

# New ARC Systems

James McClure

Advanced Research Computing  
Division of Information Technology

# Overview

“Deep Learning” System designed and acquired in partnership with the Department of Electrical and Computer Engineering

- Research computing trends
- Hardware and software overview
- System acquisition and deployment

# Research Computing Trends

- **“Big Data”** — growth rate in rate of data generation outpaces growth rate for compute

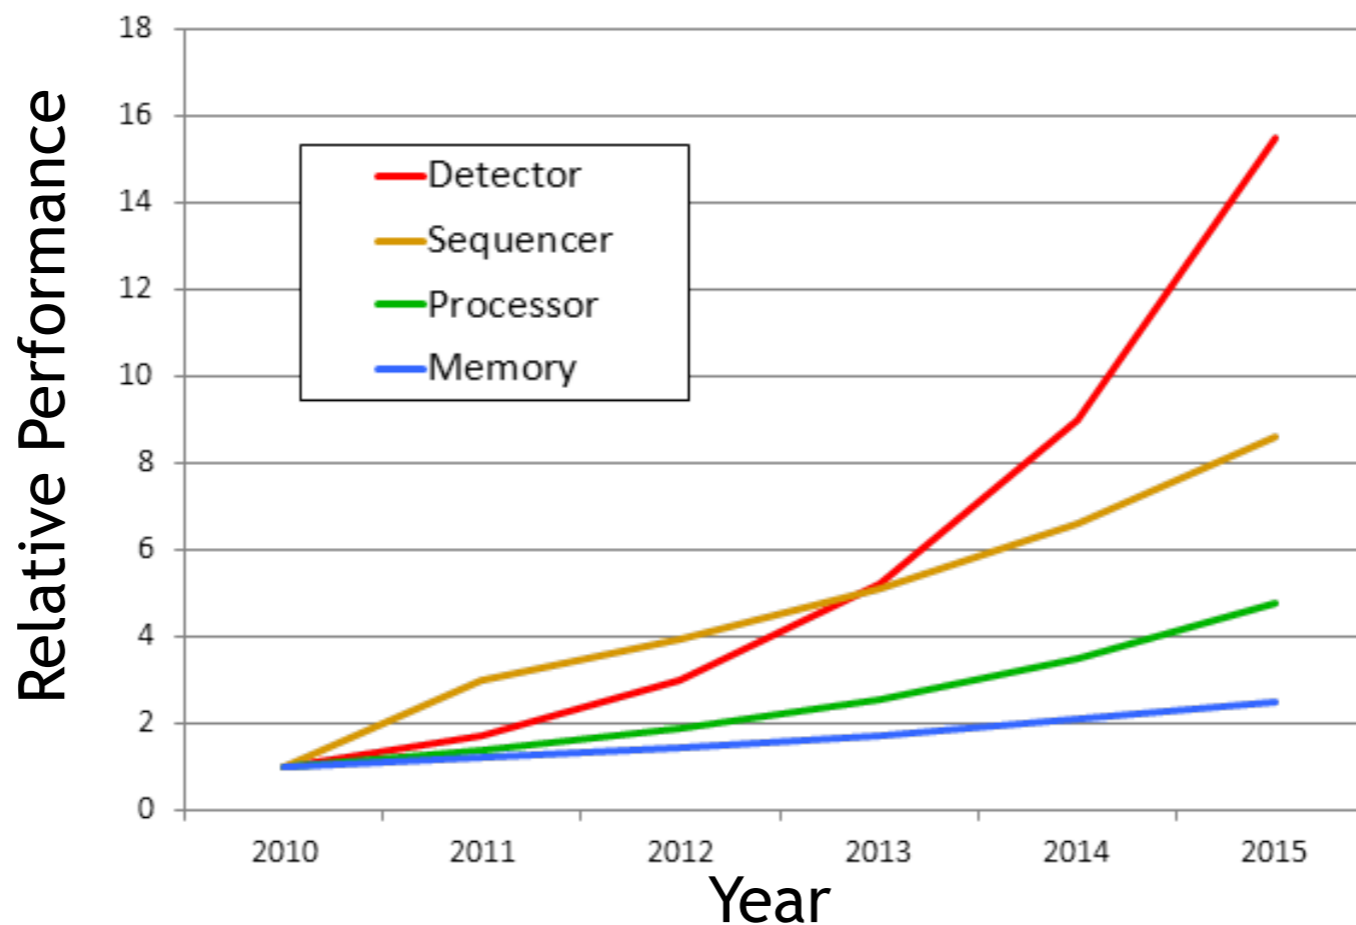


Image source: Kathy Yelick, LBNL, DOE

# Research Computing Trends

- **“Data Intensive”**— data movement constrains application performance
  - Data transfers over the network
  - I/O (bandwidth or IOPs)
  - Memory bandwidth (movement of data from memory to the processor)
- Data locality drives performance
- Both hardware and software must respect costs of moving data

# Research Computing Trends

- **“Deep Learning”** — class of algorithms built on neural networks, excel at “classification” problems
- Wide range of applications
  - image recognition
  - natural language processing
  - healthcare analytics
  - business analytics
- Algorithms are “data hungry”
- Training networks is bandwidth-intensive

# Hardware overview

## IBM S822LC Compute nodes (14)

- Dual 10-core Power8 CPU
  - 2.061 — 4.023 GHz
  - 512 GB RAM
- Four NVIDIA P100 GPU
  - 16 GB RAM per GPU
  - 720 GB/s from GPU memory
- NVLINK connects CPU and GPU memory spaces
- Mellanox EDR Infiniband interconnect (100 GB/s)
- 10GE (ethernet network)
- 1GE (management network)




Image source: nvidia

# Decision Points

- First architecture delivered by OpenPOWER alliance (IBM, NVIDIA, Mellanox)
- Hardware / software maturity a significant factor in decision process
  - Critical to identify applications in advance to support early user adoption
  - What do we know is already working in other places?

# Target Applications

- Deep Learning
- Molecular Dynamics
- Quantum Mechanics
- Computational Fluid Dynamics



Image source: nvidia



# Deep Learning Software

## IBM PowerAI Platform

### PowerAI Software Distribution

Deep  
Learning  
Frameworks

Caffe

 Caffe

IBM Caffe

 torch

 TensorFlow™

DL4J

theano

 Chainer

Supporting  
Libraries

DIGITS

OpenBLAS

Distributed  
Frameworks

Bazel

NCCL

## IBM Power System for HPC, with NVLink

Breakthrough performance for GPU accelerated applications,  
Including Deep Learning and Machine Learning



# Deployment Plan (UAS/ARC)

(PO placed in January 2017)  
(System Delivered February 6)

## Phase I (March 15)

- Basic system configuration
- Deep learning application stack

## Phase II (May 15)

- Support for general applications

# Phase I Deployment

- Provision OS — Ubuntu 16.04 (Cobbler)
- Configuration management (Ansible)
- Repository management for Debian (Aptly)
- Install GPFS (mounted over 10GB ethernet)
- Install NVIDIA drivers, CUDA, cuDNN
- Install and configure scheduler (Slurm)
- Install and test PowerAI platform
- Early users (focused on deep learning)
- Documentation for Slurm, PowerAI

# Phase II Deployment

- Install module environment
- Install OFED drivers
- Create customized build environment for debian packages
- Set up chroot to manage software installation for non-root users
- Install multiple compilers
- Build and test MPI (EDR Infiniband)
- Build and test general application software

# Current Status

- Early users are on the system
- Mid-way through Phase II deployment
- Trying to resolve as many issues as possible prior to full system release
- Build out support for application software