance communication and computation by providing a DAG for pushing updates_{halton}. For example, for the SSI workload, which is fully connected neural network, we only see a 1.1X speedup for MALT_{halton} over MALT_{all-all}. However, as the number of nodes and model sizes increase, the cost of communication begins to dominate, and using MALT_{halton} is beneficial.

Figure 13 shows the network time spent for MALT_{all-all}, MALT_{halton}, and parameter server for the entire network, for webspam workload. We find that MALT_{halton} is most network efficient. Webspam is a high-dimensional workload. MALT_{halton} only sends updates to $\log(N)$ nodes. Parameter server sends gradients and but needs to receive the whole model from the central server. We note that other optimizations such as compression, and other filters can further reduce the network costs as noted in [36]. Furthermore, when the parameter server is replicated for high-availability, there is more network traffic for additional N (asynchronous) messages for $N - way$ chain replication of the parameters.

OverallTo summarize, we find that MALT provides sending gradients (instead of sending the model) that saves network costs. Furthermore, MALT_{halton} is network efficient and achieves acceptable convergence speedup over MALT_{all-all}.

Network saturation tests: We perform infiniband network throughput tests, and measure the time to push updates in ALL case with the SVM workload. In the synchronous case, we find that all ranks operate in a log step fashion, and during the scatter phase, all machines send models at line rate (about 5GB/s). Specifically, for the webspam workload (Figure 13), we see about 5.1 GB/s (40 Gb/s) during scatter. In the asynchronous case, to saturate the network, we needed to run multiple replicas on every machine. When we run three ranks on every machine, we find that each machine sends model updates at 4.2 GB/s (33 Gb/s) for the Webspam dataset. These tests demonstrate that using a large bandwidth network is beneficial for training models with large number of parameters.

6.3 Developer Effort

We evaluate the ease of implementing parallel learning in MALT by adding support to the four applications listed in Table 3. For each application we show the amount of code we copied from suspend/resume to create checkpoint/restore modified as well as the number of new lines added. On

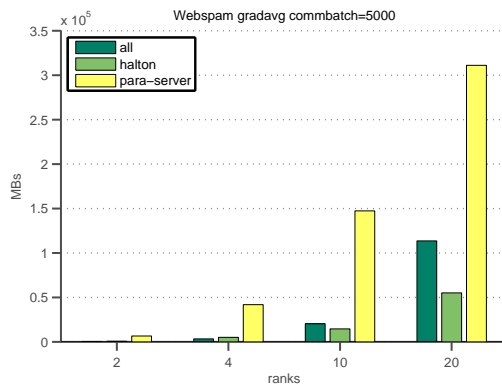


Figure 13. This figure shows the network costs for MALT_{all-all}, MALT_{halton} and parameter server for the whole network for webspam workload. We find that MALT_{halton} reduces network communication costs and provides fast convergence.

Application	Dataset	MALT annotations	
		LOC Modified	LOC Added
SVM	RCV1	105	107
Matrix Factorization	Netflix	76	82
SSI	KDD2012	82	130

Table 3. Developer effort for converting serial applications to data-parallel with MALT.

average, we moved 87 lines of code and added 106 lines, representing about 15% of overall code. In Section 4 we described the specific changes required. The new code adds support for creating MALT objects, scattering the updates, gathering and averaging the received updates. In comparison, implementing a whole new algorithm takes hundreds of lines new code assuming underlying data parsing and arithmetic libraries are provided by the processing framework. On average, we moved 87 lines of code and added 106 lines, representing about 15% of overall code.

6.4 Fault Tolerance

We evaluate the time required for convergence when a node fails. When the MALT fault monitor in a specific node receives a time-out from a failed node, it removes that node from send/receive lists. We run MALT-SVM on a 10 machine cluster to train over PASCAL-DNA [4] dataset. We inject faults on MALT jobs on one of the machines and observe recovery and subsequent convergence. We inject faults through an external script and also inject programmer errors such as divide by zero.

We find that in each case, MALT fault monitors detected the unreachable failed mode, triggered a recovery process to synchronize with the remaining nodes and continued to train. We also observe that subsequent group operations only execute on the survived node. Finally, we verify that the models converge to an acceptable accuracy in each of the failed cases. We also find that local fault monitors were able to trap processor exceptions and terminate the local training replica. We note that MALT cannot detect corruption of

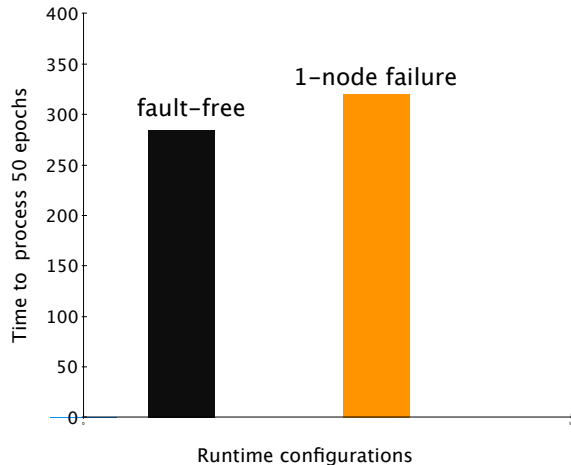


Figure 14. This figure shows the time taken to converge for DNA dataset with 10 nodes in fault-free and a single process failure case. We find that MALT is able to recover from the failure and train the model with desired accuracy.

scalar values or Byzantine failures. Figure 14 shows one instance of failure recovery, and the time to converge is proportional to the number of remaining nodes in the cluster.

7. Related work

Our work is inspired from past work on data-parallel processing, ML specific platforms, ML optimizations and RDMA based key-value stores.

Existing data-parallel frameworks: Batch processing systems based on map-reduce [23, 50] perform poorly for machine learning jobs because the iterative nature of these algorithms require frequent communication using disks. Furthermore, the data-flow model provided by map-reduce is restrictive, and limits the flexibility of expressing communication across tasks. Spark [53] provides an efficient in-memory representation to synchronize data across iterations. Spark provides copy-on-write, in-memory structures that improve performance for batch workloads that transform all-data. It also provides fault tolerance using lineage (re-computation) that enforces determinism. However, this is less efficient for [certain](#) machine learning algorithms that make fine-grained and asynchronous updates. Dryad[32] and CIEL [42] provide a more flexible data-parallel communication API to write any arbitrary data flow but share data across tasks through disks. MPI [25] provides a low-level message passing constructs, and a system similar to MALT can be built over MPI. However, MPI does not provide any fault tolerance [capabilities](#). Picollo [45] provides a distributed, master-slave key-value store, and resolves writes conflicts [with-using](#) user-defined functions. Picollo provides strong fault tolerance, consistency and determinism guarantees that can be relaxed in MALT for performance. MALT is completely asynchronous, and allows senders to transmit data with one-sided write operations. Furthermore, unlike Picollo, MALT provides a more expressive communication

API, [allowing that allows](#) the programmer to control where the updates may reside and how they may propagate.

ML frameworks: Parameter-server [36], Project Adam [15], DistBelief [22] use a master-client style communication with the parameter server, that complements MALT. MALT is a general purpose API, designed to train models in a peer-peer fashion. MALT’s peer-peer style simplifies fault-tolerance, and only requires writing code once that executes over all machines. Furthermore, most existing parallel learning frameworks require a re-write of applications and libraries. Vowpal Wabbit [52] provides data-parallel learning, where individual model replicas train in parallel and average the gradients using the LBFGS algorithm. However, it runs over Hadoop and lacks efficient shared memory semantics that MALT provides. Presto [49], has a similar goal to ours, to provide a rich developer environment for parallel learning. It provides a parallel R, since R is a common data analysis tool. Presto provides a large distributed array abstraction to shard data and removes scalability bottlenecks in R implementation. MALT library can be used to parallelize many existing learning frameworks. There are also many GPU based frameworks [18]. However, GPU speedups are limited for datasets exceeding its memory sizes (<10 GB) and training multiple-GPUs over the network incurs significant communication costs. Furthermore, programming GPUs is hard, and requires developers to worry about hiding the GPU-CPU latency apart from ML algorithm specific issues.

ML optimizations: We now discuss prior work to optimize distributed machine learning to reduce synchronization and improve convergence. Many researchers have explored improving stochastic gradient descent over distributed systems by providing *mini-batching* [19]. This reduces the amount of communication overhead since each model processes more examples before sending out updates. However, both ASGD and mini-batching show poor convergence rates and are not very computationally efficient [29]. HogWild [46] provides a single shared parameter vector and allows parallel workers to update model parameters without any locking (in a single machine). However, this method only works when updates are sparse and there is limited overlap. HogWild can also generate lots of network traffic when adopted in a distributed setting. Bounded-staleness [20] limits stale updates from stragglers by slowing down the forerunners. Iterative parameter mixing (used by MALT) has been shown to provide high-accuracy models over map-reduce [29]. Optimistic Concurrency Control [44] uses database style coordination free model updates for distributed machine learning. MALT also provides coordination free updates by allocating a per-receiver queue at sender and also avoids invoking the remote CPU by using one-sided RDMA.

RDMA systems: MALT uses one-sided RDMA writes to communicate gradients for every batch (and no reads). Recently, infiniband hardware has been used to build [transactions](#)

on objects in shared address space in FARM [24] and client-server based key-value stores with Pilaf [40] and distributed symmetric shared-memory systems in FARM [24]. These systems provide efficient PUT/GET semantics. Pilaf provides consistency using checksums, while FARM orders provides consistency by ordering DMA writes. MALT provides a finite per-sender queue at the receiver to avoid write-write conflicts. Older gradients may be over-written by a fast sender and the receiver averages (or any other user defined function) the incoming model updates for conflict resolution.

8. Conclusion

Existing map-reduce frameworks are optimized for batch processing systems and ill-suited for tasks that are iterative, fine-grained and asynchronous. Recent scalable ML platforms force developers to learn a new programming environment and rewrite their ML software. The goal of MALT is to efficiently provide data-parallelism to existing ML software. Given a list of machines and MALT library, we demonstrate that one can program ML algorithms, control the data-flow and synchrony. We provide MALT library interface for procedural (C++) and scripting (Lua) languages and demonstrate data-parallel benefits with SVM, matrix factorization and neural networks. MALT uses one-sided RDMA primitives that reduces network processing costs and transmission overhead. The new generation of RDMA protocols provide additional opportunities for optimizations. Primitives such as `fetch_and_add` can be used to perform gradient averaging in hardware and further decrease the model training costs in software.

Acknowledgments

We would like to thank our shepherd Derek Murray and the anonymous reviewers for the useful feedback. We also thank Igor Durdanovic for helping us port MALT to RAPID and Hans-Peter Graf for his support and encouragement.

References

- [1] Halton sequence. en.wikipedia.org/wiki/Halton_sequence.
- [2] Machine Learning in Python. <http://scikit-learn.org/>.
- [3] The R Project for Statistical Computing. www.r-project.org/.
- [4] PASCAL(Pattern Analysis, Statistical Modelling and Computational Learning) Large Scale Learning Challenge. <http://largescale.ml.tu-berlin.de/about/>, 2009.
- [5] A. Agarwal, O. Chapelle, M. Dudík, and J. Langford. A reliable effective terascale linear learning system. *Journal of Machine Learning Research*, 2014.
- [6] T. Alrutz, J. Backhaus, T. Brandes, V. End, T. Gerhold, A. Geiger, D. Grünewald, V. Heuveline, J. Jägersküpper, A. Knüpfer, et al. Gaspi—a partitioned global address space programming interface. In *Facing the Multicore-Challenge III*, pages 135–136. Springer, 2013.
- [7] K. Bache and M. Lichman. UCI machine learning repository, 2013.
- [8] B. Bai, J. Weston, D. Grangier, R. Collobert, K. Sadamasa, Y. Qi, O. Chapelle, and K. Weinberger. Supervised semantic indexing. In *Proceedings of the 18th ACM conference on Information and knowledge management*. ACM, 2009.
- [9] C. M. Bishop et al. *Pattern recognition and machine learning*, volume 1. springer New York, 2006.
- [10] L. Bottou. Large-scale machine learning with stochastic gradient descent. In Y. Lechevallier and G. Saporta, editors, *Proceedings of the 19th International Conference on Computational Statistics (COMPSTAT'2010)*, August 2010. Springer.
- [11] L. Bottou. Stochastic gradient descent tricks. In *Neural Networks: Tricks of the Trade*, pages 421–436. Springer, 2012.
- [12] J. Canny and H. Zhao. Butterfly mixing: Accelerating incremental-update algorithms on clusters. In *SDM'13*, 2013.
- [13] M. Castro and B. Liskov. Proactive recovery in a byzantine-fault-tolerant system. In *Proceedings of the 4th conference on Symposium on Operating System Design & Implementation-Volume 4*. USENIX Association, 2000.
- [14] T. D. Chandra, R. Griesemer, and J. Redstone. Paxos made live: an engineering perspective. In *Proceedings of the twenty-sixth annual ACM symposium on Principles of distributed computing*. ACM, 2007.
- [15] T. Chilimbi, Y. Suzue, J. Apacible, and K. Kalyanaraman. Project adam: Building an efficient and scalable deep learning training system. In *OSDI*, 2014.
- [16] C. Chu, S. K. Kim, Y.-A. Lin, Y. Yu, G. Bradski, A. Y. Ng, and K. Olukotun. Map-reduce for machine learning on multicore. *Advances in neural information processing systems*, 19:281, 2007.
- [17] J. Cipar, Q. Ho, J. K. Kim, S. Lee, G. R. Ganger, G. Gibson, K. Keeton, and E. Xing. Solving the straggler problem with bounded staleness. *HotOS. USENIX Association*, 2013.
- [18] A. Coates, B. Huval, T. Wang, D. Wu, B. Catanzaro, and N. Andrew. Deep learning with cots hpc systems. In *Proceedings of The 30th International Conference on Machine Learning*, 2013.
- [19] A. Cotter, O. Shamir, N. Srebro, and K. Sridharan. Better mini-batch algorithms via accelerated gradient methods. In *Advances in Neural Information Processing Systems*, 2011.
- [20] H. Cui, J. Cipar, Q. Ho, J. K. Kim, S. Lee, A. Kumar, J. Wei, W. Dai, G. R. Ganger, P. B. Gibbons, et al. Exploiting bounded staleness to speed up big data analytics. 2014.
- [21] W. Dai, J. Wei, X. Zheng, J. K. Kim, S. Lee, J. Yin, Q. Ho, and E. P. Xing. Petuum: A framework for iterative-convergent distributed ml. *arXiv preprint arXiv:1312.7651*, 2013.
- [22] J. Dean, G. Corrado, R. Monga, K. Chen, M. Devin, Q. V. Le, M. Z. Mao, M. Ranzato, A. W. Senior, P. A. Tucker, et al. Large scale distributed deep networks. In *NIPS*, 2012.
- [23] J. Dean and S. Ghemawat. Mapreduce: simplified data processing on large clusters. *Communications of the ACM*, 51(1):107–113, 2008.

- [24] A. Dragojević, D. Narayanan, O. Hodson, and M. Castro. Farm: fast remote memory. In *Proceedings of the 11th USENIX Symposium on Networked Systems Design and Implementation (NSDI 14)*, 2014.
- [25] E. Gabriel, G. E. Fagg, G. Bosilca, T. Angskun, J. J. Dongarra, J. M. Squyres, V. Sahay, P. Kambadur, B. Barrett, A. Lumsdaine, et al. Open mpi: Goals, concept, and design of a next generation mpi implementation. In *Recent Advances in Parallel Virtual Machine and Message Passing Interface*, pages 97–104. Springer, 2004.
- [26] R. Gemulla, E. Nijkamp, P. J. Haas, and Y. Sismanis. Large-scale matrix factorization with distributed stochastic gradient descent. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 2011.
- [27] [GPI2: Programming Next Generation Supercomputers Benchmarks](#).
- [28] A. Halevy, P. Norvig, and F. Pereira. The unreasonable effectiveness of data. *Intelligent Systems, IEEE*, 24(2):8–12, 2009.
- [29] K. B. Hall, S. Gilpin, and G. Mann. Mapreduce/bigtable for distributed optimization. In *NIPS LCCC Workshop*, 2010.
- [30] M. Hilbert and P. López. The worlds technological capacity to store, communicate, and compute information. *Science*, 332(6025):60–65, 2011.
- [31] IDIAP Research Institute, New York University and NEC Laboratories America. Torch7: Scientific computing for LuaJIT. <http://torch.ch/>.
- [32] M. Isard, M. Budi, Y. Yu, A. Birrell, and D. Fetterly. Dryad: distributed data-parallel programs from sequential building blocks. In *Proceedings of the 2nd ACM european conference on Computer Systems*, 2007.
- [33] Kaggle. Tencent 2012 KDD Cup. <https://www.kddcup2012.org>, 2012.
- [34] [A. Kalia, M. Kaminsky, and D. G. Andersen. Using RDMA efficiently for key-value services. In Proceedings of the 2014 ACM conference on SIGCOMM. ACM, 2014.](#)
- [35] Y. Koren, R. Bell, and C. Volinsky. Matrix factorization techniques for recommender systems. *Computer*, 42(8):30–37, 2009.
- [36] M. Li, D. Andersen, A. Smola, J. Park, A. Ahmed, V. Josifovski, J. Long, E. Shekita, and B.-Y. Su. Scaling distributed machine learning with the parameter server. In *OSDI*, 2014.
- [37] M. Li, D. G. Andersen, and A. Smola. Distributed delayed proximal gradient methods. In *NIPS Workshop on Optimization for Machine Learning*, 2013.
- [38] M. Li, T. Zhang, Y. Chen, and A. J. Smola. Efficient mini-batch training for stochastic optimization. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '14. ACM, 2014.
- [39] R. McDonald, K. Hall, and G. Mann. Distributed training strategies for the structured perceptron. In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*. Association for Computational Linguistics, 2010.
- [40] C. Mitchell, Y. Geng, and J. Li. Using one-sided rdma reads to build a fast, cpu-efficient key-value store. In *Proceedings of the 2013 USENIX Annual Technical Conference, USENIX ATC*, volume 13, 2013.
- [41] D. G. Murray, F. McSherry, R. Isaacs, M. Isard, P. Barham, and M. Abadi. Naiad: a timely dataflow system. In *Proceedings of the Twenty-Fourth ACM Symposium on Operating Systems Principles*. ACM, 2013.
- [42] D. G. Murray, M. Schwarzkopf, C. Smowton, S. Smith, A. Madhavapeddy, and S. Hand. Ciel: A universal execution engine for distributed data-flow computing. In *NSDI*, volume 11, 2011.
- [43] NEC Laboratories America. MiLDE: Machine Learning Development Environment. <http://www.nec-labs.com/research-departments/machine-learning/machine-learning-software/Milde>.
- [44] X. Pan, J. E. Gonzalez, S. Jegelka, T. Broderick, and M. Jordan. Optimistic concurrency control for distributed unsupervised learning. In *Advances in Neural Information Processing Systems*, 2013.
- [45] R. Power and J. Li. Piccolo: Building fast, distributed programs with partitioned tables. In *OSDI*, 2010.
- [46] B. Recht, C. Re, S. Wright, and F. Niu. Hogwild: A lock-free approach to parallelizing stochastic gradient descent. In *Advances in Neural Information Processing Systems*, 2011.
- [47] J. R. Santos, Y. Turner, and G. Janakiraman. End-to-end congestion control for infiniband. In *INFOCOM 2003. Twenty-Second Annual Joint Conference of the IEEE Computer and Communications. IEEE Societies*, volume 2. IEEE, 2003.
- [48] L. G. Valiant. A bridging model for parallel computation. *Communications of the ACM*, 33(8):103–111, 1990.
- [49] S. Venkataraman, E. Bodzsar, I. Roy, A. AuYoung, and R. S. Schreiber. Presto: distributed machine learning and graph processing with sparse matrices. In *Proceedings of the 8th ACM European Conference on Computer Systems*. ACM, 2013.
- [50] T. White. *Hadoop: The definitive guide*. ” O’Reilly Media, Inc.”, 2009.
- [51] W. Xu. Towards optimal one pass large scale learning with averaged stochastic gradient descent. *arXiv preprint arXiv:1107.2490*, 2011.
- [52] Yahoo Corp. Vowpal Wabbit. <http://hunch.net/vw/>.
- [53] M. Zaharia, M. Chowdhury, T. Das, A. Dave, J. Ma, M. McCauley, M. J. Franklin, S. Shenker, and I. Stoica. Resilient distributed datasets: A fault-tolerant abstraction for in-memory cluster computing. In *Proceedings of the 9th USENIX conference on Networked Systems Design and Implementation*. USENIX Association, 2012.
- ~~M. Zaharia, M. Chowdhury, M. J. Franklin, S. Shenker, and I. Stoica. Spark: cluster computing with working sets. In Proceedings of the 2nd USENIX conference on Hot topics in cloud computing, 2010.~~