ASCR Research Division Long Range Plan (LRP)

(For 2018 – 2030)

William Harrod

DOE / ASCR

December 20, 2016

U.S. DEPARTMENT OF ENERGY | Office of Science

Purpose of Long Range Plan (LRP)

- Provide the necessary new vision for ASCR Research to support DOE in an era of profound technological change in the computer and semiconductor industries.
- Computing is fundamental to DOE's missions in scientific discovery, energy prosperity and environmental health, and national security.
- Establish ambitious goals and associated research initiatives necessary to enable current and future DOE missions
- Enhance U.S. leadership and support the research community in furthering the national scientific endeavor

Long Range Planning Team
Development Team

William J. Harrod
ASCR, Research Division Director

- Frank Alexander, PhD,
  Los Alamos National Laboratory
- David L. Brown, PhD,
  Lawrence Berkeley National Laboratory
- Richard Carlson,
  ASCR, Next Generation Networking
- Scott Collis, PhD,
  Sandia National Laboratory
- Lori A. Diachin, PhD,
  Lawrence Livermore National Laboratory
- Adolfy Hoisie, PhD,
  Pacific Northwest National Laboratory
- Kerstin Kleese van Dam,
  Brookhaven National Laboratory
- Steven Lee, PhD,
  ASCR, Applied Mathematics
- Barney Maccabe, PhD,
  Oak Ridge National Laboratory
- Lucy Nowell, PhD,
  ASCR, Computer Science
- Marc Snir, PhD,
  Argonne National Laboratory
- Ceren Susut, PhD,
  ASCR, Computational Partnerships

LRP Timeline & Overall Process

- LRP Document Development
  - Nov 2015 – Initiated discussions on LRP development process
  - Jan 2016 – First meeting of LRP Development Team
  - Dec 2016 – Preliminary version of LRP document
- Workshops
  - Data Convergence – Spring 2016
  - Computing Beyond 2025 – Summer 2016
  - Research Initiatives – (TBD) 2017
- Strategic Plan
  - Based on LRP document
  - Developed by DOE HQ, ASCR

LRP Outline

ASCR Research

- ASCR Research Mission
  - Create, develop, and deploy the computational and networking capabilities that underlie DOE's ability to fulfill its missions and that support its ability to meet future challenges
  - Maximize the utility of ASCR facilities for analyzing, modeling, simulating and predicting complex phenomena critical to DOE's missions.
- ASCR Research Vision
  - Along with simulation science, ASCR Research will lead progress in data-driven advanced computing for scientific discovery and science-based policy-making across the DOE mission space

LRP Goals

1. Improve responsiveness of foundational research programs to DOE missions
2. Maximize the return from DOE investments in computing hardware and software
3. Enhance scientific productivity: computationally-enabled scientific and engineering results from theoretical and experimental studies
4. Capitalize on new and emerging computational technologies
5. Maintain and grow the HPC workforce
6. Increase ASCR's contribution to US competitiveness and policy making

Major Research Initiatives

1. Drive the Convergence of Data and Computing
2. Bring the Power of Artificial Intelligence to Scientific Discovery
3. Advance New Hardware/Software Computing Paradigms to Glean Unexploited Efficiencies
4. Foster Ubiquitous Use of Advanced Computing for Scientific Discovery and Decision Support

Indicators of Success

The extent to which the following capabilities are achieved:

- Revolutionize and automate scientific processes;
- Shield DOE scientists and engineers from changes in the underlying implementations while continuing delivery of effective performance;
- Provide DOE scientists and engineers with greater confidence in results and improved timeliness for science and decisions; and
- Give policy makers greater confidence in scientific results and greater understanding of the uncertainty associated with those results

Goal 1: Responsive Research

Improve the Responsiveness of Foundational Research in Applied Mathematics, Computer Science, and Networking to DOE missions

Strategy: *Create foundational research programs in applied mathematics, computer science, and networking as enabling technologies for scientific discovery and decision-support for government policy makers.*

Key Concepts and Terms

- Algorithms and techniques
- System characteristics
- Complex physical systems
- Uncertainty

Goal 2: Development/Execution Productivity

Maximize the Return from DOE Investments in Computing Hardware and Software (Development/Execution Productivity)

Strategy: *Create and develop enabling technologies to best utilize DOE's computational and scientific facilities of the future.*

Key Concepts and Terms

- Development infrastructure
- Computer system management
- Transparency, convenience

Goal 3: Scientific Productivity

Enhance Computationally-Enabled Scientific and Engineering Results from Theoretical and Experimental Studies (Scientific Productivity)

Strategy: *Create and develop innovative computational and networking capabilities that enable scientists to effortlessly exploit the full potential of emerging DOE SC computational and experimental facilities.*

Key Concepts and Terms

- Collaboration
- Data infrastructure
- Workflow management

Goal 4: Sustainable Computing Advances

Capitalize on Emerging Computational Technologies

Strategy: *Invest in the research required to create a new curve of sustainable advances in computational capabilities that will continue the reliable progress maintained over the past several decades.*

Key Concepts and Terms

- CMOS and beyond
- New algorithms
- Prototype system stacks

Goal 5: Increased HPC Workforce

Maintain and Grow the HPC Workforce

Strategy: *Actively grow and sustain the DOE laboratory and university workforce in HPC-relevant applied mathematics, computer science, and networking through recruitment and incentives.*

Key Concepts and Terms

- HPC-ready
- Educational curricula
- Diverse workforce

Goal 6: Expanded Scope of Support

Increase Contribution to US Competitiveness and to Policy Making

Strategy: *Facilitate the use of advanced computing in support of DOE's applied missions, U.S. industrial competitiveness, and national policy decisions.*

Key Concepts and Terms

- "First to benefit"
- "What-if" scenarios
- HPC ecosystem
- Provenance, V&V, trustworthiness

Research Thrust 1: Drive the convergence of data and computing for DOE missions

- Historically, DOE's HPC facilities focused primarily on supporting large-scale numerical simulations for theoretical investigations. However, the need for HPC support in analyzing increasingly massive experimental results, analyzing intermediate simulation results, and steering experiments has been steadily growing.

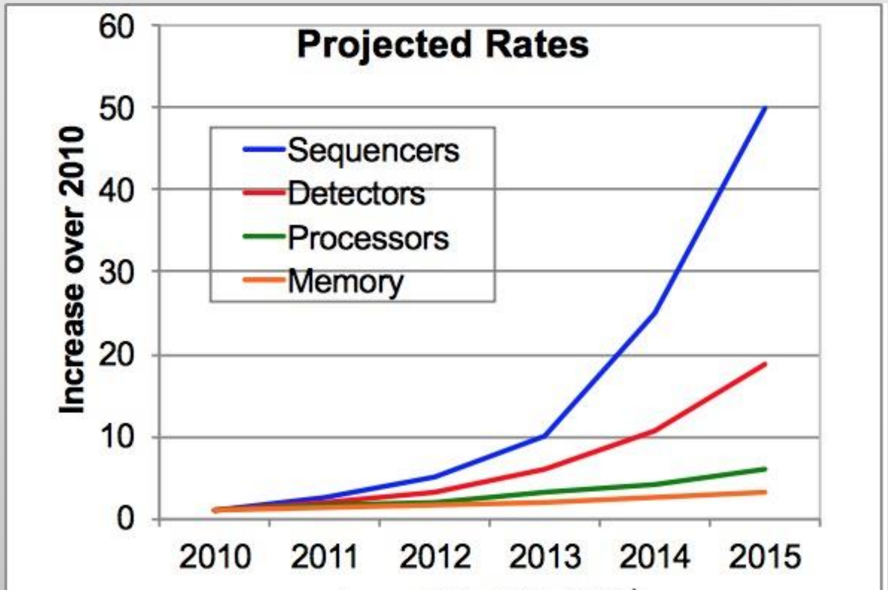

Image: Kathy Yelick (LBNL)

Reference Documents

- Experimental and Observational Data workshop Findings and Recommendations, Bethel
- Synergistic Challenges In Data-Intensive Science And Exascale Computing, ASCAC/Sarkar
- National Strategic Computing Initiative Strategic Plan
- LRP Data Convergence Report (draft), Bell
- A Superfacility for Data Intensive Science, Yelick
- Workflow Report
- Exascale Computing and Big Data, By Daniel A. Reed, Jack Dongarra Communications of the ACM, Vol. 58 No. 7, Pages 56 - 68

Research Thrust 1: Unified Data and Computing – Research Directions

- *ASCR Research will accelerate the pace, quality, and accuracy of both scientific discovery and use of science in policy-making by "tightening the loop" between simulations and experiments.*
- *Research Directions*
    1. New mathematical methods and algorithms to integrate HPC modeling and simulation with experimental data
    2. Software environments to support integrated data analysis and simulations: providing interoperability between runtime systems and virtualization technologies
    3. Infrastructure for composing workflows for distributed experiments
    4. Feedback mechanisms between data analysis and simulations for steering experiments and validating codes
    5. Using artificial intelligence (AI) to improve data analytics
    6. Network infrastructure to support escalating data traffic and communications
    7. Distributed data management

Research Thrust 2: Bring the power of artificial intelligence to scientific discovery and policy making

- Despite major advances in artificial intelligence (AI) over the past decade, their potential benefits have not yet been extended to DOE's mission space, especially in areas dependent on or related to HPC.

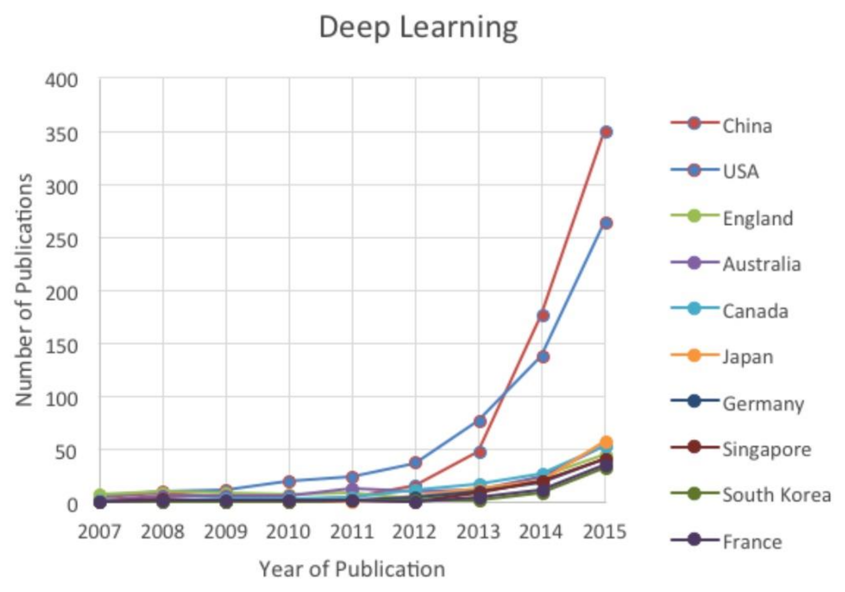

Figure 1: Journal articles mentioning "deep learning" or "deep neural network", by nation.[62]

Reference Documents

- National Artificial Inteligence Research and Development Strategic Plan
- Preparing for the Future of Artifical Intelligence
- National Strategic Computing Initiative Strategic Plan
- Machine Learning and Understanding for Intelligent Extreme Scale Scienti c Computing and Discovery
- Neuromorphic Computing Architectures, Models, and Applications

Research Thrust 2: Using AI – Research Directions

- *ASCR Research will develop, adapt, and extend artificial intelligence (AI) approaches to advance knowledge generation; create more flexible, reliable computing infrastructures; and support decision making.*
- *Research Directions*
  1. Representations for automated analysis, human comprehension, and explanation
  2. Natural language understanding of scientific literature
  3. Machine Learning (ML) for finding rare events and discovering patterns in data, text, and publications; and dynamic control of computing resources
  4. Multi-modal data analysis for analyzing heterogeneous datasets at multiple levels of abstraction/resolution
  5. Introspective, adaptive systems for optimizing HPC operation and managing HPC centers

Research Thrust 3: Advance new computing paradigms in hardware and software to create new pathways of continuing performance gains

- The imminent of end of Moore's Law/Dennard Scaling, with no clear replacement technology on the horizon, creates unique opportunities (through U.S. government/industry partnerships) to develop mutually-beneficial disruptive technologies that can drive the future of mainstream computing.
- ASCR Research must set the course of new hardware and software technology development to provide the foundations for future HPC systems (beyond exascale) and to establish new pathways of sustainable progress for future DOE missions.

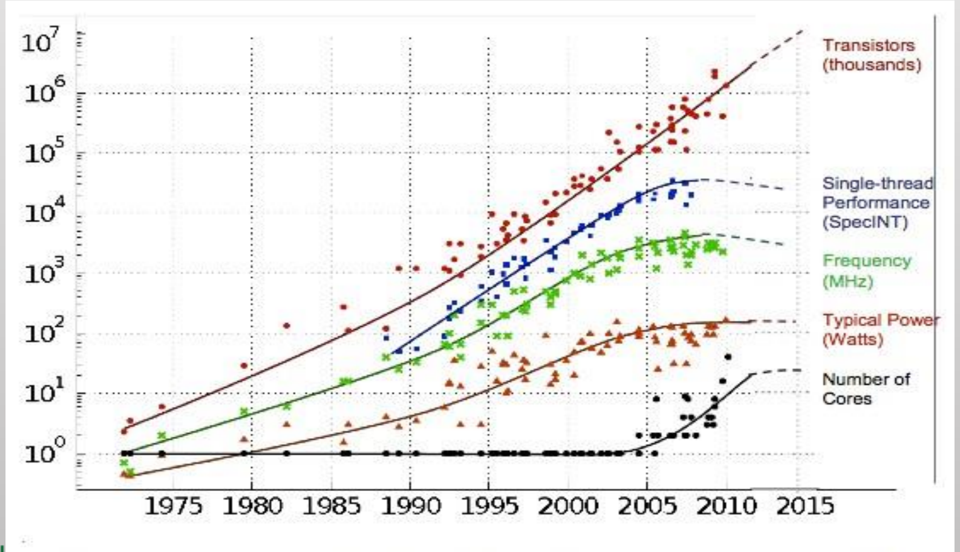

Original data collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond and C. Batten
Dotted line extrapolations by C. Moore

Reference Documents

- The Future of Computing Performance
- Advancing Quantum Information Science: National Challenges and Opportuities
- Neuromorphic Computing: From Materials to Systems Architecture
- Computing Beyond 2025 Report (TBD)
- Research and Education in Computational Science and Engineering

Research Thrust 3: New computing paradigms– Research Directions

- *ASCR Research will deliver essential increases in effective performance and functionality of computing systems to meet the needs of DOE and the nation.*
- *Research Directions*
  1. Extending CMOS: develop highly-concurrent, energy-efficient architectures
  2. Theory of computation: programming models to utilize future hybrid architectures
  3. Scalable, adaptive algorithms that exploit the increased capabilities of emerging computing architectures
  4. Software methodology, including theory of computation, formal methods, and computational complexity analyses applied to scientific codes and collaboration
  5. Modeling and simulation: for predicting system, application and algorithm performance

Research Thrust 4: Enable ubiquitous computing for science-based decision support

- DOE has a unique and growing role in using HPC systems to support decision making in a wide range of settings, including driving engineering solutions, accelerating the pace of scientific discovery, and supporting science-based policy decisions.
- As advanced computing becomes increasingly important in driving consequential decisions at shorter time scales, with explosively increasing data, issues of correctness, provenance, uncertainty quantification, stability, etc. must be resolved.

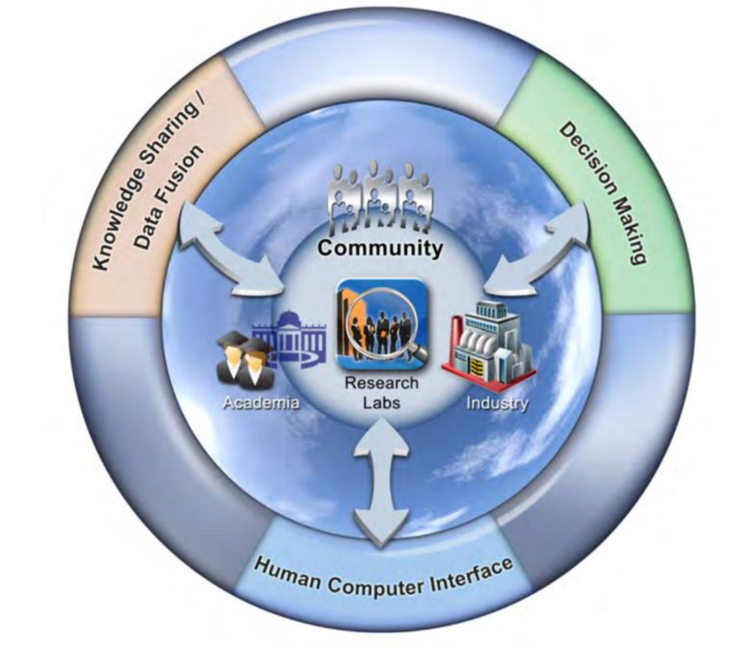

Reference Documents

Accelerating Scientific Knowledge Discovery

ASCR Cybersecurity for Scientific Computing Integrity

Management, Visualization, and Analysis of Experimental and Observational Data (EOD) The Convergence of Data and Computing Workshop Final Report

The Future of Scientific Workflows

Research and Education in Computational Science and Engineering

Research Thrust 4: Ubiquitous computing

- *ASCR Research will develop efficient and high-confidence systems for science, evidence-based policy, and operational support at the pace, complexity, and impact of advancing technology.*
- *Research Directions*
  1. "Vanguard" programming models: tools, environments and methodologies
  2. Ensuring correctness and integrity of software design and operation
  3. Foundations for human and machine decision support
  4. Information retrieval and archiving methods
  5. Collaborative environments for scientifically diverse, geographically-distributed researchers
  6. Human-machine partnerships for augmenting human capabilities in science

Crosscutting Directives
Topics

- ***Complexity:*** with the proliferation of heterogeneous, potentially-incompatible subsystems, components, and technologies that will need to be unified, complexity threatens to overwhelm HPC hardware/software systems and applications
- ***Productivity:*** increasing HPC system complexity threatens to decrease realizable system performance gains, efficiency of code development, system accessibility, and the ability to exploit the full system potential
- ***Software sustainability:*** without serious forethought to future evolution of HPC architectures, future software upgrade needs, requirements for seamless incorporation into DOE HPC facilities, maintenance, etc., the life-cycle of HPC software threatens to be labor-intensive, expensive to maintain, and incompatible with other software.

BACKUP

Long Range Plan will enable:

- Increased accuracy in modeling systems and analyzing experiments involving processes taking place across a wide range of time and length scales;
- Greater ability to validate application/workflow behavior & integrity
- Globally distributed resources and facilities to be rigorously combined into composable systems with well-defined mathematical properties;
- Efficient, effective achievement of the full the potential of today's and future HPC systems and advanced networks for science and engineering applications;
- Software required to make effective use of future-generation supercomputers and the future scientific user facilities.

LPR Development Results

- New mathematics that enable increasing accuracy in modeling systems and analyzing experiments involving processes taking place across a wide range of time and length scales;
- New theories, models, simulations, and experiments that describe and validate application/workflow behavior over composable systems;
- Software tools and services that enable globally distributed resources and facilities to be rigorously combined into composed systems with well-defined mathematical properties;
- Computer science and algorithm innovations that increase the productivity, energy efficiency, and resiliency of future-generation supercomputers;
- Software, tools, SMART services, and middleware to harness more efficiently and effectively the potential of today's high performance computing systems and advanced networks for science and engineering applications;
- Networking and collaboration tools to make scientific resources readily available to a broad spectrum of scientists, including those in national laboratory, university, and industrial settings; and
- Software in the areas of numerical and non-numerical algorithms, operating systems, data management, analyses, representation model development, user interfaces, and other tools required to make effective use of future-generation supercomputers and the data sets from current and future scientific user facilities.

Goal 1 and Objectives

- **Improve the Responsiveness of Foundational Research in Applied Mathematics, Computer Science, and Networking to DOE missions**
- **Objective 1.1:** Lay the groundwork for optimizing the performance, energy efficiency, and robustness of algorithms operating on future HPC systems
- **Objective 1.2:** Enable the exploration and analysis of increasingly complex physical systems on future HPC facilities
- **Objective 1.3:** Develop decision support methodologies for management and control of future increasingly complex HPC systems

Goal 2 and Objectives

- **Maximize the Return from DOE Investments in Computing Hardware and Software (Developer Productivity)**
- **Objective 2.1:** Foster widespread use of HPC facilities by software application developers through transparent user environments
- **Objective 2.2:** Make HPC facilities run more efficiently and effectively

Goal 3 and Objectives

- **Enhance Computationally-Enabled Scientific and Engineering Results from Theoretical and Experimental Studies (Scientific Productivity)**
- **Objective 3.1:** Develop collaborative environments that can support dynamically changing, geographically-distributed, integrated theoretical/experimental teams
- **Objective 3.2:** Enable more efficient, effective data retrieval and archiving
- **Objective 3.3:** Improve management of complexity to enhance user productivity in HPC facilities

Goal 4 and Objectives

- **Capitalize on Emerging Computational Technologies**
- **Objective 4.1:** Extract as-yet unexploited performance gains from CMOS
- **Objective 4.2:** Investigate and develop component and system designs to assess how emerging technologies can be adapted for DOE target applications
- **Objective 4.3:** Pursue research advances in mathematical libraries and frameworks for emerging computational systems
- **Objective 4.4:** Investigate and develop prototype software stacks that enable high productivity for emerging computing technologies

Goal 5 and Objectives

- **Maintain and Grow the HPC Workforce**
- **Objective 5.1:** Develop and maintain a highly qualified and appropriately diverse workforce for computational facilities and mission application areas

Goal 6 and Objectives

- **Increase Contribution to US Competitiveness and to Policy Making**
- **Objective 6.1:** Make significant contributions to the HPC ecosystem
- **Objective 6.2:** Make significant contributions to the development of the science-based computational discovery and policy-making environment