Accelerating Hyperscale Graph Analysis and Machine Learning with D4M on the MIT SuperCloud

Vijay Gadepally, Jeremy Kepner, Peter Michaleas, Lauren Milechin
vijayg@ll.mit.edu

MIT LINCOLN LABORATORY
SUPERCOMPUTING CENTER

Outline

- **Introduction**
- **D4M and MIT SuperCloud**
- **Cyber Network Analysis**
- **Signal Processing on Networks**
- **Summary**

Trend: sophisticated cyber threats

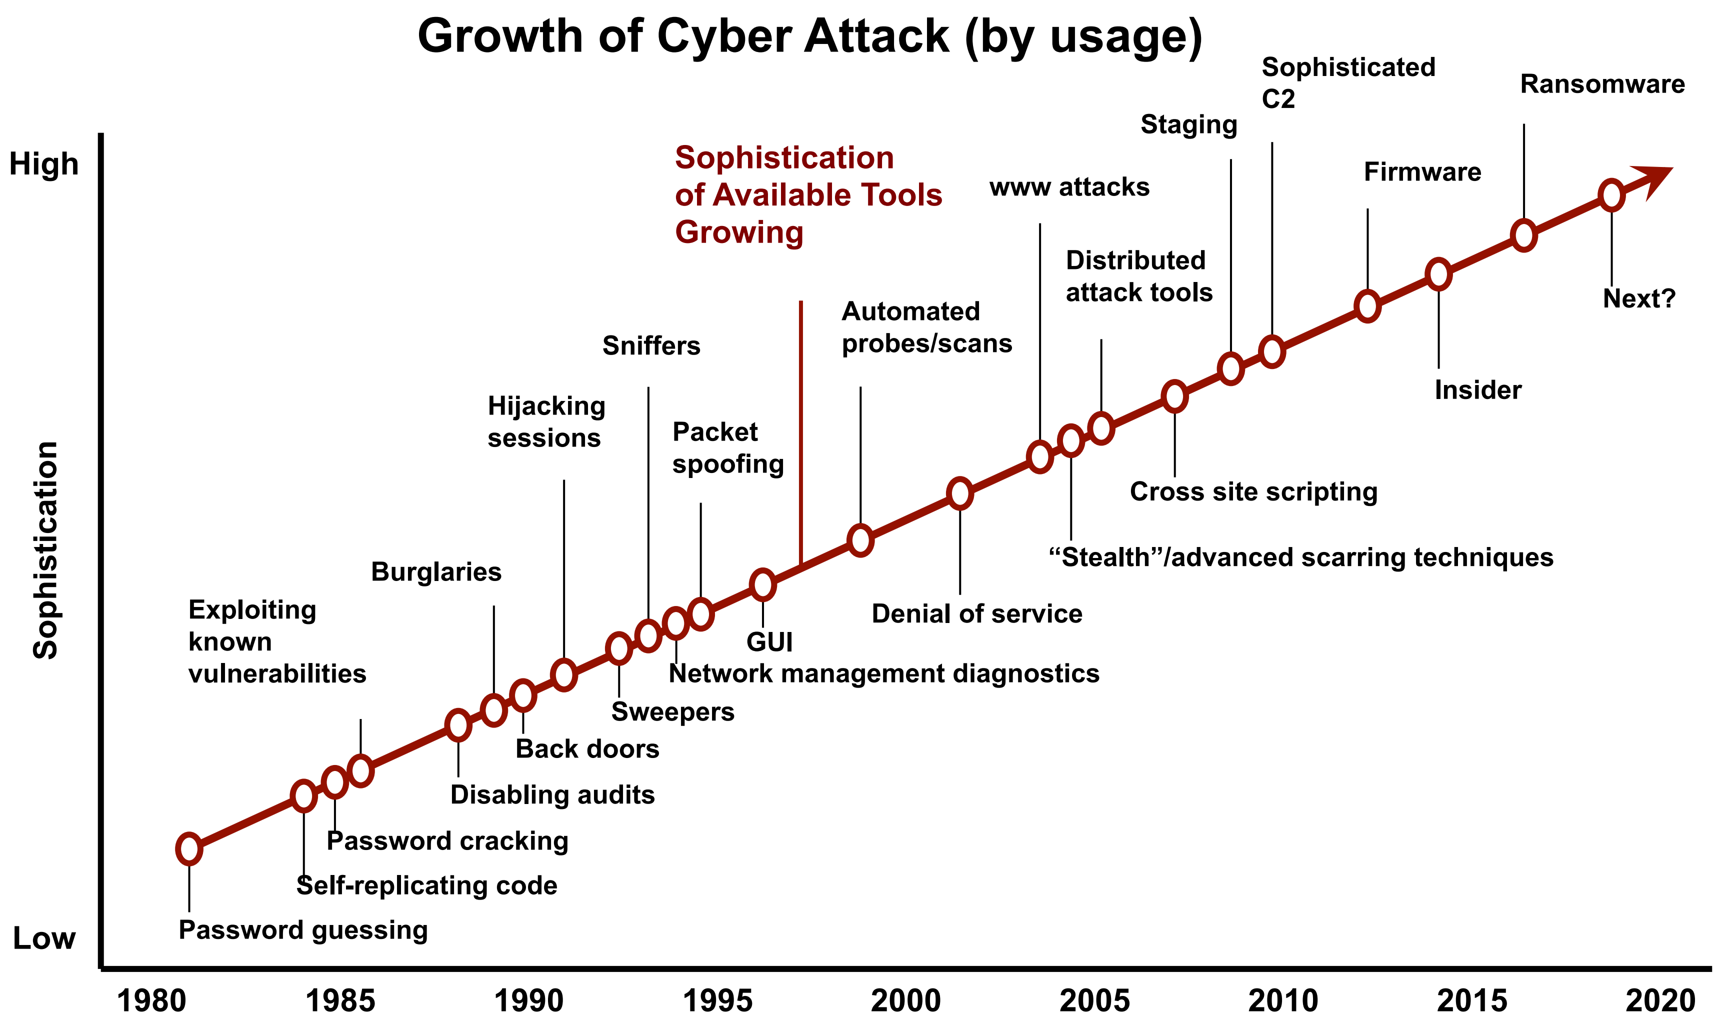

NOTEWORTHY FACTS

- 250K new malware programs are registered each day
- There were 357M new email malware variants in 2016 - 36% more new variants than in 2014.
- There were 463M new variants of ransomware in 2016 - 36% more new variants than in 2015.
- 99 days to detect compromise - adversary gains access in 3
- Internet of Things and Cloud are hot targets (e.g. Mirai botnet) – 2 min to compromise
- Projected cyber attack costs in 2019: $2.1T

Sources: https://www.symantec.com/security-center/threat-report
http://expandedramblings.com/index.php/cybersecurity-statistics/
http://www.nato.int/docu/Review/2016/Also-in-2016/cyber-defense-nato-security-role/EN/index.htm
https://www.fireeye.com/blog/threat-research/2017/03/m-trends-2017.htm
https://www.ag-test.org/en/statistics/malware

Trend: volume of data outpacing storage and processing

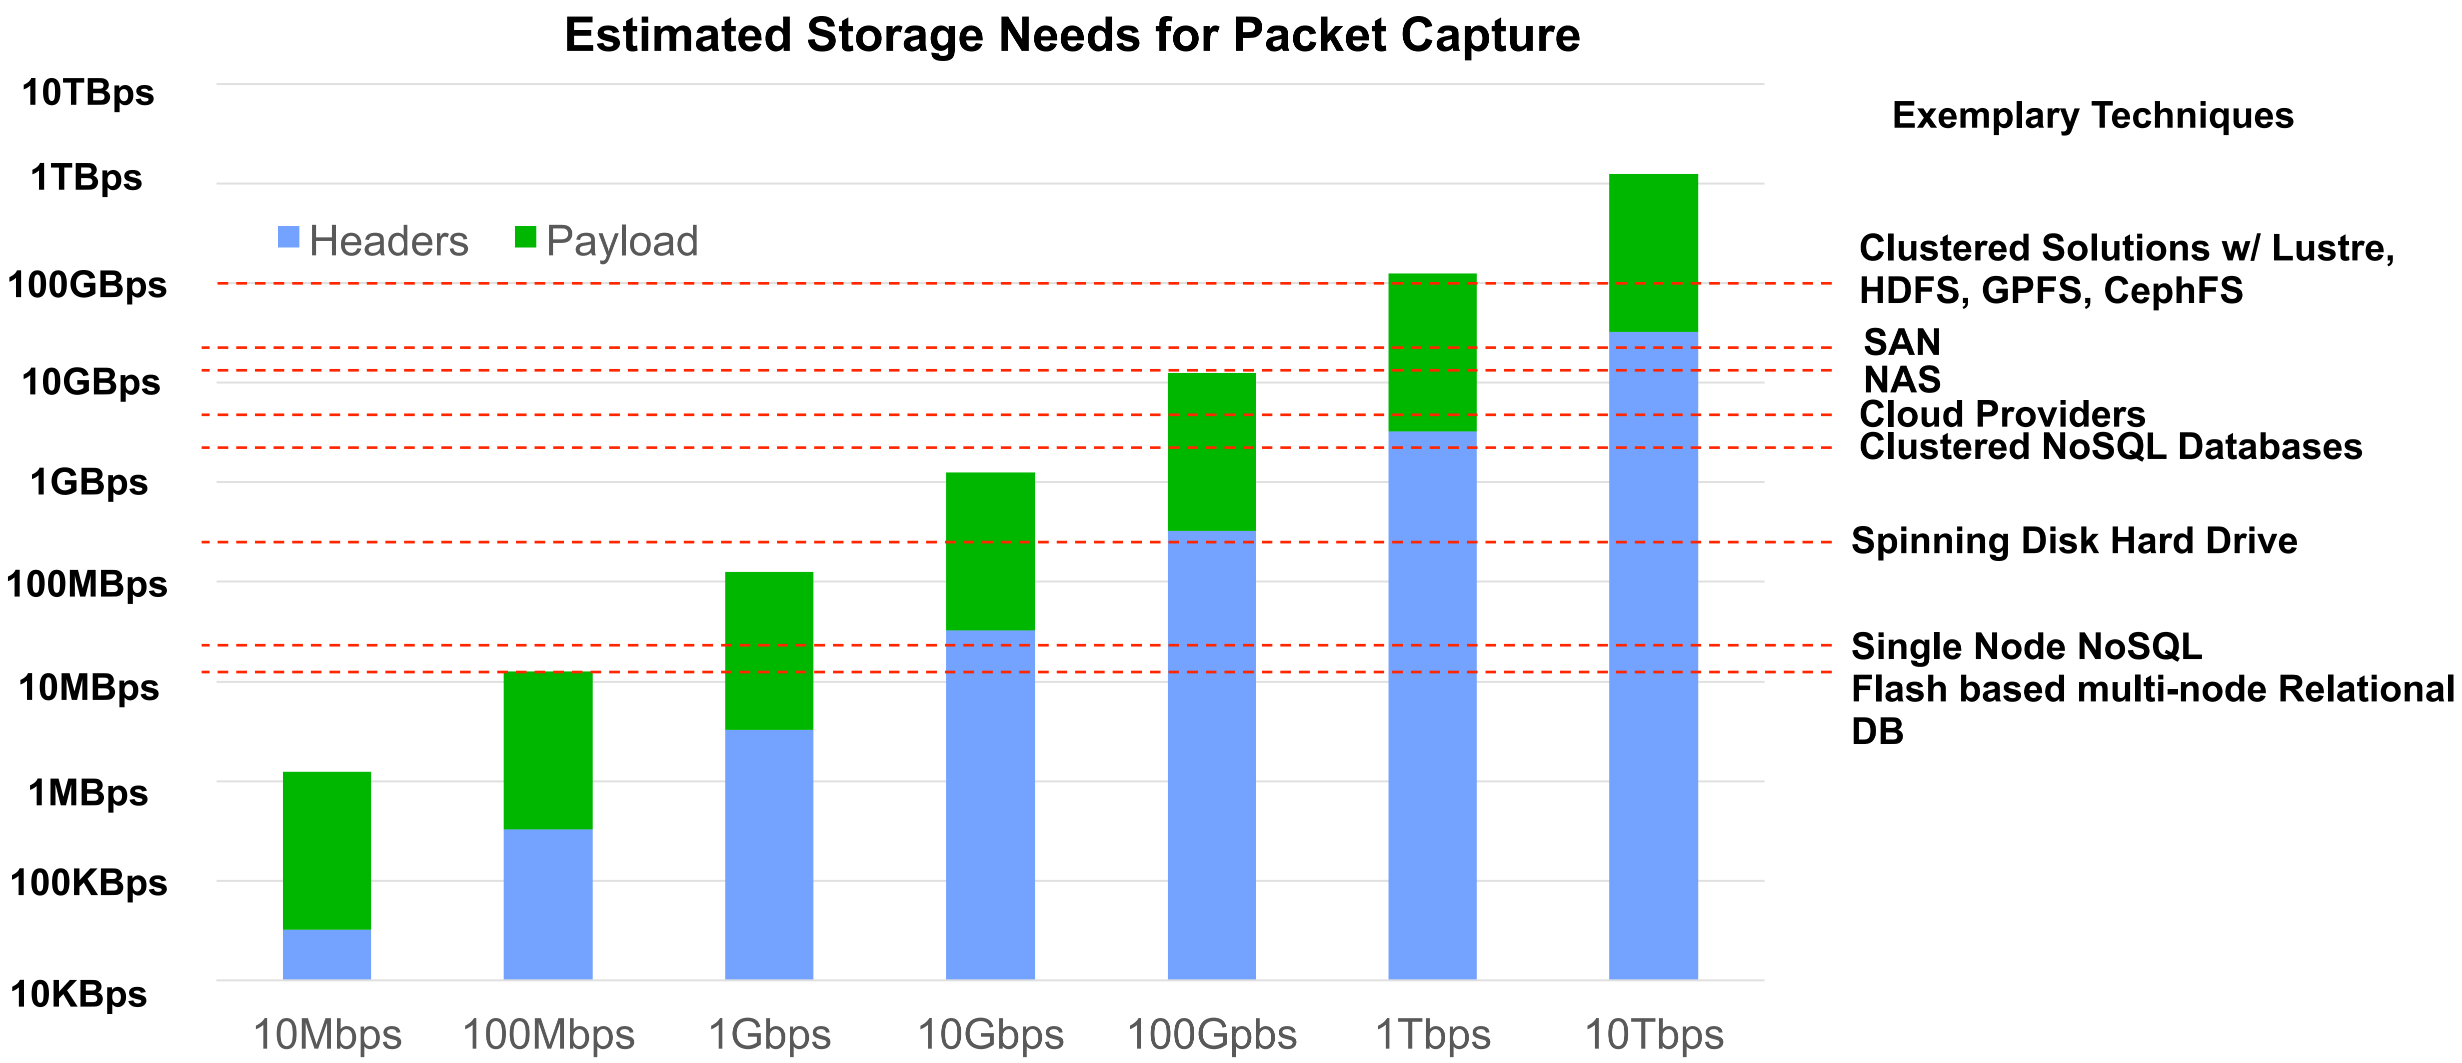

Assumptions: 600 bytes/packet, 40 Byte headers

An Analogy

1910 - 1940

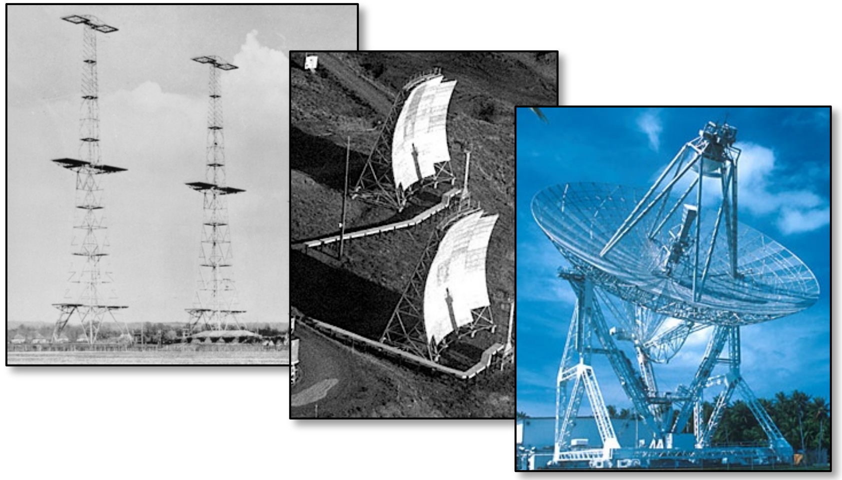

1940 -

1990 - 20XX

?

20XX -

Analogy … cont'd

Sensor Data:

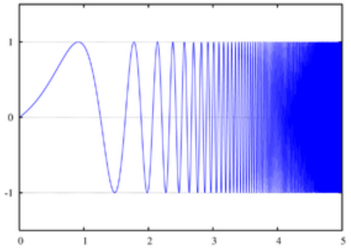

structured dense

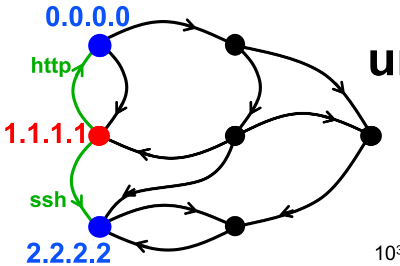

unstructured sparse

Background Distribution:

Gaussian

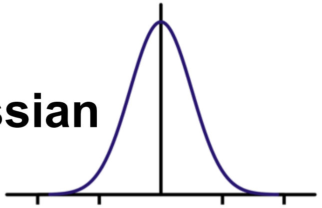

power law

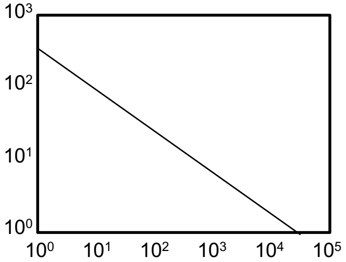

Mathematics:

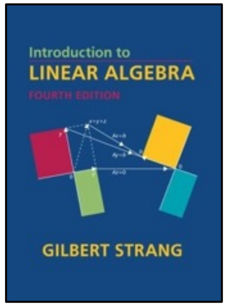

linear algebra

D4M

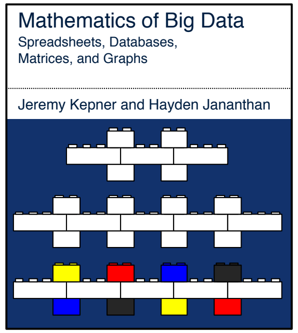

associative array algebra

Compute:

dense vector processor

graph processor

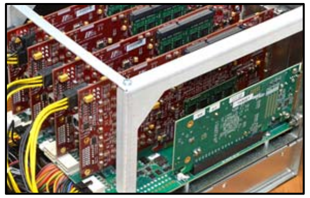

or

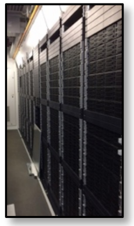

super computer

Major Challenge:
Surpassing the Big Data Analytics Scalability Wall

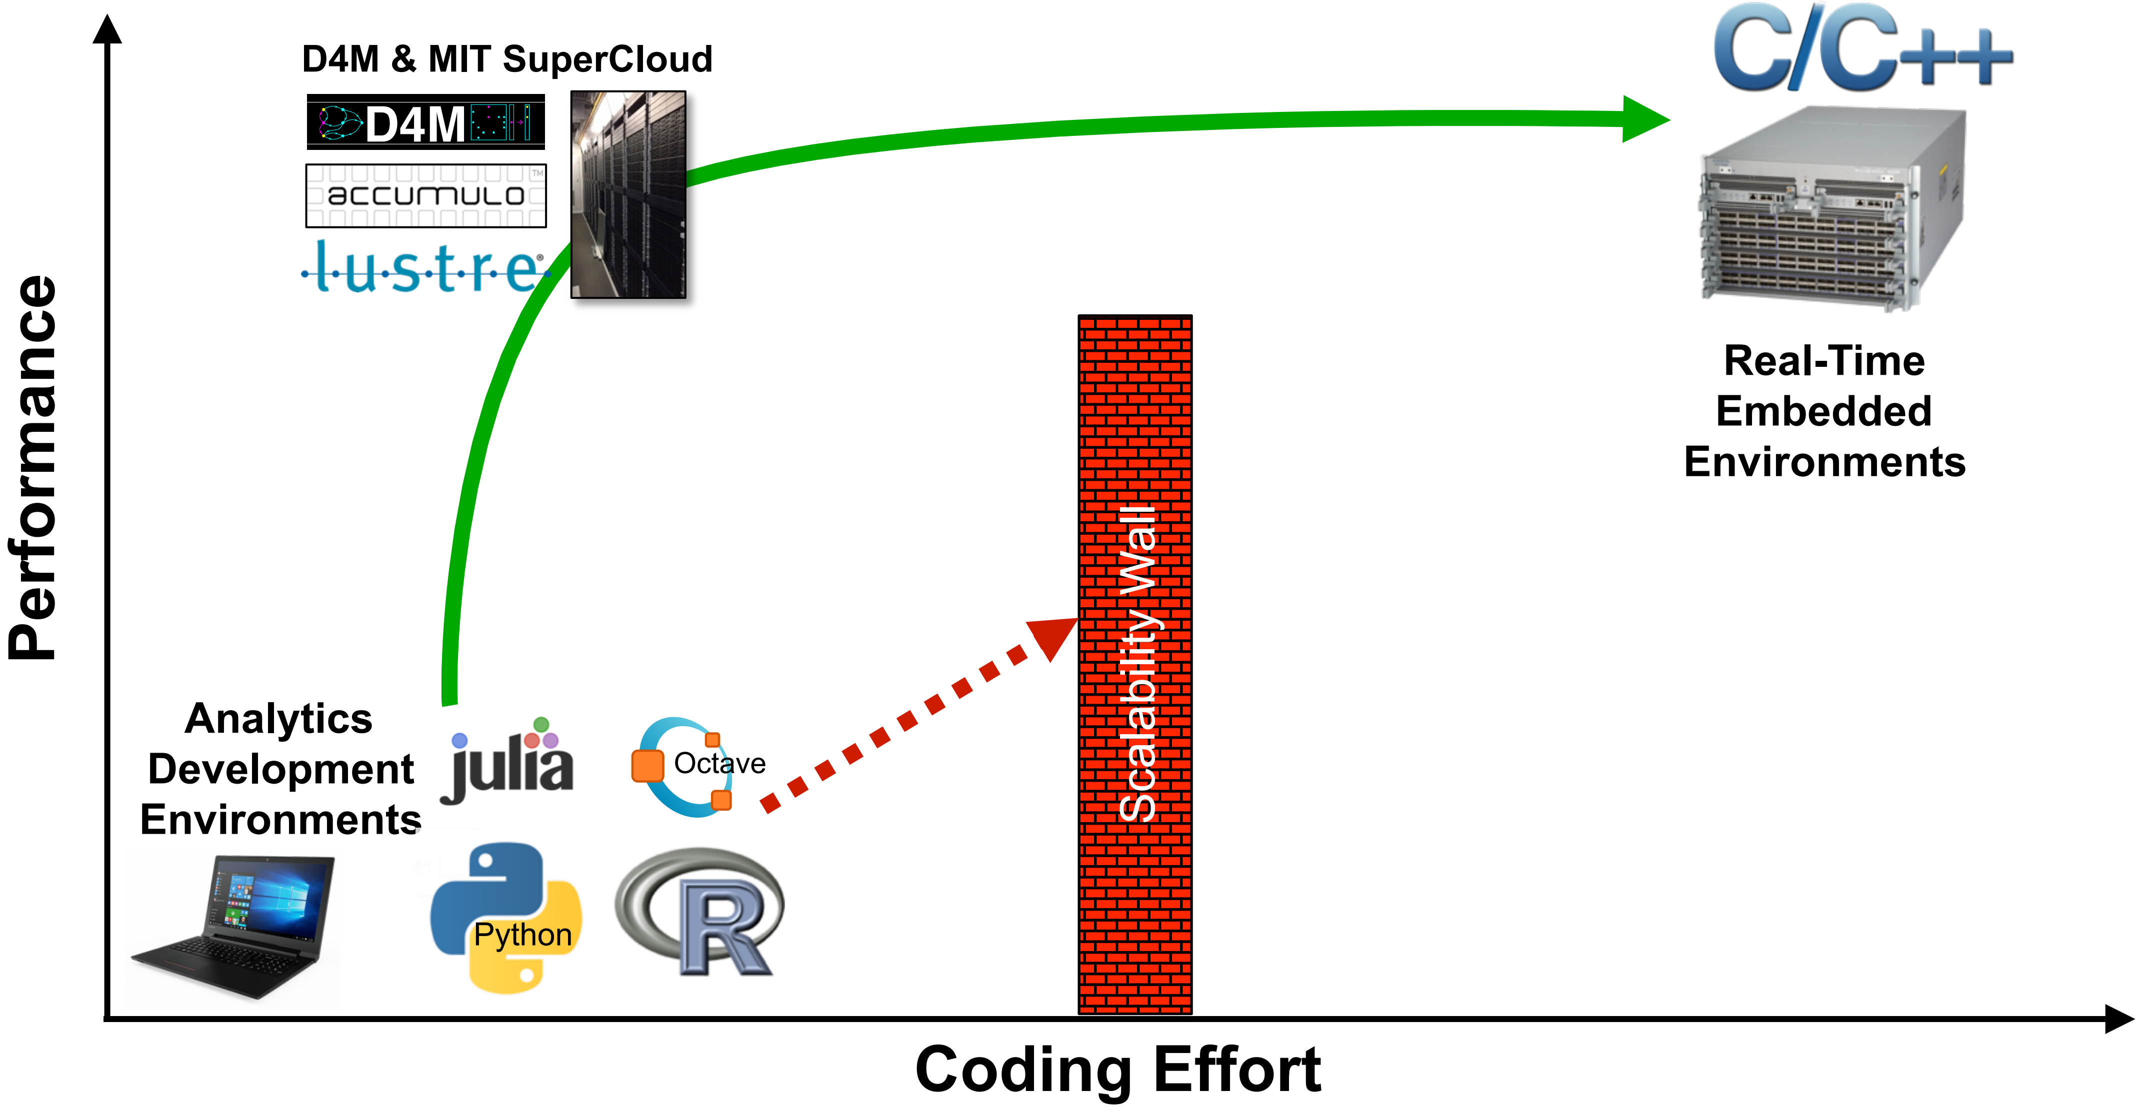

Exemplary Packet Capture Pipeline

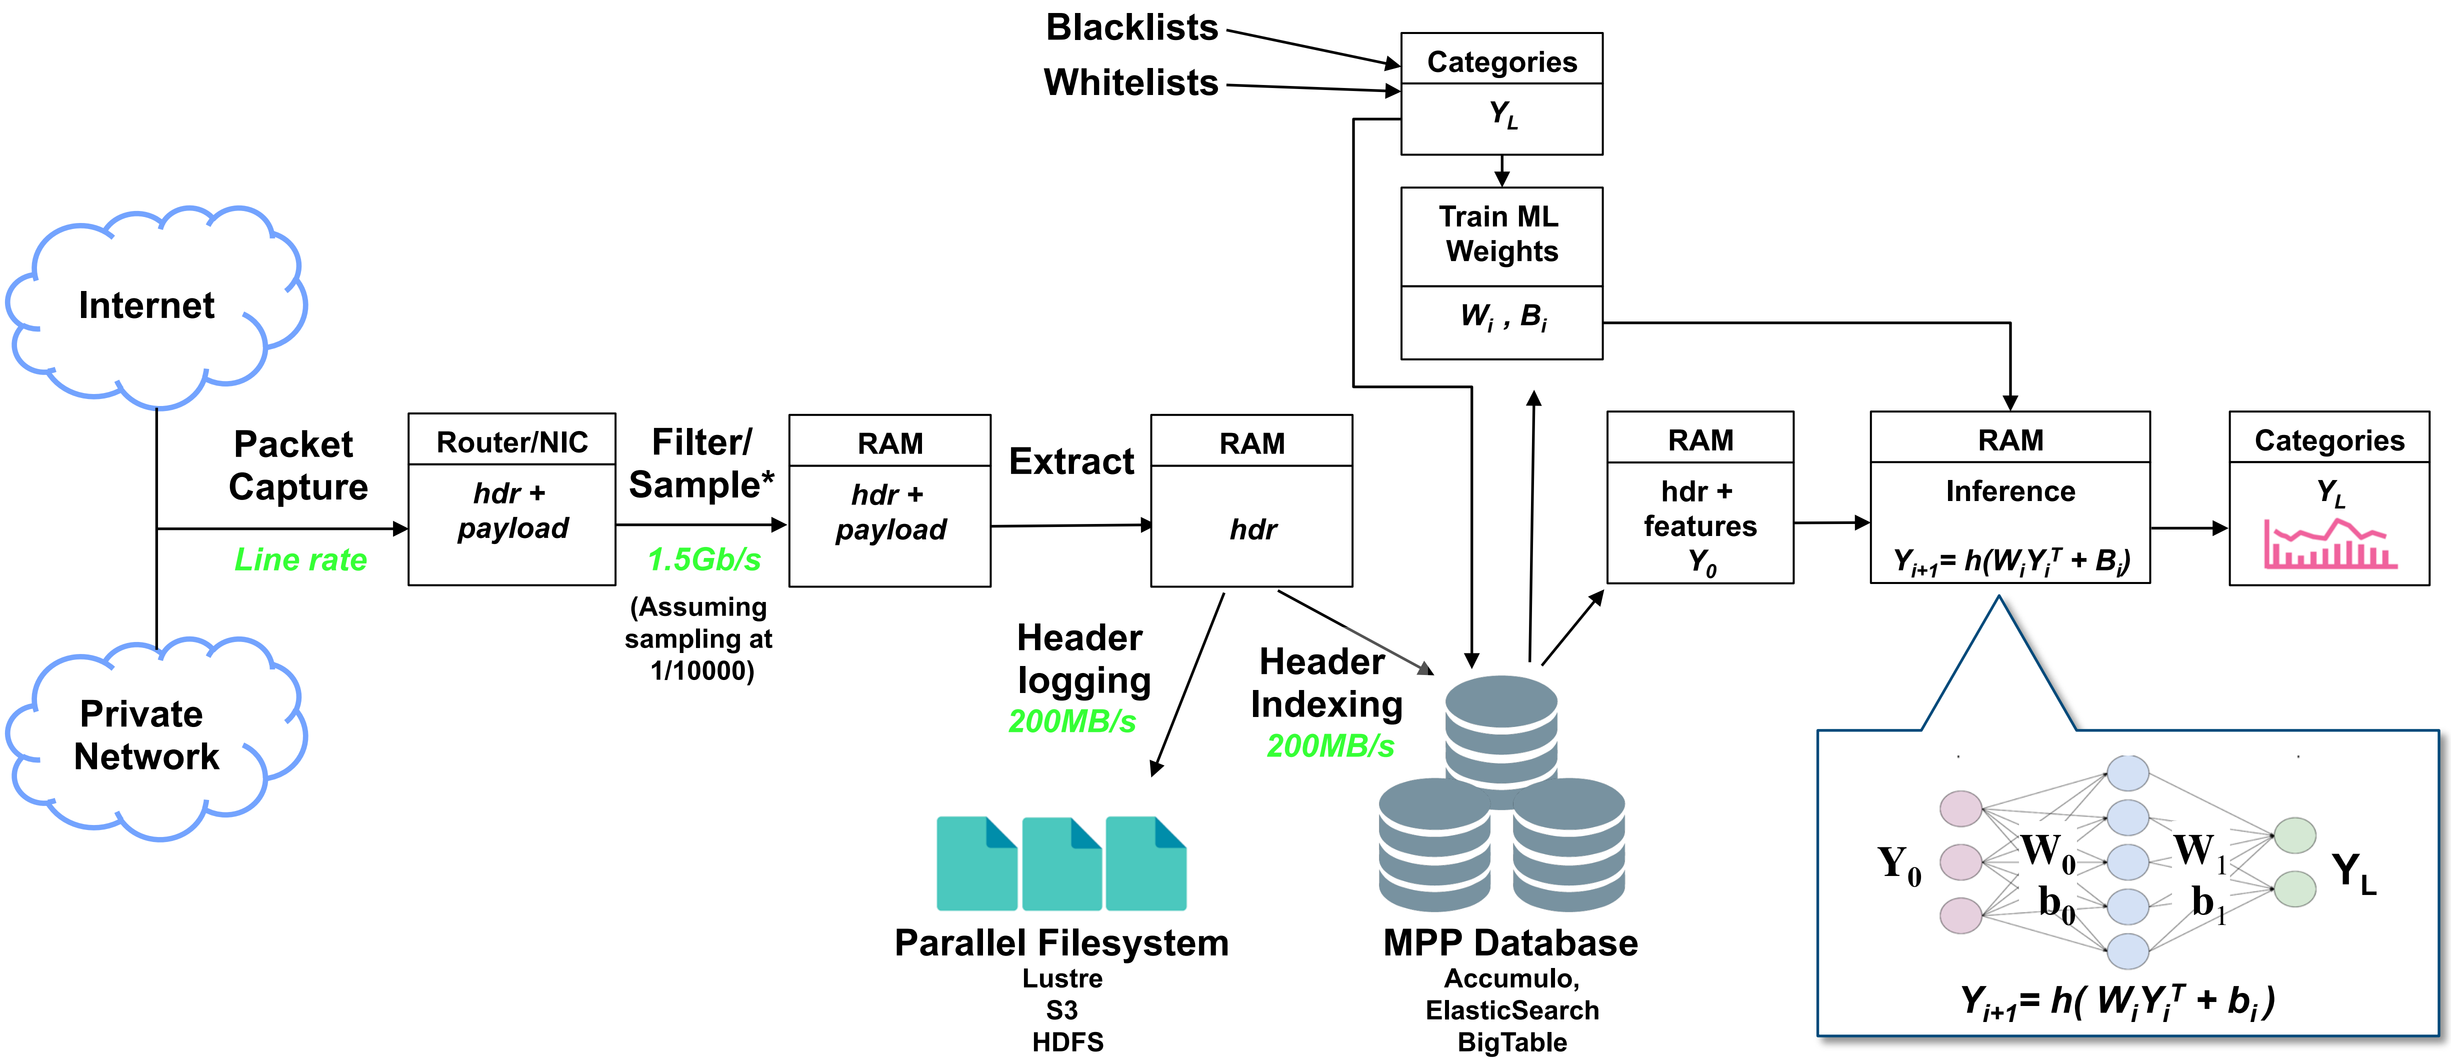

MPP=Massively Parallel Processing
ML= Machine Learning

* Gadepally, et al. "Sampling operations on big data." *Asilomar Signals, Systems and Computers 2015*

Outline

- Introduction
- **D4M and MIT SuperCloud**
- Cyber Network Analysis
- Signal Processing on Networks
- Summary

What is D4M?

- The Dynamic Distributed Dimensional Data Model
  - Support for mathematical foundation – associative arrays
  - A schema to represent most unstructured data as associative arrays
  - Library of software tools to connect with variety of databases such as Apache Accumulo, SciDB, mySQL, PostgreSQL, …
- Software tools currently implemented in MATLAB/ Octave, Julia and Python*
- Connect to databases via JDBC (relational), SHIM (SciDB) or custom Java API (Accumulo)

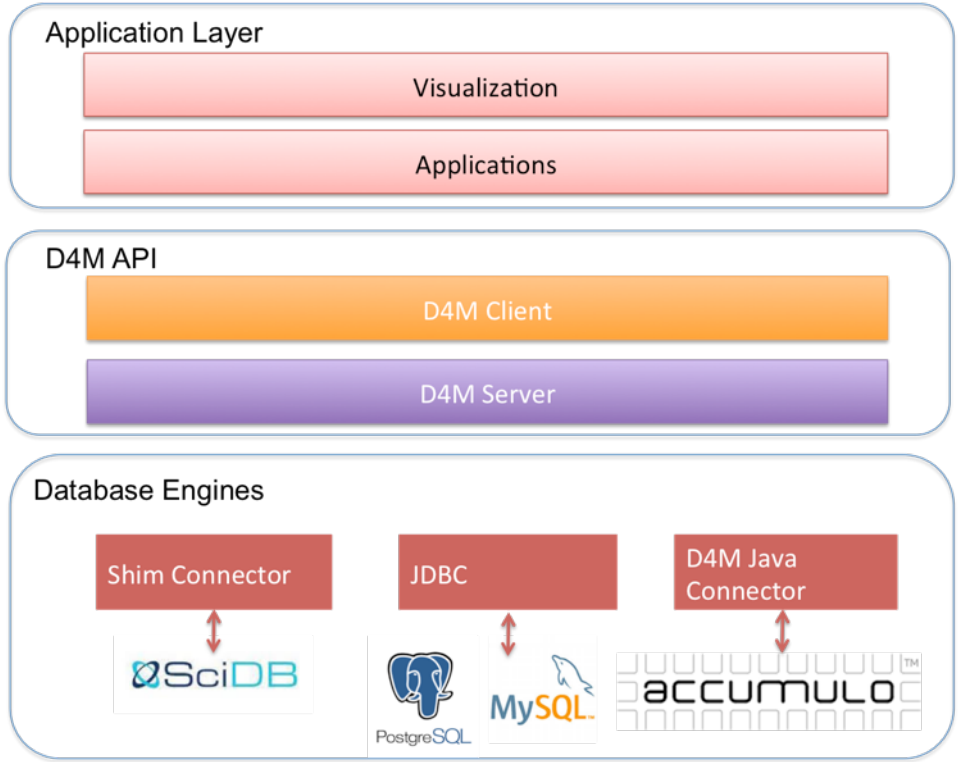

* In development

D4M Associative Arrays

SQL Set Operations	NoSQL Graph Operations	NewSQL Linear Algebra

	src	link	dest
001	1.1.1.1	http	0.0.0.0
002	0.0.0.0	udp	1.1.1.1
003	1.1.1.1	ssh	2.2.2.2

SELECT 'src' FROM T
WHERE 'dest=1.1.1.1'

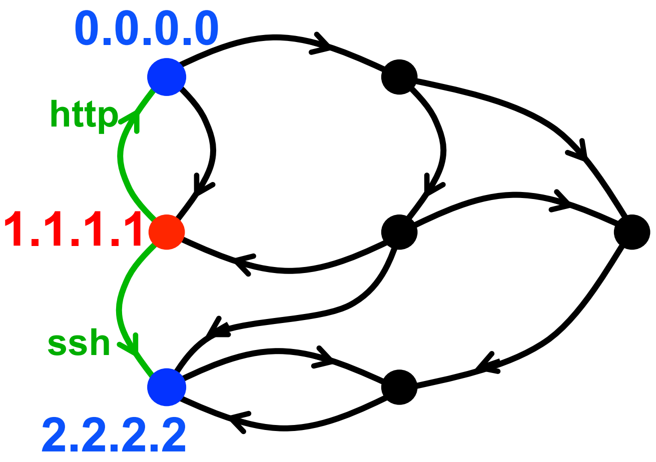

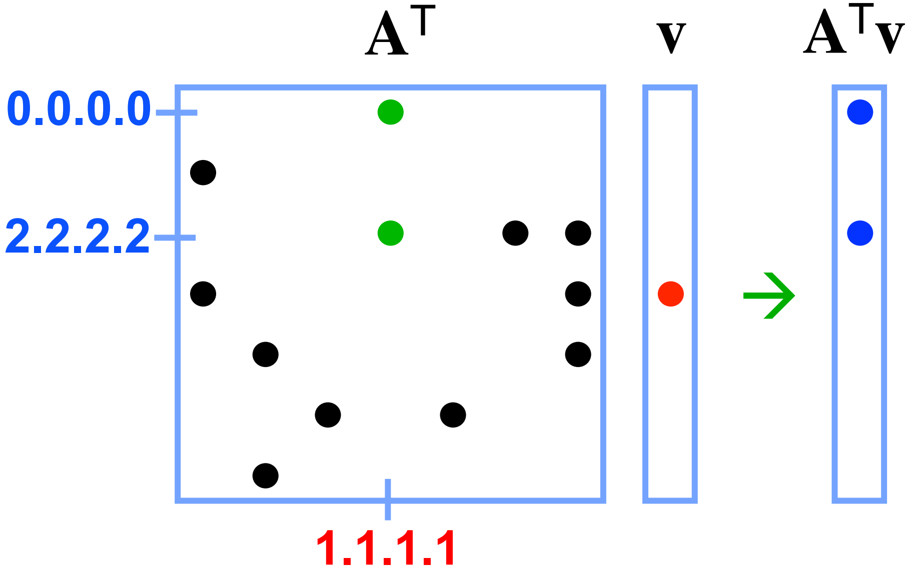

Operation: finding Alice's nearest neighbors

Associative Array Algebra Provides a Unified Mathematics for SQL, NoSQL, NewSQL

$A = \mathbb{S}^{N \times M}(k_1, k_2, v, \oplus)$ $\quad (k_1, k_2, v) = A$ $\quad C = A^T$ $\quad C = A \oplus B$ $\quad C = A \otimes C$ $\quad C = A\,B = A \oplus.\otimes B$

Operations in all representations are equivalent and are linear systems

D4M = Dynamic Distributed Dimensional Data Model (d4m.mit.edu)
Mathematics of Big Data, Kepner & Jananthan, MIT Press 2018

Associative Arrays for Deep Neural Networks -based on the GraphBLAS standard-

- **Increased abstraction at deeper layers**

$$\mathbf{y}_{i+1} = h(\mathbf{W}_i \mathbf{y}_i + \mathbf{b}_i)$$

requires a non-linear function, such as

$$h(\mathbf{y}) = \max(\mathbf{y},0)$$

- **Matrix multiply $\mathbf{W}_i \mathbf{y}_i$ dominates compute**

Remark: can rewrite using GraphBLAS as

$$\mathbf{y}_{i+1} = \mathbf{W}_i \mathbf{y}_i \otimes \mathbf{b}_i \oplus 0$$

where $\oplus = \max()$ and $\otimes = +$

DNN oscillates over two linear semirings

$$S_1 = (\mathbb{R} , + ,x, 0,1)$$

$$S_2 = (\{-\infty \cup \mathbb{R}\},\max,+,-\infty,0)$$

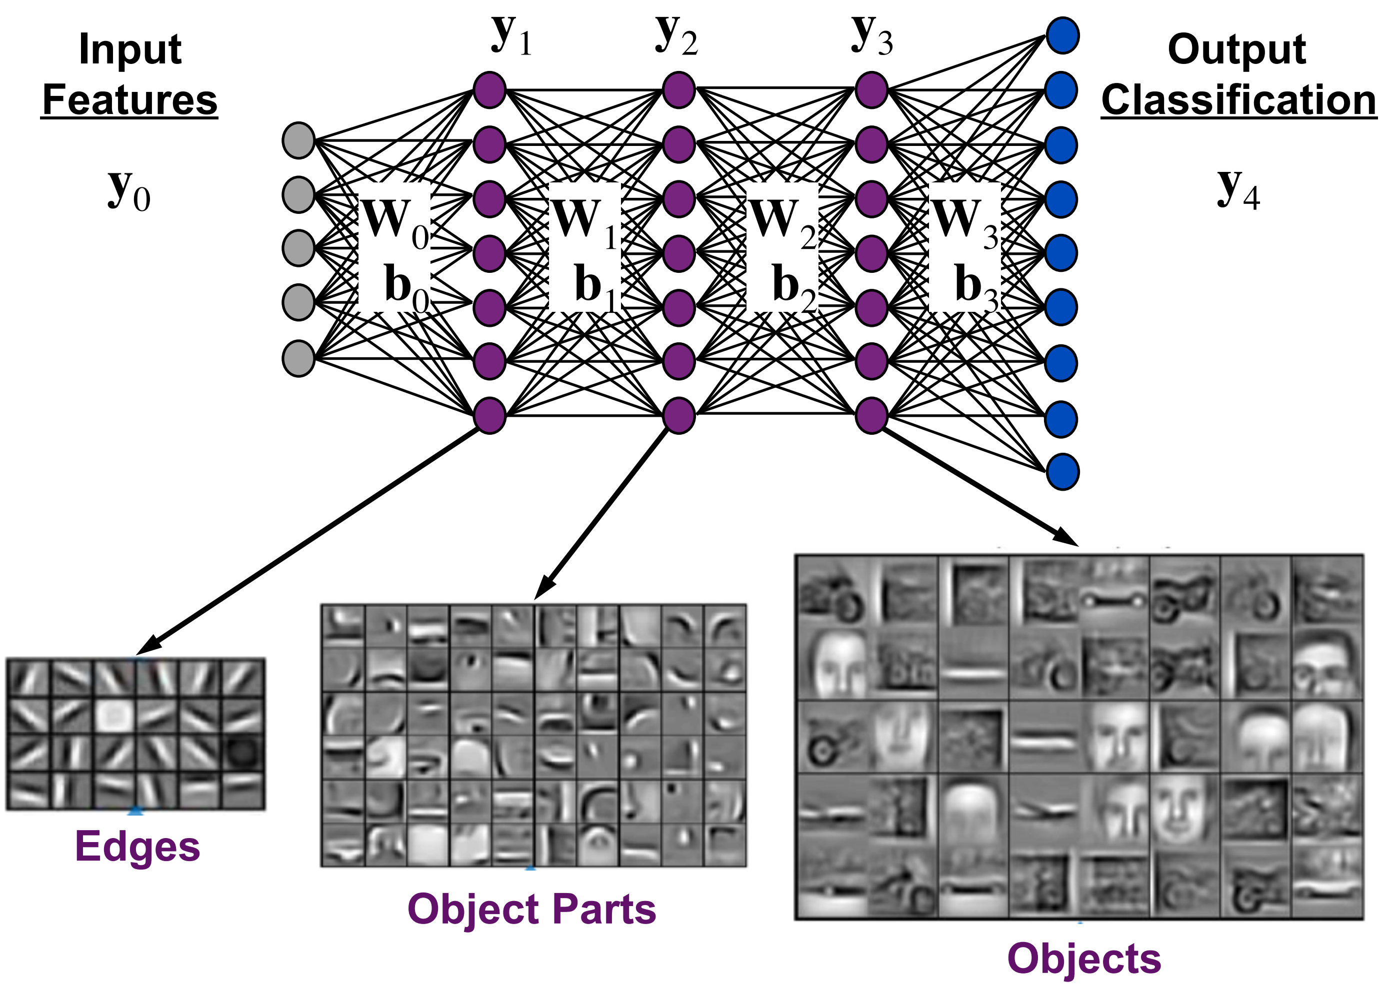

CMD – Computing on Masked Associative Array Data

Computing on Masked Data (CMD) combines concepts from:

- Cryptography
- Advanced database technologies

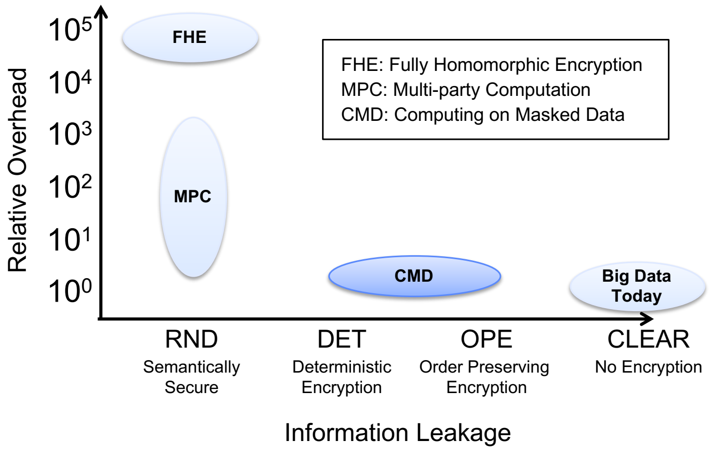

Design criteria:

- Performance
- Security

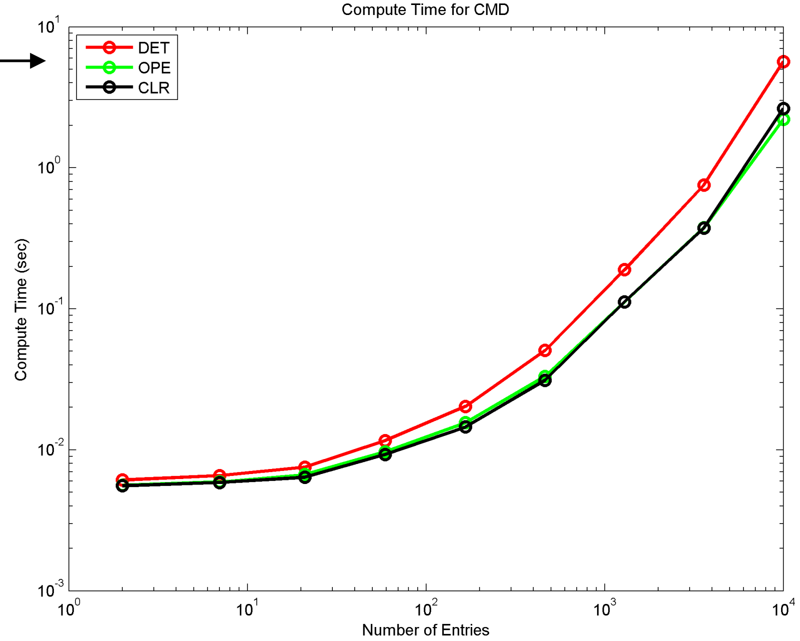

Raw source data

log_id	src_ip	dest_ip
001	128.0.0.1	208.29.69.138
002	192.168.1.2	157.166.255.18
003	128.0.0.1	74.125.224.72

Using CMD

Masked associative array

	BGDJBEAB…	PJDMJPCGG…	RSTPWRQQI…	SWVUZZVZJ…
EQKRP…	SKASEMIC		SKASEMIC	
BYZZO…		SKASEMIC		
CJYTG…	SKASEMIC		SKASEMIC	SKASEMIC

DET=Deterministic Encryption
OPE=Order Preserving Encryption

MIT SuperCloud
-petascale infrastructure-

	TX-Green Upgrade
Processor	Intel Xeon 64 Core
Total Cores	41,472
Peak Petaflops	1.724
Top500 Petaflops	1.025 (measured)
Total Terabytes	124
Network Link	Intel OmniPath 25 GB/s

Based on Nov 2016 Top500.org list

- #1 in Massachusetts
- #1 in New England
- #2 in the Northeast
- #3 at a US University
- #3 at a University in the Western Hemisphere
- #43 in the United States
- #106 in the World

Only zero carbon emission system in Top500

Manycore system sustains MIT's leadership position in interactive supercomputing

- Compatible with all existing LLSC software
- Provides processing (6x) and bandwidth (20x) for big data and machine learning applications

Interactive High Performance Machine Learning (HPML)
- Interactive Launch on 32,000+ Cores -

- **Machine Learning models require**
  - **High level programming environments for building models**
  - **Rapid interaction with analyst**
- **Standard approaches take minutes to hours to launch on thousands of cores**
- **MIT SuperCloud optimizes every aspect of HPML system to enable**
  - **Launching hundreds of machine learning models in seconds**
  - **32,000+ cores (512 64-core Xeon nodes)**
  - **Truly interactive machine learning**

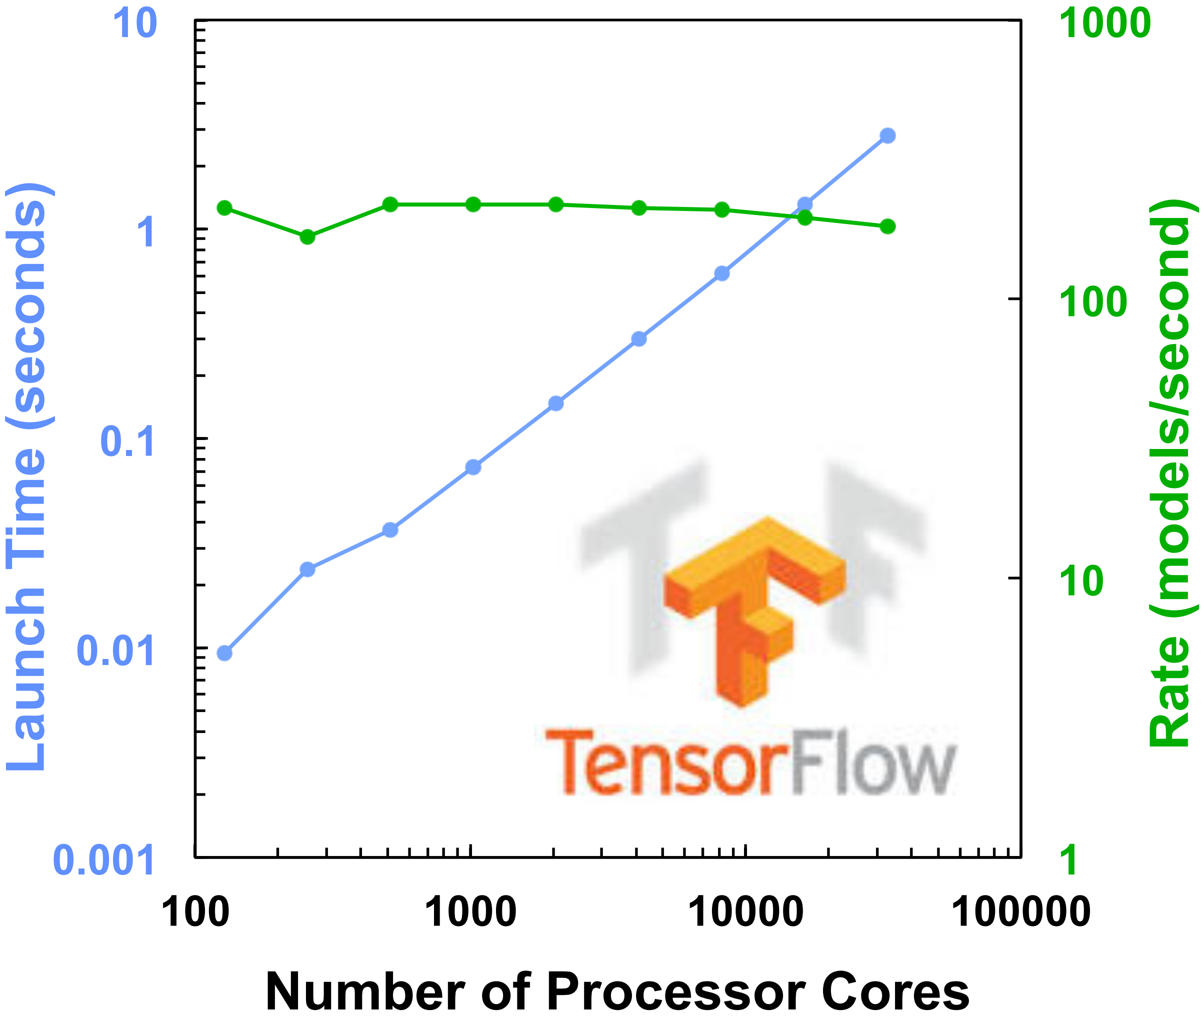

High Performance HPDA Launch on 32,000+ Cores

- **High Performance Data Analysis (HPDA) requires**
  - High level programming environments
  - Rapid interaction and fast turnaround
- **Standard approaches take minutes to hours to launch on thousands of cores**
- **MIT SuperCloud optimizes every aspect of HPDA system to enable**
  - Launching 32,000+ HPDA environments
  - 32,000+ cores (512 x 64-core Xeon nodes)
  - Launched in 25 seconds
  - 1000+ launches/second
  - 500x faster than standard approaches[1]
  - Truly interactive supercomputing

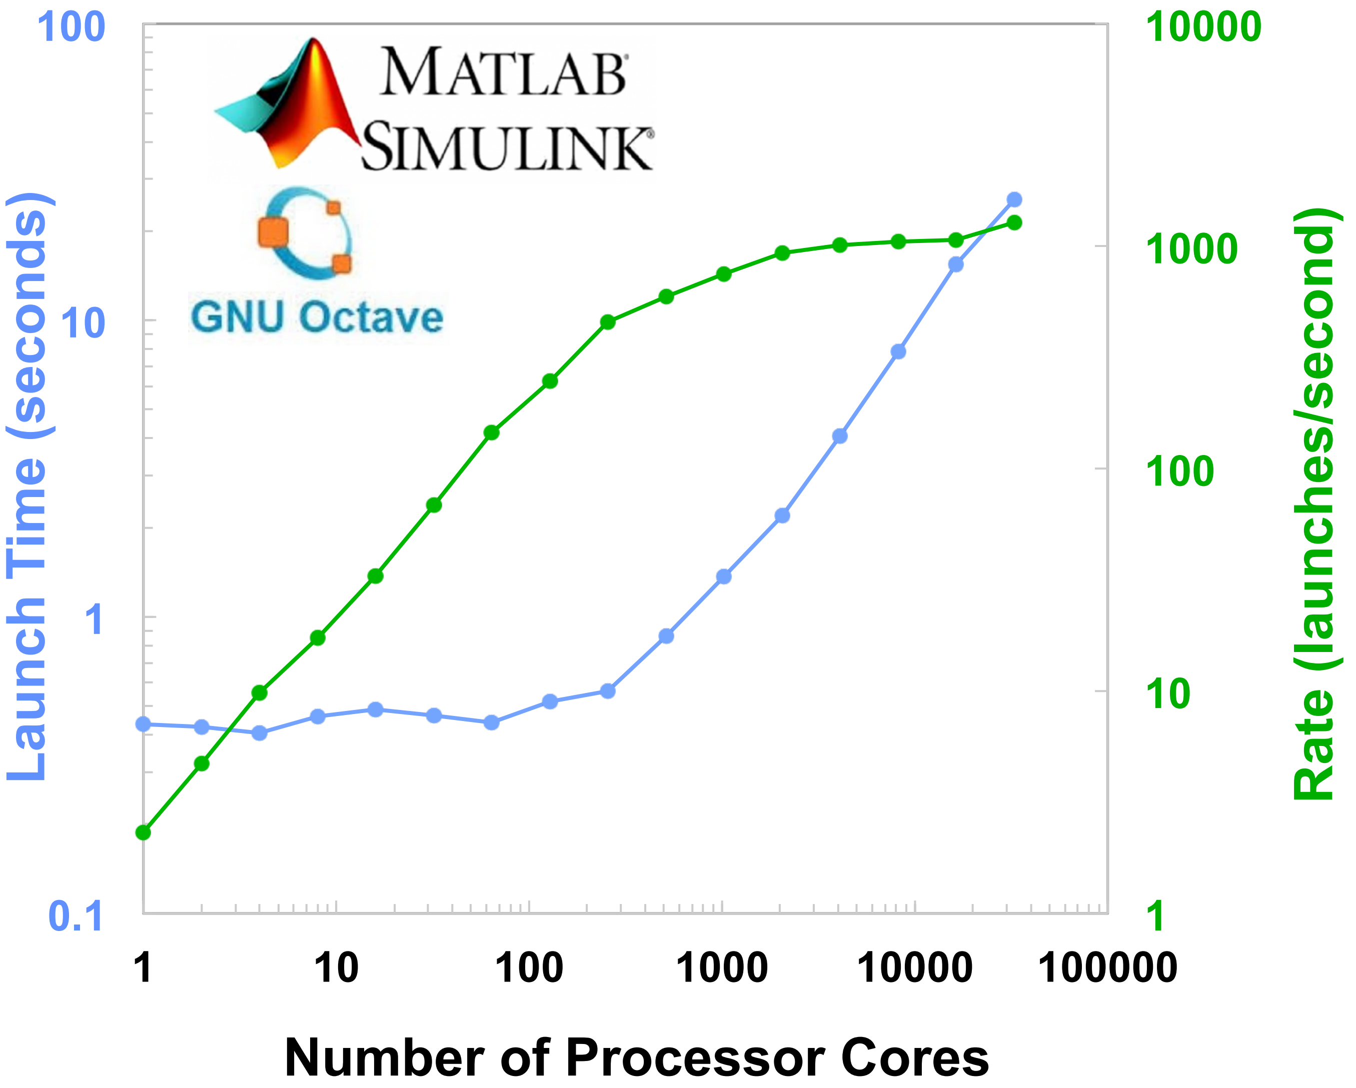

[1] *Scalable System Scheduling for HPC and Big Data*, Reuther et al, Journal of Parallel and Distributed Computing, 2017

SuperCloud Analytics Environment

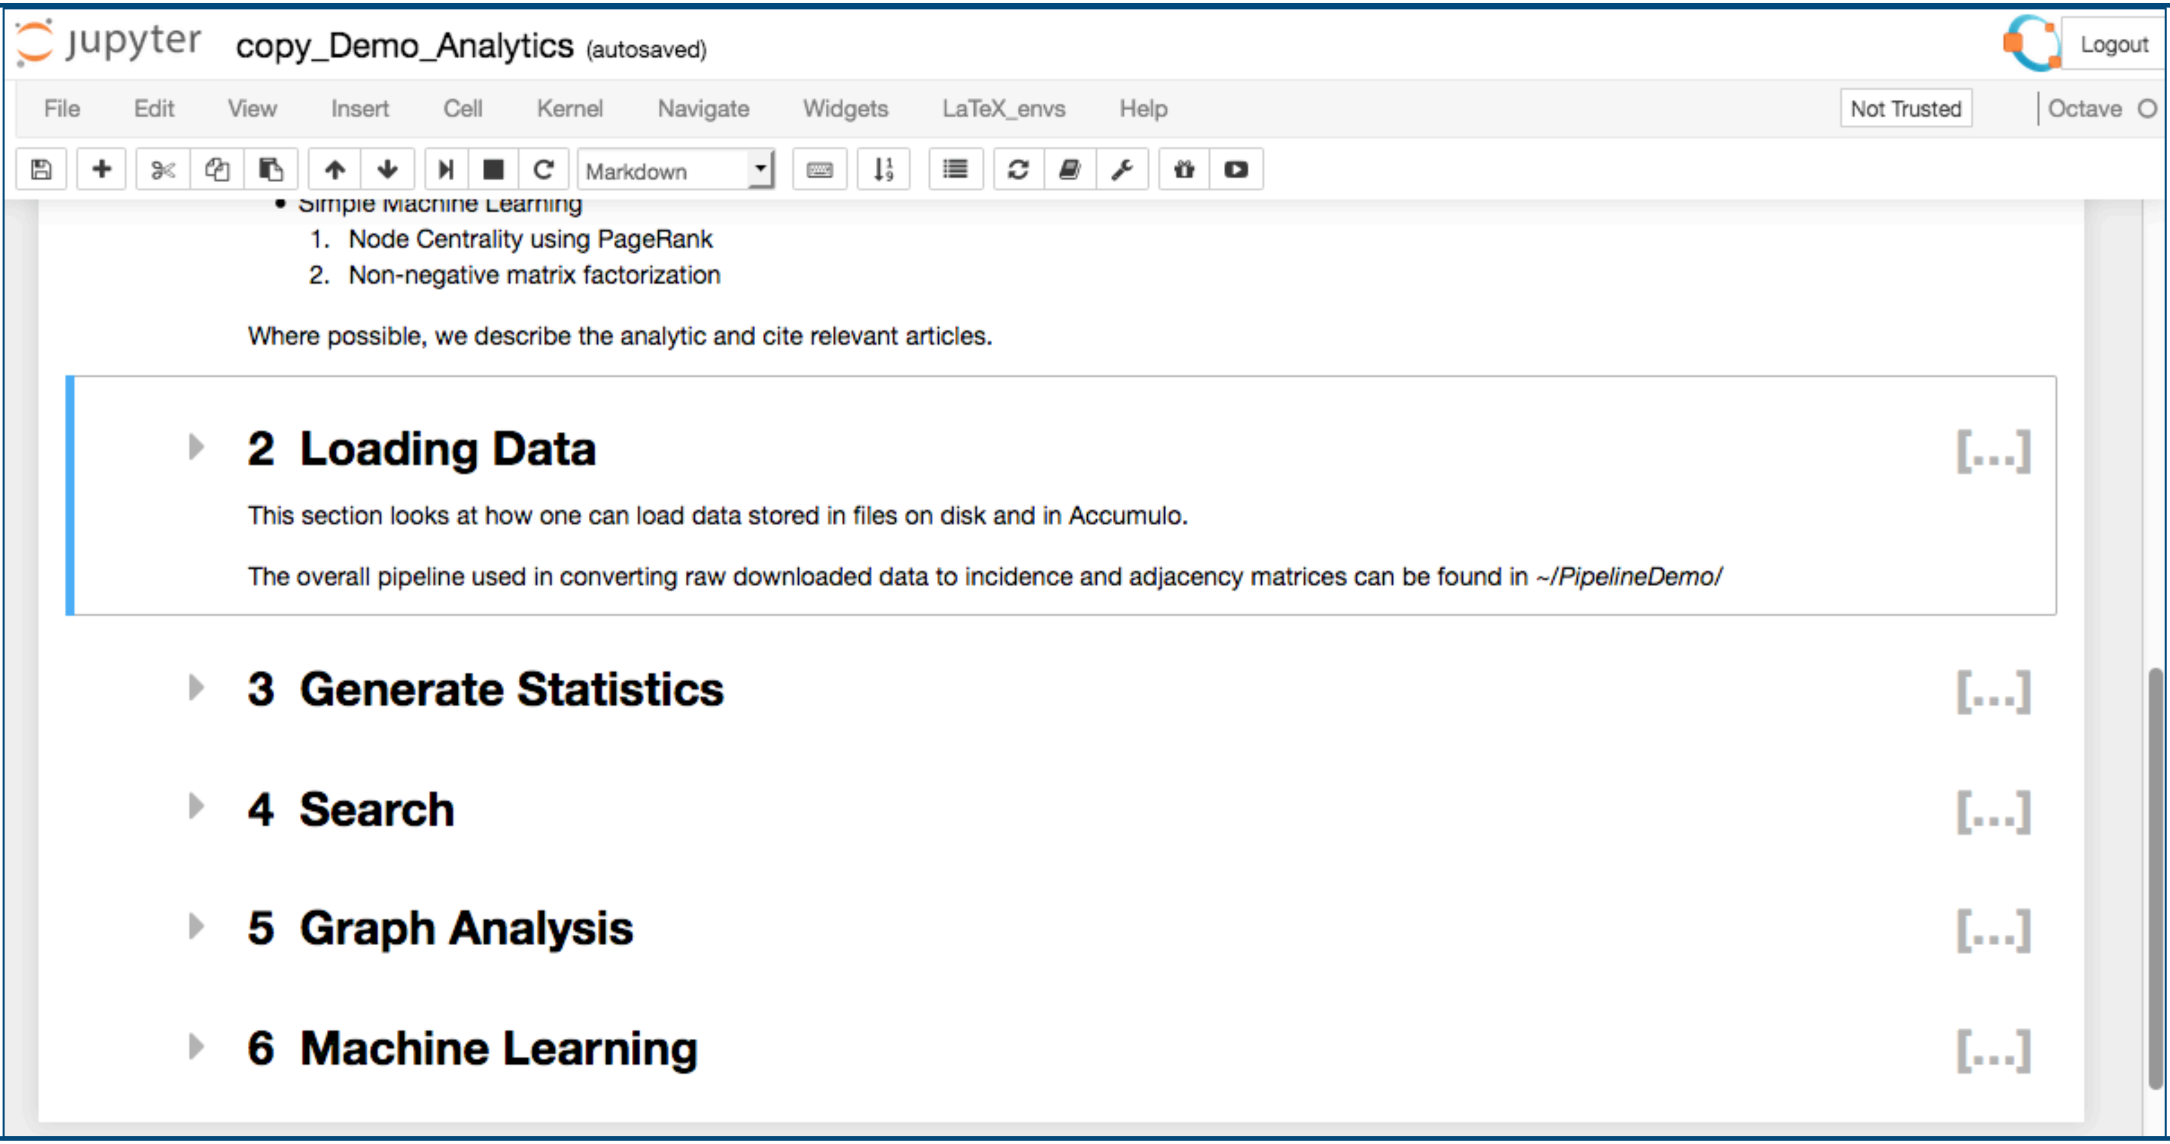

Outline

- Introduction
- D4M and MIT SuperCloud
- Cyber Network Analysis
- Signal Processing on Networks
- Summary

Data: Largest Public Internet Packet Capture (PCAP)

- **Measurement and Analysis of Wide-area Internet (MAWI) working group**
  - Day-in-the-Life of the Internet 2015 (http://mawi.wide.ad.jp/mawi/ditl/ditl2015)
  - Day-in-the-Life of the Internet 2017 (http://mawi.wide.ad.jp/mawi/ditl/ditl2017)
  - 2x48=96 hours of 1 Gigabit packet capture (PCAP) headers collected in Tokyo
  - 0.7 TB compressed; 20 TB in analysts friendly form
    - Normalized, sorted, indexed, and read optimized
  - IP addressed deterministically anonymized *within* each collect
    - Network analysis is still valid

- **Center for Applied Internet Data Analysis (CAIDA) working group**
  - 2016 (https://data.caida.org/datasets/passive-2016/equinix-chicago)
  - 4x1=4 hours of 10 Gigabit packet capture (PCAP) headers collected in Chicago
  - 0.4 TB compressed; 10 TB in analysts friendly form
    - Normalized, sorted, indexed, and read optimized
  - IP addressed deterministically anonymized *within* each collect
    - Network analysis is still valid

Algorithm Development Pipeline

D4M

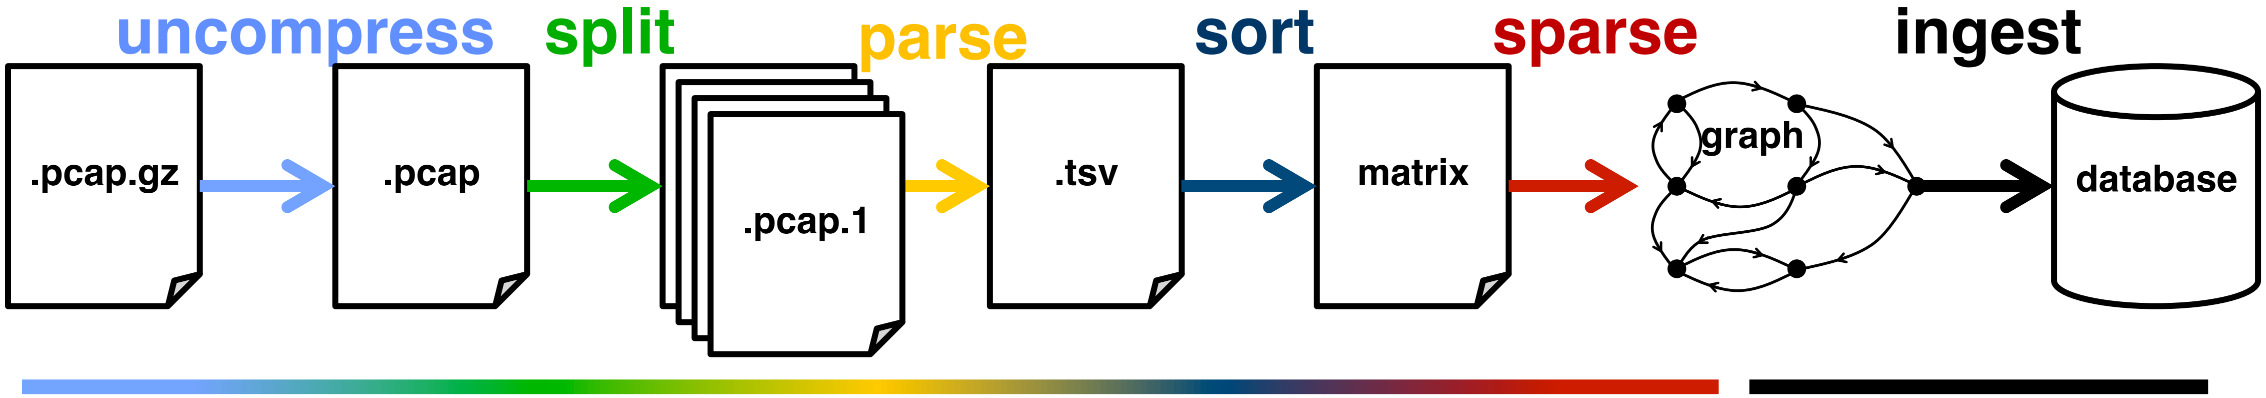

lustre®

accumulo™

Packet Capture Analysis Performance
- Algorithm Development Pipeline on 20,000+ Cores -

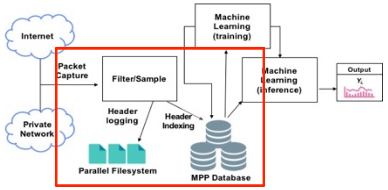

- **Development of novel computer network traffic analytics requires**
  - **High level programming environments**
  - **Massive packet capture (PCAP) data**
  - **Diverse data products for "at scale" algorithm pipeline development**
- **Benchmarked processing 96 hours of Gigabit PCAP[1] data with MIT SuperCloud**
  - **Provides scaling performance for designing real analytic development systems**

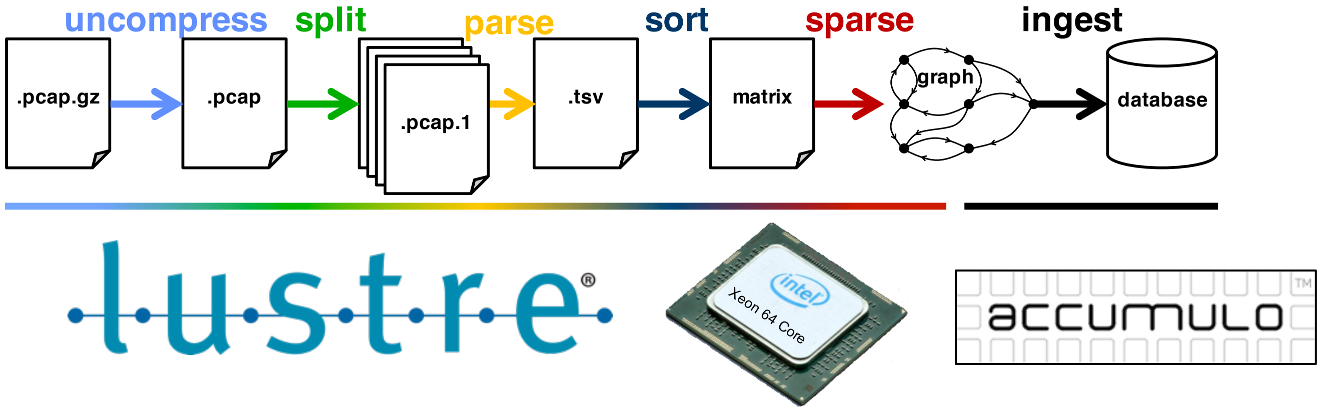

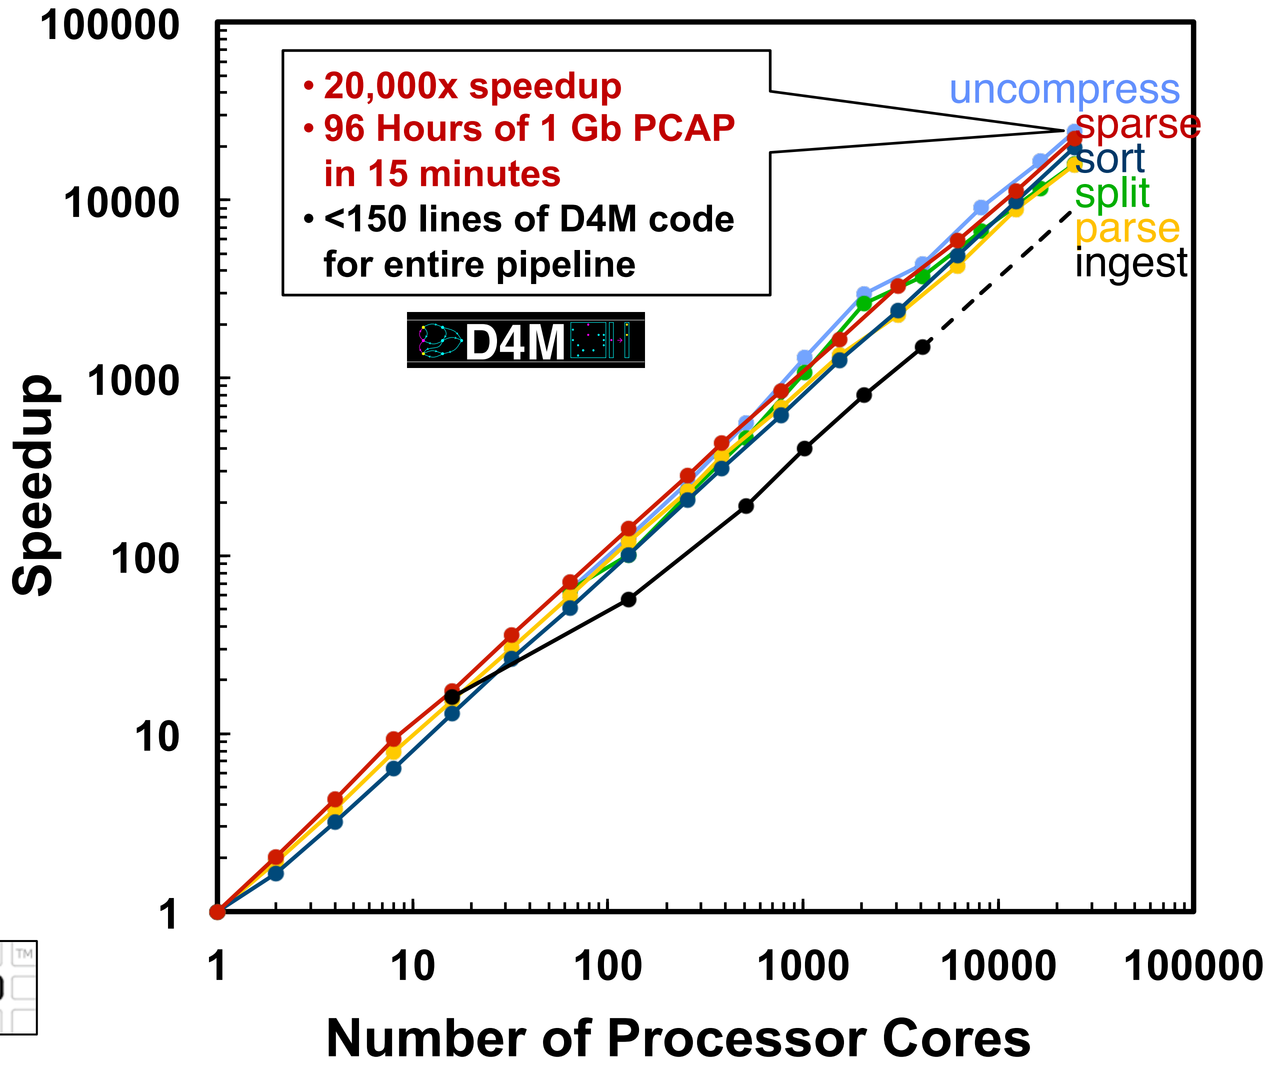

[1]Measurement and Analysis of Wide-area Internet (MAWI) working group

Outline

- Introduction
- MIT SuperCloud
- Hyperscale Analysis Pipeline
- **Signal Processing on Networks**
- Summary

Signal Processing Refresher

Signal Processing on Networks:

- Detecting signal from noise
- Operations on multi-dimensional structured or unstructured data

- Uses background model to separate signal from noise
- Traditionally, dealt with sound, images, video
  - D4M allows us to extend these techniques to unstructured data

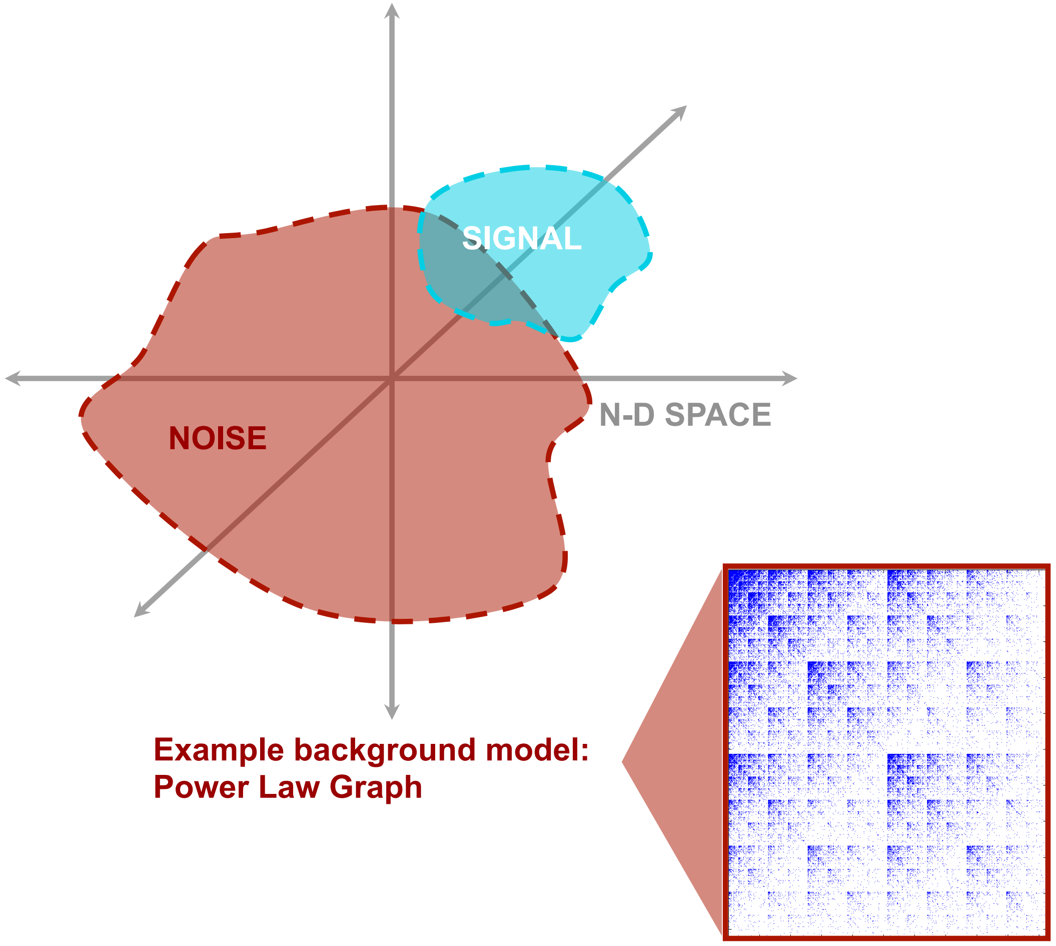

Signal Processing uses a background model to distinguish signal from noise

100,000 & 10,000,000 Edge Equalized Network Streams

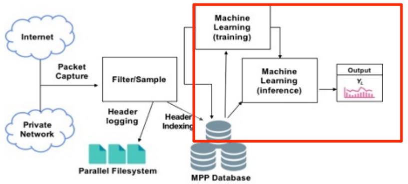

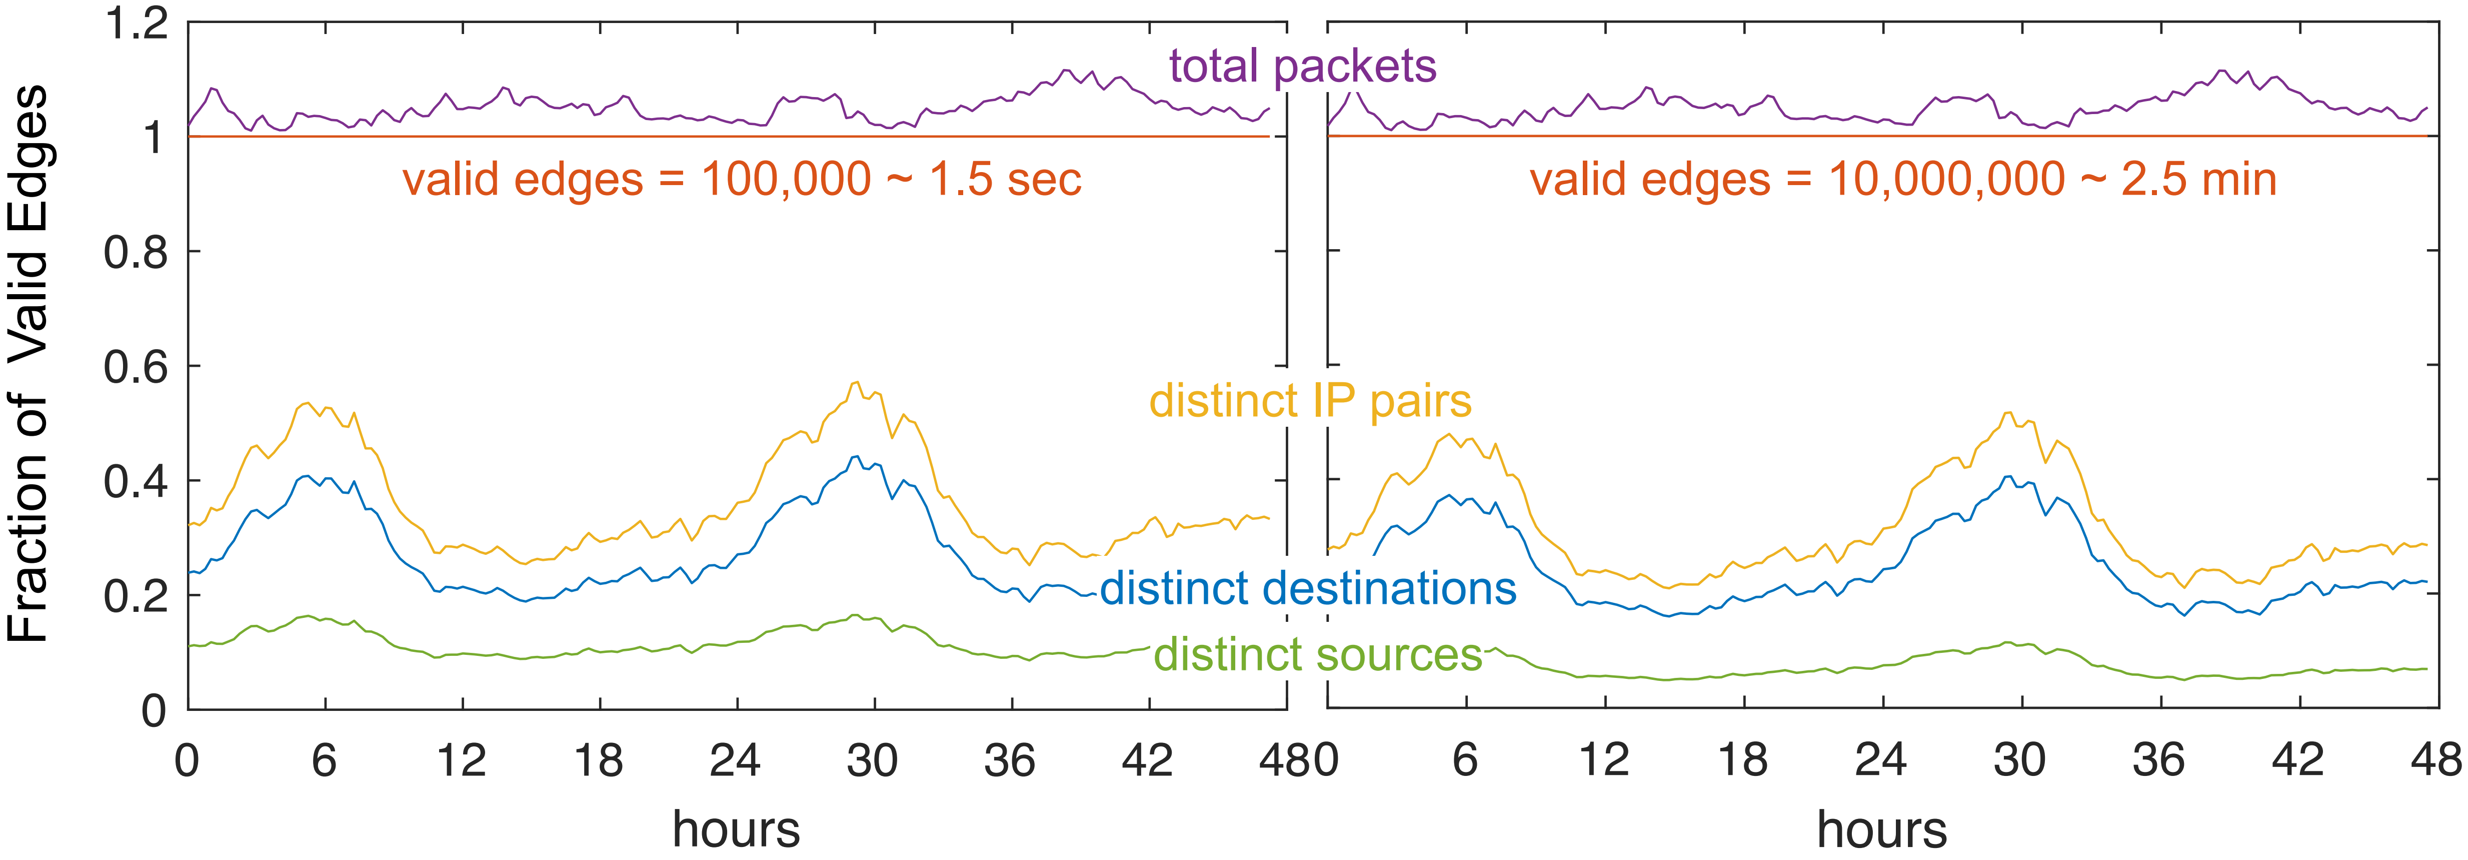

- **Key to signal processing on networks: understanding background behavior**
- **Edge equalization shows clear diurnal pattern independent of edge size**

Background: Five Standard Power Law Networks

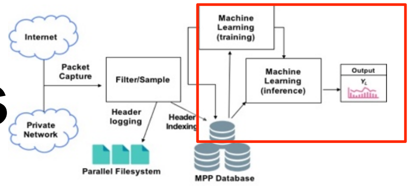

- **Edge equalization, $\log_2$ binning, and probability normalization produces consistent low variance power laws**
- **Clear baseline for anomaly detection**

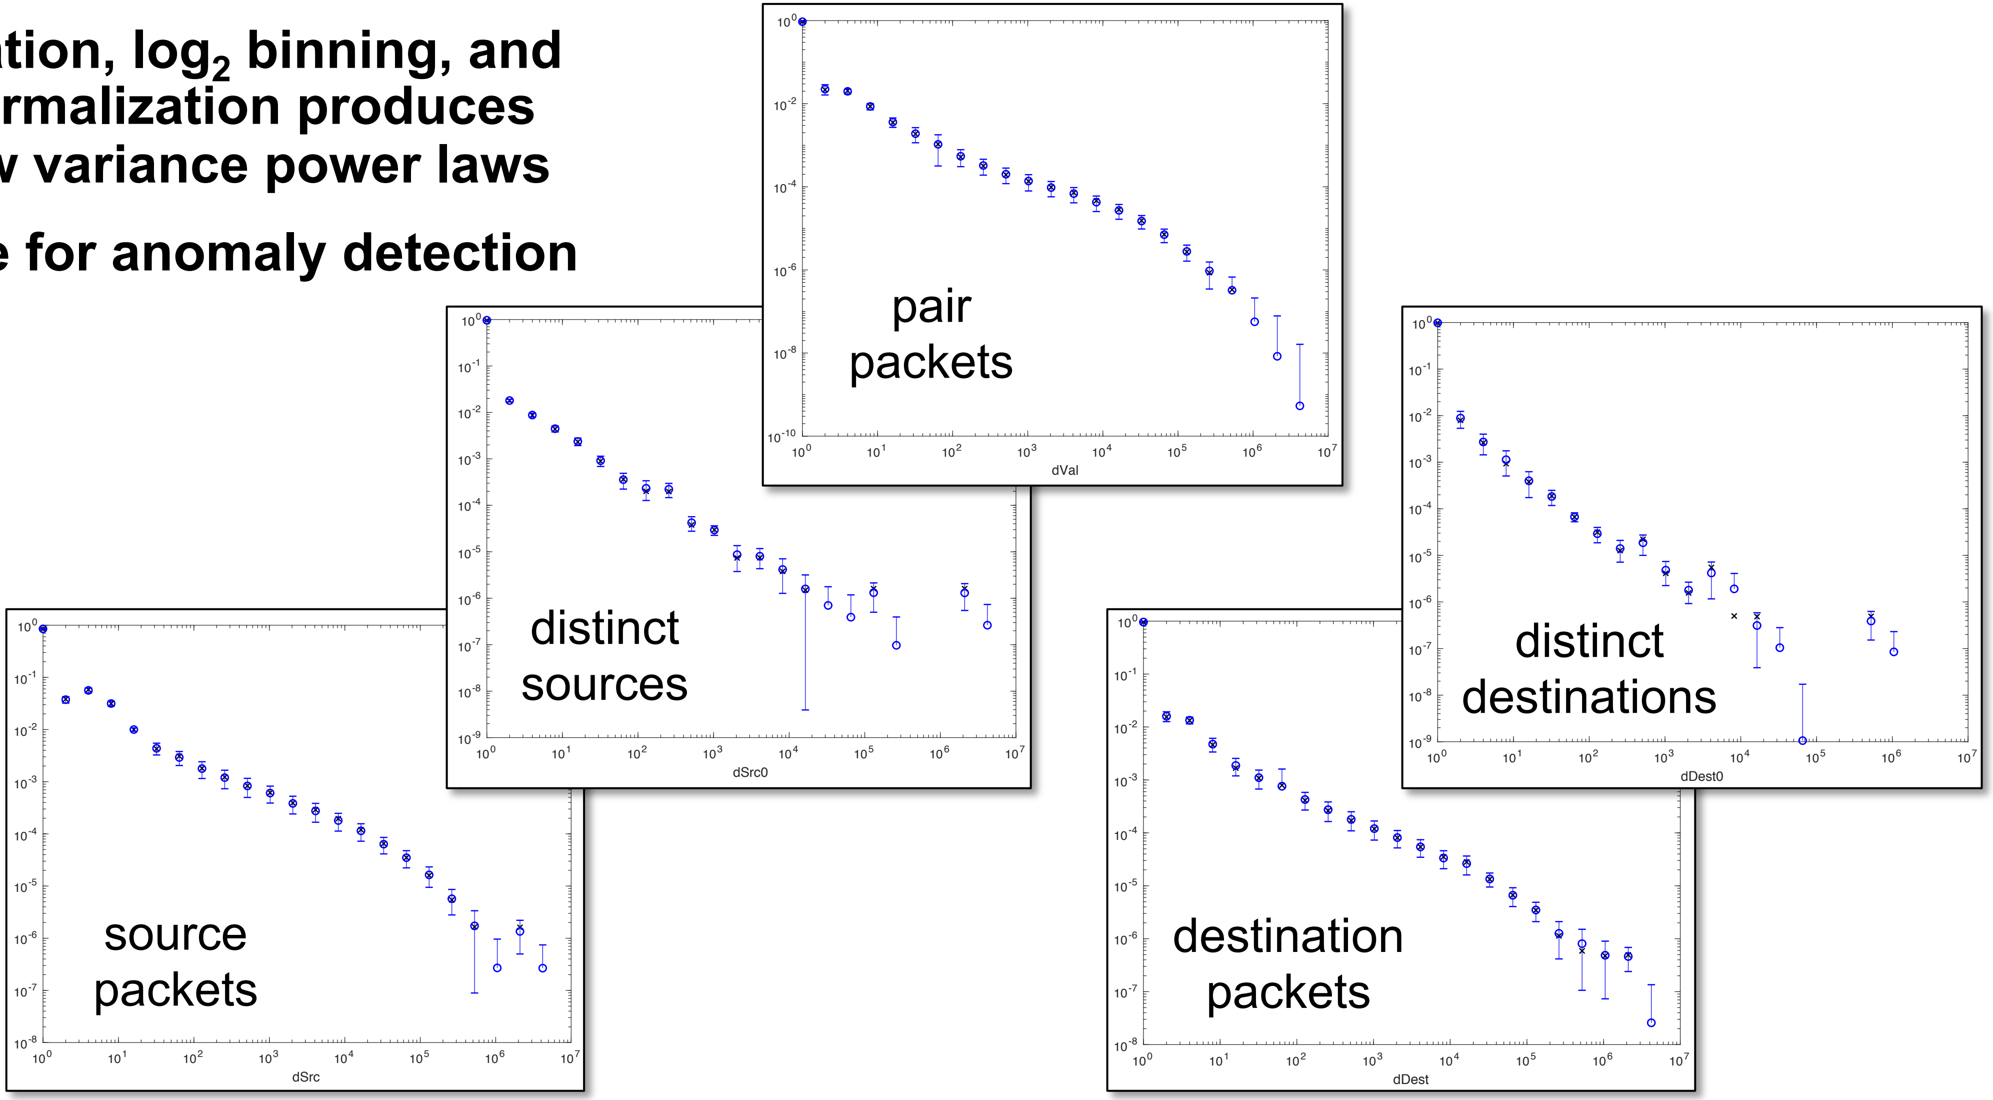

Internet “Tides”

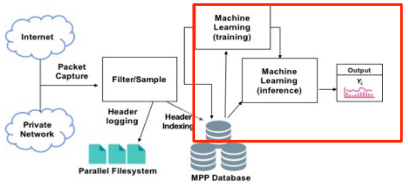

- **Low variance in power law distributions allow tracking over time**
- **Diurnal behavior oscillates between lower and higher fraction of**
  - **Single source nodes**

    vs
  - **Higher sources nodes**

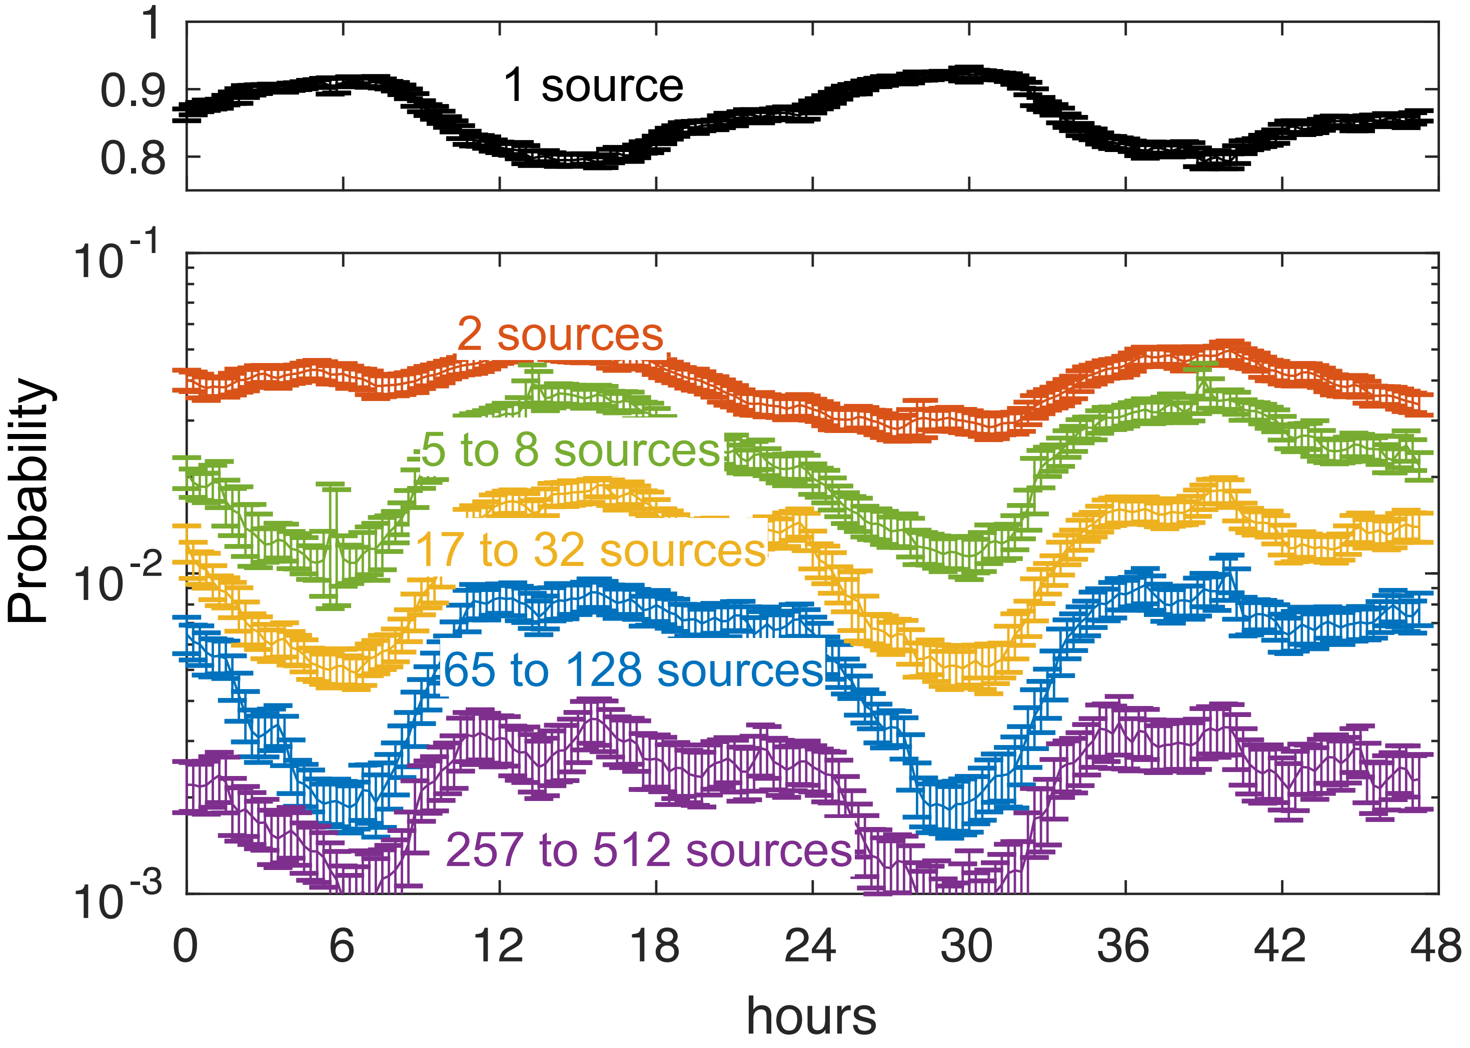

Diurnal Power Law Envelope

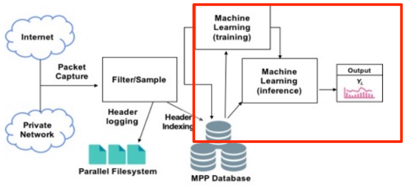

- **Diurnal variations in power law oscillates between two extremes**
- **Clearly defines normal behavior of Internet in this data**

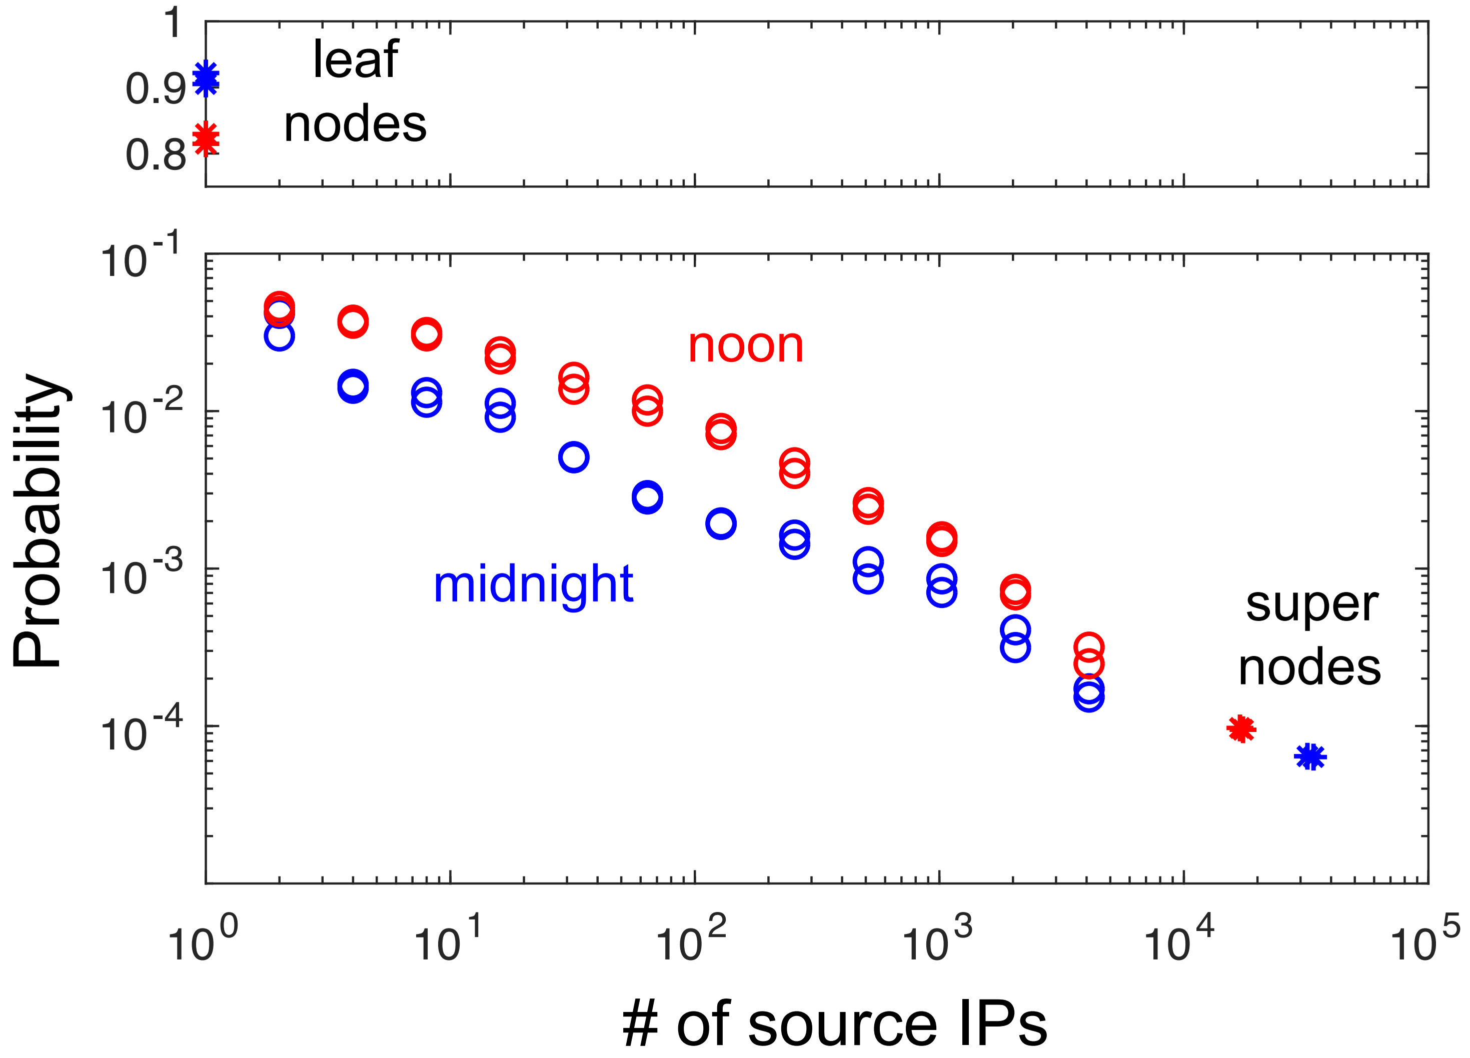

-40 dB Internet Network Anomaly Detection

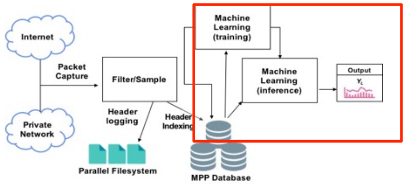

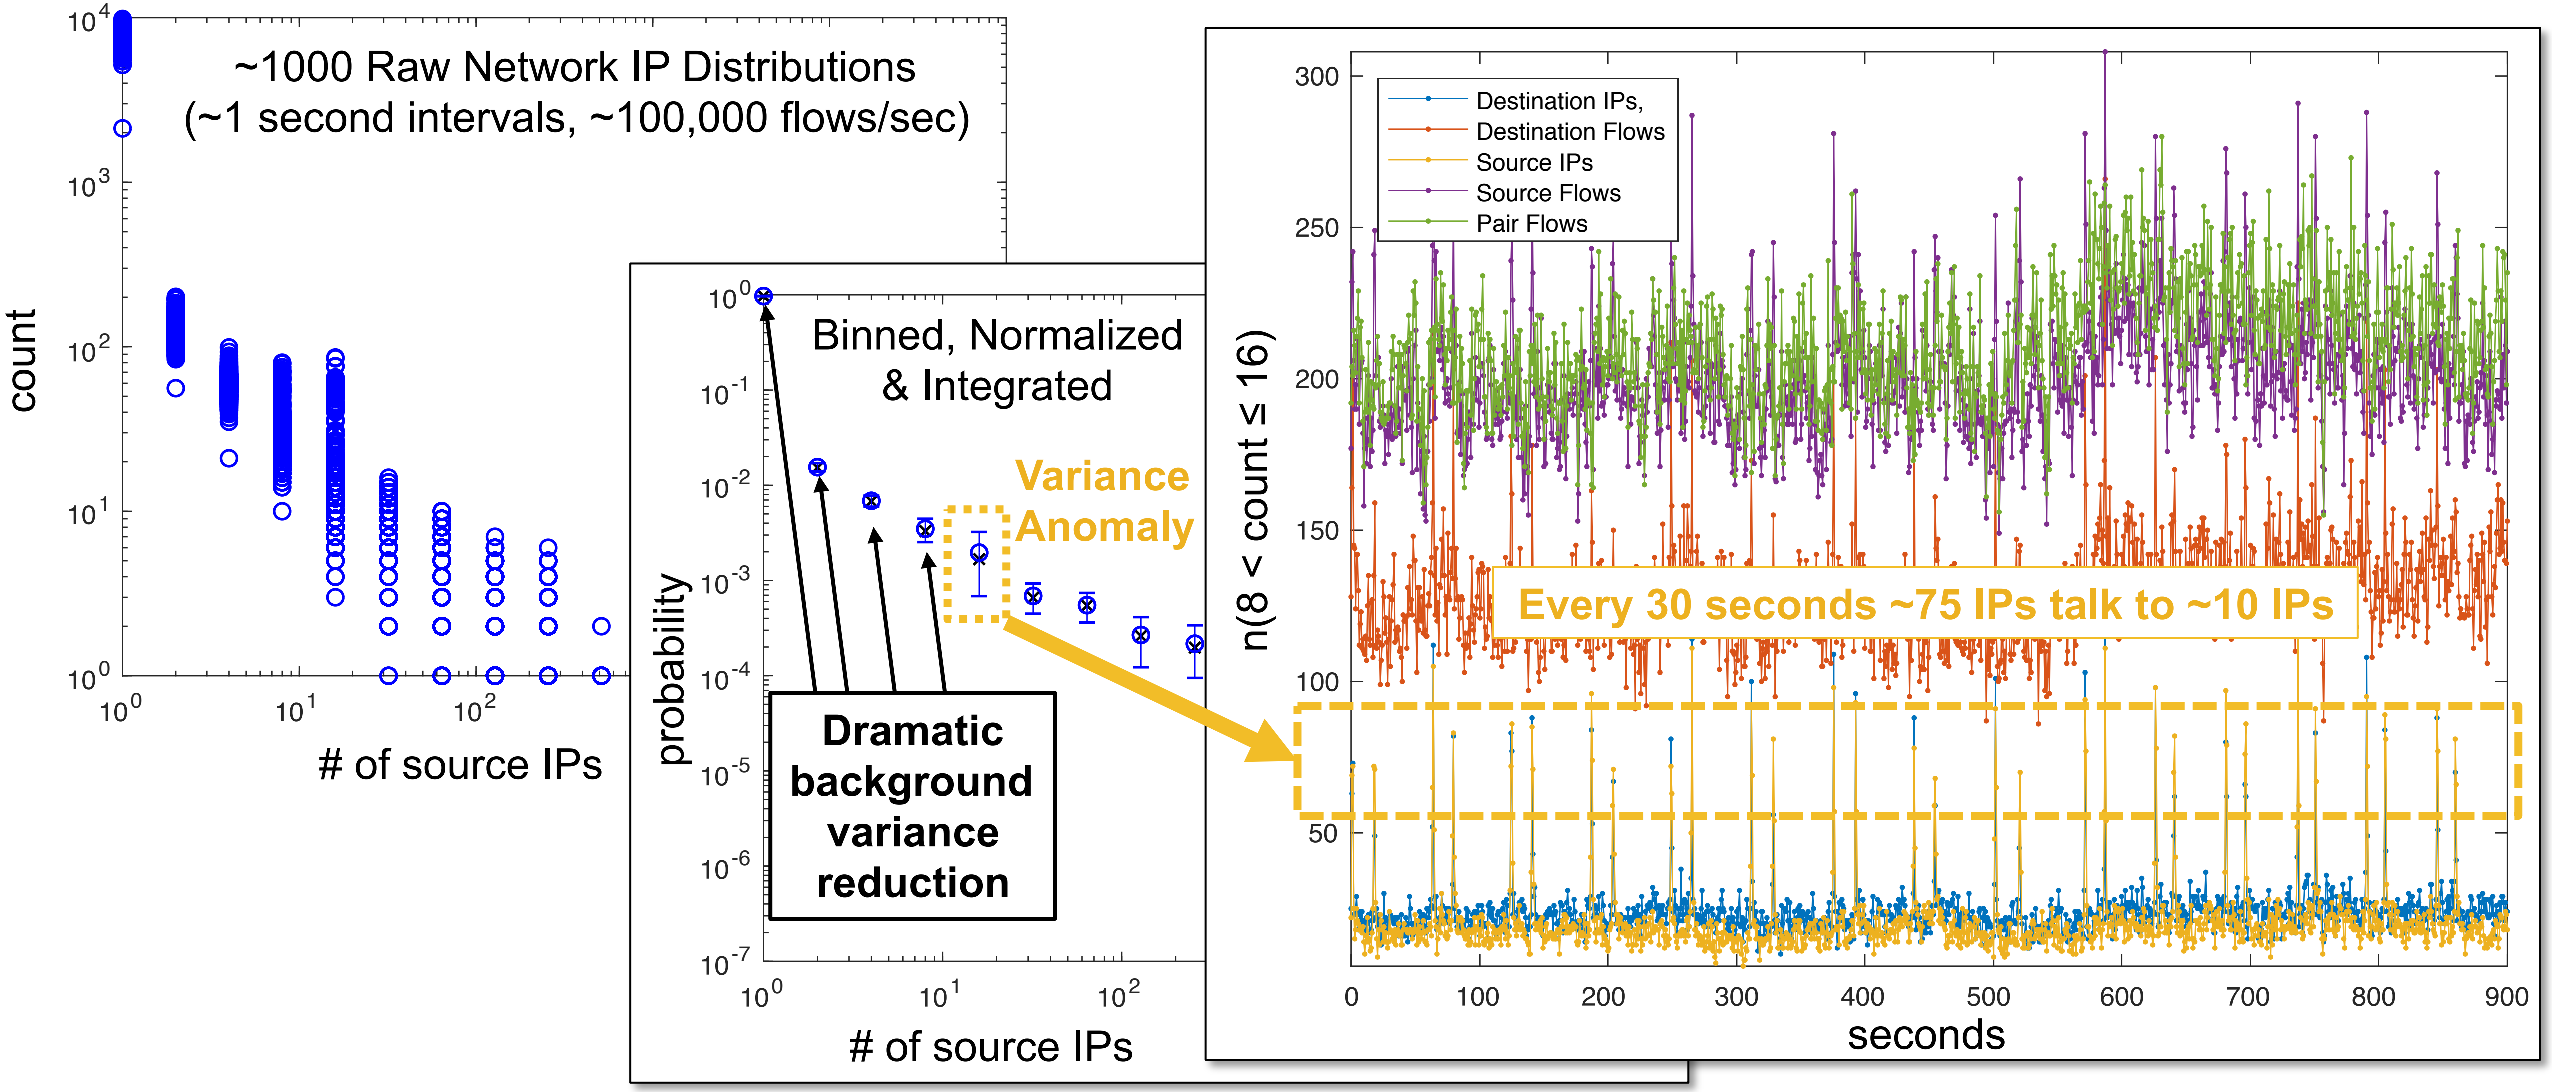

Outline

- Introduction
- MIT SuperCloud
- Hyperscale Analysis Pipeline
- Signal Processing on Networks
- **Summary**

22nd IEEE HPEC Conference
September 25-27, 2018 (ieee-hpec.org)

- **Premiere conference on High Performance Extreme Computing**
  - Largest computing conference in New England (250+ people)
- **Invited Speakers (2018)**
  - Ms. Barbara Helland (DOE)
  - Dr. Rich Linderman (DoD)
  - Prof. Suzanne Sze (MIT)
- **Special sessions on**
  - Amazon/IEEE Graph Challenge
  - Quantum Computing
  - Big Data
  - GPU & FPGA Computing

IEEE HPEC

Platinum Co-Sponsors	Gold Sponsors
SEAGATE Hewlett Packard Enterprise	intel DELL EMC

Silver Sponsors	Cooperating Society	Technical Organizer
MITRE NVIDIA	siam	

- **6-Page Papers Due May 18, 2018**

GPU = Graphics Processing Unit
FPGA = Field Programmable Gate Array

Summary

- **Internet analysis requires methods for detecting faint signals that can leverage signal processing theory**

- **D4M analytics environment and associative array mathematics can be used to prototype complex algorithms that describe internet phenomenology**

- **MIT SuperCloud allows analysts to interactively test algorithms on 10,000+ cores in their preferred environments (Jupyter) and programming languages (Python, Julia, Matlab, Octave, …)**

Backup

Exemplary Packet Capture Pipeline

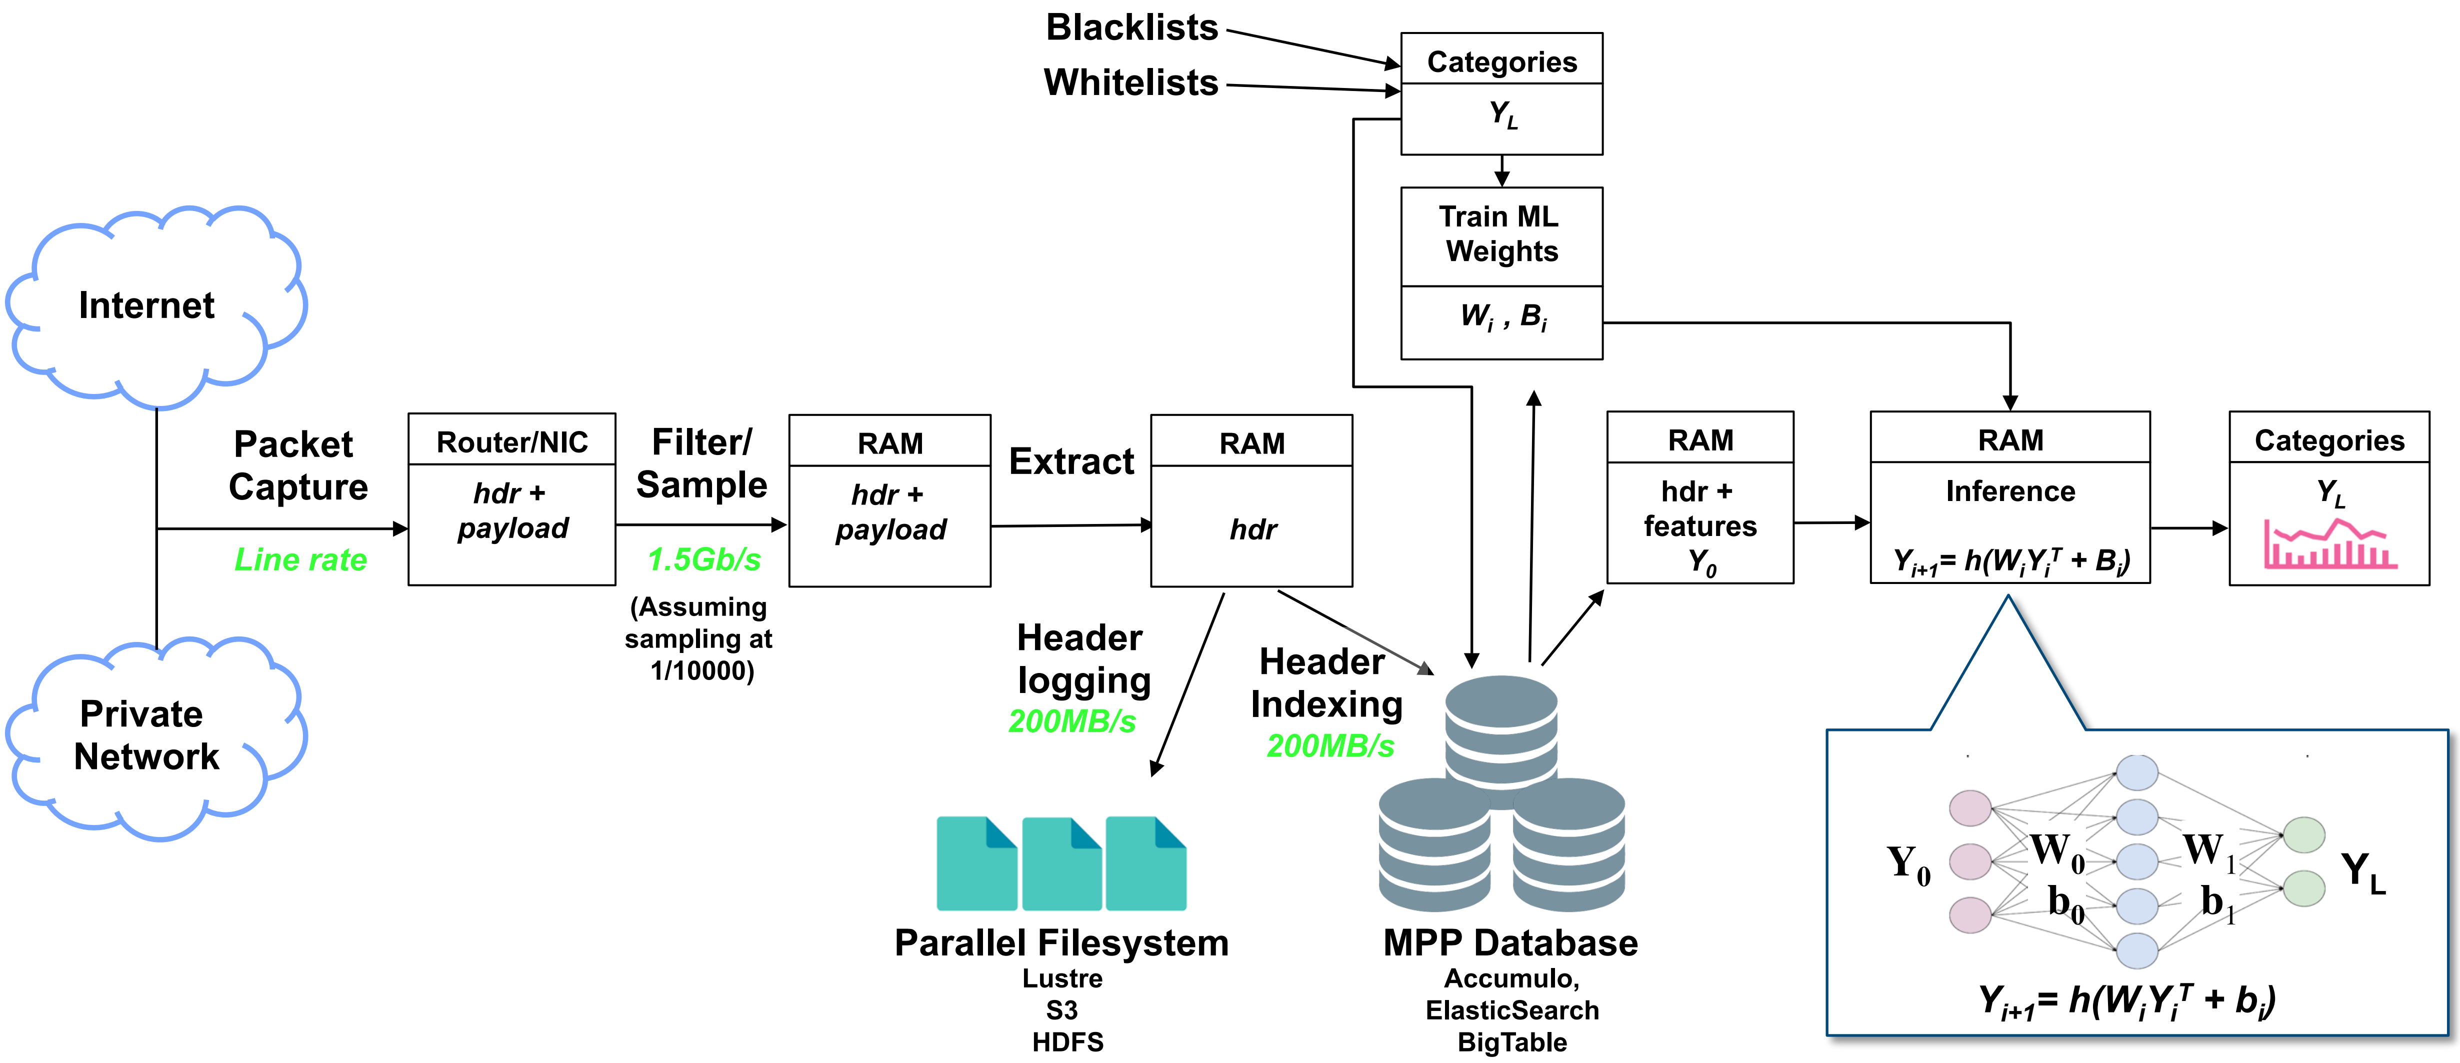

MPP=Massively Parallel Processing
ML= Machine Learning

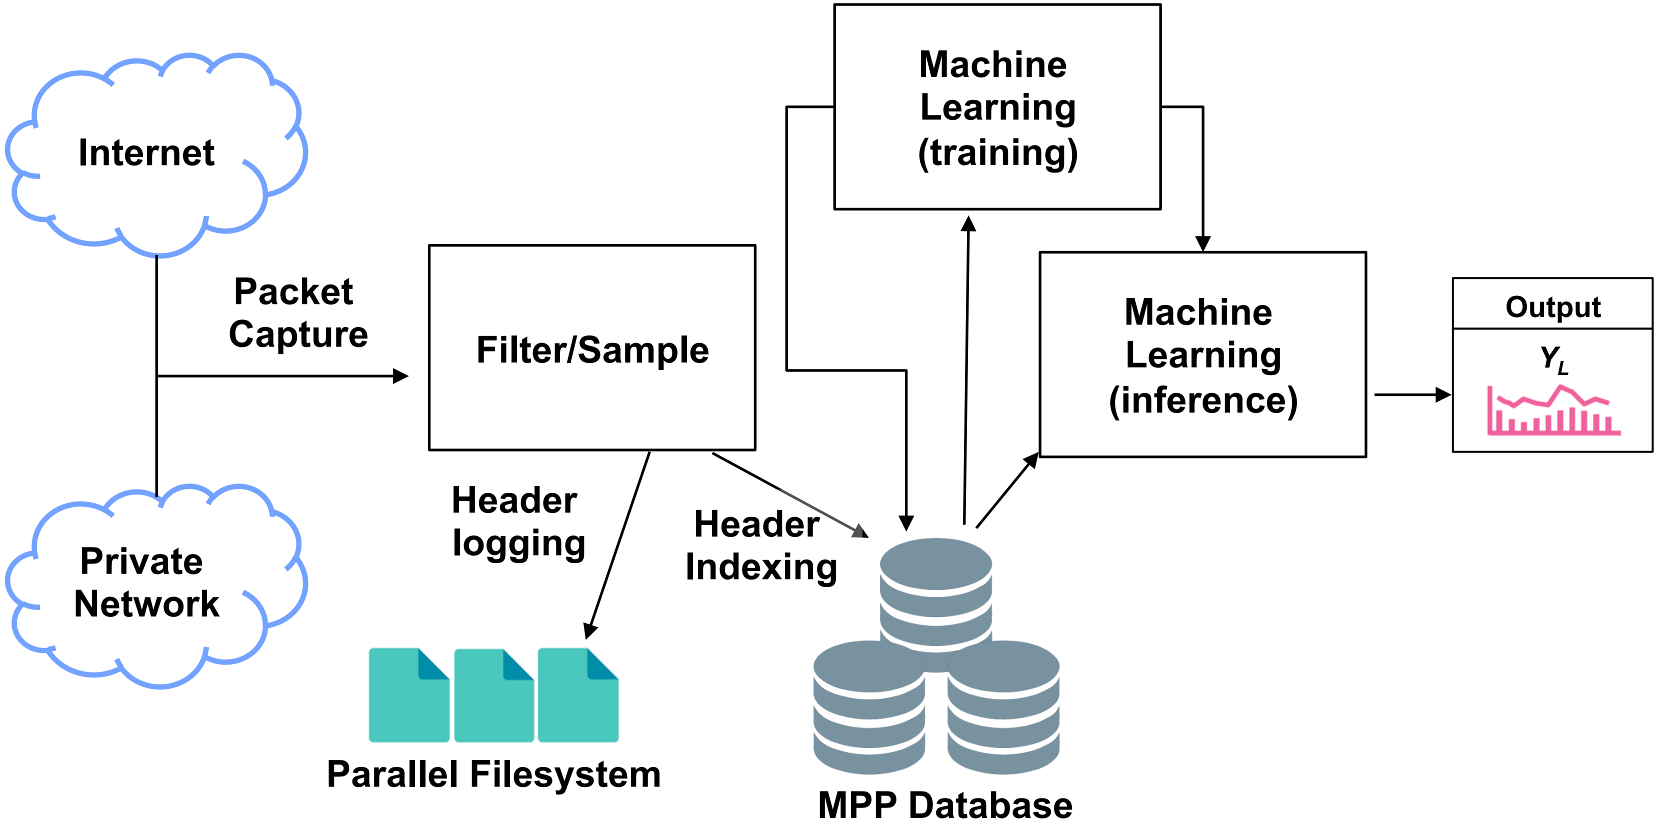
Internet
Private Network
Packet Capture
Filter/Sample
Header logging
Parallel Filesystem
Header Indexing
MPP Database
Machine Learning (training)
Machine Learning (inference)
Output
Y_L