
RESEARCH	Large Scale Machine Learning in Shared Computing Environments	
EDUCATION	Carnegie Mellon University , Pittsburgh, PA.	
	Ph.D., Electrical and Computer Engineering, May, 2019	
	Master of Science., Electrical and Computer Engineering, May 2016	
	<ul style="list-style-type: none"> • Thesis: Improving ML Applications in Shared Computing Environments • Advisors: Greg Ganger, Phil Gibbons 	
	Northeastern University , Boston, Massachusetts	
	B.S., Electrical and Computer Engineering, May 2014	
	<ul style="list-style-type: none"> • Summa Cum Laude 	
RESEARCH EXPERIENCE	Research Assistant Parallel Data Lab, Carnegie Mellon University	Aug 2014 to May 2019

PipeDream: Pipeline Parallelism for DNN Training. PipeDream is a new distributed training system for deep neural networks (DNNs) that partitions ranges of DNN layers among machines, and aggressively pipelines computation and communication. Today’s pervasive use of data-parallel training performs well for DNNs of up to 10–20 million model parameters, but inter-machine communication dominates for models that are even 10x larger (e.g., up to 85% of time training the VGG16 model is spent on communication) – it seems likely that models will only get larger in the future. PipeDream’s pipelined approach reduces communication by over 95% for the same frequency of synchronization and allows complete overlapping of communication and computation. PipeDream’s design efficiently handles the systematic splitting of work into pipeline stages, model versioning, coordination of the forward and backward passes, and the other consistency challenges associated with pipelined DNN training. As a result, provides a 3x or more improvement in “time to target accuracy” compared to efficient data-parallel training for large models like VGG16, without reducing the performance of training smaller models like Inception-BN.

Work published at SysML ’18.

Tributary: spot-dancing for elastic services with latency SLOs. Tributary, is an elastic control system that embraces the uncertain nature of transient cloud resources, such as AWS spot instances, to manage elastic services with latency SLOs more robustly and more cost-effectively. Such resources are available at lower cost, but with the proviso that they can be preempted en masse, making them risky to rely upon for business-critical services. Tributary creates models of preemption likelihood and exploits the partial independence among different resource offerings, selecting collections of resource allocations that satisfy SLO requirements and adjusting them over time, as client workloads change. Although Tributary’s collections are often larger than required in the absence of preemptions, they are cheaper because of both lower spot costs and partial refunds for preempted resources. At the same time, the often-larger sets allow unexpected workload bursts to be absorbed without SLO violation. Over a range of web service workloads, we find that Tributary reduces cost for achieving a given SLO by 81–86% compared to traditional scaling on non-preemptible resources, and by 47–62% compared to the high-risk approach of the same scaling with spot resources.

Work published at Usenix ATC ’18.

Proteus: agile ML elasticity through tiered reliability in dynamic resource markets. Many shared computing clusters allow users to utilize excess idle resources at lower cost or priority, with the proviso that some or all may be taken away at any time. But, exploiting such dynamic resource availability and the often fluctuating markets for them requires agile elasticity and effective acquisition strategies. Proteus aggressively exploits such transient revocable resources to do machine learning (ML) cheaper and/or faster. Its parameter server framework, TierML, efficiently adapts to bulk additions and revocations of transient machines, through a novel 3-stage active-backup approach, with minimal use of more costly non-transient resources. Its BidBrain component adaptively allocates resources from multiple EC2 spot markets to minimize average cost per work as transient resource availability and cost change over time. Our experiments show that Proteus reduces cost by 85% relative to non-transient pricing, and by 43% relative to previous approaches, while simultaneously reducing run-times by up to 37%.

Work published at EuroSys '17.

FlexRR: Addressing the straggler problem for iterative convergent parallel ML. FlexRR provides a scalable, efficient solution to the straggler problem for iterative machine learning (ML). The frequent (e.g., per iteration) barriers used in traditional BSP-based distributed ML implementations cause every transient slowdown of any worker thread to delay all others. FlexRR combines a more flexible synchronization model with dynamic peer-to-peer re-assignment of work among workers to address straggler threads. Our experiments with real straggler behavior observed on Amazon EC2 and Microsoft Azure, as well as injected straggler behavior stress tests, confirm the significance of the problem and the effectiveness of FlexRR's solution. Using FlexRR, we consistently observe near-ideal run-times (relative to no performance jitter) across all real and injected straggler behaviors tested.

Work published at SoCC '16.

INTERNSHIPS

Microsoft Research

May 2017 to August 2017

Research Intern: Worked on a novel machine learning training system (PipeDream) for deep neural networks. Responsibilities included developing research ideas and system implementation.

Spectral Sciences Incorporated

January 2013 to June 2013

Co-Op Software Engineer: Created graphical user interface for MODTRAN program in C, including plotting interface. Wrote binary test code for MODTRAN code in Fortran and created a wrapping script in PERL. Updated MODTRAN code from Fortran 77 to Fortran 90

Motorola Mobility

January 2012 to June 2012

Co-Op Software Engineer: Worked as part of the BSR sustaining team. Fixed customer related issues and developed new features including the Early Patching System and IPV6 Neighbor Discovery customer interface.

Charles River Development

January 2011 to August 2011

SQA CoOp Engineer: Worked as part of the Automaton / Infrastructure Group. Created components status check system which automated the process of checking the statuses of virtual machines, and posting them on-line. Rebuilt and centralized the process of starting regression testing, which decreased the time necessary to complete the process by over 75%. Also, responsible for bitmap testing and worked on several other projects including Log Backup, QA Wiki Page, and Result Reporting.

Parametric Technology Corporation

May 2010 to August 2010

QA Intern: Performed final installation checks before shipping, conducted export tests, wrote code for tools, created a customer-filed bug database, set up local area networks,

assembled plotter printers, and performed plot checks.

CONFERENCE
TALKS

Tributary: spot-dancing for elastic services with latency SLOs.
In *Usenix ATC 2018*, Boston, MA, July 2018

Proteus: agile ML elasticity through tiered reliability in dynamic resource markets.
In *EuroSys 2017*, Belgrade, Serbia, May 2017

Addressing the Straggler Problem for Iterative Convergent Parallel ML.
In *SoCC 2016*, Santa Clara, CA, Oct 2016

PUBLICATIONS

Aaron Harlap, Andrew Chung, Alexey Tumanov, Gregory Ganger and Phillip Gibbons
Tributary: spot-dancing for elastic services with latency SLOs.
In *USENIX Annual Technical Conference (ATC '18)*, 2018.

Aaron Harlap, Deepak Narayanan, Amar Phanishayee, Vivek Seshadri,
Gregory R. Ganger, Phillip B. Gibbons
PipeDream: Pipeline Parallelism for DNN Training.
In *SysML Conference (SysML '18)*, 2018.

Aaron Harlap, Alexey Tumanov, Andrew Chung, Gregory Ganger and Phillip Gibbons
Proteus: agile ML elasticity through tiered reliability in dynamic resource markets.
In *European Conference on Computer Systems (EuroSys '17)*, 2017.

Kevin Hsieh, **Aaron Harlap**, Nandita Vijaykumar, Dimitris Konomis,
Gregory R. Ganger, Phillip B. Gibbons, Onur Mutlu
Gaia: Geo-Distributed Machine Learning Approaching LAN Speeds
In *Symposium on Networked Systems Design and Implementation (NSDI' 17)*, 2017

Aaron Harlap, Henggang Cui, Wei Dai, Jinliang Wei, Gregory R. Ganger, Phillip B.
Gibbons, Garth A. Gibson, and Eric P. Xing.
Addressing the Straggler Problem for Iterative Convergent Parallel ML.
In *ACM Symposium on Cloud Computing (SoCC'16)*, 2016.

AWARDS

Student Awards — Northeastern University
Northeastern University Honors Program
Northeastern University Dean's List (x8)

Travel Awards
SoCC 2016, Santa Clara, CA, Oct 2016

COURSES

Carnegie Mellon University (GPA 3.89/4)
18847 Machine Learning Infrastructure, Fall 2017.

15719 Advanced Cloud Computing, Spring 2016.

15721 Advanced Database Systems, Spring 2016. Course project:
Lock free database index.

18749 Building Reliable Distributed Systems, Fall 2015. Course project:
Elastic machine learning on temporarily available/affordable resources

18601 Entrepreneurship Innovation Technology, Spring 2015

18746 Storage Systems, Spring 2015. Course project:
Cloudfs, a hybrid file system integrating solid-state devices and cloud storage

15712 Advanced Topics in Operating Systems, Fall 2014. Course project:
Straggler mitigation in parallel machine learning

10701 Machine Learning, Fall 2014. Course project:
Censorship Detection at Global Scale

TEACHING
EXPERIENCE

Teaching Assistant

18746/15746 Storage Systems

Fall 2016 and 2017

SKILLS

Computer Programming

C++, Python, AWS EC2, Shell

Software Systems

Git, Linux, LaTeX, Pytorch, Tensorflow, Caffe