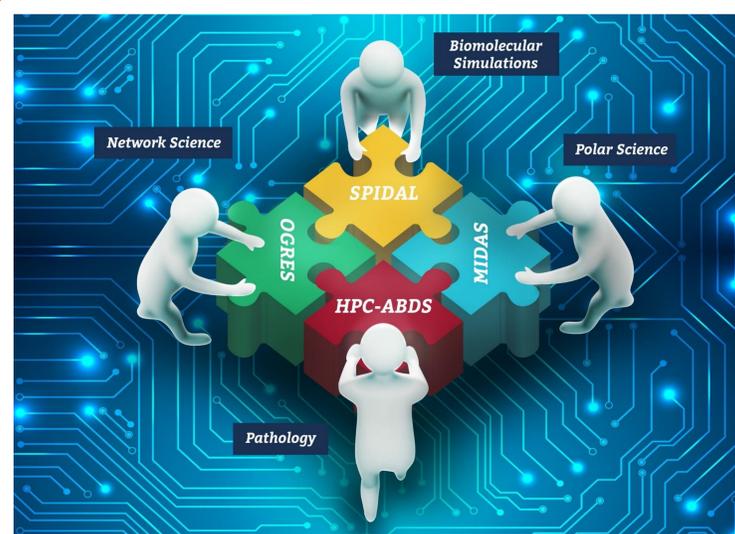


Community Driven High-Performance Big Data

for bio-physical applications based on HPC, distributed systems, network science, GIS and machine learning



MIDAS High Performance Middleware

for Data-Intensive Analytics and Science

Harp HP Machine Learning

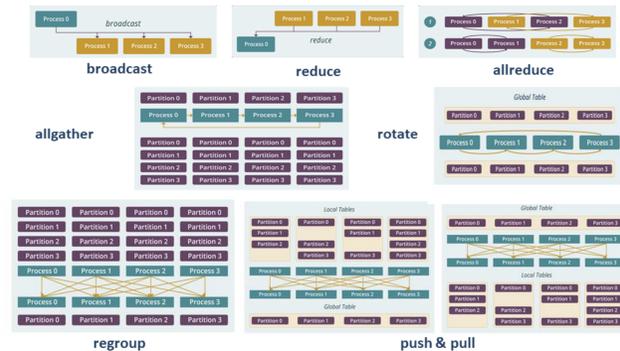
Harp-DAAL: A HPC-Cloud convergence framework for productivity and performance

- Harp is an Hadoop plug-in offering highly optimized **Map collective** operations plus support of computation models covering much machine learning
- Harp-DAAL seamlessly integrates Intel's DAAL Data Analytics Library
- DAAL: SubGraph Mining, K-means, Matrix Factorization, Recommender System (ALS), Singular Value Decomposition (SVD), QR Decomposition, Neural Network, Covariance, Low Order Moments, Naive Bayes Reduce, Linear Regression, Ridge Regression, Principal Component Analysis (PCA)
- **No DAAL yet:** DA-MDS Multi-dimensional Scaling (DA is Deterministic Annealing), Directed Force Dimension, Irregular DAVS Clustering, DA Semimetric Clustering, Multi-class Logistic Regression, SVM, Latent Dirichlet Allocation, Random Forest

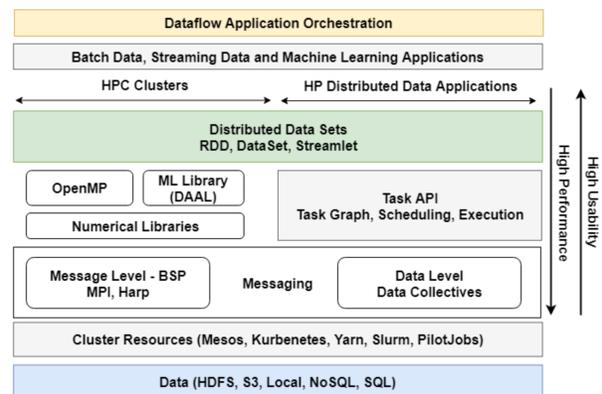


Expand the Applicability of MapReduce to more classes of Applications on HPC-Cloud Platforms (Multicore, Manycore, other Accelerators, ...)

Map Collective Run time merges MapReduce and HPC



Twister2 HP Big Data Architecture



Deployment of MIDAS and SPIDAL

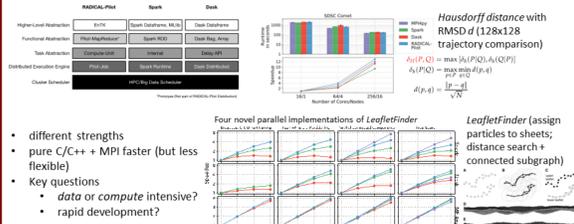
- Goals of deployment
- Add to existing frameworks as high performance components (scikit-learn, Spark, TensorFlow)
- Deploy SPIDAL algorithms on big data environments
- Combine SPIDAL algorithms into complete applications (dataflow orchestration) including both streaming and data pipelines (pre-processing and post processing)
- Be user friendly - easy deployments and integrations
- SPIDAL algorithms using a common set of tools (Harp, MPI, OpenMP)
- Common data abstraction for Input/Output
- Deploy in native fashion using Twister2
- Full applications including streaming and data processing components
- Release the algorithms both in binary format, source code and SaaS
- <https://github.com/orgs/hpcanalytics>
- Comprehensive documentation including user videos
- Web portal for testing
- Deploy Harp as Apache Project; algorithms use Apache processes to build communities

SPIDAL Libraries and Communities

Biomolecular Simulation Community

Biomolecular Simulations application of high-performance data analytics to problems in the biomolecular simulation community (Rutgers, Utah, Arizona State)

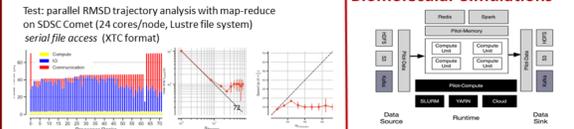
BENCHMARKING: Evaluation of task-parallel frameworks



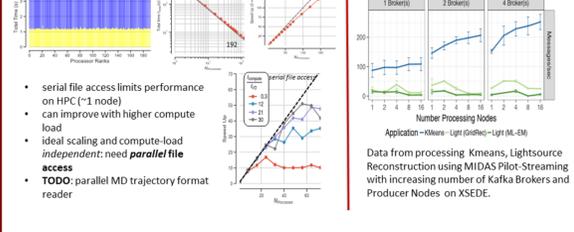
NEW ALGORITHMS AND PRODUCTION-GRADE IMPLEMENTATIONS

PMDA Python library: parallel MDAnalysis <https://github.com/hpcanalytics/pmda>
 cpptraj: trajectory analysis tool <https://github.com/Amber-MD/cpptraj>
 Simple decorators/classes for writing parallel analysis with MDAnalysis and dask
 parallel LeafletFinder
 parallel PSA/Hausdorff RMSD, contacts, g(r), ...

IDENTIFYING AND OVERCOMING BARRIERS TO PROGRESS: I/O in parallel trajectory analysis

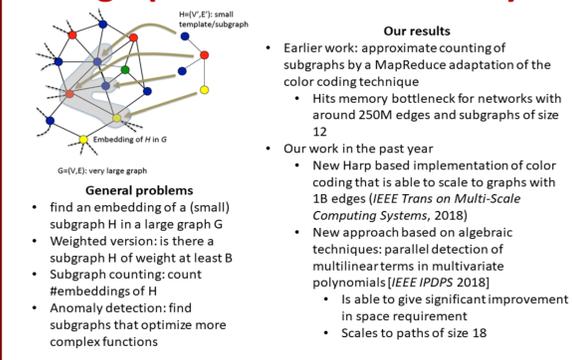


Streaming and Biomolecular Simulations

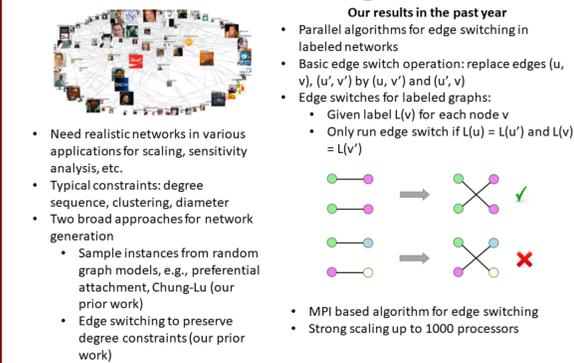


Network Science Community

Subgraph detection and analysis



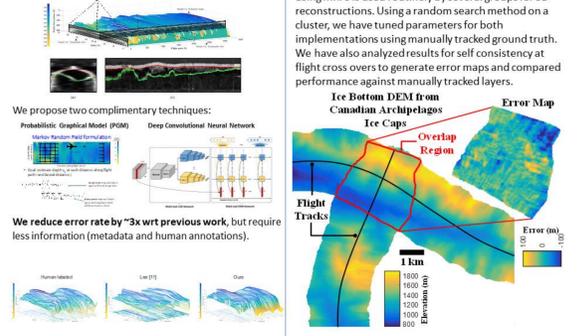
Parallel network generation



Remote Sensing / Polar Community

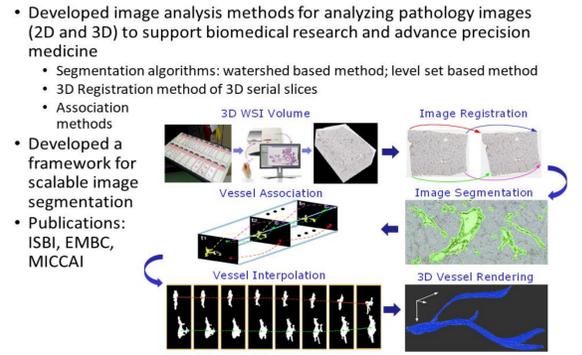
Inferring structure of polar ice sheets through radar informatics

We reconstruct the 3d surface of ice-bedrock layers, using echograms from ground-penetrating radar.



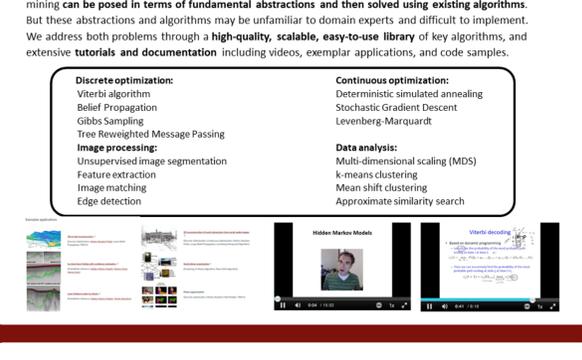
Pathology Community

Pathology image analysis



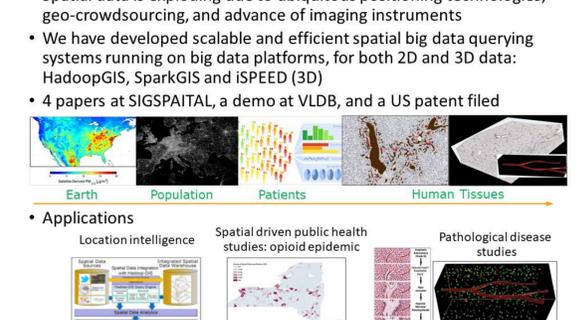
SPIDAL Components

Image Processing and Machine Learning Toolkit



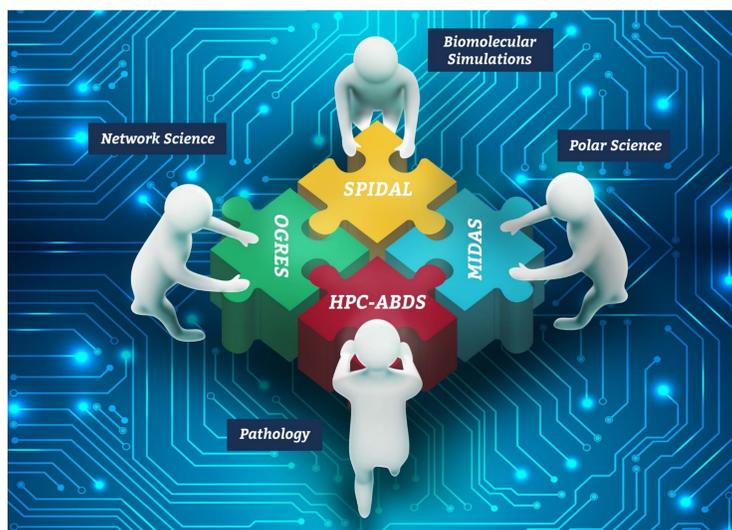
Spatial GIS Community

Spatial big data querying systems



Community Driven High-Performance Big Data

for bio-physical applications based on HPC, distributed systems, network science, GIS and machine learning



Components of NSF 1443054

- Active:**
 - Design and build SPIDAL – a Scalable Parallel Interoperable Data Analytics Library
 - Domain specific libraries in communities
 - Core Machine Learning Libraries
 - High Performance Data Middleware MIDAS based on HPC-ABDS systems Harp and Twister2
 - Packaging SPIDAL and MIDAS; building community
- Completed:**
 - MIDAS Pilot Jobs HPC Task Management
 - High Performance for Java/C++ for Machine Learning
 - NIST Big Data Application Analysis: Features of data intensive Applications deriving 64 Convergence Diamonds
 - HPC-ABDS: Cloud-HPC Interoperable software with performance of HPC and rich functionality of commodity Apache Stack
 - Implementations: HPC and Clouds with DevOps

Deployment of MIDAS and SPIDAL

- Goals of deployment
 - Add to existing frameworks as high performance components (scikit-learn, Spark, TensorFlow)
 - Deploy SPIDAL algorithms on big data environments
 - Combine SPIDAL algorithms into complete applications (dataflow orchestration) including both streaming and data pipelines (pre-processing and post processing)
 - Be user friendly - easy deployments and integrations
- SPIDAL algorithms using a common set of tools (Harp, MPI, OpenMP)
- Common data abstraction for Input/Output

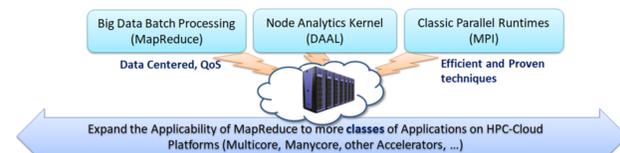
MIDAS High Performance Middleware

for Data-Intensive Analytics and Science

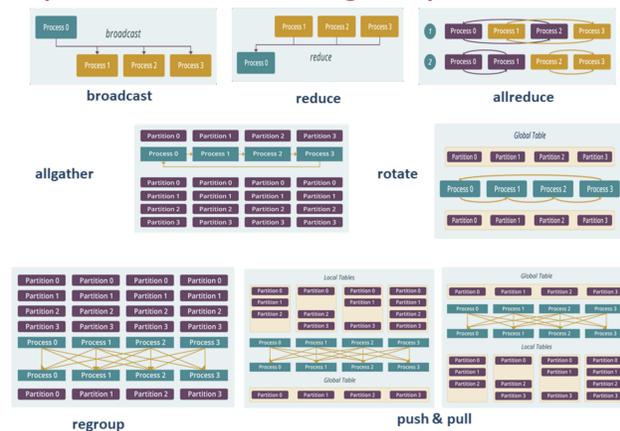
Harp HP Machine Learning

Harp-DAAL: A HPC-Cloud convergence framework for productivity and performance

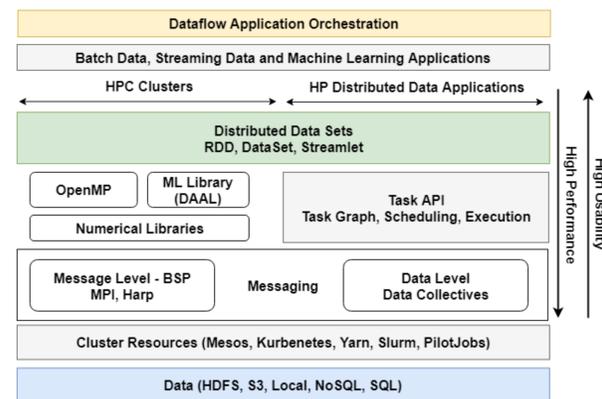
- Harp** is an Hadoop plug-in offering highly optimized **Map collective** operations plus support of computation models covering much machine learning
- Harp-DAAL** seamlessly integrates Intel's DAAL Data Analytics Library
- DAAL**: SubGraph Mining, K-means, Matrix Factorization, Recommender System (ALS), Singular Value Decomposition (SVD), QR Decomposition, Neural Network, Covariance, Low Order Moments, Naive Bayes Reduce, Linear Regression, Ridge Regression, Principal Component Analysis (PCA)
- No DAAL yet**: DA-MDS Multi-dimensional Scaling (DA is Deterministic Annealing), Directed Force Dimension, Irregular DAVS Clustering, DA Semimetric Clustering, Multi-class Logistic Regression, SVM, Latent Dirichlet Allocation, Random Forest



Map Collective Run time merges MapReduce and HPC



Twister2 HP Big Data Architecture



- Deploy in native fashion using Twister2
 - Full applications including streaming and data processing components
- Release the algorithms both in binary format, source code and SaaS
 - <https://github.com/orgs/hpcanalytics>
- Comprehensive documentation including user videos
- Web portal for testing
- Deploy Harp as Apache Project; algorithms use Apache processes to build communities

SPIDAL Libraries and Communities

Biomolecular Simulation Community

Biomolecular Simulations application of high-performance data analytics to problems in the biomolecular simulation community (Rutgers, Utah, Arizona State)

BENCHMARKING: Evaluation of task-parallel frameworks

NEW ALGORITHMS AND PRODUCTION-GRADE IMPLEMENTATIONS

Simple decorators/classes for writing parallel analysis with `MDAnalysis` and `dask`

parallel `LeafletFinder` for writing parallel analysis

parallel `PSA/Hausdorff` with `MDAnalysis` and `dask`

IDENTIFYING AND OVERCOMING BARRIERS TO PROGRESS: I/O in parallel trajectory analysis

Test: parallel RMSD trajectory analysis with map-reduce on SDSComet (24 cores/node, Lustre file system) serial file access (XTC format)

Streaming and Biomolecular Simulations

serial file access limits performance on HPC (~1 node)

can improve with higher compute load

ideal scaling and compute-load independent: need **parallel file access**

TODD: parallel MD trajectory format reader

Network Science Community

Subgraph detection and analysis

Our results

- Earlier work: approximate counting of subgraphs by a MapReduce adaptation of the color coding technique
 - Hits memory bottleneck for networks with around 250M edges and subgraphs of size 12
- Our work in the past year
 - New Harp based implementation of color coding that is able to scale to graphs with 1B edges (*IEEE Trans on Multi-Scale Computing Systems*, 2018)
 - New approach based on algebraic techniques: parallel detection of multilinear terms in multivariate polynomials (*IEEE PDPS* 2018)
 - Is able to give significant improvement in space requirement
 - Scales to paths of size 18

General problems

- find an embedding of a (small) subgraph H in a large graph G
- Weighted version: is there a subgraph H of weight at least B
- Subgraph counting: count #embeddings of H
- Anomaly detection: find subgraphs that optimize more complex functions

Parallel network generation

Our results in the past year

- Parallel algorithms for edge switching in labeled networks
- Basic edge switch operation: replace edges $(u, v), (u', v')$ by (u, v') and (u', v)
- Edge switches for labeled graphs:
 - Given label $L(v)$ for each node v
 - Only run edge switch if $L(u) = L(v)$ and $L(v') = L(v)$

Need realistic networks in various applications for scaling, sensitivity analysis, etc.

Typical constraints: degree sequence, clustering, diameter

Two broad approaches for network generation

- Sample instances from random graph models, e.g., preferential attachment, Chung-Lu (our prior work)
- Edge switching to preserve degree constraints (our prior work)

MPI based algorithm for edge switching

Strong scaling up to 1000 processors

Remote Sensing / Polar Community

Inferring structure of polar ice sheets through radar informatics

We reconstruct the **3d surface of ice-bedrock layers**, using echograms from ground-penetrating radar.

We propose two complementary techniques:

- Probabilistic Graphical Model (PGM)
- Deep Convolutional Neural Network

We reduce error rate by ~3x wrt previous work, but require less information (metadata and human annotations).

Our HPC implementation of the PGM technique is now used routinely to process 2D radar data collected by CREISIS (50-100 TB/year), and our HPC implementation using MRFs is used routinely by several groups for 3D reconstructions. Using a random search method on a cluster, we have tuned parameters for both implementations using manually tracked ground truth. We have also analyzed results for self consistency at flight cross overs to generate error maps and compared performance against manually tracked layers.

Ice Bottom DEM from Canadian Archipelago

Ice Caps

Overlap Region

Flight Tracks

Error Map

Pathology Community

Pathology image analysis

- Developed image analysis methods for analyzing pathology images (2D and 3D) to support biomedical research and advance precision medicine
 - Segmentation algorithms: watershed based method; level set based method
 - 3D Registration method of 3D serial slices
 - Association methods
- Developed a framework for scalable image segmentation
- Publications: ISBI, EMBC, MICCAI

SPIDAL Components

Image Processing and Machine Learning Toolkit

Many problems in scientific applications of computer vision, image processing, machine learning, and data mining can be posed in terms of fundamental abstractions and then solved using existing algorithms. But these abstractions and algorithms may be unfamiliar to domain experts and difficult to implement. We address both problems through a high-quality, scalable, easy-to-use library of key algorithms, and extensive tutorials and documentation including videos, exemplar applications, and code samples.

Discrete optimization: Viterbi algorithm Belief Propagation Gibbs Sampling Tree-Weighted Message Passing	Continuous optimization: Deterministic simulated annealing Stochastic Gradient Descent Levenberg-Marquardt
Image processing: Unsupervised image segmentation Feature extraction Image matching Edge detection	Data analysis: Multi-dimensional scaling (MDS) k-means clustering Mean shift clustering Approximate similarity search

Spatial GIS Community

Spatial big data querying systems

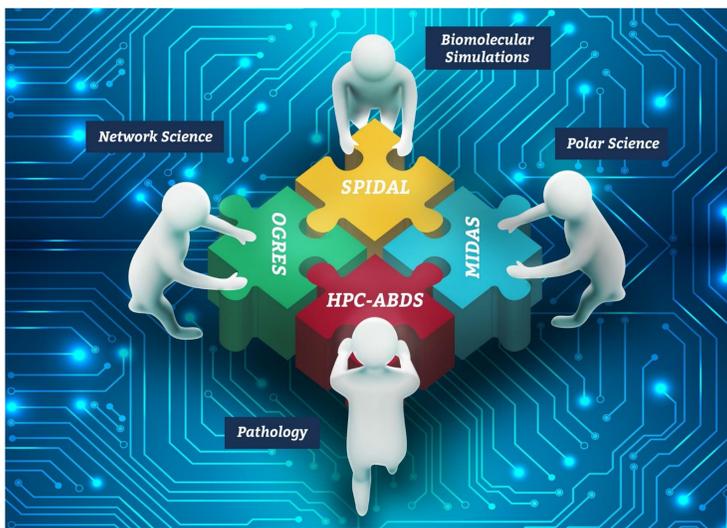
- Spatial data is exploding due to ubiquitous positioning technologies, geo-crowdsourcing, and advance of imaging instruments
- We have developed scalable and efficient spatial big data querying systems running on big data platforms, for both 2D and 3D data: HadoopGIS, SparkGIS and iSPEED (3D)
- 4 papers at SIGSPATIAL, a demo at VLDB, and a US patent filed

Applications

- Location intelligence
- Spatial driven public health studies: opioid epidemic
- Pathological disease studies

Community Driven High-Performance Big Data

for bio-physical applications based on HPC, distributed systems, network science, GIS and machine learning



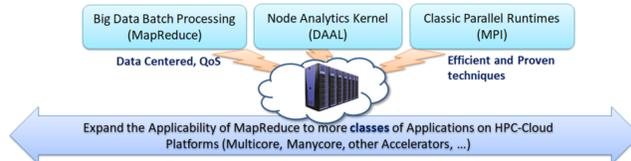
MIDAS High Performance Middleware

for Data-Intensive Analytics and Science

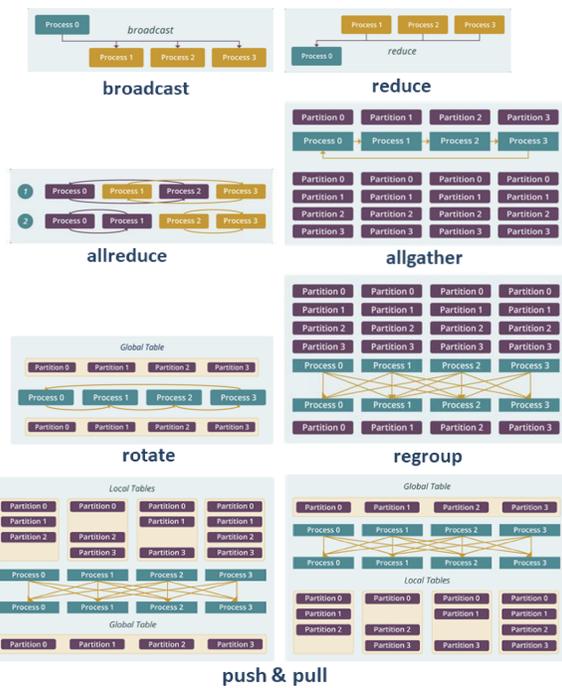
Harp HP Machine Learning

Harp-DAAL: A HPC-Cloud convergence framework for productivity and performance

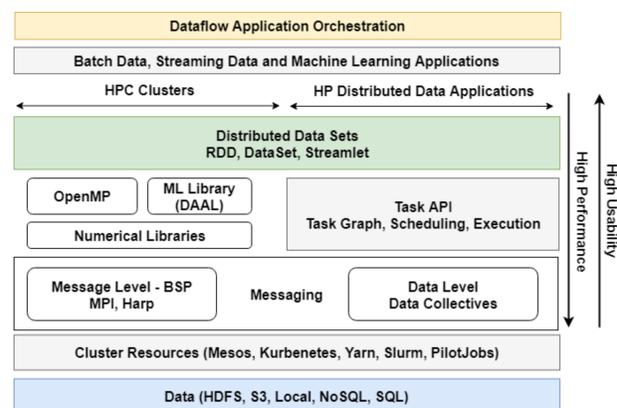
- Harp is an Hadoop plug-in offering highly optimized **Map collective** operations plus support of computation models covering much machine learning
- Harp-DAAL seamlessly integrates Intel's DAAL DAAL node Data Analytics Library
- DAAL: SubGraph Mining, K-means, Matrix Factorization, Recommender System (ALS), Singular Value Decomposition (SVD), QR Decomposition, Neural Network, Covariance, Low Order Moments, Naive Bayes Reduce, Linear Regression, Ridge Regression, Principal Component Analysis (PCA)
- No DAAL yet:** DA-MDS Multi-dimensional Scaling (DA is Deterministic Annealing), Directed Force Dimension, Irregular DAVS Clustering, DA Semimetric Clustering, Multi-class Logistic Regression, SVM, Latent Dirichlet Allocation, Random Forest



Map Collective Run time merges MapReduce and HPC



Twister2 HP Big Data Architecture



SPIDAL Libraries and Communities

Biomolecular Simulation Community

Biomolecular Simulations application of high-performance data analytics to problems in the biomolecular simulation community (Rutgers, Utah, Arizona State)

BENCHMARKING: Evaluation of task-parallel frameworks

NEW ALGORITHMS AND PRODUCTION-GRADE IMPLEMENTATIONS

- parallel LeafletFinder for writing parallel analysis
- parallel PSA/Hausdorff with MDAnalysis and dask
- RMSD, contacts, g(r), ...

IDENTIFYING AND OVERCOMING BARRIERS TO PROGRESS: I/O in parallel trajectory analysis

Test: parallel RMSD trajectory analysis with map-reduce on SDS5 Comet (24 cores/node, Lustre file system) serial file access (XTC format)

Streaming and Biomolecular Simulations

Extended MIDAS to integrate Streaming into HPC simulation and analysis workflows.

Application: $\text{MDS} \rightarrow \text{Light} \rightarrow \text{Light} \rightarrow \text{Light} \rightarrow \text{EM}$

Data from processing: Kmeans, Lightsources Reconstruction using MIDAS Pilot-Streaming with increasing number of Kafka Brokers and Producer Nodes on XSEDE.

Remote Sensing / Polar Community

Inferring structure of polar ice sheets through radar informatics

Our HPC implementation of the PGM technique is now used routinely to process 2D radar data collected by CryoSat (50-100 TB/year), and our HPC implementation using MRFs is used routinely by several groups for 3D reconstructions. Using a random search method on a cluster, we have tuned parameters for both implementations using manually tracked ground truth. We have also analyzed results for self consistency at flight cross overs to generate error maps and compared performance against manually tracked layers.

We reconstruct the 3d surface of ice-bedrock layers, using echograms from ground-penetrating radar.

We propose two complimentary techniques: Probabilistic Graphical Model (PGM) and Deep Convolutional Neural Network.

We reduce error rate by ~3x wrt previous work, but require less information (metadata and human annotations).

Pathology Community

Pathology image analysis

- Developed image analysis methods for analyzing pathology images (2D and 3D) to support biomedical research and advance precision medicine
- Segmentation algorithms: watershed based method; level set based method
- 3D Registration method of 3D serial slices
- Association methods
- Developed a framework for scalable image segmentation
- Publications: ISBI, EMBC, MICCAI

Components of NSF 1443054

- Active:**
- Design and build SPIDAL – a Scalable Parallel Interoperable Data Analytics Library
 - Domain specific libraries in communities
 - Core Machine Learning Libraries
 - High Performance Data Middleware MIDAS based on HPC-ABDS systems Harp and Twister2
 - Packaging SPIDAL and MIDAS; building community
- Completed:**
- MIDAS Pilot Jobs HPC Task Management
 - High Performance for Java/C++ for Machine Learning
 - NIST Big Data Application Analysis: Features of data intensive Applications deriving 64 Convergence Diamonds
 - HPC-ABDS: Cloud-HPC Interoperable software with performance of HPC and rich functionality of commodity Apache Stack
 - Implementations: HPC and Clouds with DevOps

Deployment of MIDAS and SPIDAL

- Goals of deployment
 - Add to existing frameworks as high performance components (scikit-learn, Spark, TensorFlow)
 - Deploy SPIDAL algorithms on big data environments
 - Combine SPIDAL algorithms into complete applications (dataflow orchestration) including both streaming and data pipelines (pre-processing and post processing)
 - Be user friendly - easy deployments and integrations
- SPIDAL algorithms using a common set of tools (Harp, MPI, OpenMP)
- Common data abstraction for Input/Output
- Deploy in native fashion using Twister2
 - Full applications including streaming and data processing components
- Release the algorithms both in binary format, source code and SaaS
 - https://github.com/orgs/hpcanalytics
- Comprehensive documentation including user videos
- Web portal for testing
- Deploy Harp as Apache Project; algorithms use Apache processes to build communities

Network Science Community

Subgraph detection and analysis

Our results

- Earlier work: approximate counting of subgraphs by a MapReduce adaptation of the color coding technique
 - Hits memory bottleneck for networks with around 250M edges and subgraphs of size 12
- Our work in the past year
 - New Harp based implementation of color coding that is able to scale to graphs with 1B edges (IEEE Trans on Multi-Scale Computing Systems, 2018)
 - New approach based on algebraic techniques: parallel detection of multilinear terms in multivariate polynomials (IEEE IPDPS 2018)
 - Is able to give significant improvement in space requirement
 - Scales to paths of size 18

General problems

- find an embedding of a (small) subgraph H in a large graph G
- Weighted version: is there a subgraph H of weight at least B
- Subgraph counting: count #embeddings of H
- Anomaly detection: find subgraphs that optimize more complex functions

Parallel network generation

Our results in the past year

- Parallel algorithms for edge switching in labeled networks
- Basic edge switch operation: replace edges (u, v) , (u', v') by (u, v') and (u', v)
- Edge switches for labeled graphs:
 - Given label $L(v)$ for each node v
 - Only run edge switch if $L(u) = L(u')$ and $L(v) = L(v')$

Need realistic networks in various applications for scaling, sensitivity analysis, etc.

Typical constraints: degree sequence, clustering, diameter

Two broad approaches for network generation

- Sample instances from random graph models, e.g., preferential attachment, Chung-Lu (our prior work)
- Edge switching to preserve degree constraints (our prior work)

MPI based algorithm for edge switching

Strong scaling up to 1000 processors

SPIDAL Components

Image Processing and Machine Learning Toolkit

Many problems in scientific applications of computer vision, image processing, machine learning, and data mining can be posed in terms of fundamental abstractions and then solved using existing algorithms. But these abstractions and algorithms may be unfamiliar to domain experts and difficult to implement. We address both problems through a high-quality, scalable, easy-to-use library of key algorithms, and extensive tutorials and documentation including videos, exemplar applications, and code samples.

Discrete optimization:

- Viterbi algorithm
- Belief Propagation
- Gibbs Sampling
- Tree Reweighted Message Passing

Image processing:

- Unsupervised image segmentation
- Feature extraction
- Image matching
- Edge detection

Continuous optimization:

- Deterministic simulated annealing
- Stochastic Gradient Descent
- Levenberg-Marquardt

Data analysis:

- Multi-dimensional scaling (MDS)
- k-means clustering
- Mean shift clustering
- Approximate similarity search

Spatial GIS Community

Spatial big data querying systems

- Spatial data is exploding due to ubiquitous positioning technologies, geo-crowdsourcing, and advance of imaging instruments
- We have developed scalable and efficient spatial big data querying systems running on big data platforms, for both 2D and 3D data: HadoopGIS, SparkGIS and iSPEED (3D)
- 4 papers at SIGSPATIAL, a demo at VLDB, and a US patent filed

Applications

- Location intelligence
- Spatial driven public health studies: opioid epidemic
- Pathological disease studies