U.S. DEPARTMENT OF ENERGY | Office of Science

ASCAC SUBCOMMITTEE ON SYNERGISTIC CHALLENGES IN DATA-INTENSIVE SCIENCE AND EXASCALE COMPUTING

Vivek Sarkar, Ph.D.

Rice University

Subcommittee chair

DOE ASCAC Subcommittee Report
March 5, 2013

Subcommittee Members

Last name	First name	Affiliation
Chen (*)	Jacqueline	Sandia
Choudhary	Alok	Northwestern U.
Feldman	Stuart	Google
Hendrickson	Bruce	Sandia
Johnson	Chris	U. Utah
Mount	Richard	SLAC
Sarkar (**)	Vivek	Rice U.
White (*)	Victoria	FermiLab
Williams (*)	Dean	LLNL

(*) ASCAC member, (**) Subcommittee chair

Our Charge

Department of Energy
Office of Science
Washington, DC 20585

Office of the Director

July 25, 2012

Professor Roscoe Giles, ASCAC Chair
Department of Electrical & Computer Engineering
Boston University
8 St. Mary's Street
Boston, MA 02215

Dear Professor Giles:

Thank you for the recent Advanced Scientific Computing Advisory Committee (ASCAC) report on the Computational Sciences Graduate Fellowship. The report was thorough, informative and very timely.

Overcoming the challenges of managing data rates and movement of data in an exascale computing environment will likely require significant research investments. In addition to the challenges and opportunities of exascale computing, the Office of Science is facing related challenges from data-intensive research activities, such as the growing volumes of data generated at our next generation scientific user facilities and by the new genomics-based technologies that are enabling a revolution in systems biology research. The Linac Coherent Light Source, for example, currently generates several petabytes of data each year and the National Synchrotron Light Source II, currently under construction and scheduled to begin operations later this decade, is expected to generate hundreds of petabytes of data each year. In order to maximize the return on our limited federal resources, we need to understand the similarities among and differences between these data challenges and the potential to leverage research investments to address issues spanning both exascale and data-intensive science.

By this letter, I am charging the ASCAC to assemble a subcommittee to examine the potential synergies between the challenges of data-intensive science and exascale. The subcommittee should take into account the Department's mission needs, which define the Office of Science's unique role in data-intensive science vis-a-vis other agencies. The subcommittee should specifically address what investments are most likely to positively impact both our exascale goals and our data-intensive science research programs, including data management at our next generation facilities.

I would appreciate the committee's preliminary comments by November 2012 and a final report by March 30, 2013. I appreciate ASCAC's willingness to undertake this important activity.

If you have any questions regarding this matter, please contact either Daniel Hitchcock, the Associate Director of the Office of Science for ASCR or Christine Chalk, the Designated Federal Official for the ASCAC.

Sincerely,

W. F. Brinkman
Director, Office of Science

Interpretation of Charge

- There are several federal and commercial initiatives under way to address the challenges of big data
  - Scope of this study was restricted to the intersection of big data challenges and exascale computing in the context of data-intensive science applications
- Timeframe assumed was ~ 2022 (when exascale capability is expected)
  - Nearer term considerations also received attention since they will influence the migration path to exascale computing.
- Charge asked for identification of investments that are most likely to positively impact both data-intensive science research goals and exascale computing goals.
  - Since Facilities is the focus of another ASCAC study, this study focused on investments that can leverage synergies between data-intensive science and exascale computing rather than on facilities.

Context for our charge

- Several recent ASCR workshops have focused on exascale & data-intensive computing

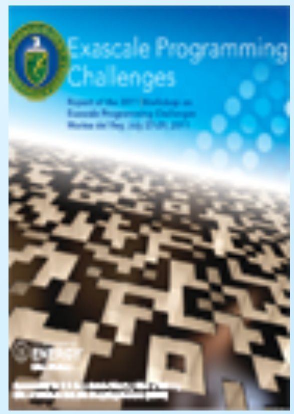

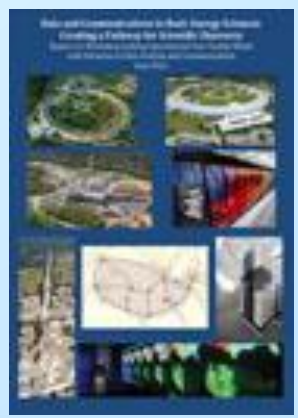

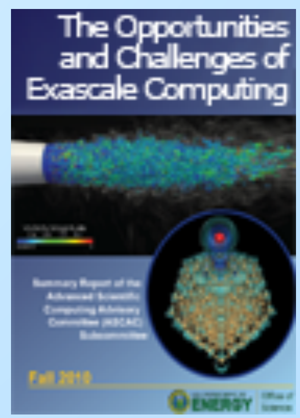

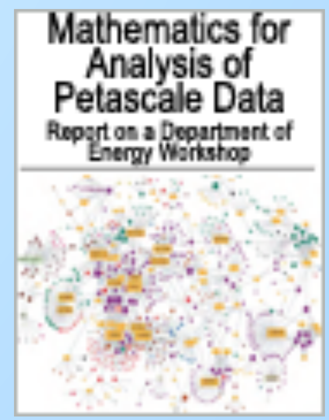

Big Data and the Fourth Paradigm

- Beyond theory, experiments, and simulations
- ***Data-driven science*** ➔ use information in massive datasets to drive scientific discovery
- Challenges arise from increasing velocity, heterogeneity, and volume of data generation

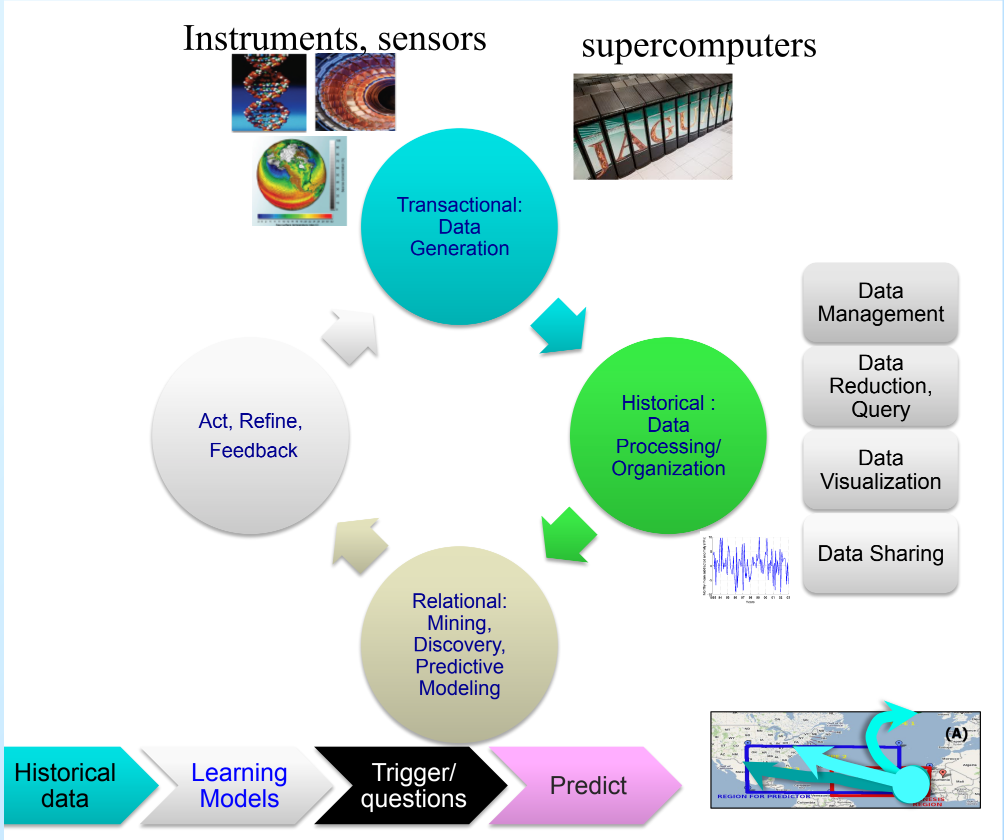

Data Challenges in Science

Overall trend: most science domains will become data-intensive in the exascale timeframe (and many well before then)

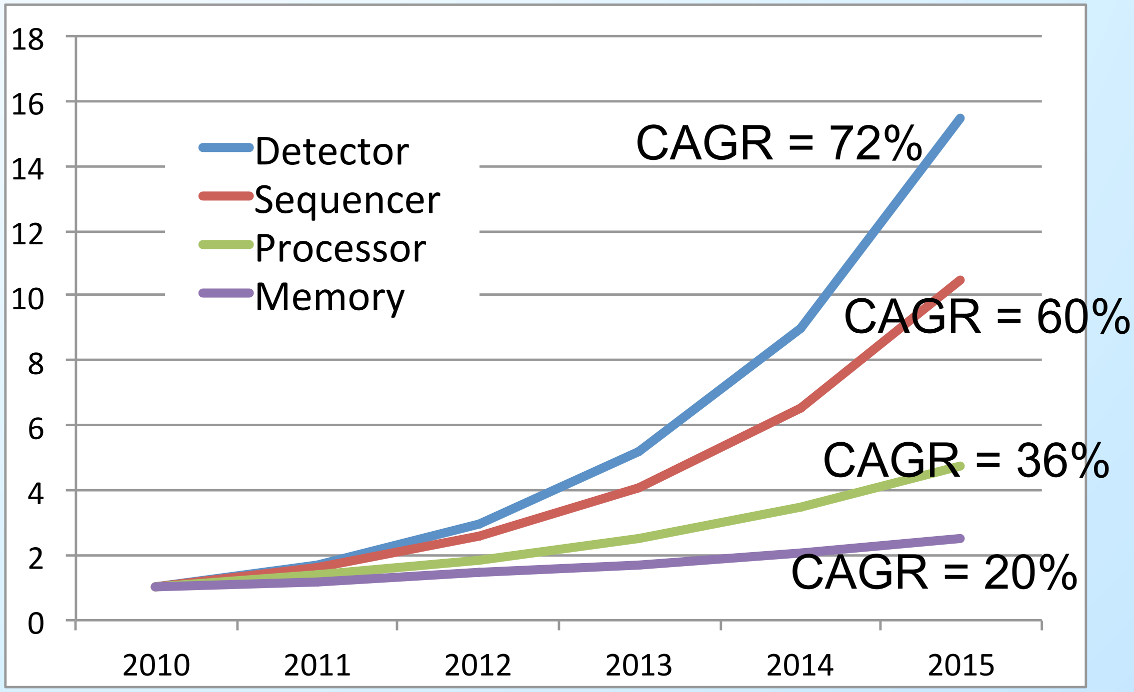

Source: notional figure, courtesy of Kathy Yelick

Genomics

Data Volume increases to 10 PB in FY21

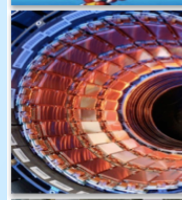

High Energy Physics (Large Hadron Collider)

15 PB of data/year

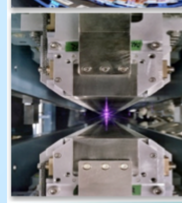

Light Sources

Approximately 300 TB/day

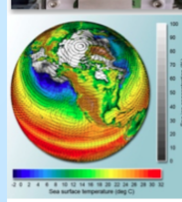

Climate

Data expected to be hundreds of 100 EB

Source: Bill Harrod, SC12 plenary presentation

Data Challenges in High Energy Physics: Large Hadron Collider exemplar

- ATLAS and CMS detectors generate analog data at rates equivalent to 1PB/second
- Output rate after *data reduction* is 1GB/second ~ 10PB/year
- Storage of cumulative derived data, simulated data, replicated data is currently ~ 100PB, and is rapidly increasing
- Workflow: homogeneous community of physicists access read-only shared data using the Worldwide LHC Computing Grid (WLCG)

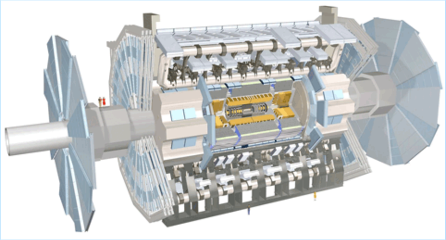

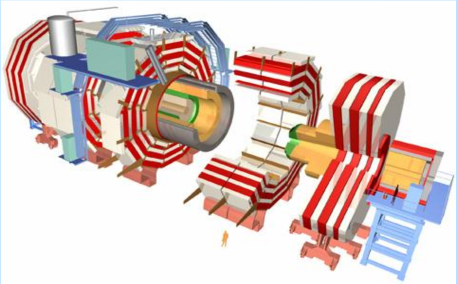

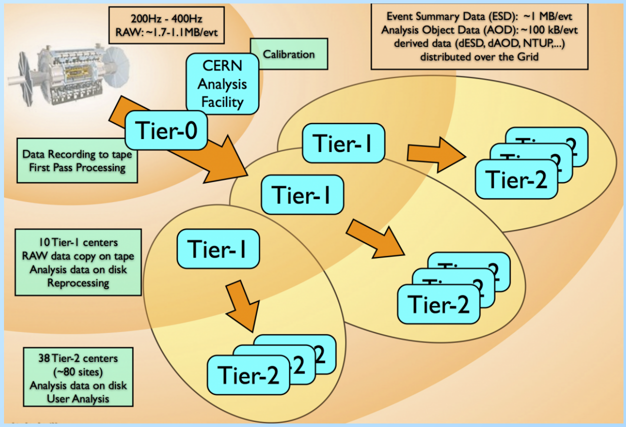

Data Challenges in Climate Science

Federated data enterprise system with significant challenges

- Velocity: distributing live data streams and large volume data movement quickly and efficiently
- Volume: analyzing large-volume data in-place for big data analytics
- Heterogeneous workflows: on-demand data products for heterogeneous communities (scientists, policy makers, farmers, insurance industry, …)

Earth System Grid Federation (ESGF) manages several petabytes of data

Simulation

Observation

ESGF
Earth System Grid Federation

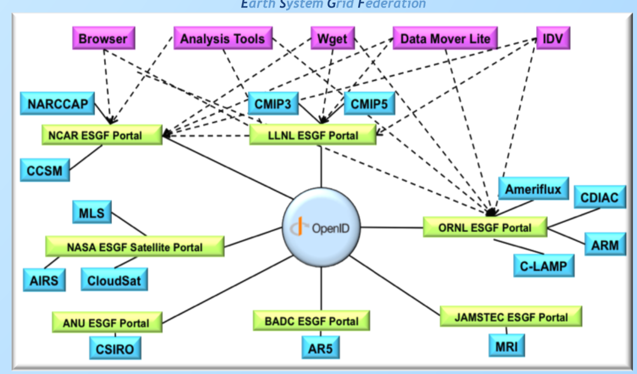

Data Challenges in Large-Scale Simulations: S3D Combustion code exemplar

- Goal: simulate turbulence-chemistry interaction at conditions that are representative of realistic systems
  - High pressure
  - Turbulence intensity
  - Turbulent length scales
  - Sufficient chemical fidelity to differentiate effects of fuels
- Exascale simulation will require 3PB of memory, and will generate 400PB of raw data (1PB every 30 minutes)
- Workflow challenges include co-design for simulation and in-situ analyses

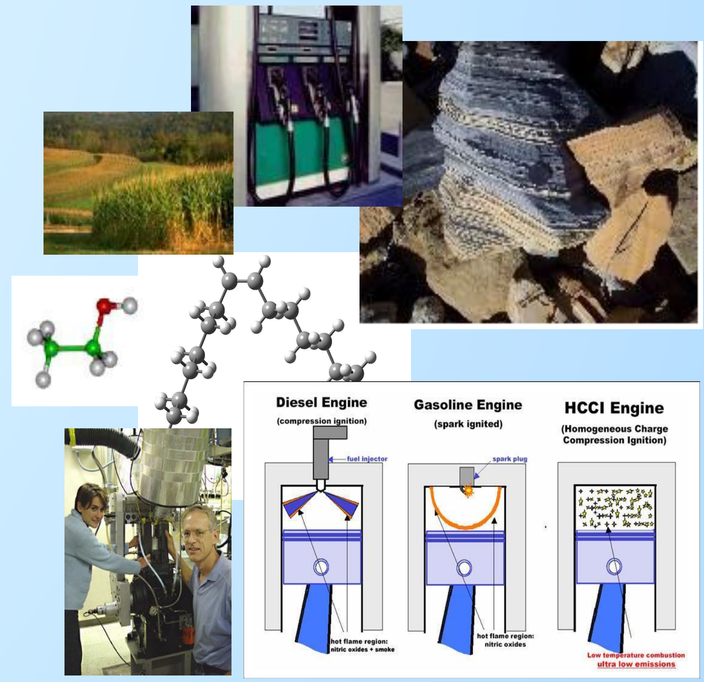

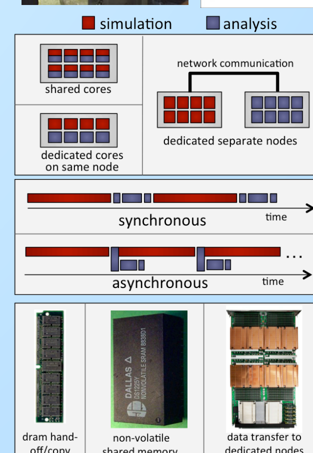

Data Challenges in Biology and Genomics: KBase exemplar

- Data-intensive challenges include
  - Biophysical simulations of cellular environments
  - Cracking the 'signaling code' of the genome across the tree of life (reconstruction of cellular networks across species)
  - Reverse engineering the human brain

- KBase center currently manages about 2 petabytes of data (plant genomes, process data) for genomics research; workflow based on a service-based infrastructure
- Significant differences between data characteristics in Kbase and other domains (lots of integer data, random access, large intermediate data size during computations, poor locality in cross-correlation, …)

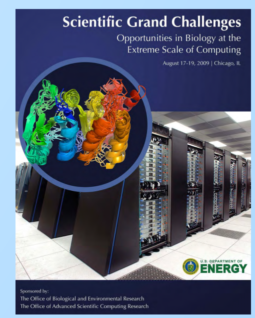

KBase
PREDICTIVE BIOLOGY
DOE Systems Biology Knowledgebase

Data Challenges in Light Sources: APS and LCLS exemplars

Advanced Photon Source (APS)

- Includes about 65 beam lines, with ~ 1TB of data generated per day
  - Future light sources are expected to generate data at the rate of 1TB per second
- GridFTP and GlobusOnline services help some APS users with their workflow, but many others bring their own storage devices and perform manual analysis of their data

Linac Coherent Light Source (LCLS)

- Provides users access to ~ 2.5PB storage facility via LCLS portal, where data is stored for 2 years, and an on-line cache of ~ 50TB, where data is stored for 5 days.
- These volumes are expected to increase dramatically in the future

Data Analysis and Visualization:
From Big Data to "Big Information"

"Hence a wealth of information creates a poverty of attention and a need to allocate that attention efficiently among the overabundance of information sources that might consume it."
--- Herbert Simon, Designing Organizations for an Information-Rich World

- Widening gap between I/O and computational rates will make *in-situ* analysis & visualization a necessity for exascale

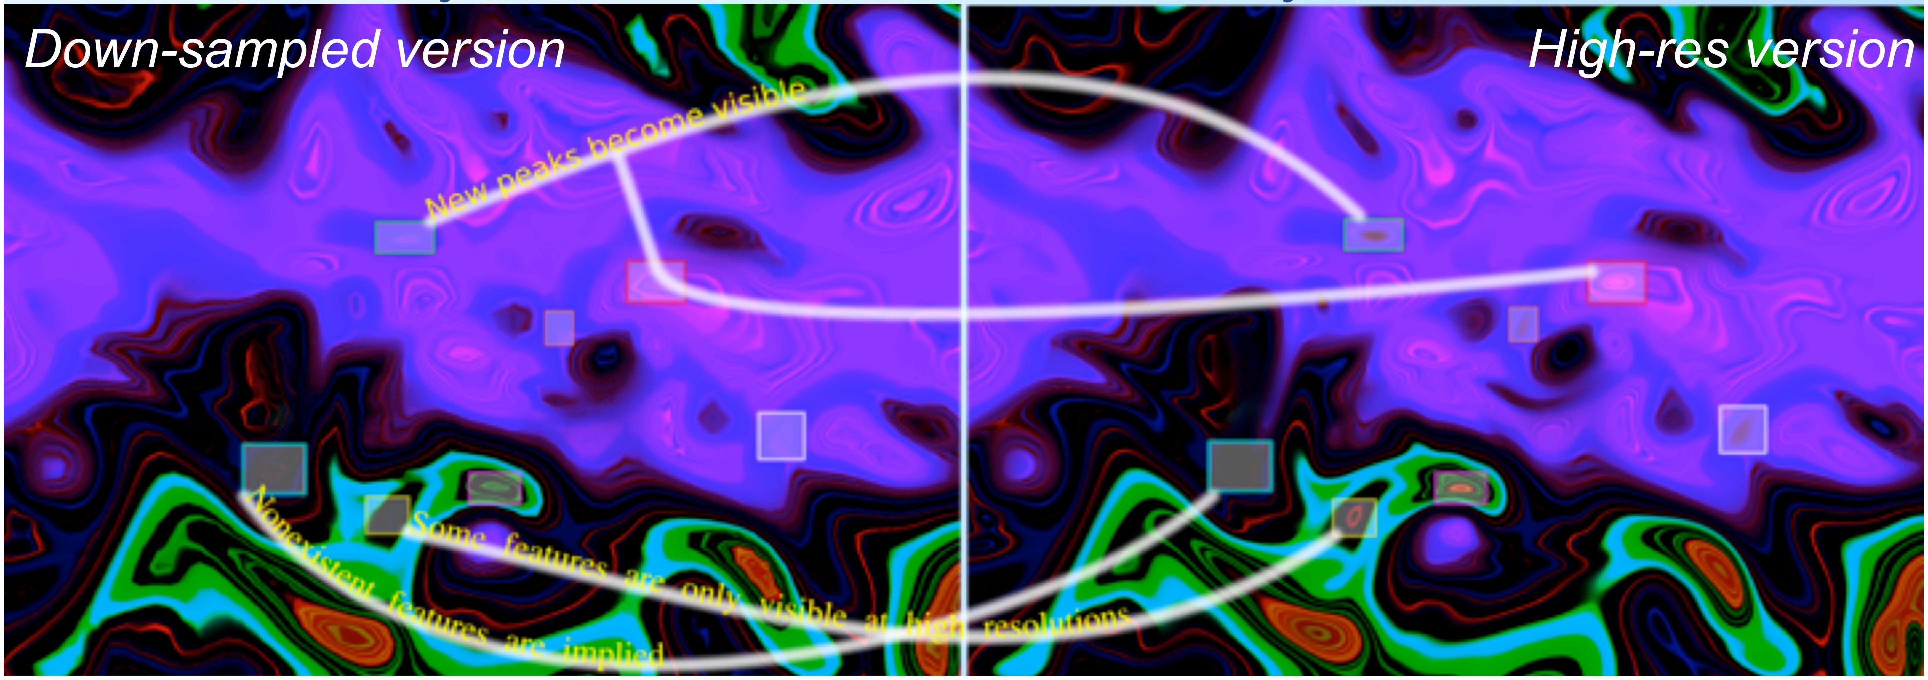

Topological Analysis & Volume Visualization of Combustion simulation

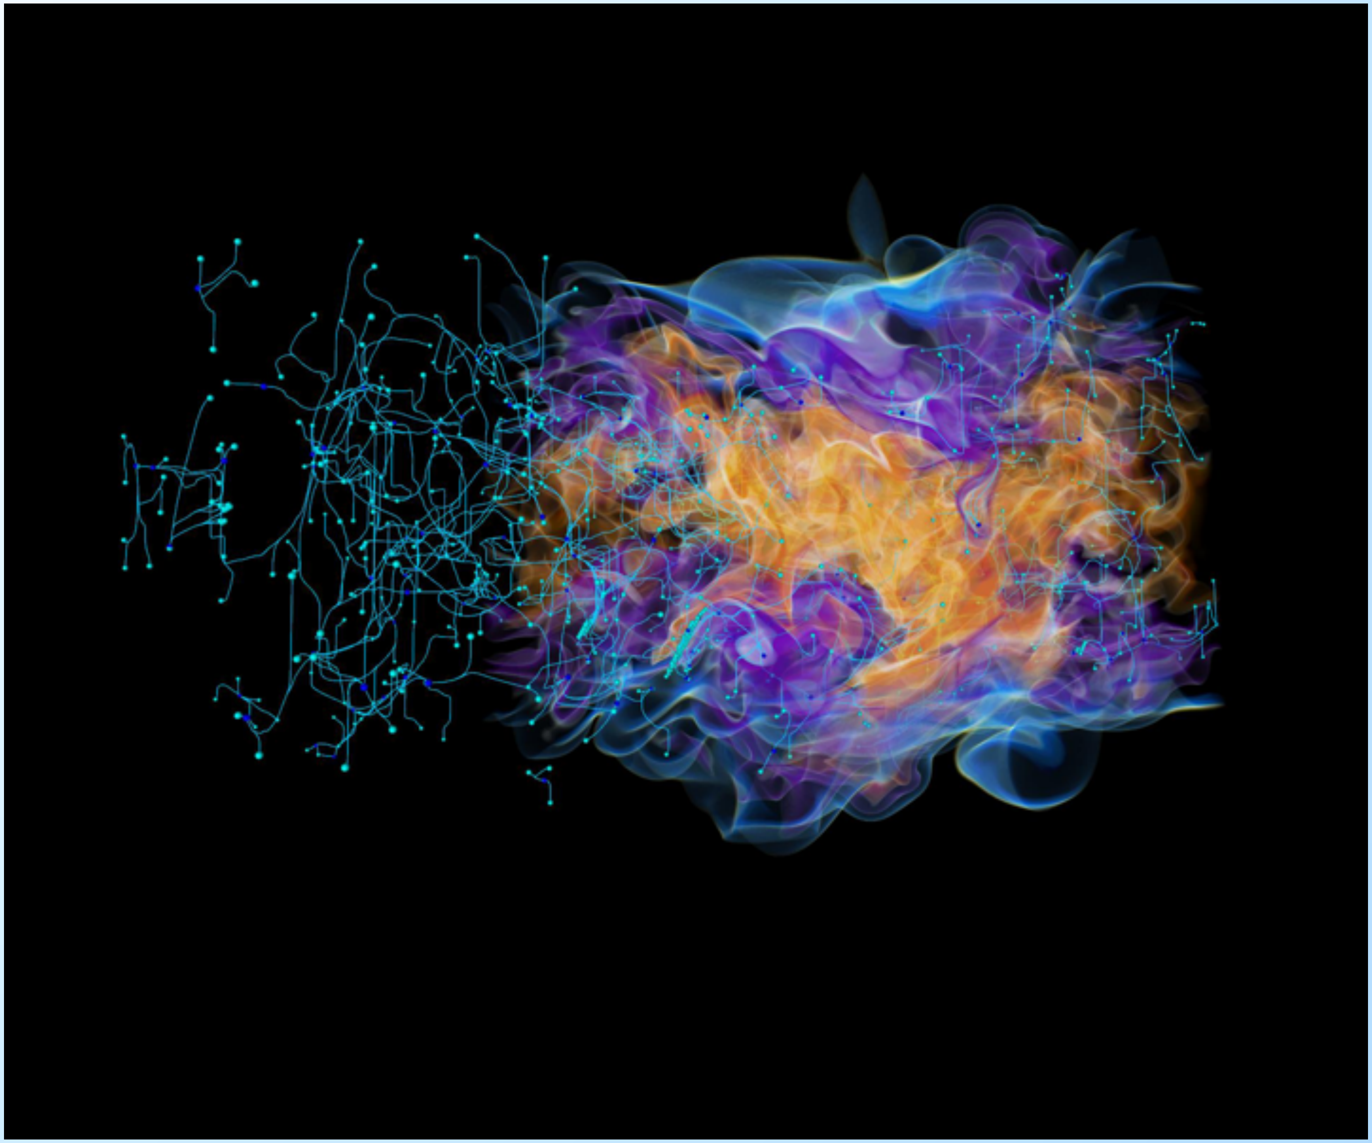

Custom Visualizations for Biological Data Analysis

M. Meyer et al., EuroVis 2010. Pathline

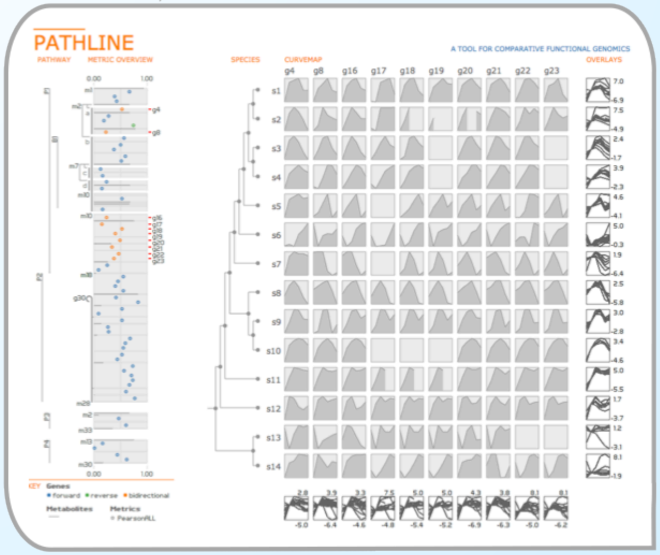

MizBee *M. Meyer et al., InfoVis 2009.*

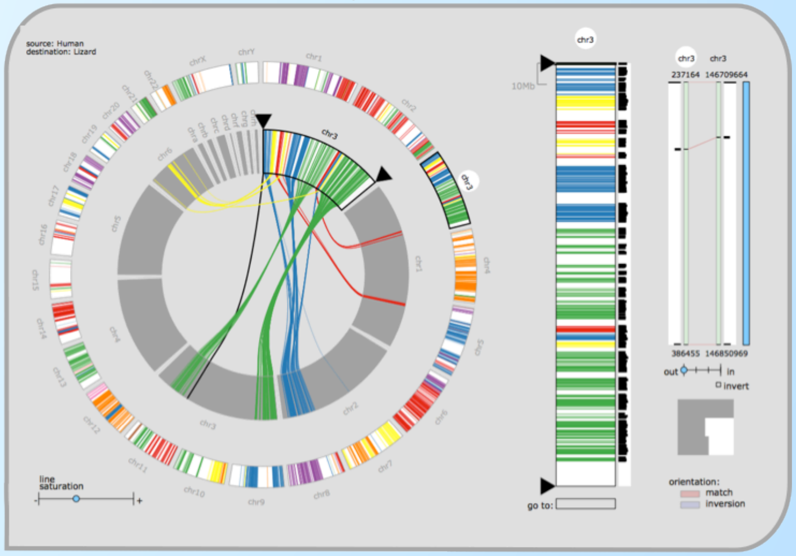

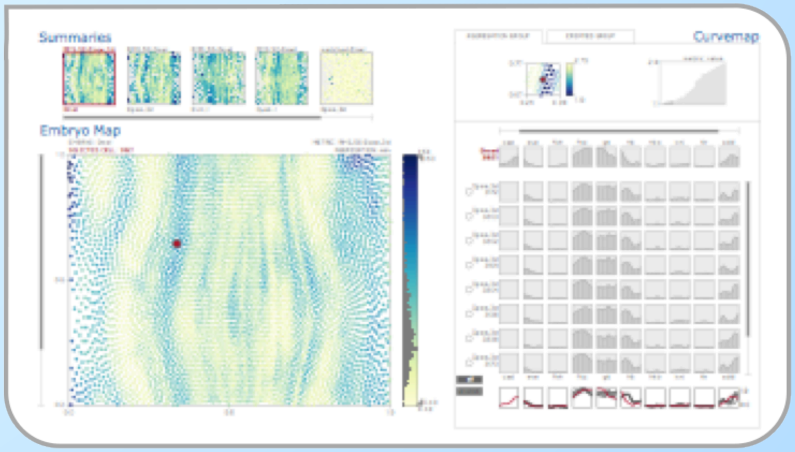

InSite

MulteeSum *M. Meyer et al., InfoVis 2010.*

Data Streaming and Near-Sensor Computing

Data Streaming exemplar

- ORNL Spallation Neutron Source (SNS)
- Challenge: reduction and visualization of some of the large SNS data sets take hours after data has been collected
- ADARADA streaming data system provides in-situ reduction of data as it is generated from the instrument
  - Challenges in in-situ reduction is synergistic with data movement challenges in exascale computing

Near-sensor computing exemplars

- HEP, radio telescopes, light sources, …
- Triggers detect events of interest to be recorded
- Filters reduce data as close to the instrument as possible
- After data has been reduced by triggers and filters, it is curated and archived for re-processing and re-analysis

Intertwined requirements for Big Data and Exascale Computing (Big Compute)

- Big Data will be analyzed by Big Compute
  - Big Data generated by the data-driven paradigm will need to be analyzed by Big Compute (exascale or extreme-scale) systems
    - “Extreme-scale systems” refer to all classes of systems built using exascale technologies
- Big Compute will generate Big Data
  - Data-intensive simulations on Big Compute (exascale) systems will generate volumes of Big Data comparable to data generated by the largest science experiments
- Data-driven (big data) and data-intensive (big compute) approaches have evolved somewhat independently of each other
  - Important for each to learn lessons from the other because their fates are intertwined

Cross-cutting Issues

- Data lifecycle – retention, preservation, sharing
- Software challenges
- Technology disruptions e.g., new analysis algorithms
- Provenance, metadata, security, privacy
- Expertise and skills gap

18

Findings

Finding 1: There are opportunities for investments that can benefit both data-intensive science and exascale computing.

- Data-intensive science relies on the collection, analysis and management of massive volumes of data, which will have to be performed by exascale systems or, more generally, by "extreme-scale" components of exascale systems.
- Extreme-scale components will include innovative memory hierarchies and data movement optimizations that will be essential for analysis components of data-intensive science workflows.
- Additional synergy between algorithms for near-sensor computing in experimental facilities and algorithms for in-situ analysis in simulations.

Findings (contd)

Finding 2: Integration of data analytics with exascale simulations represents a new kind of workflow that will impact both data-intensive science and exascale computing.

- Exascale simulations require in-situ analysis and visualization, thus necessitating a new kind of workflow for scientists.
- In-situ analysis will impact the workloads that high-end computers have traditionally been designed for e.g., increases in integer and branch operations than before.
- Tighter integration of simulation and analytics in the science workflow will impact co-design of these systems for future workloads, and will require development of new classes of proxy applications to capture their combined characteristics.

Findings (contd)

Finding 3: There is an urgent need to simplify the workflow for data-intensive science.

- Workflow needs to integrate best computational algorithms with the best interactive techniques and interface, while reducing complexity.
- Workflow should be able to transparently support decisions such as when to move data to computation or computation to data.
- Recent proposal for a Virtual Data Facility (VDF) will significantly help in this regard.

Findings (contd)

Finding 4: There is a need to increase the pool of computer and computational scientists trained in both exascale and data-intensive computing.

- Earlier workflow models allowed for a separation of concerns between computation and analytics that is no longer possible
  - This separation of concerns allowed science to progress with personnel that may be experts in computation or in analysis, but not both.
- This approach is not sustainable in data-intensive science where the workflow for computation and analysis will have to be co-designed.
- Need for increase in number of computer and computational scientists trained in both exascale and data-intensive computing.

Recommendations

Recommendation 1: The DOE Office of Science should give higher priority to investments that can benefit both data-intensive science and exascale computing so as to leverage their synergies.

- For science domains that need exascale simulations, commensurate investments in exascale computing capabilities and data infrastructure are necessary.
- In other domains, extreme-scale components of exascale systems are necessary for near-sensor computing and other tiers of data analysis.
- Innovations in algorithms to address fundamental challenges in concurrency, data movement, and resilience will benefit data analysis and computational techniques for both data-intensive science and exascale computing.

Recommendations (contd)

Recommendation 2: DOE ASCR should give higher priority to investments that simplify the science workflow and improve the productivity of scientists involved in data-intensive and exascale computing.

- Must pay greater attention to simplifying human-in-the-loop workflows for data-intensive science.
  - Virtual Data Facility (VDF) will provide a simpler portal for data services than current systems.
- Recommend development of libraries of scalable data analytics and data mining algorithms and software components for use in workflows.
- Recommend creation of new classes of proxy applications to capture the combined characteristics of simulation and analytics to feed into future design/co-design activities.

Recommendations (contd)

Recommendation 3: DOE ASCR should adjust investments in programs such as fellowships, career awards, and funding grants, to increase the pool of computer and computational scientists trained in both exascale and data-intensive computing.

- There is a significant gap between the number of current computational and computer scientists trained in both exascale and data-intensive computing and the future needs for this combined expertise in support of DOE's science missions.
- ASCR investments such as fellowships, career awards, and funding grants should look to increase the pool of computer and computational scientists trained in both exascale and data-intensive computing.

Conclusions

- This study reviewed current practice and future plans in multiple science domains in the context of the Big Data and the Exascale Computing challenges that they will face in the future, drawing from public presentations, workshop reports and expert testimony.
- Data-intensive research activities are increasing in all domains of science, and exascale computing is a key enabler of these activities.
- The report includes key findings and recommendations from the perspective of identifying investments that are most likely to positively impact both data-intensive science goals and exascale computing goals.