Enabling Scale-Up, Scale-Out, and Scale-Deep for Big Data

Dr. Jeremy Kepner

Jan 2017

MIT LINCOLN LABORATORY
SUPERCOMPUTING CENTER

This material is based upon work supported by the Assistant Secretary of Defense for Research and Engineering under Air Force Contract No. FA8721-05-C-0002 and/or FA8702-15-D-0001. Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Assistant Secretary of Defense for Research and Engineering.

© 2016 Massachusetts Institute of Technology.

Delivered to the U.S. Government with Unlimited Rights, as defined in DFARS Part 252.227-7013 or 7014 (Feb 2014). Notwithstanding any copyright notice, U.S. Government rights in this work are defined by DFARS 252.227-7013 or DFARS 252.227-7014 as detailed above. Use of this work other than as specifically authorized by the U.S. Government may violate any copyrights that exist in this work.

Outline

- Introduction
- Approach
- Future
- Summary

Big Data Challenges

<u>Volume</u>

- **Challenge: Scale of data beyond what current approaches can handle**

<u>Velocity</u>

- **Challenge: Rate of data beyond what current approaches can handle**

<u>Variety</u>

- **Challenge: Diversity of data beyond what current approaches can handle**

Big Data Hardware Solutions

Volume

- Challenge: Scale of data beyond what current approaches can handle
- Hardware Solution: Scale-out, more servers per data center (hyperscale)

Velocity

- Challenge: Rate of data beyond what current approaches can handle
- Hardware Solution: Scale-up, more transistors per server (accelerators)

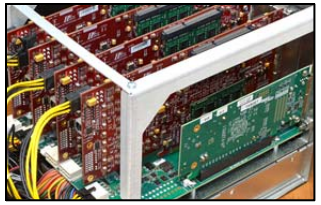

Variety

- Challenge: Diversity of data beyond what current approaches can handle
- Hardware Solution: Scale-deep, more customizable processors (FPGAs, ...)

Requires mathematically rigorous approaches to insulate users from scaling

Our Approach

- **Deliver database capabilities to a much broader set of data**
  - Sorting, indexing, search, ..., for all data
  - Provide declarative, mathematically rigorous interfaces that enable hardware scale-out

- **Integrate computational capabilities into a much broader set of databases**
  - Graph processing, matrix mathematics, machine learning, ..., inside the database
  - Provide declarative, mathematically rigorous interfaces that enable hardware scale-up

- **BigDAWG interface to allow diverse data to move seamlessly across databases**
  - SQL, NoSQL, NewSQL, ..., in the same application
  - Provide declarative, mathematically rigorous interfaces that enable hardware scale-deep

"Bring the power of databases to all data." –Jim Held (Intel Fellow)

Outline

- Introduction
- Approach
- Future
- Summary

Modern Database Paradigm Shifts

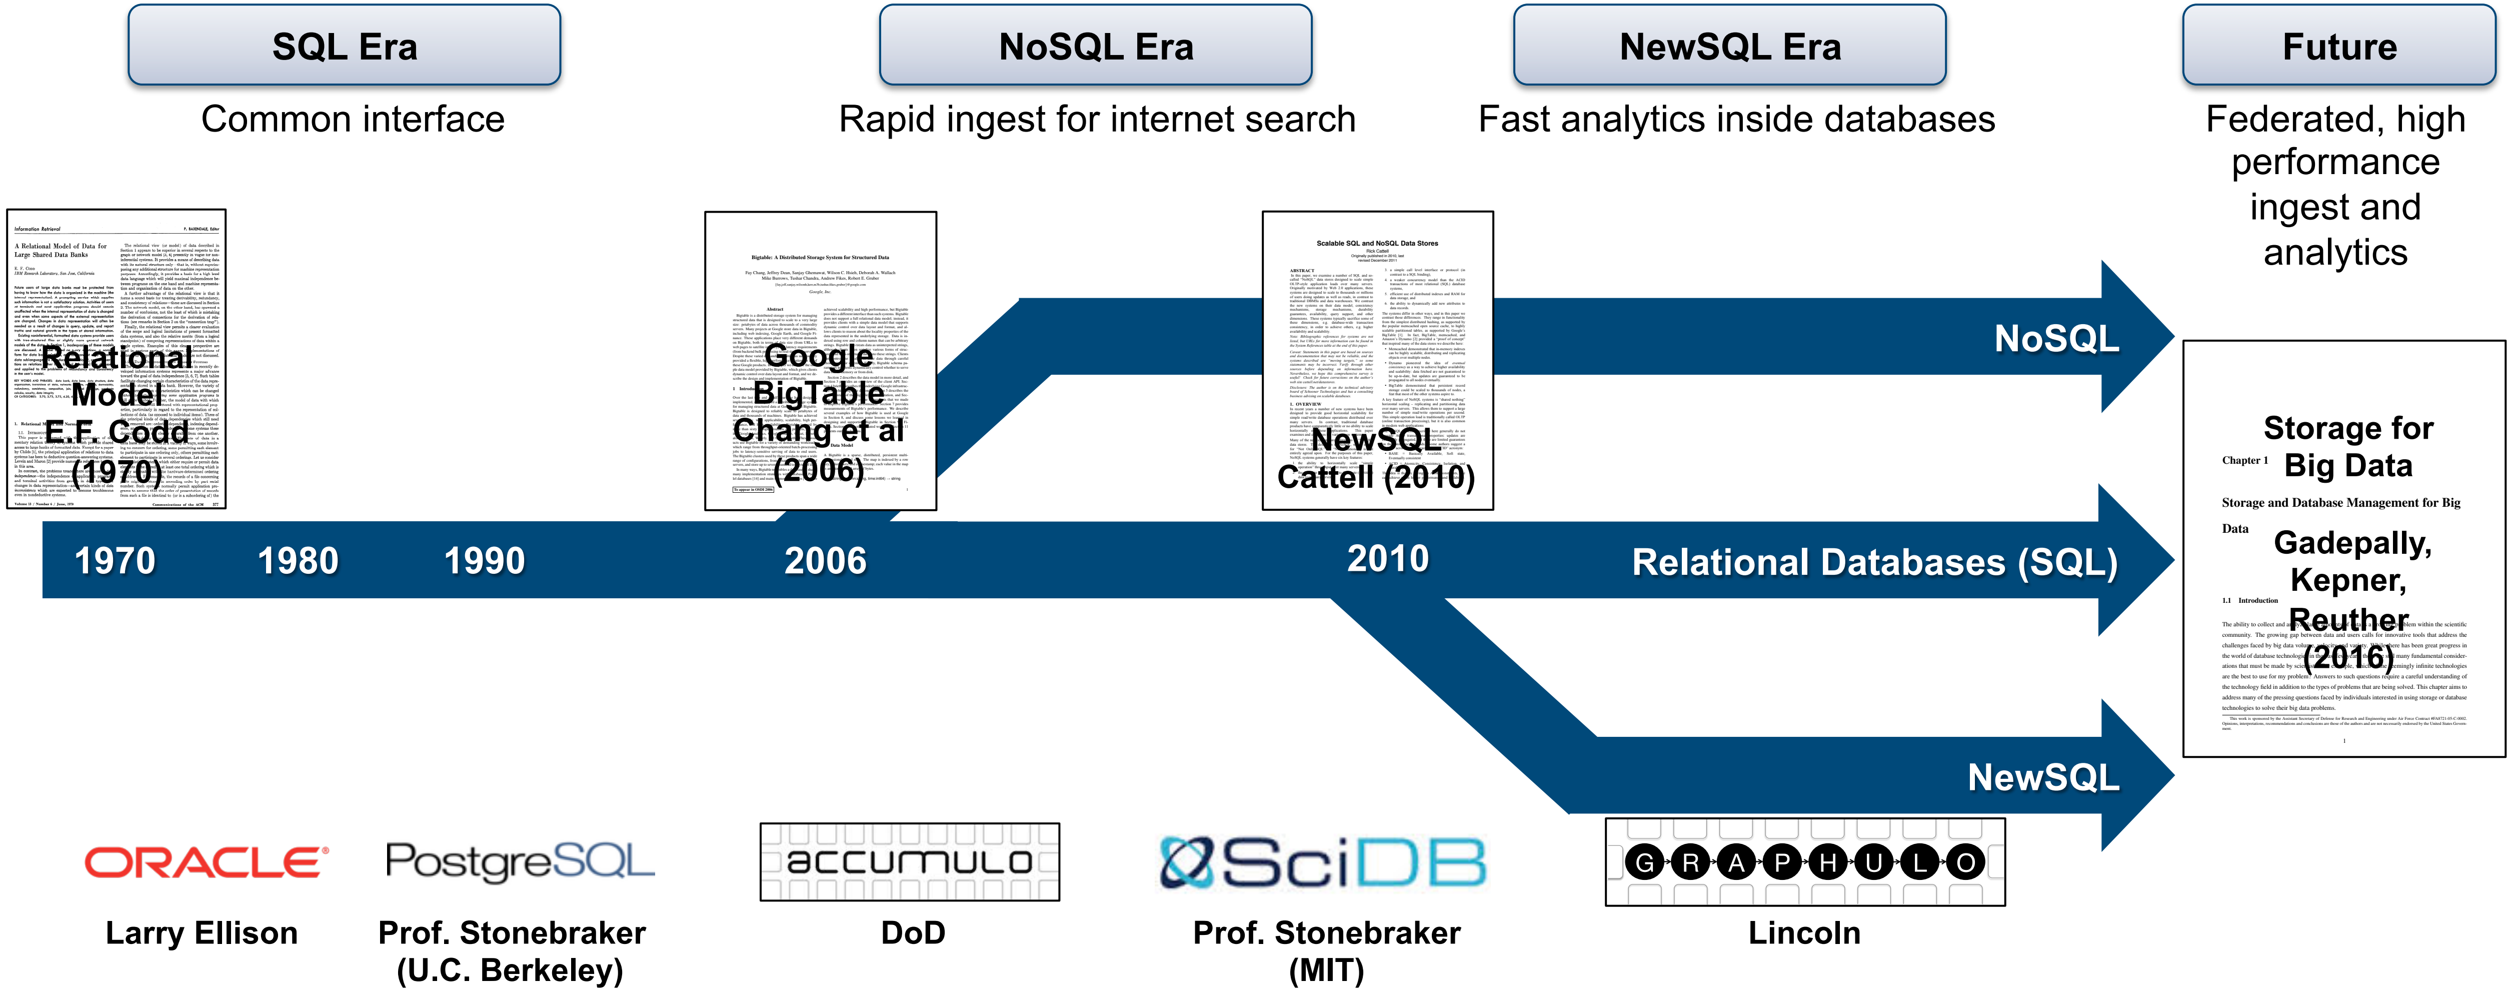

SQL = Structured Query Language
NoSQL = Not only SQL

Declarative, Mathematically Rigorous Interfaces

SQL
Set Operations

	from	link	to
001	alice	cited	bob
002	bob	cited	alice
003	alice	cited	carl

SELECT
WHERE from=alice

NoSQL
Graph Operations

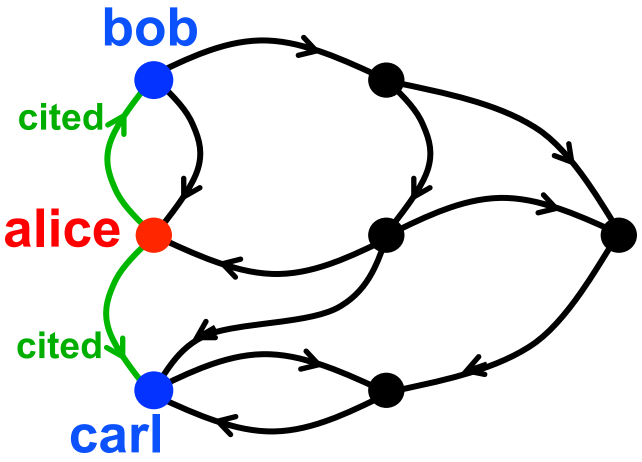

NewSQL
Linear Algebra

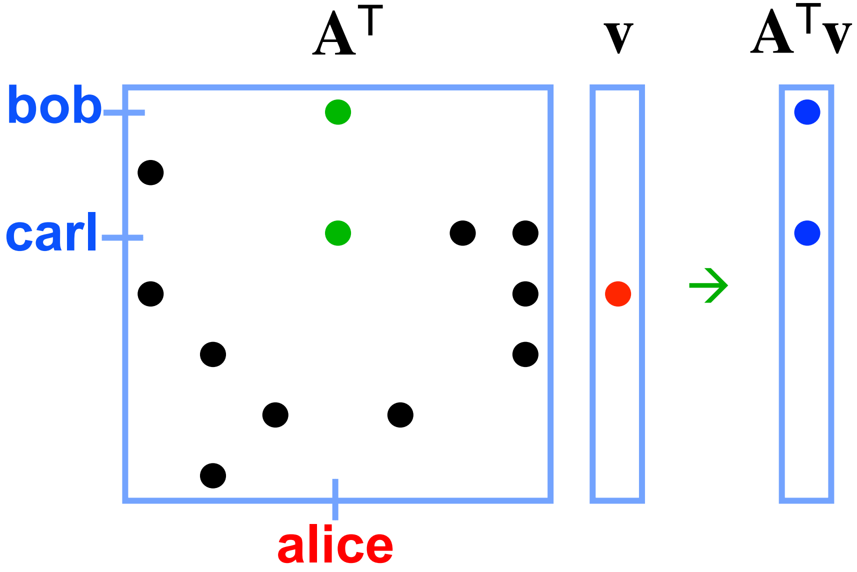

Associative Array Algebra Provides a Unified Mathematics for SQL, NoSQL, NewSQL

$\mathbf{A} = \mathbb{S}^{N\times M}(\mathbf{k}_1,\mathbf{k}_2,\mathbf{v},\oplus)$ $\quad (\mathbf{k}_1,\mathbf{k}_2,\mathbf{v}) = \mathbf{A}$ $\quad \mathbf{C} = \mathbf{A}^{\mathsf{T}}$ $\quad \mathbf{C} = \mathbf{A} \oplus \mathbf{B}$ $\quad \mathbf{C} = \mathbf{A} \otimes \mathbf{C}$ $\quad \mathbf{C} = \mathbf{A}\,\mathbf{B} = \mathbf{A} \oplus.\!\otimes \mathbf{B}$

Operations in All Representations are Equivalent

Advanced Computing Landscape

Mission Drivers

High performance data analysis: machine learning, graph analytics, sequence analysis, ...

Physical systems modeling: robotic vehicles, electronic devices, ...

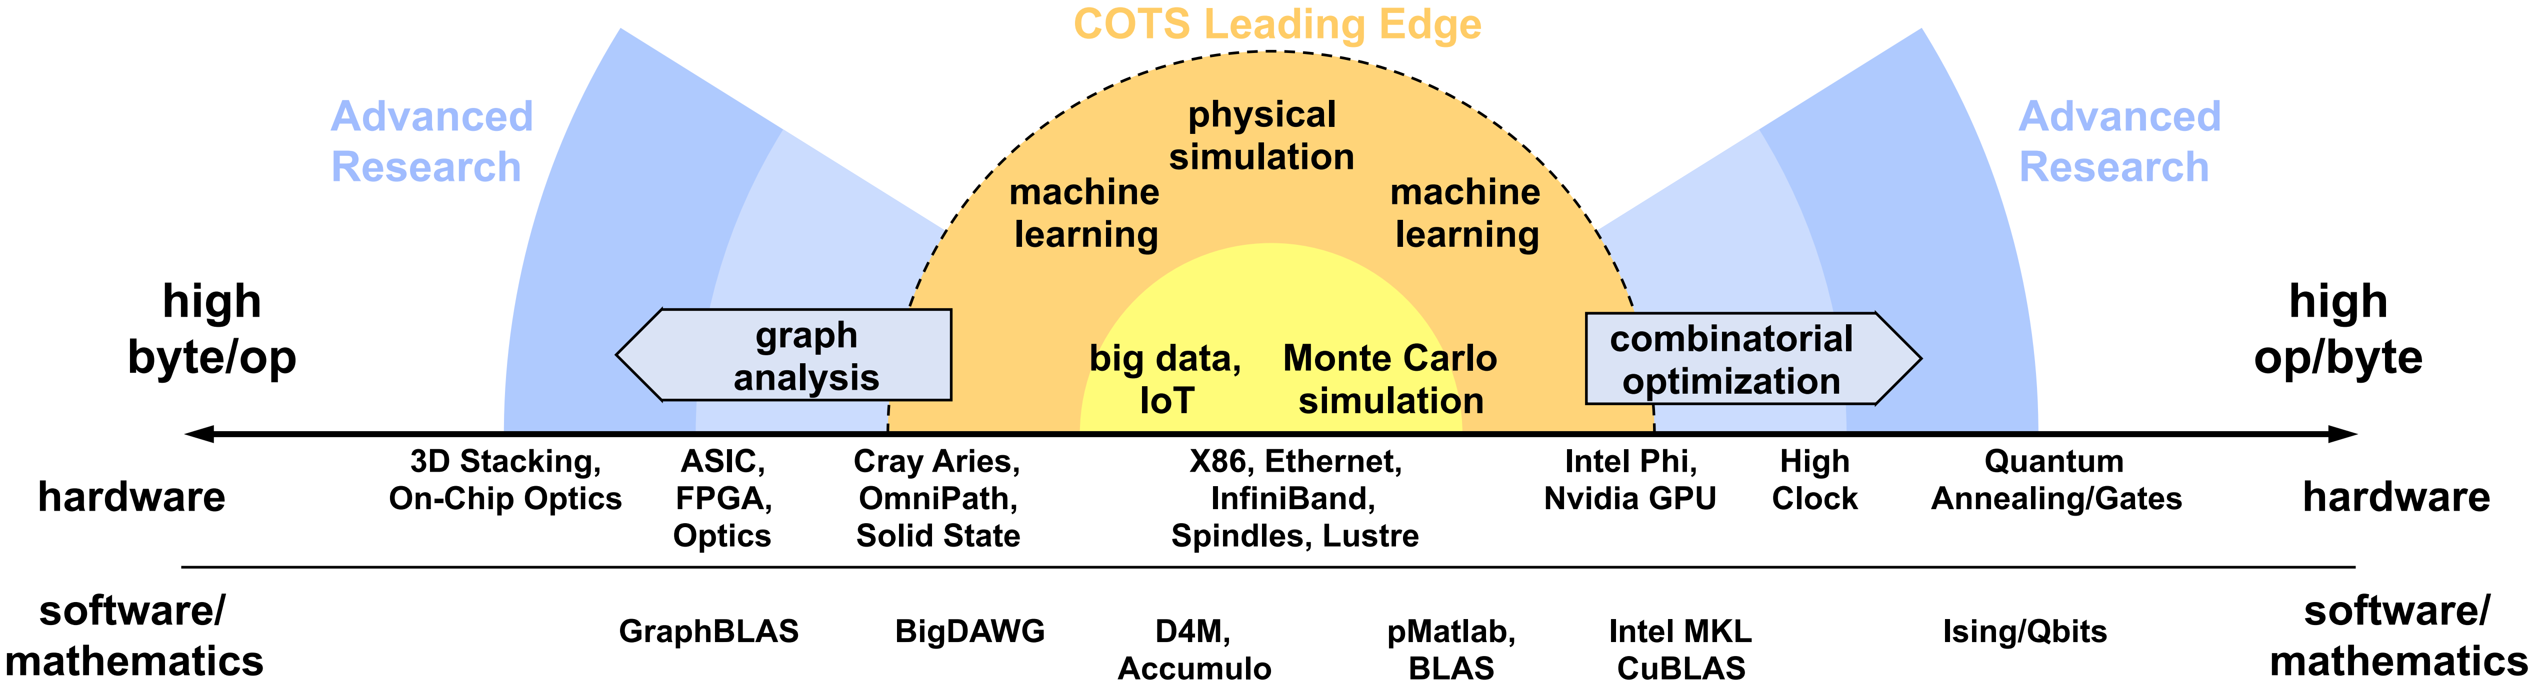

Lincoln Laboratory Supercomputing Center

LLSC Advanced Computing Landscape

Mission Drivers

machine learning, graph analytics, sequence analysis, ...

tic vehicles, electronic devices, ...

Scale-Up: Intel KNL
Parallel Matlab ~2 Teraflops

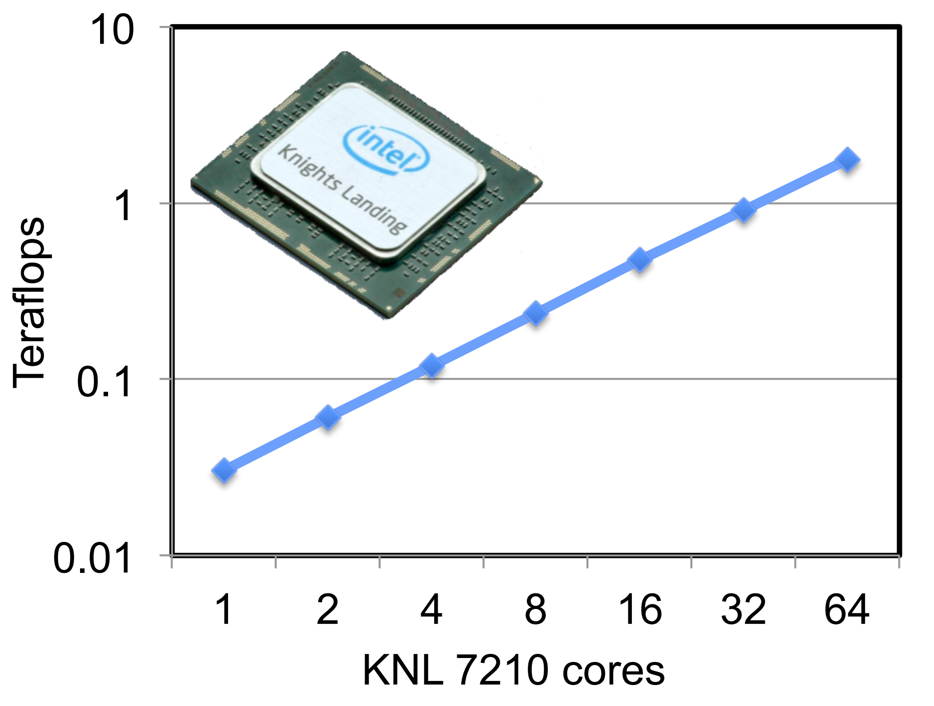

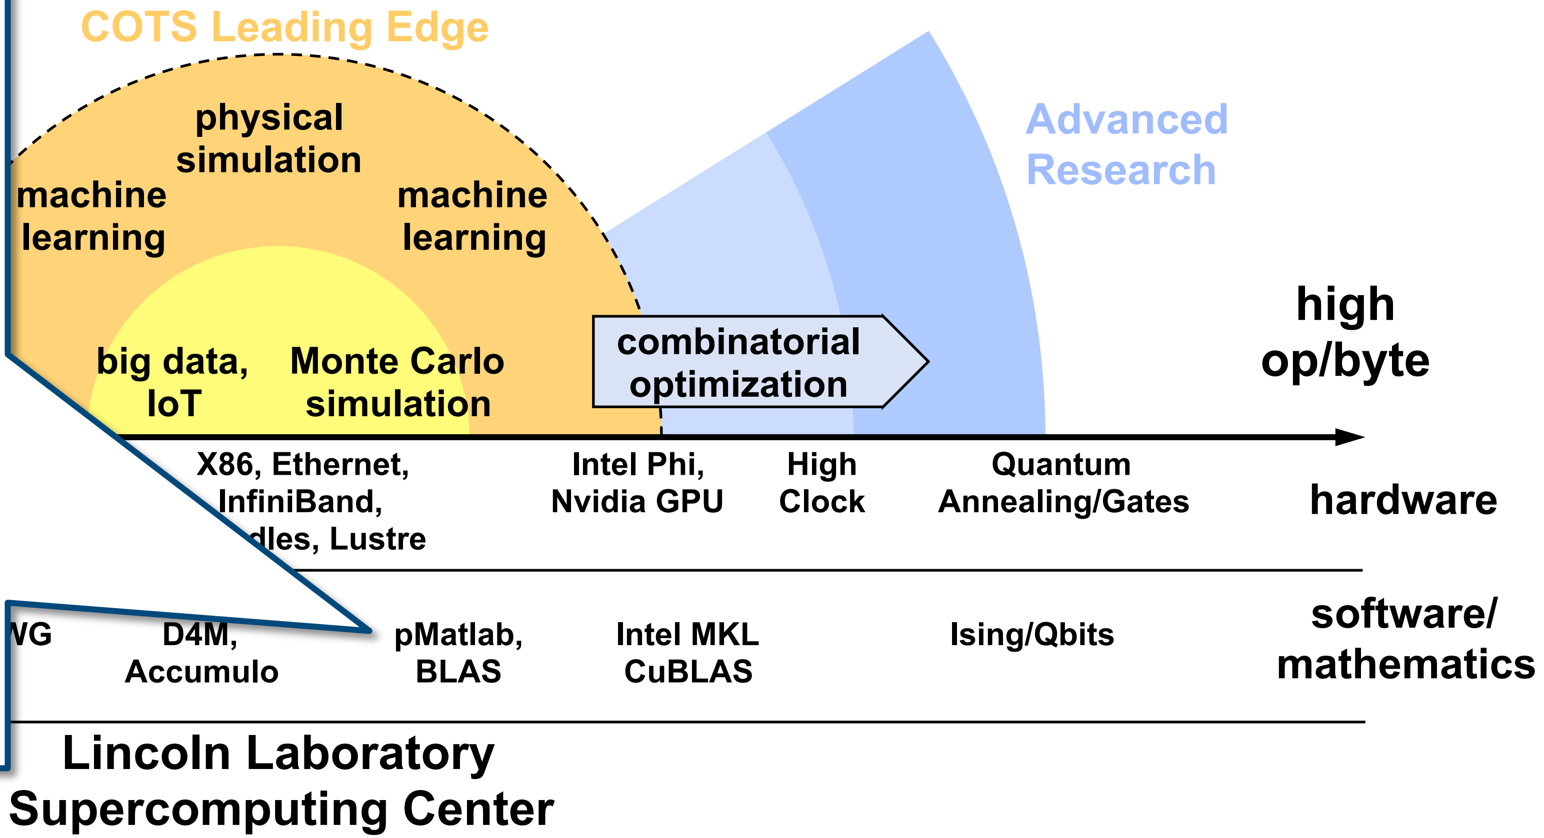

LLSC Advanced Computing Landscape

Mission Drivers

High performance data analysis: machine learning, grap

Physical systems modeling: robotic vehicles, electronic

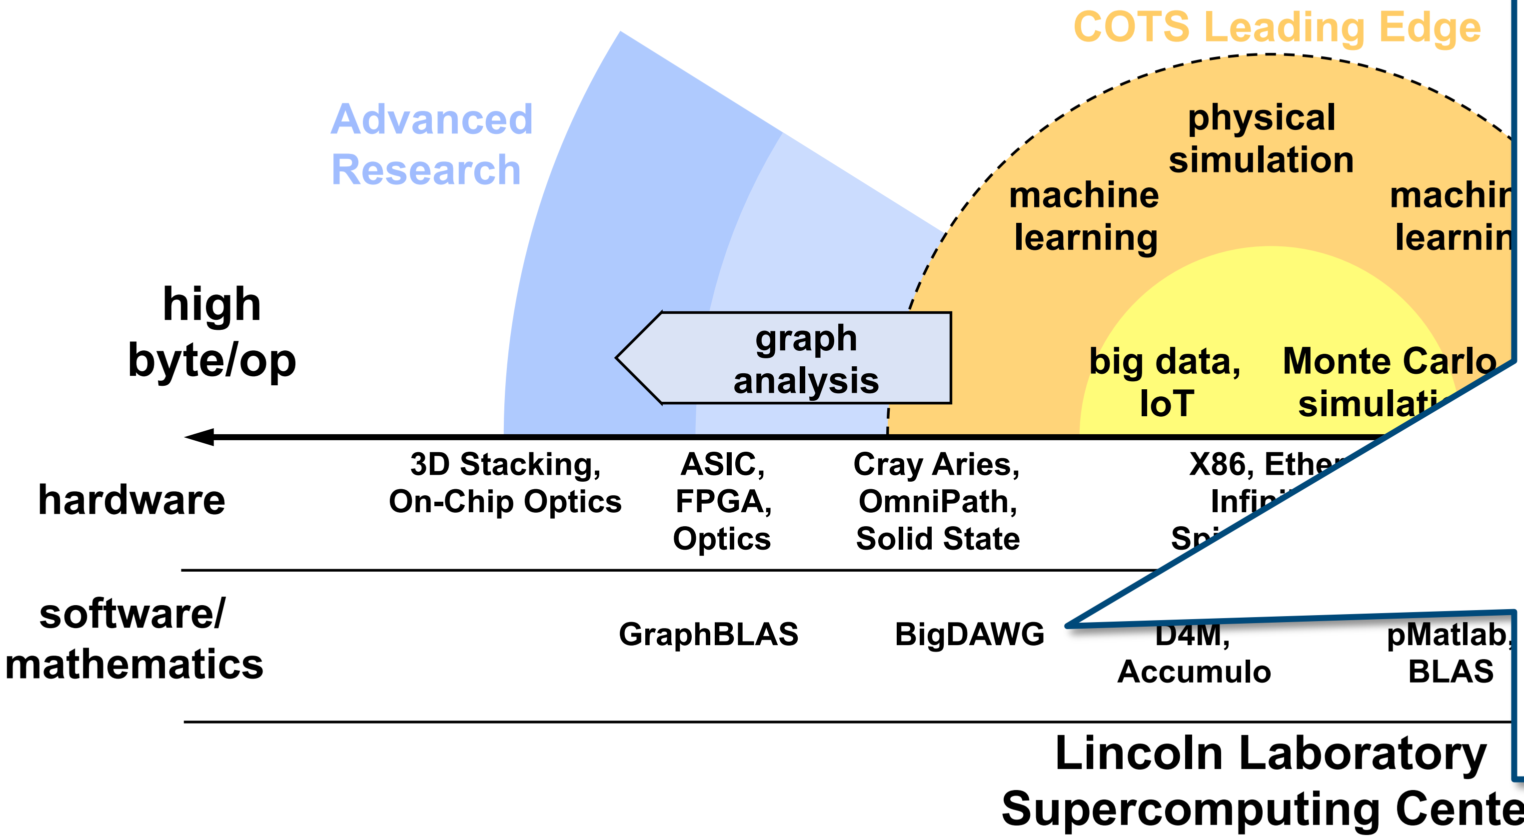

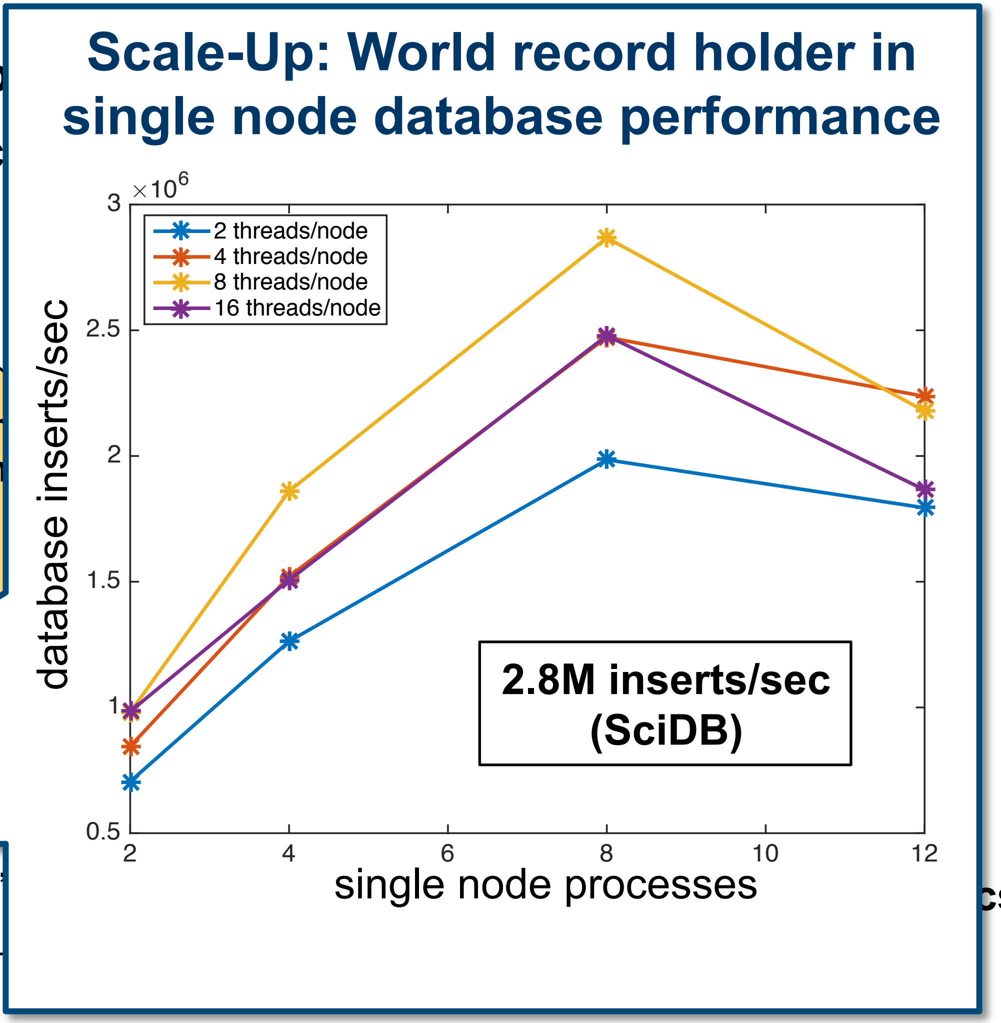

D4M and Large Array Databases for Management and Analysis of Large Biomedical Imaging Data, Samsi et al., NEDB 2016

LLSC Advanced Computing Landscape

Mission Drivers

High performance data analysis: machine learning, grap

Physical systems modeling: robotic vehicles, electronic

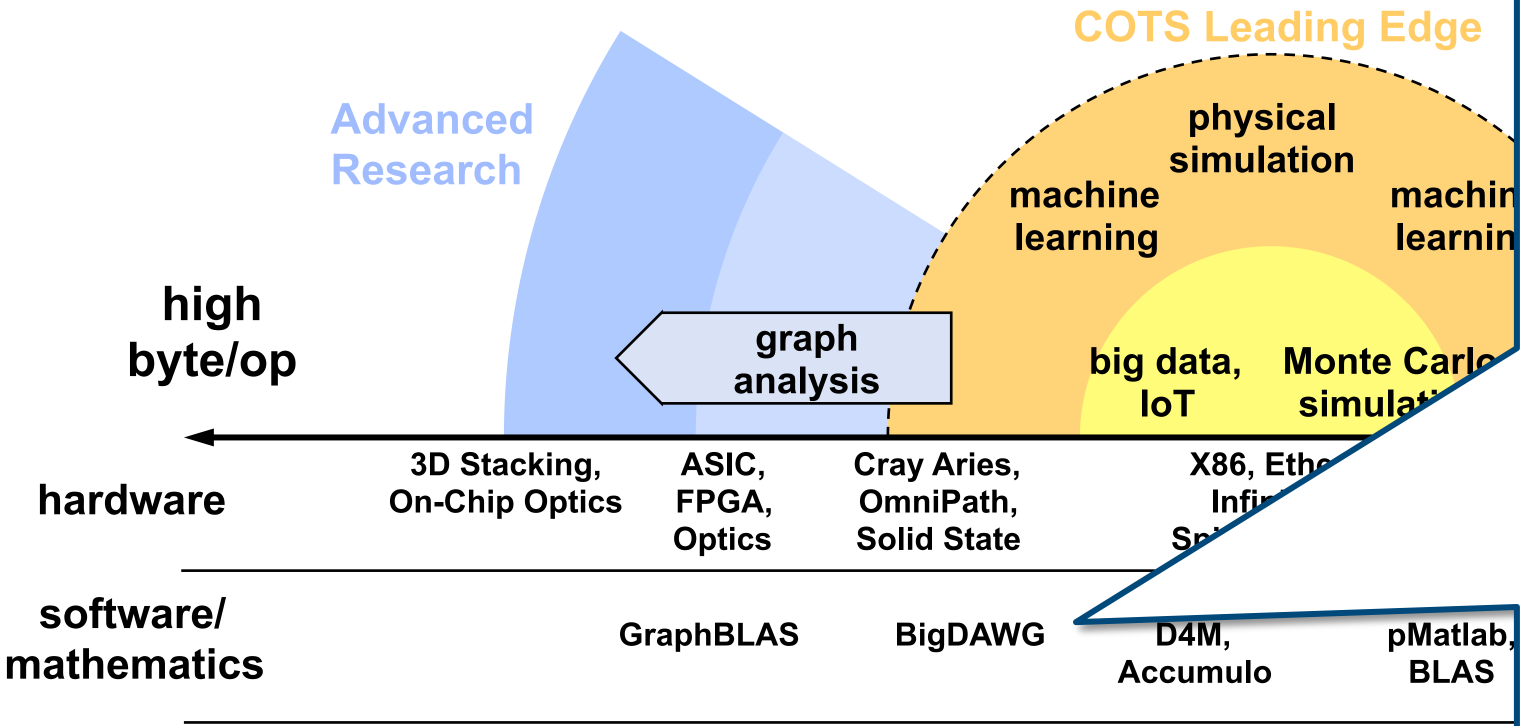

Lincoln Laboratory Supercomputing Center

Scale-Out: World record holder in system wide database performance

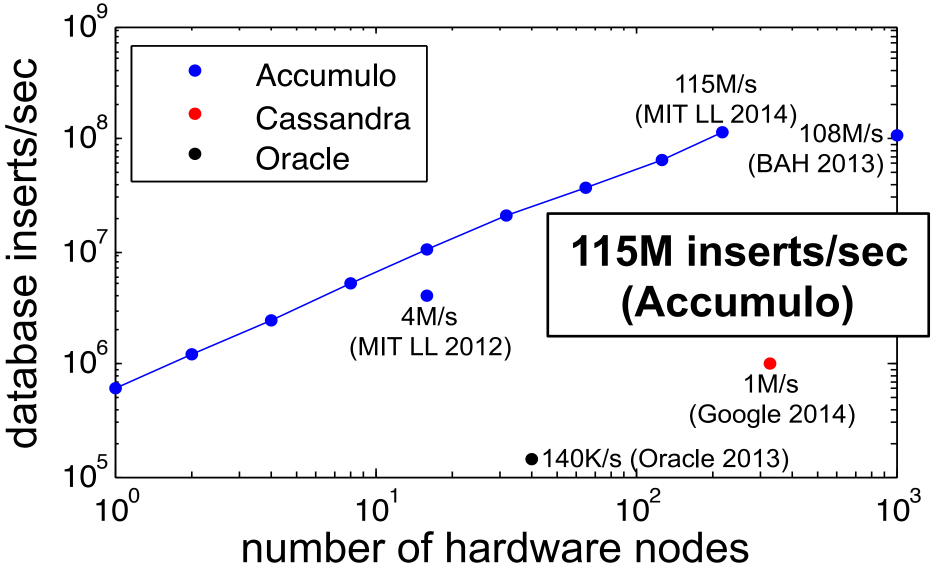

Achieving 100,000,000 database inserts per second using Accumulo and D4M, IEEE HPEC 2014

LLSC Advanced Computing Landscape

Mission Drivers

High performance data analysis: mac

Physical systems modeling: robotic v

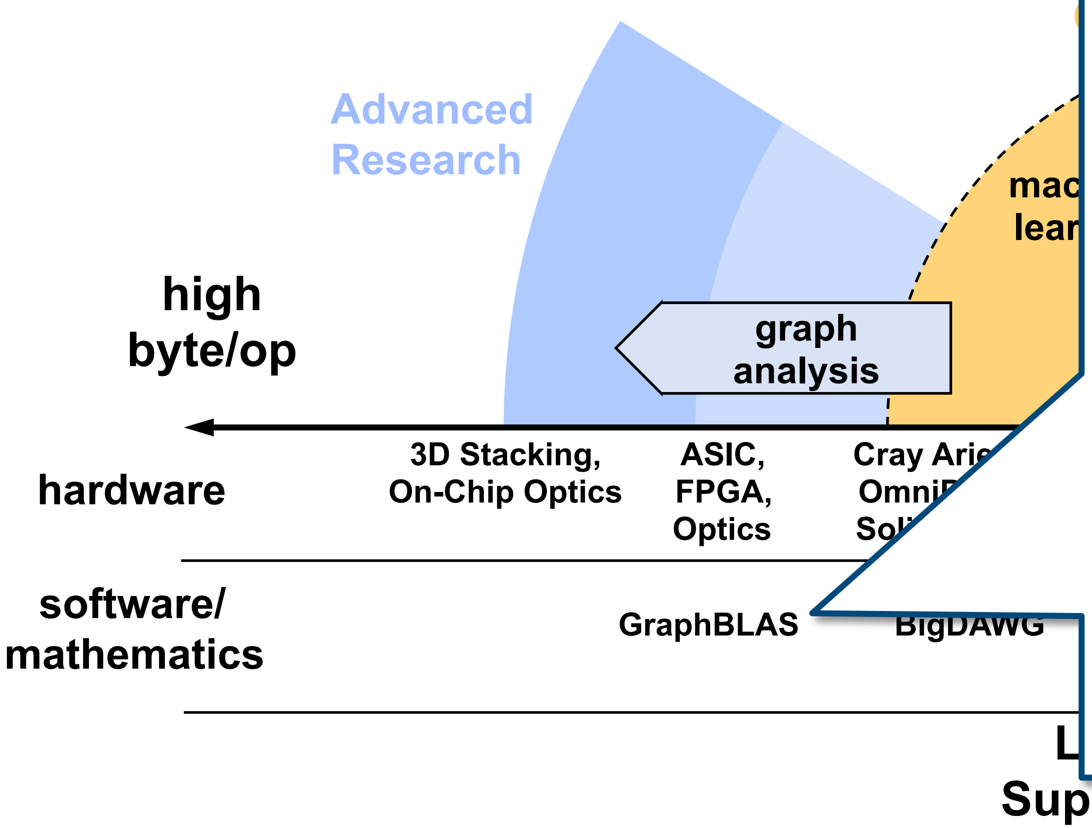

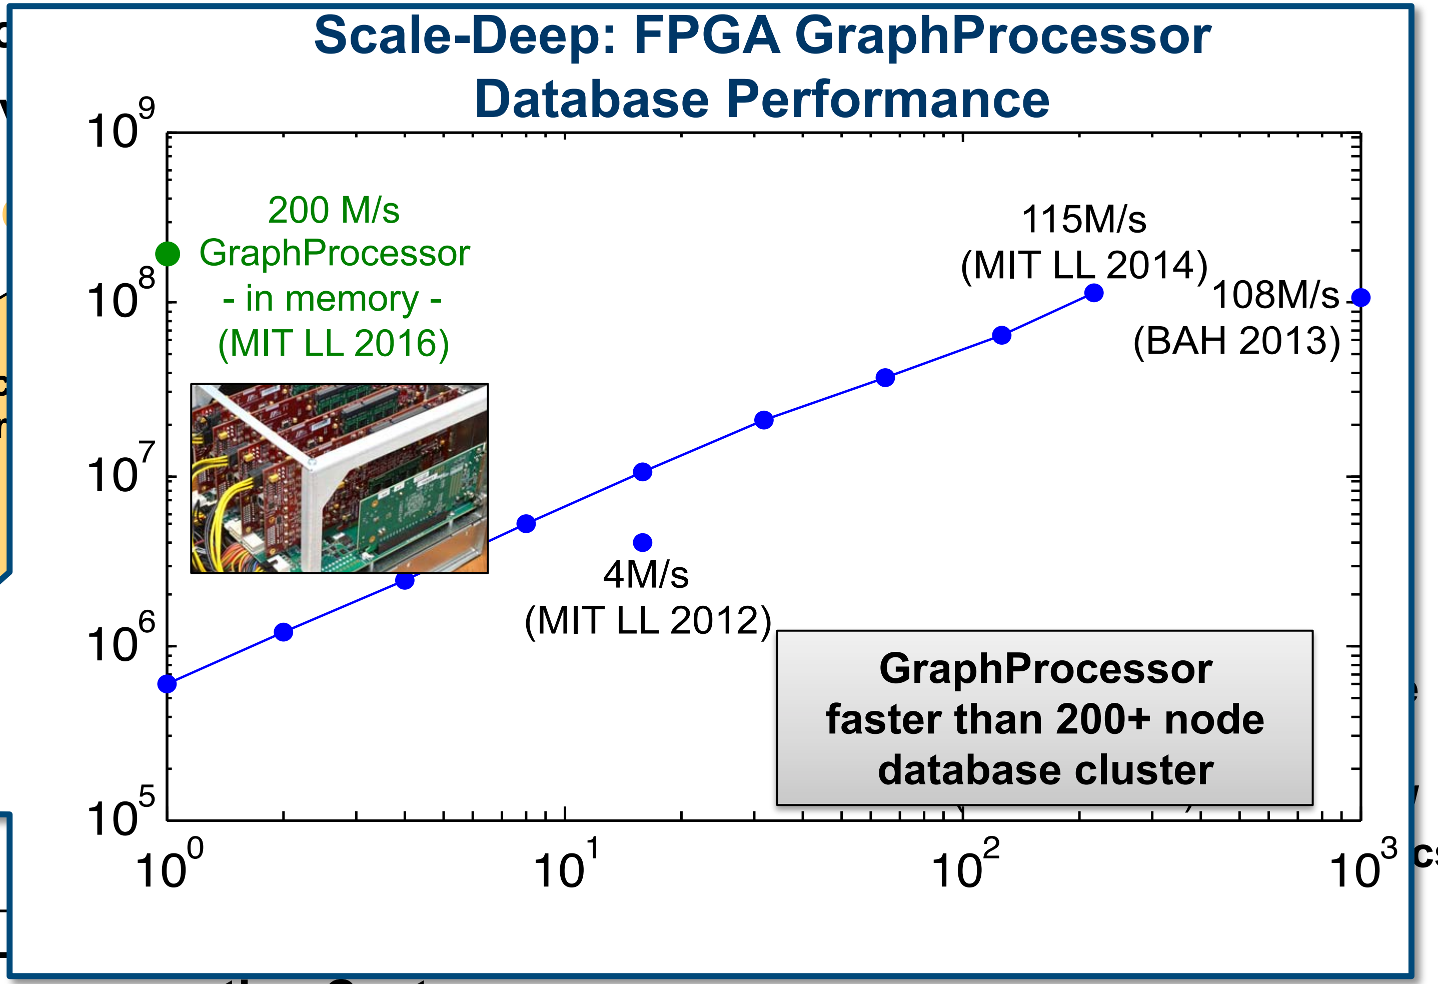

Novel Graph Processor Architecture, Prototype System, and Results, Song, et al., IEEE HPEC 2016

Outline

- Introduction
- Approach
- Future
- Summary

Challenge: New OS for a New Era of Computers

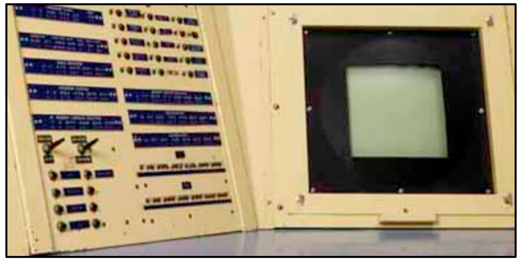

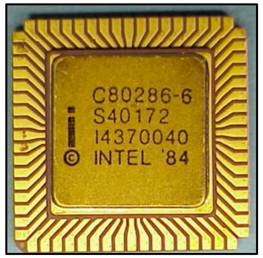

Computing Past

Serial

Local

Homogeneous

Deterministic

OS Managed

Processes

Memory

Files

Communications

Security

Computing Present

Massively Parallel

Distributed

Heterogeneous

Non-Deterministic

User Managed

Processes

Memory

Files

Communications

Security

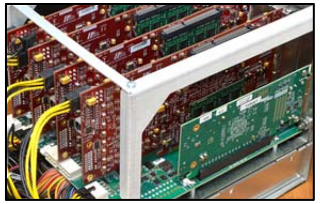

TabularOSA: Tabular Operating System Architecture

Potential Organizing Principles

- Faster, simpler, easier-to-use, provides the data to know what happened when
- Database, data analysis, machine learning are first order operations
- OS designed to analyze itself
- Graduate Unix's "everything is a file" philosophy to "everything is a table”
- Rigorously defined mathematical interfaces and properties

Books

2009 (Scale-Up, Scale-Out)

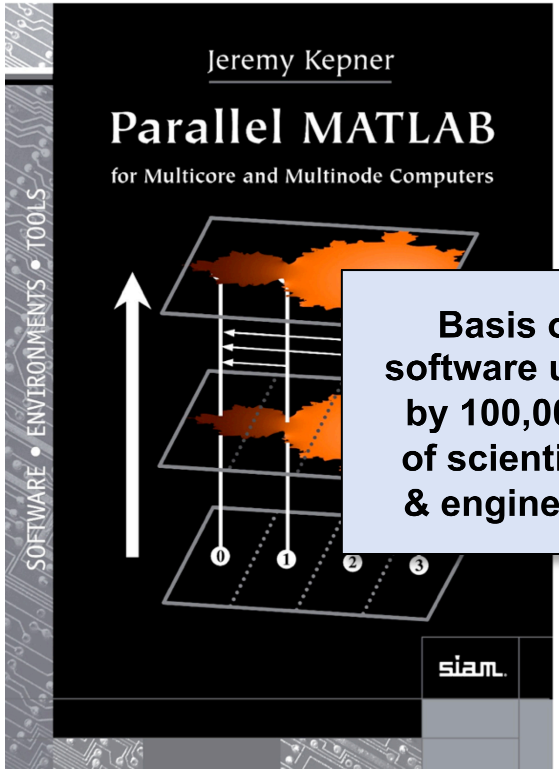

Basis of software used by 100,000s of scientists & engineers

2011 (Scale-Deep)

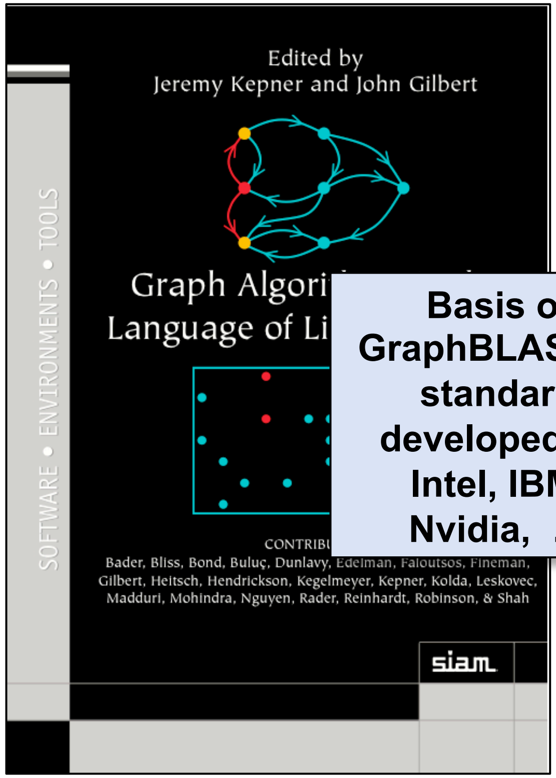

Basis of GraphBLAS.org standard developed by Intel, IBM, Nvidia, …

2017 (Scale-Deep)

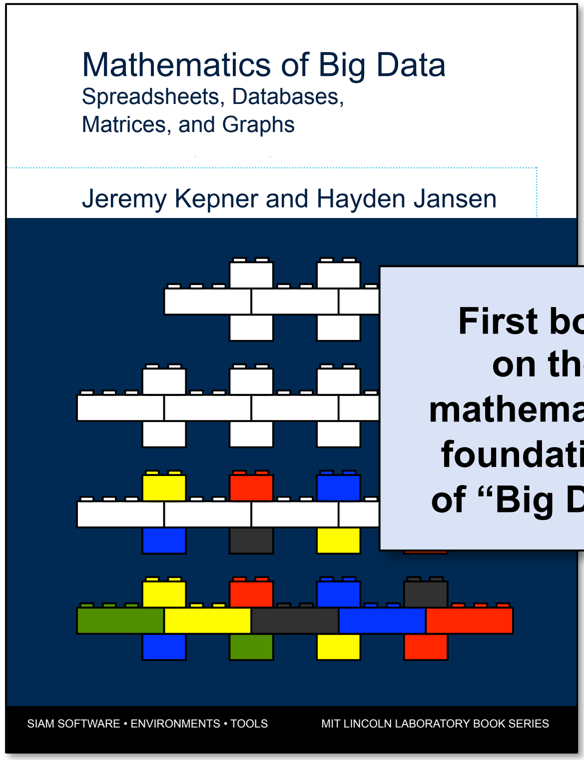

First book on the mathematical foundations of "Big Data"

Based on strong collaborations with MIT faculty, researchers, and students: Parallel Matlab (Prof. Edelman), Graph Algorithms (Prof. Gilbert), Big Data (Profs. Madden & Stonebraker)

IEEE-HPEC.org

IEEE
Advancing Technology for Humanity

2017 IEEE High Performance Extreme Computing Conference (HPEC '17)

Twenty-first Annual HPEC Conference

12 - 14 September 2017
Westin Hotel, Waltham, MA USA

IEEE HPEC
High Performance Extreme Computing

Welcome to HPEC 2017

2016 Sponsors

Platinum Co-Sponsors

HPEC is the largest computing conference in New England and is the premier conference in the world on the convergence of High Performance and Embedded Computing. We are passionate about performance. Our community is interested in computing hardware, software, systems and applications where performance matters. We welcome experts and people who are new to the field.

Need to Know!

Call for Papers 2017

Submission Deadline: May 19, 2017

GraphChallenge
(coming soon)

Summary

Volume

- Challenge: Scale of data beyond what current approaches can handle
- Hardware Solution: Scale-out, more servers per data center (hyperscale)

Velocity

- Challenge: Rate of data beyond what current approaches can handle
- Hardware Solution: Scale-up, more transistors per server (accelerators)

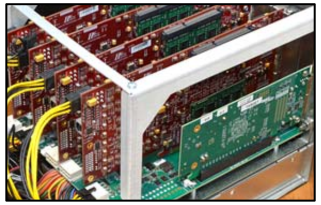

Variety

- Challenge: Diversity of data beyond what current approaches can handle
- Hardware Solution: Scale-deep, more customizable processors (FPGAs, ...)

Requires mathematically rigorous approaches to insulate users from scaling