

Intelligent HPC Cloud

Illustrated by Harp and Harp-DAAL at Indiana University

IEEE Cloud Computing Conference

June 26, 2017

Judy Qiu

Intelligent Systems Engineering Department, Indiana University

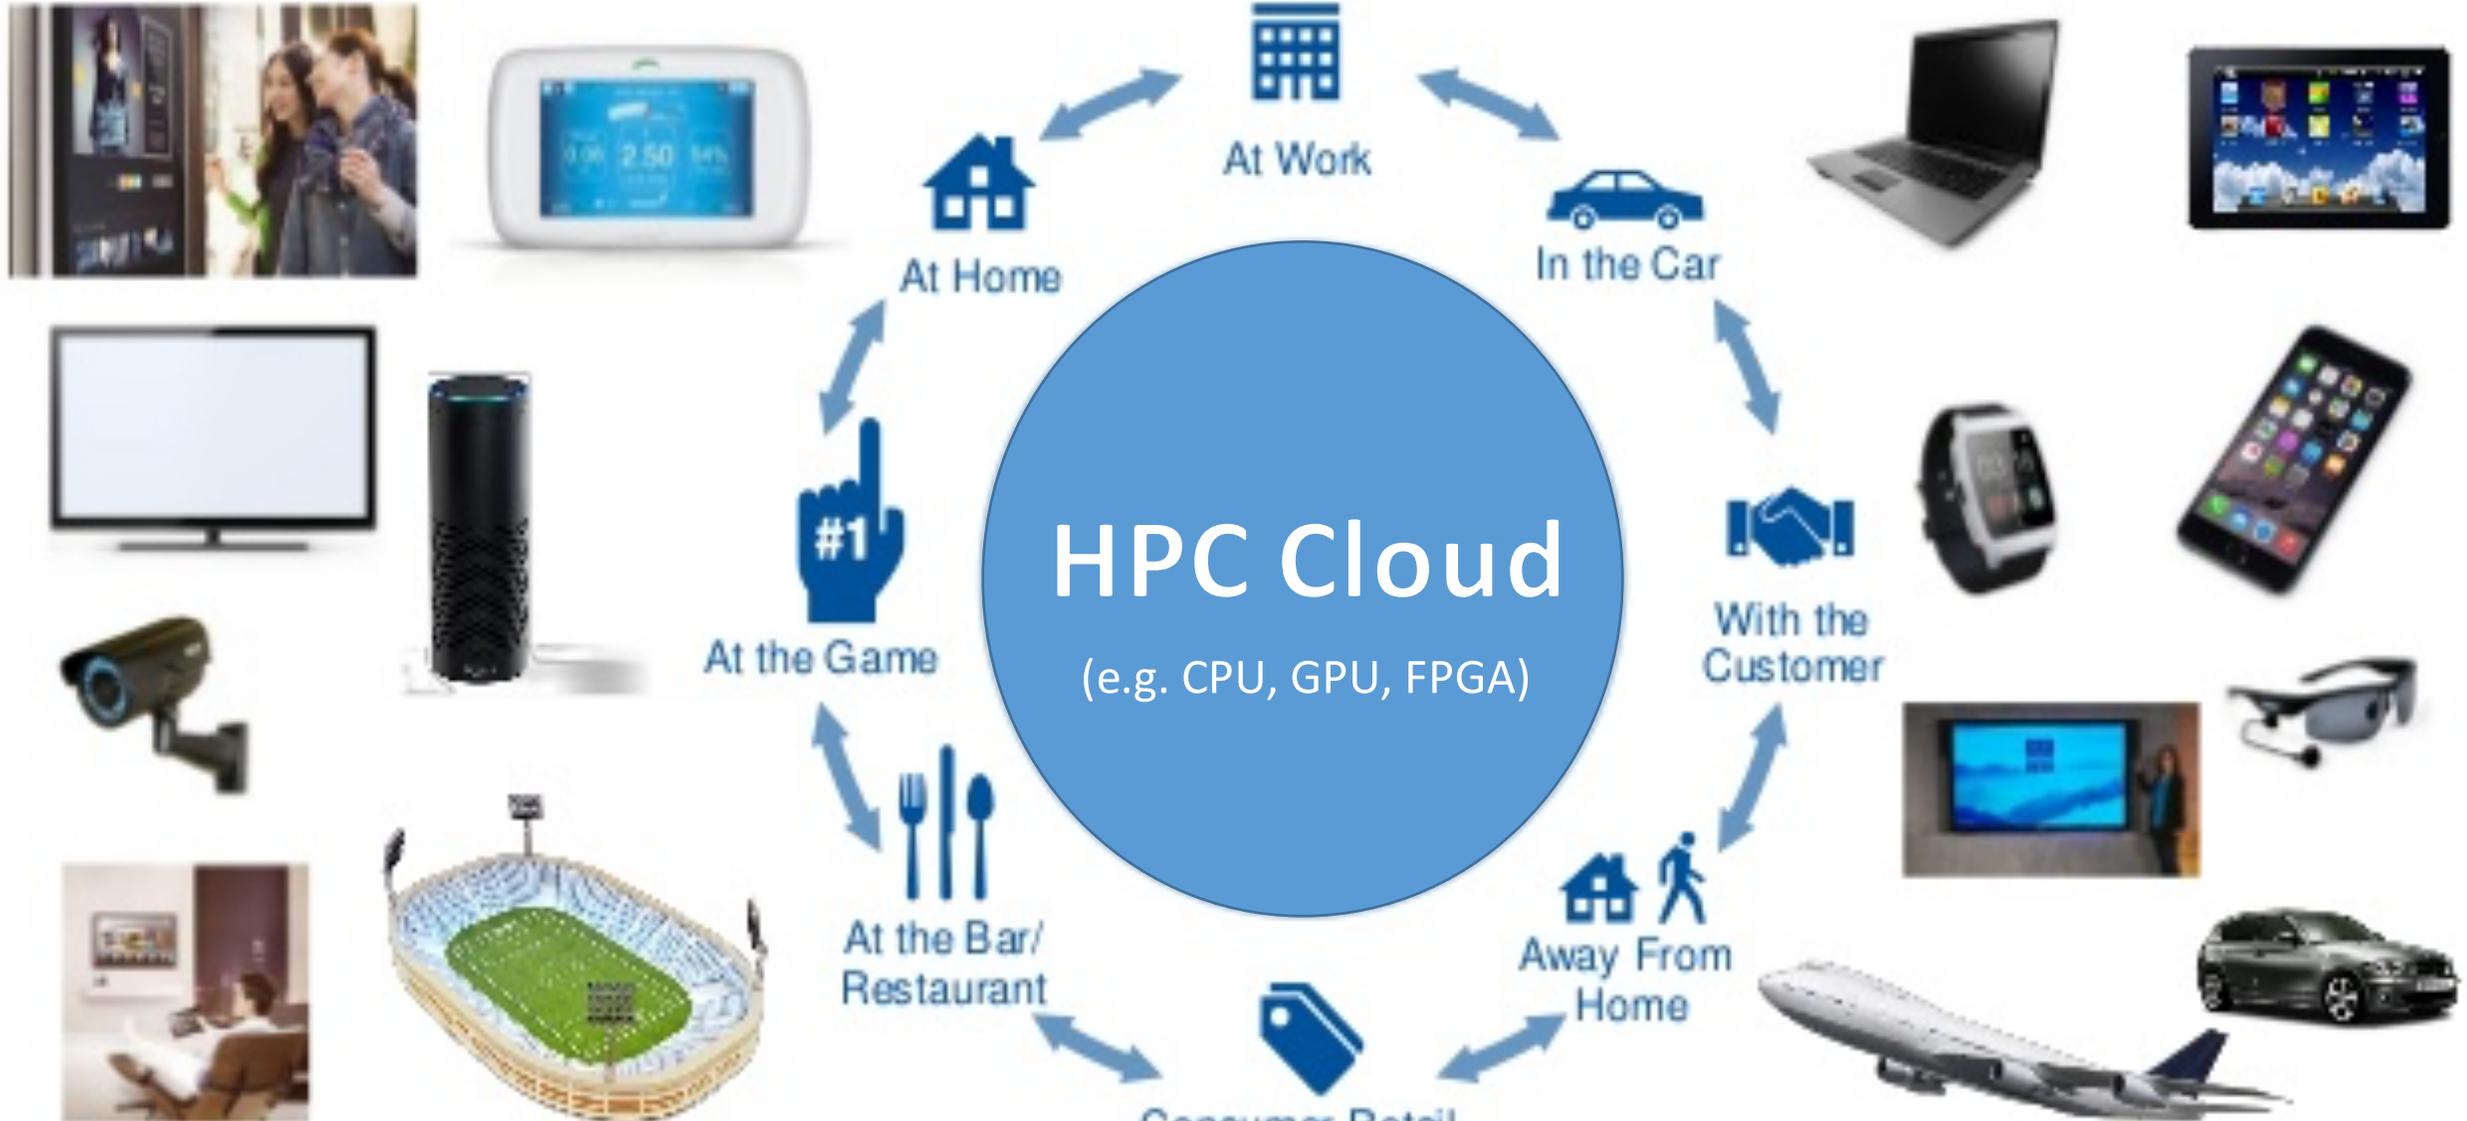
Email: xqiu@indiana.edu

Intelligent HPC Cloud

- 1. General Purpose Machine Intelligence requires both Cloud and HPC Technologies**
- 2. HPC-Apache Big Data Stack supports AI and IoT**
- 3. Illustration of Harp (Hadoop plug in) and Intel's High Performance Data Analytics**

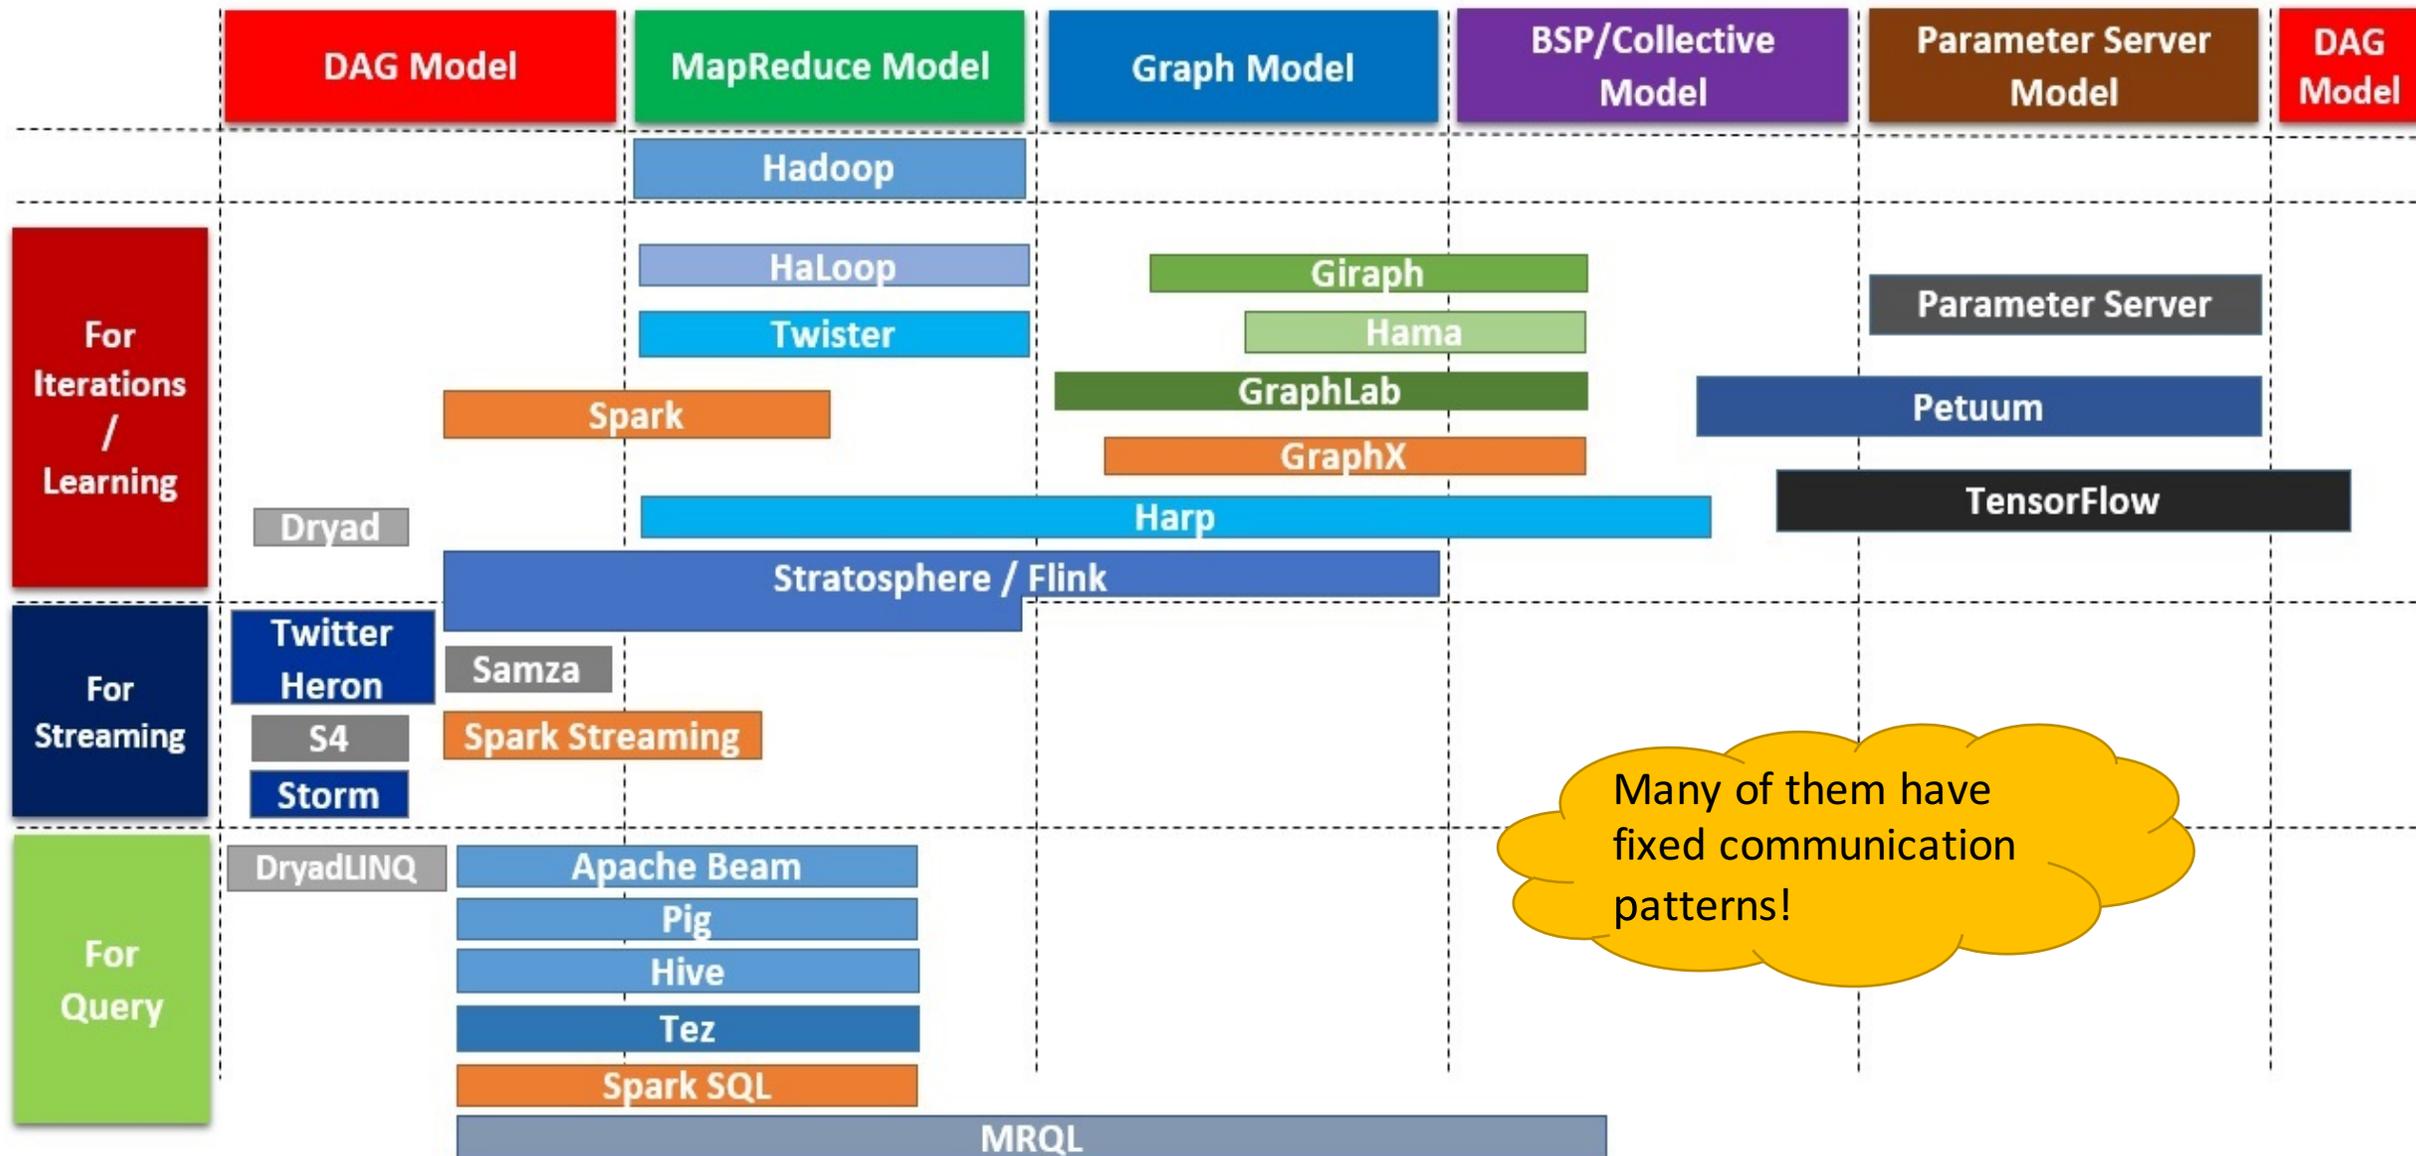
50 Billion Devices by 2020

World Popular will be 7.6 billion by 2020



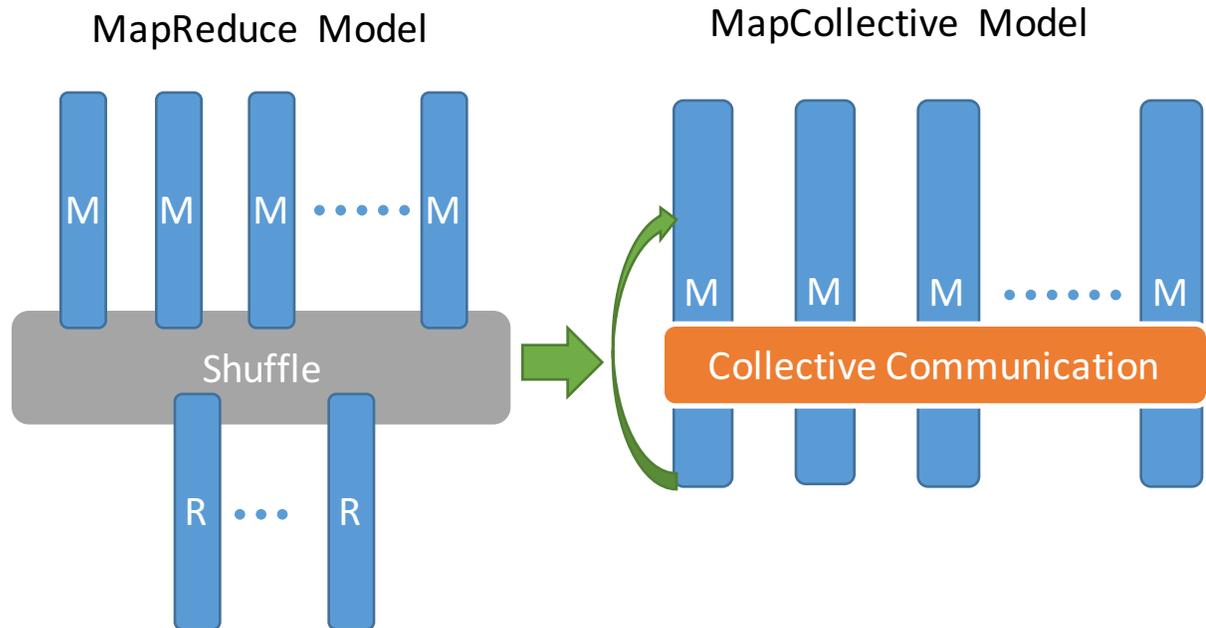
High Performance – Apache Big Data Stack

HPC-ABDS as Cloud-HPC interoperable software with performance of HPC (High Performance Computing) and the rich functionality of the commodity Apache Big Data Stack was a bold idea developed.

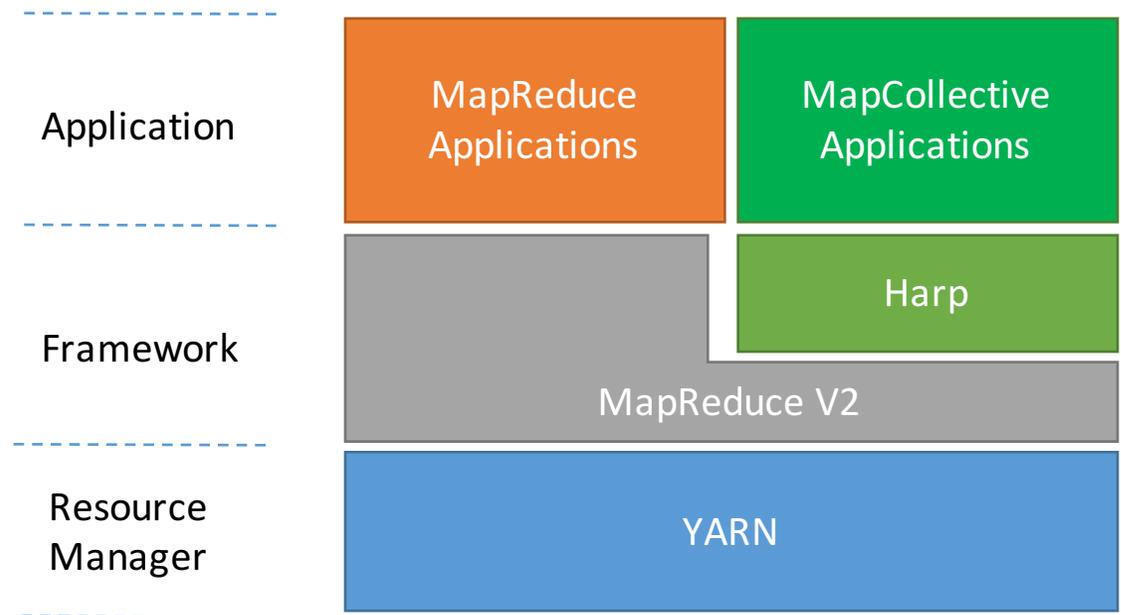


The Concept of Harp Plug-in for Hadoop

Parallelism Model



Architecture



Harp is an open-source project developed at Indiana University [6], it has:

- MPI-like **collective communication** operations that are highly optimized for big data problems.
- Harp has efficient and innovative **computation models** for different machine learning problems.

[6] Harp project. Available at <https://dsc-spidal.github.io/harp>



Data Management

Data sources
Data dictionaries
Data model
Numeric tables & matrices
Compression

Algorithms

Analysis
Training
Prediction

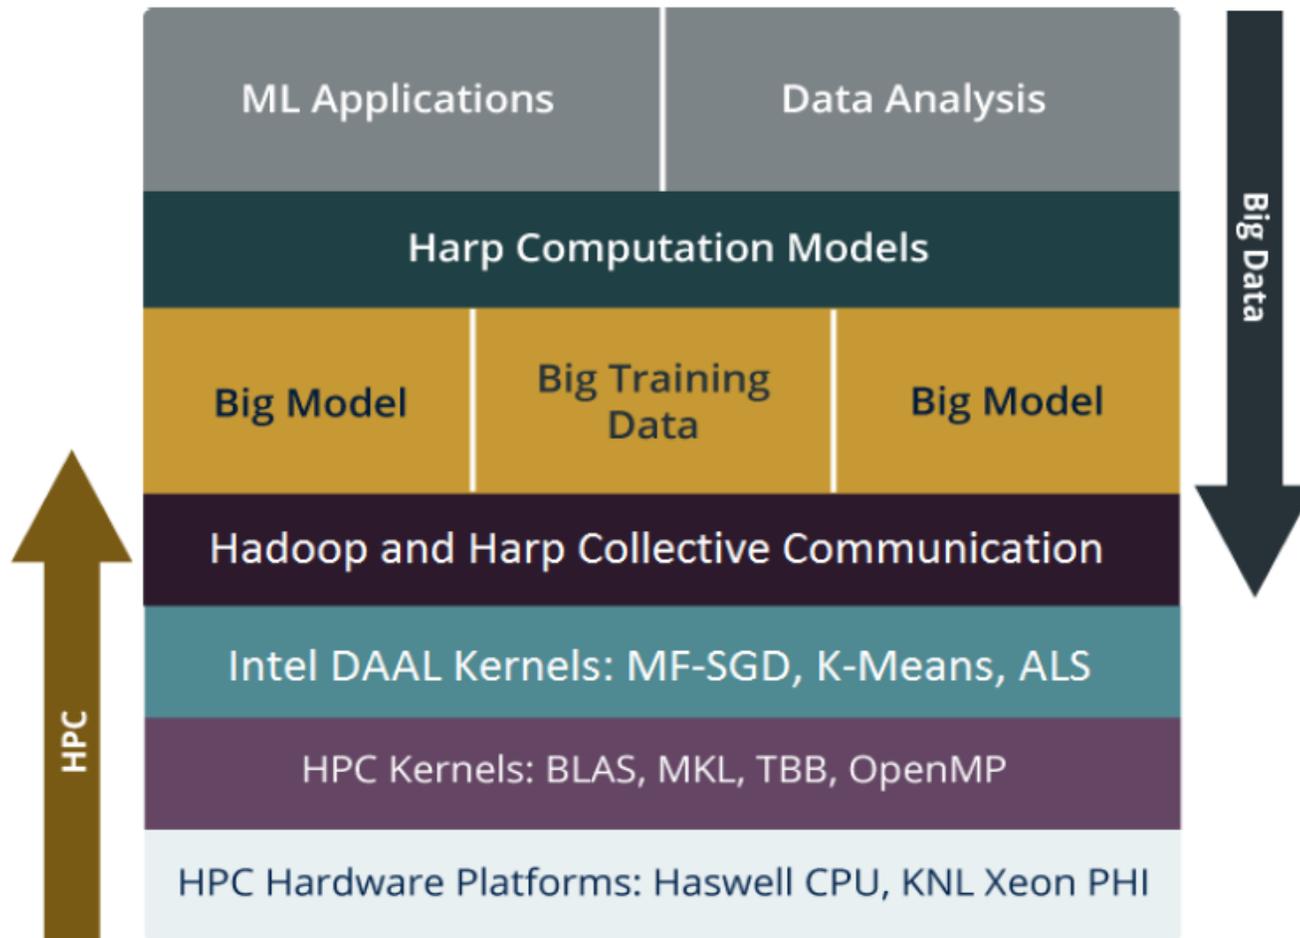
Services

Memory allocation
Error handling
Collections
Shared pointers

DAAL is an open-source project that provides:

- Algorithms Kernels to Users
 - Batch Mode (Single Node)
 - **Distributed Mode (multi nodes)**
 - Streaming Mode (single node)
- Data Management & APIs to Developers
 - Data structure, e.g., Table, Map, etc.
 - HPC Kernels and Tools: MKL, TBB, etc.
 - Hardware Support: Compiler

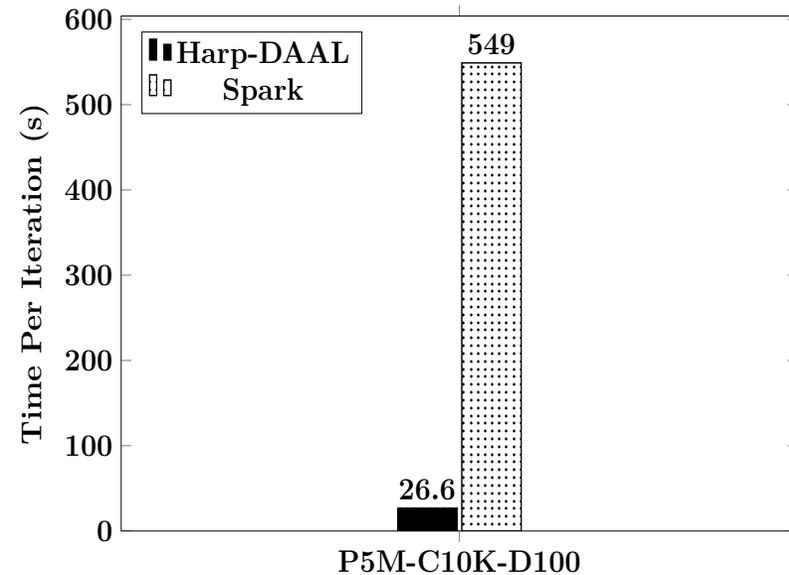
Harp-DAAL enable faster Machine Learning Algorithms with Hadoop Clusters on Multi-core and Many-core architectures



- Bridge the gap between HPC hardware and Big data/Machine learning Software
- Support Iterative Computation, Collective Communication, Intel DAAL and native kernels
- Portable to new many-core architectures like Xeon Phi and run on Haswell and KNL clusters

Performance of Harp-DAAL on KNL Single Node

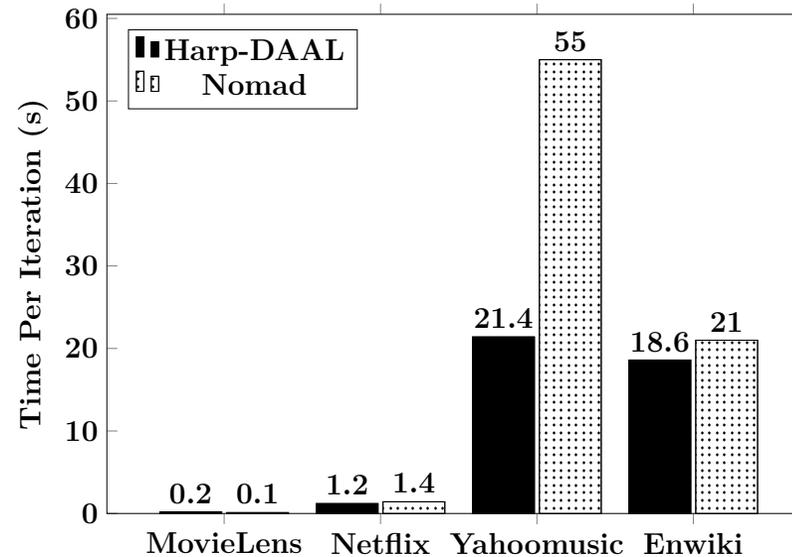
Harp-DAAL vs. Spark vs. NOMAD



Harp-DAAL-Kmeans vs. Spark-Kmeans:

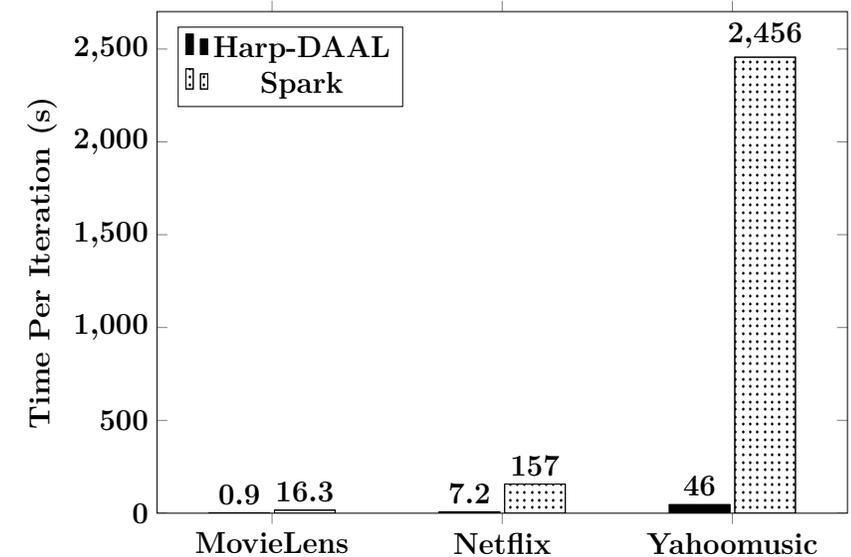
~ 20x speedup

- 1) Harp-DAAL-Kmeans invokes MKL matrix operation kernels at low level
- 2) Matrix data stored in contiguous memory space, leading to regular access pattern and data locality



Harp-DAAL-SGD vs. NOMAD-SGD

- 1) Small dataset (MovieLens, Netflix): comparable perf
- 2) Large dataset (YahooMusic, Enwiki): **1.1x to 2.5x**, depending on data distribution of matrices



Harp-DAAL-ALS vs. Spark-ALS

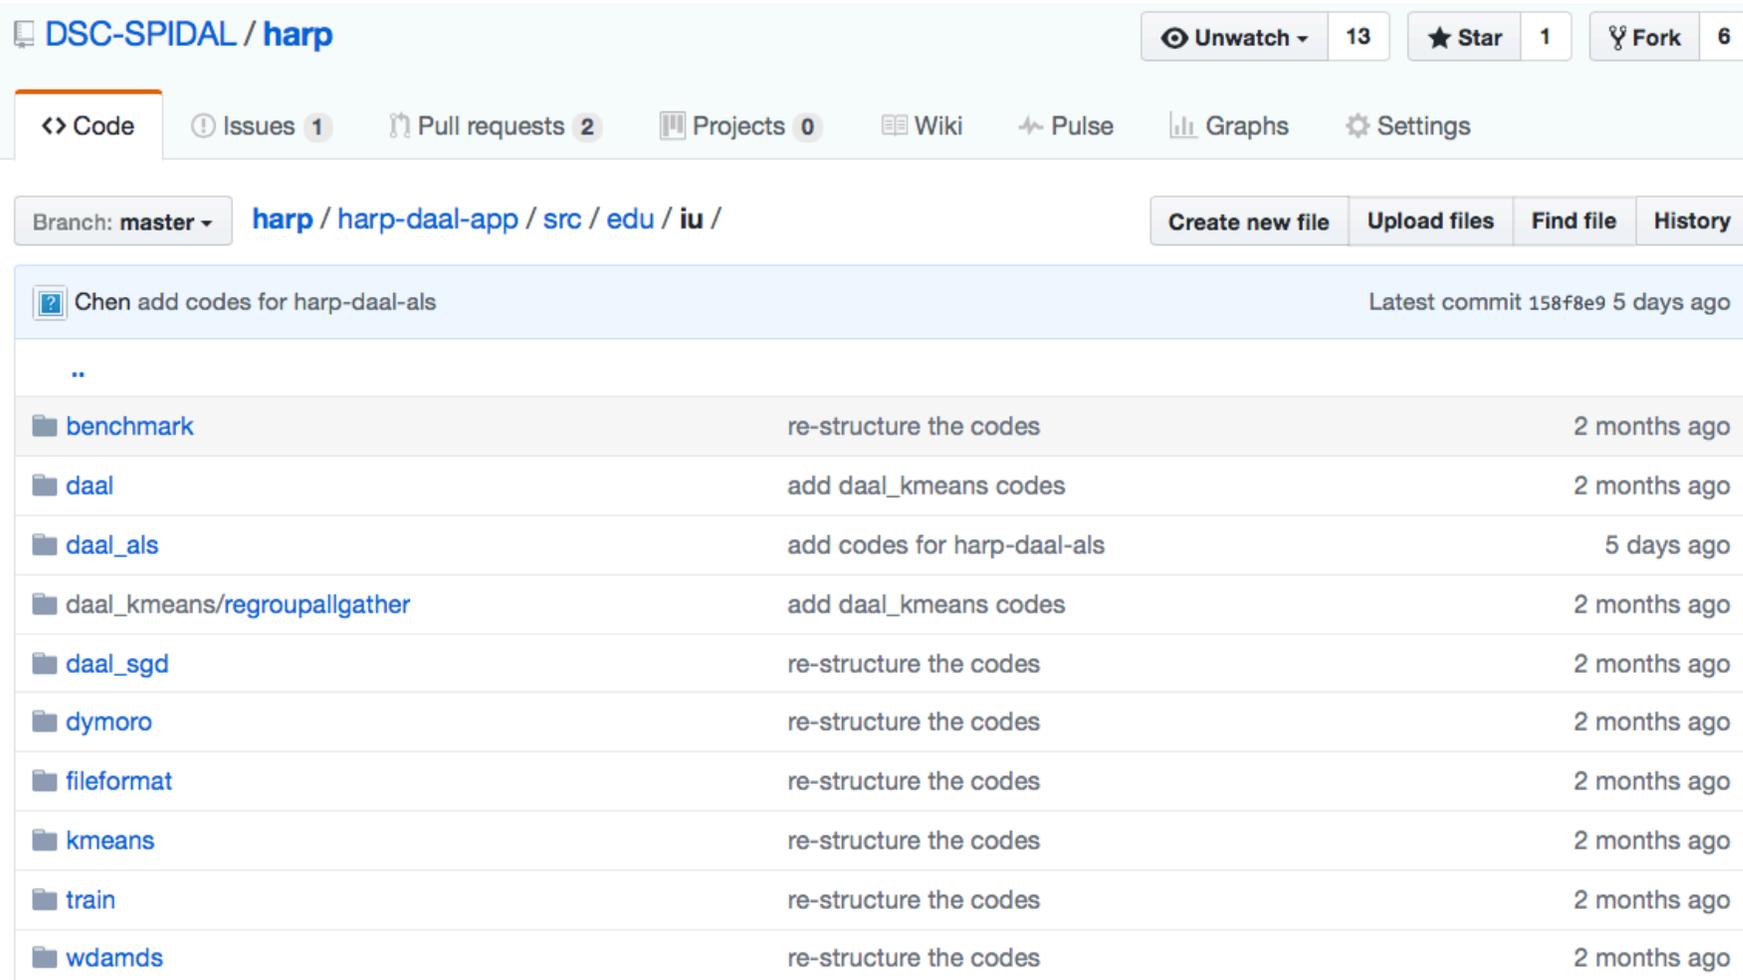
20x to 50x speedup

- 1) Harp-DAAL-ALS invokes MKL at low level
- 2) Regular memory access, data locality in matrix operations

Hadoop/Harp-DAAL: Prototype and Production Code

Source codes became available on Github at [Harp-DAAL project](#) in February, 2017.

- Harp-DAAL follows the same standard of DAAL's original codes
- Six Applications
 - Harp-DAAL Kmeans
 - Harp-DAAL MF-SGD
 - Harp-DAAL MF-ALS
 - Harp-DAAL SVD
 - Harp-DAAL PCA
 - Harp-DAAL Neural Networks



DSC-SPIDAL / **harp** Unwatch 13 Star 1 Fork 6

<> Code Issues 1 Pull requests 2 Projects 0 Wiki Pulse Graphs Settings

Branch: **master** [harp](#) / [harp-daal-app](#) / [src](#) / [edu](#) / [iu](#) / Create new file Upload files Find file History

[Chen add codes for harp-daal-als](#) Latest commit 158f8e9 5 days ago

..		
benchmark	re-structure the codes	2 months ago
daal	add daal_kmeans codes	2 months ago
daal_als	add codes for harp-daal-als	5 days ago
daal_kmeans/regroupallgather	add daal_kmeans codes	2 months ago
daal_sgd	re-structure the codes	2 months ago
dymoro	re-structure the codes	2 months ago
fileformat	re-structure the codes	2 months ago
kmeans	re-structure the codes	2 months ago
train	re-structure the codes	2 months ago
wdamds	re-structure the codes	2 months ago

Scalable Algorithms implemented using Harp

Algorithm	Category	Applications	Features	Computation Model	Collective Communication
K-means	Clustering	Most scientific domain	Vectors	AllReduce	allreduce, regroup+allgather, broadcast+reduce, push+pull
				Rotation	rotate
Multi-class Logistic Regression	Classification	Most scientific domain	Vectors, words	Rotation	regroup, rotate, allgather
Random Forests	Classification	Most scientific domain	Vectors	AllReduce	allreduce
Support Vector Machine	Classification, Regression	Most scientific domain	Vectors	AllReduce	allgather
Neural Networks	Classification	Image processing, voice recognition	Vectors	AllReduce	allreduce
Latent Dirichlet Allocation	Structure learning (Latent topic model)	Text mining, Bioinformatics, Image Processing	Sparse vectors; Bag of words	Rotation	rotate, allreduce
Matrix Factorization	Structure learning (Matrix completion)	Recommender system	Irregular sparse Matrix; Dense model vectors	Rotation	rotate
Multi-Dimensional Scaling	Dimension reduction	Visualization and nonlinear identification of principal components	Vectors	AllReduce	allgather, allreduce
Subgraph Mining	Graph	Social network analysis, data mining, fraud detection, chemical informatics, bioinformatics	Graph, subgraph	Rotation	rotate
Force-Directed Graph Drawing	Graph	Social media community detection and visualization	Graph	AllReduce	allgather, allreduce