The AI/BD/HPC Technology Intersection & Precision Medicine

Dimitri Kusnezov
CEA/TeraTec
March 6, 2018

In the past several years, we have taken advantage of a number of opportunities to push the intersection of next generation high-performance computing technologies and our approaches to precision medicine. Today we are in the throes of piecing together what is likely the most unique intersection of medical data and computer technologies. But more deeply, the traditional paradigm of computer simulation needs fundamental revision. This is the time for a number of reasons. I will review what the drivers are, why now, how this has been approached over the past 2-3 years, and where we are heading.

Disclaimer: These views are mine and not of the DOE.

Prediction is part of your day to day lives…

- First prediction of solar eclipse by Thales of Miletus *(585BC)*
- Will PSG beat Real Madrid today without Neymar?
- Tide tables at Mont St Michel
- Lifetime of first excited state of hydrogen
- Laminar-/turbulent-flow drag of a thin plate,
  aligned with flow direction
- Performance prediction for a new aircraft
- Will the Seine flood above 6m soon?
- Climate prediction
- Astrology

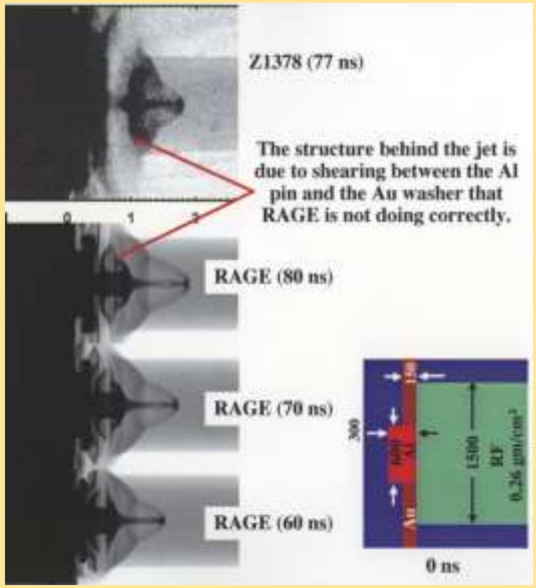

… But consequences of poor predictions are not all the same.

Today we use traditional HPC to make science based predictions to inform increasingly serious problems.

Why are we thinking about this?

Precision Medicine

Department of Energy (DOE) and the Life Sciences

A long history rooted in the development of extensive radiation biology programs. Expertise in the biosciences remains core to our responsibilities. We now turn to biology to tackle problems from clean energy production, to determining genomic properties, molecular and regulatory mechanisms, and the resulting functional potential of microbes, plants, and biological communities central to DOE missions.

Perhaps the best known contribution to life-sciences is our role in sequencing the human genome. Last year we celebrated the 25th anniversary of its start – something that has been transformative in domains from the economy to precision medicine.

- In 1986, DOE announced the Human Genome Initiative
- In 1990 DOE & NIH joined forces and launch the Human Genome Project (HGP)
- In 2000 we completed the working draft with the HGP team and Celera
- In 2003, memorably, the finished human genome was delivered on the 50th anniversary of the Watson-Crick paper

The willingness to attack complex and significant problems is a hallmark of DOE. Since the completion of the first sequence, the cost to do this has come down a factor of one million – something reflective of the robust ecosystem that has been created in this field due to this initiative.

A few remarks on DOE and Computing

Simulation is at the core of our mission success

- *Many of the mission areas, from energy and emergency response/preparedness, to nuclear security and science, involve classes of problems that cannot be tested other than virtually.*
- *DOE and the labs have equities in outcomes; Context of the missions provides intellectual headroom and career long commitments to the outcomes*
- *'Go to' agency for informing urgent decisions*

This is the only way to approach problems that cannot be instrumented

- *Built around Uncertainty Quantification, Verification & Validation*
- *Many successes. Billions in deferred costs. Likely more in the future*
- *Innovation in our missions will require virtual tools to explore ideas*
- *Rigor in prediction - a place we must continue to pioneer.*
- *Problems typically have a cost for inaction.*

Sector-specific agency for Energy Infrastructure *(PDD-21)*

- *Resiliency' against natural and man-made conditions*
- *Cross-sector dependencies*

Opportunity for broader societal impacts

- *Extreme power requirements becomes a barrier to innovation.*

We announced our next steps towards “exascale” (10^{18} floating operations/second) in late 2014…

Next systems: Sierra (LLNL), Summit (ORNL), Aurora (ANL)

Department of Energy

Department of Energy Awards $425 Million for Next Generation Supercomputing Technologies

NOVEMBER 14, 2014

Home » Department of Energy Awards $425 Million for Next Generation Supercomputing Technologies

WASHINGTON — U.S. Secretary of Energy Ernest Moniz today announced two new High Performance Computing (HPC) awards to put the nation on a fast-track to next generation exascale computing, which will help to advance U.S. leadership in scientific research and promote America’s economic and national security.

Secretary Moniz announced $325 million to build two state-of-the-art supercomputers at the Department of Energy’s Oak Ridge and Lawrence Livermore National Laboratories. The joint Collaboration of Oak Ridge, Argonne, and Lawrence Livermore (CORAL) was established in early 2014 to leverage supercomputing investments, streamline procurement processes and reduce costs to develop supercomputers that will be five to seven times more powerful when fully deployed than today’s fastest systems in the U.S. In addition, Secretary Moniz also announced approximately $100 million to further develop extreme scale supercomputing technologies as part of a research and development program titled FastForward 2.

MOBILE

IBM and Nvidia win $425M to build two monstrous supercomputers for the Department of Energy

DEAN TAKAHASHI @DEANTAK NOVEMBER 14, 2014 7:00 AM

The Department of Energy has awarded $425 million in federal funding to IBM, Nvidia, and other companies who will build two giant supercomputers many times more powerful than today’s most powerful machines.

Ernest Moniz, U.S. Secretary of Energy, announced Friday that the government will fund two high-performance computing supercomputers as part of a drive toward exascale computing, or a machine capable of one exaflop, or 10 to the 18th power floating point operations per second. Moniz said that the investments will help ensure the U.S. will have scientific, economic, and national security for future generations. The unstated goal: We must have more computing power for scientific research than other countries can marshal, because data is power.

About $325 million of the money will go toward two state-of-the-art supercomputers at the Department of Energy’s Oak Ridge and Lawrence Livermore National Laboratories. Those computers are expected to be five to seven times more powerful than today’s fastest systems in the U.S. Those systems are needed because, well, big data is getting bigger.

…as well as post Moore's law investments

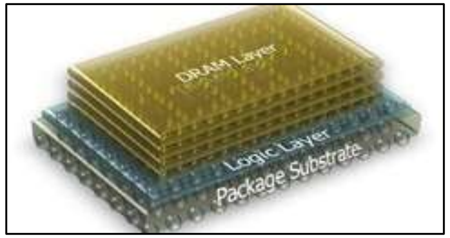

Memory intensive architectures

Neuromorphic learning systems

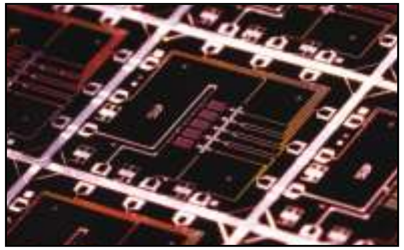

Quantum information systems

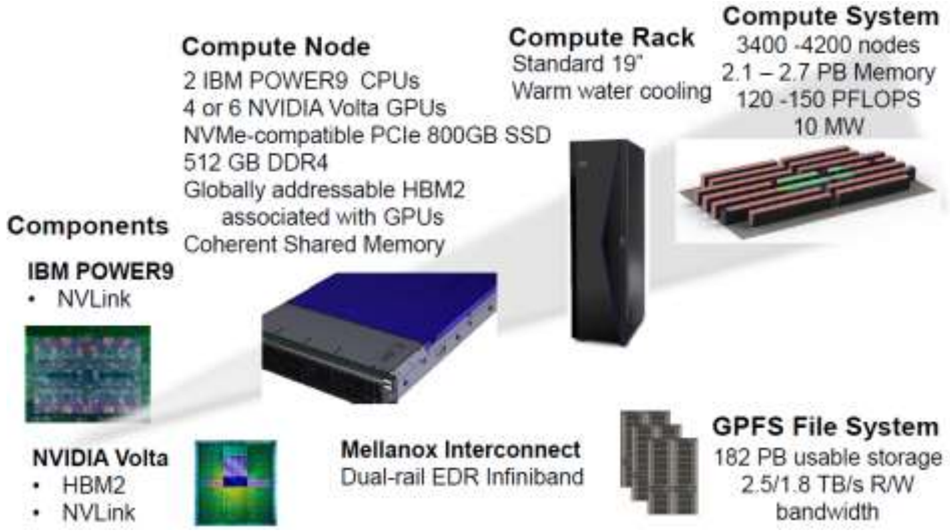

16 million neurons and 4 billion synapses and consume the energy equivalent of a hearing-aid battery – a mere 2.5 watts of power.

Recognized the node coherence on the CORAL nodes provided a means to rethink fundamentally the architectures in early 2015 based on reinventing UQ:

Started to recognize a growing set of opportunities:

- Data beyond what we can handle is around the corner
- Beyond Moore’s Law technology progress
- Node coherence and PCIe slots on CORAL that was part of early delivery in mid-2015
- No methods for Uncertainty Quantification (UQ) and AI

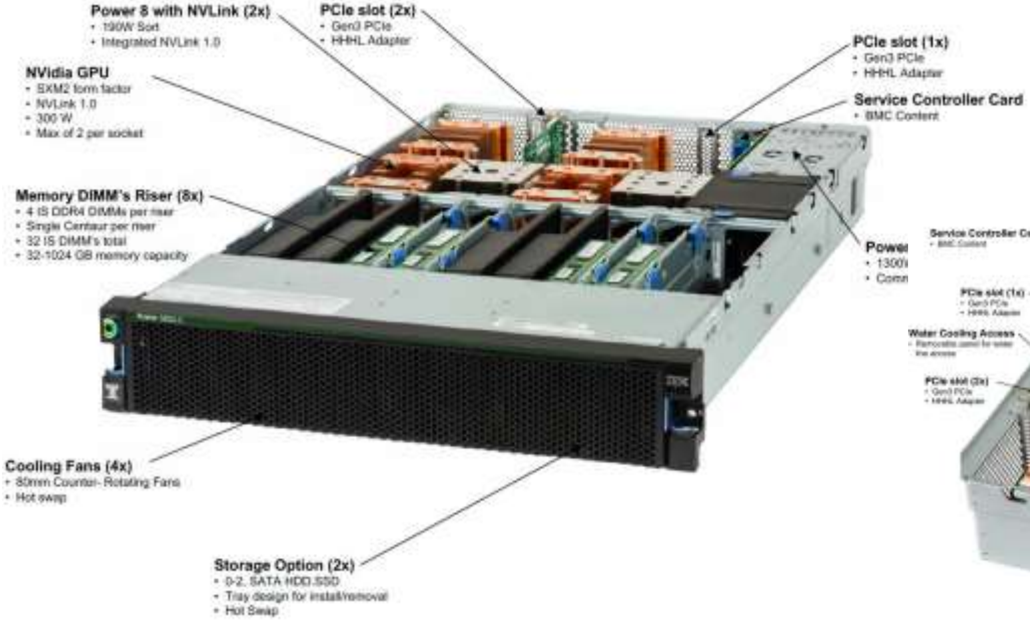

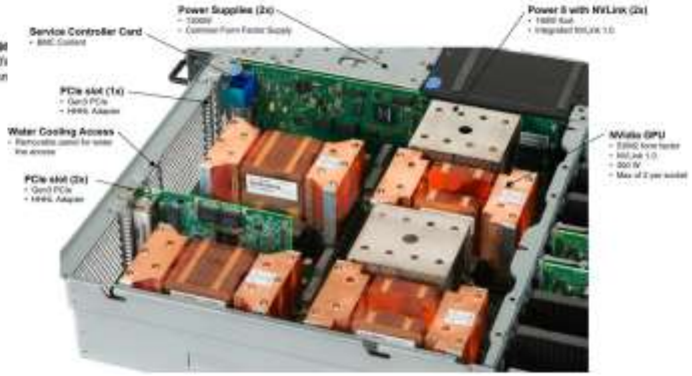

16 million neurons and 4 billion synapses and consume the energy equivalent of a hearing aid battery – a mere 2.5 watts of power.

Some of my concerns with Uncertainty Quantification today

- My focus has been on pragmatic solutions since we have missions to deliver
- But today I believe that our model is dated

- Usually start in a space of discretized space of guiding equations
- Run a deterministic code end to end
- Bayesian approaches, Latin hypercube,…
- Post-hoc add-on of tools to: accrue uncertainty, better sample initial conditions, parameters,…

- UQ is also a hard problem. Identifying the underlying dynamical equation from any amount of experimental data, however precise is a provably computationally hard problem (it is NP hard), both for classical and quantum mechanical systems.

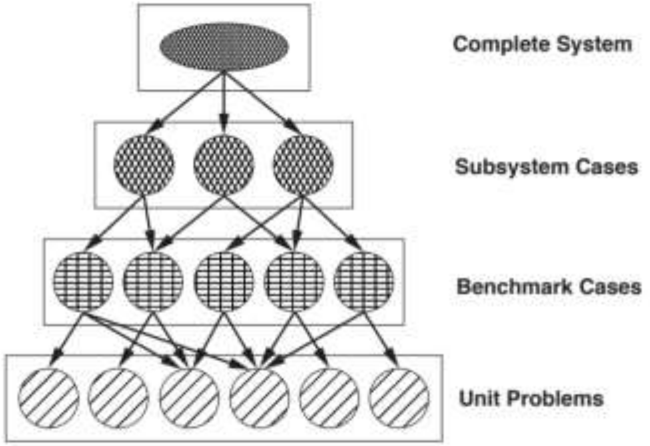

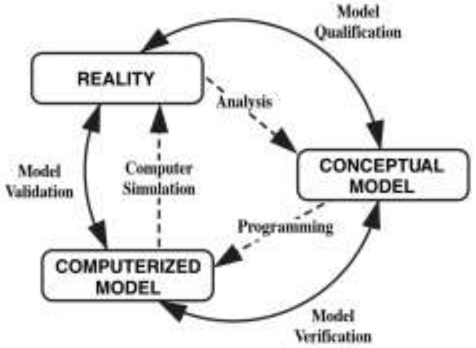

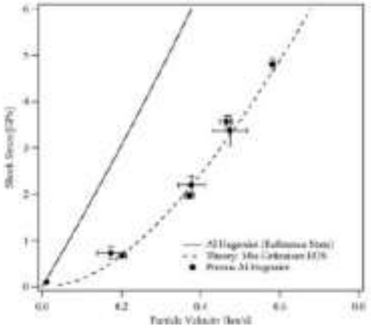

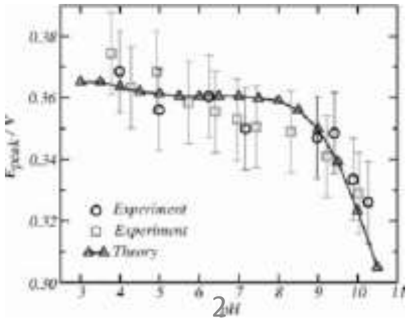

But we were challenged with mistiming

- DOE’s missions, from nuclear and national security, to energy and science, will be are not yet at the point where this is visibly urgent. I expect this in the 2020s
- However, industry today is making technology choices that are starting to define their architectural directions.

How do we find ways to innovate in next generation supercomputing in domains rich with data but short in technology?

Use potential attractors as a strategy to draw in broader thinking and resources

- Force the rethinking of traditional paradigms by challenging them with qualitatively new classes of prediction and a richness of data
- Use the qualities of data to change how we think of many of our traditional approaches from architectures to Uncertainty Quantification (UQ) to codes
- Align with where next economic drivers could provide most amplification
- Use Codesign as a philosophy
- Bring in DOE’s multiphysics labs
- Use ESNet
- Public Private partnerships
- Tech sector engagement: IBM, Intel, GE, Nvidia…
- Use legislative support: 21st Century CURES Act
- Co-lead of Data/Tech Track for VP

Pathways opened up to try to drive innovation and shared risk

THE WHITE HOUSE

Office of the Press Secretary

For Immediate Release July 29, 2015

EXECUTIVE ORDER

\- - - - - - -

CREATING A NATIONAL STRATEGIC COMPUTING INITIATIVE

By the authority vested in me as President by the Constitution and the laws of the United States of America, and to maximize benefits of high-performance computing (HPC) research, development, and deployment, it is hereby ordered as follows:

Section 1. Policy. In order to maximize the benefits of HPC for economic competitiveness and scientific discovery, the United States Government must create a coordinated Federal strategy in HPC research, development, and deployment. Investment in HPC has contributed substantially to national economic prosperity and rapidly accelerated scientific discovery. Creating and deploying technology at the leading edge is vital to ... spurring innovati... National Strategi... whole-of-governme... multi-agency stra... executed in colla... maximize the bene...

Over the pas... been maintained t... and deployment of... performance on ap... industry, and aca... coming decades wi... increasing demand... challenges and op... and competition w... require a cohesiv... Government and a...

Sec. 2. Objectives. Executive departments, agencies, and offices (agencies) participating in the NSCI shall pursue five strategic objectives:

(1) Accelerating delivery of a capable exascale computing system that integrates hardware and software capability to deliver approximately 100 times the performance of current 10 petaflop systems across a range of applications representing government needs.

(2) Increasing coherence between the technology base used for modeling and simulation and that used for data analytic computing.

(3) Establishing, over the next 15 years, a viable path forward for future HPC systems even after the limits of current semiconductor technology are reached (the "post-Moore's Law era").

(3) The United States must adopt a whole-of-government approach that draws upon the strengths of and seeks cooperation among all executive departments and agencies with significant expertise or equities in HPC while also collaborating with industry and academia.

technologies ... scientific ... advances are, to the greatest extent, shared between the United States Government and industrial and academic sectors.

(2) The Uni... collaboration, relying on the respective strengths of government, industry, and academia to maximize the ...

The White House

Office of the Press Secretary

For Immediate Release January 30, 2015

FACT SHEET: President Obama's Precision Medicine Initiative

Building on President Obama's announcement in his State of the Union Address, today the Administration is unveiling details about the Precision Medicine Initiative, a bold new research effort to revolutionize how ... improve health and treat disease. Launched with a $215 million investment in the President's 2016 Budget, the Precision Medicine Initiative will pioneer a new model of patient-powered research tha... promises to accelerate biomedical discoveries and provide clinicia... new tools, knowledge, and therapies to select which treatments wi... best for which patients.

Most medical treatments have been designed for the "average pat... a result of this "one-size-fits-all-approach," treatments can be very successful for some patients but not for others. This is changing w... emergence of precision medicine, an innovative approach to disea... prevention and treatment that takes into account individual differe... people's genes, environments, and lifestyles. Precision medicine ... clinicians tools to better understand the complex mechanisms und... patient's health, disease, or condition, and to better predict which treatments will be most effective.

Ecosystem:
1 W/ 1TF wearables
100TF wall-socket systems
1PF rack
…

The White House

Office of the Press Secretary

For Immediate Release January 28, 2016

Memorandum -- White House Cancer Moonshot Task Force

January 28, 2016

MEMORANDUM FOR THE HEADS OF EXECUTIVE DEPARTMENTS AND AGENCIES

SUBJECT: White House Cancer Moonshot Task Force

Cancer is a leading cause of death, and cancer incidence is expected to increase worldwide in the coming decades. But today, cancer research is on the cusp of major breakthroughs. It is of critical national importance that we accelerate progress towards prevention, treatment, and a cure -- to double the rate of progress in the fight against cancer -- and put ourselves on a path to achieve in just 5 years research and treatment gains that otherwise might take a decade or more. To that end, I hereby direct the following:

Section 1. White House Cancer Moonshot Task Force. There is established, within the Office of the Vice President, a White House Cancer Moonshot Task Force (Task Force), which will focus on making the most of Federal investments, targeted incentives, private sector efforts from industry and philanthropy, patient engagement initiatives, and other mechanisms to support cancer research and enable progress in treatment and care. The Vice President shall serve as Chair of the Task Force.

(a) Membership of the Task Force. In addition to the Vice President, the Task Force shall consist of the heads of the executive branch departments, agencies, and offices listed below:

(i) the Department of Defense;

(ii) the Department of Commerce;

(iii) the Department of Health and Human Services;

(iv) the Department of Energy;

(v) the Department of Veterans Affairs;

Some key NSCI threads:

- DOE is a lead agency
- Broad federal deployment
- All of government; holistic
- Public/Private

SOTU: 1/12/16, 9pm

Some early successes under President Trump. Currently on DOE's web page…

The Department of Veterans Affairs (VA) and the Department of Energy (DOE) are partnering to drive technology innovation and transform health care delivery for Veterans. The partnership brings together VA's unparalleled and vast array of healthcare and genomic data with DOE's world class high performance computing (HPC), artificial intelligence and data analytics. By combining expertise, we can push the frontiers of data analytics, next-generation computing, precision health, genomic sciences, and health care delivery. This partnership supports:

- **Innovation tied to design and development of DOE's next generation supercomputing that will merge Big Data (BD), Artificial Intelligence (AI) and High-Performance Computing (HPC) as well as innovation in population science using complex health system and genomic data for knowledge generation.**
- **Better Healthcare via using supercomputing to inform when and how to treat our Veterans to improve outcomes and reduce cost.**
- **Better Science via a cadre of researchers and clinicians who specialize in healthcare with the DOE experts in HPC, AI &BD.**
- **Better Government via interagency collaborations bringing to bear the full capabilities and expertise within and public private partnerships.**

The starting point for the DOE-VA partnership is MVP CHAMPION (Million Veterans Program Computational Health Analytics for Medical Precision to Improve Outcomes Now). Under MVP CHAMPION, VA and DOE will establish a scientific computing environment that will not only house

More Info

• Find
Progra

• Cont
Inform
free at

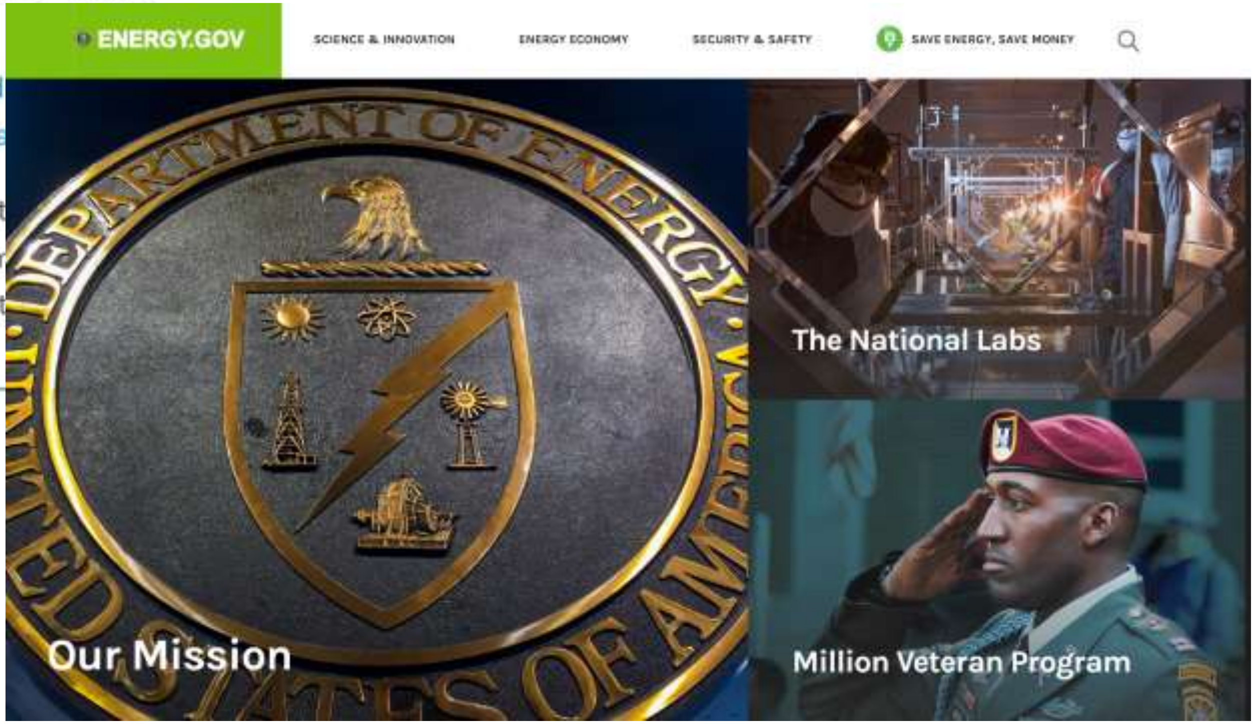

Piecing together the most unique data and next generation computational tools under one roof

Partnerships to date:

Key labs and vital people:
Argonne: Stevens, Madduri…
Livermore: Streitz, Brase, Paragas…
Los Alamos: McMahon, McCabe…
Oak Ridge: Tourassi, Begoli,…
…

- Federal Agencies
  - National Cancer Institute - early 2015
  - Department of Veterans Affairs – early 2016 & early 2017
- Pharma – late 2015
  - GlaxoSmithKline (GSK)
- Tech Sector - ongoing
  - IBM, Intel, GE, Nvidia,…
- Foreign Governments – late 2015
  - Norway - Oslo Cancer Cluster

Our DOE network

It also takes data – very hard to get data. A data snapshot today:

- ~ 24 M Veterans health records
- 620,000 genotype; ~10% with deep sequence
- 2M failed drug compounds and associated data, records and tests
- Norway: 1,728,336 unique screening results, 10,753,752 individual records
- SEER (Surveillance, Epidemiology, and End Results) cancer registry data from US states. So far Louisiana, Georgia,…
- TRACK-TBI data and associated brain scans and information
- Many small scale experiments on PDX, RAS on membranes,…

Creating the cycle of improving what we measure, how we measure and why we measure is important in the development of our methods.

Joint Design of Advanced Computing Solutions for Cancer

JDACS4C

Exascale technologies driving advances

NCI
National Cancer Institute

DOE
Department of Energy

Cancer driving computing advances

Initiatives Supported NSCI and PMI

NIH NATIONAL CANCER INSTITUTE

Argonne National Laboratory

Oak Ridge National Laboratory

Lawrence Livermore National Laboratory

Los Alamos National Laboratory

Frederick National Lab for Cancer Research

Integrated Precision Oncology

Molecular | Pre-clinical | Population

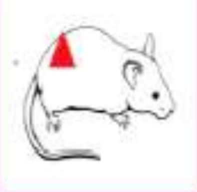

Pre-clinical Domain – Improved predictive models

Computational/hybrid predictive models of drug response

Improved experimental design

Clinical Domain – Precision oncology surveillance

Expanded SEER database information capture

Modeling patient health trajectories

Molecular Domain – Multiscale biological models

Models for RAS-RAS complex interactions

Insight into RAS related cancers

CANcer Distributed Learning Environment (CANDLE)

Uncertainty Quantification (UQ)

Patient Derived Xenograft Models

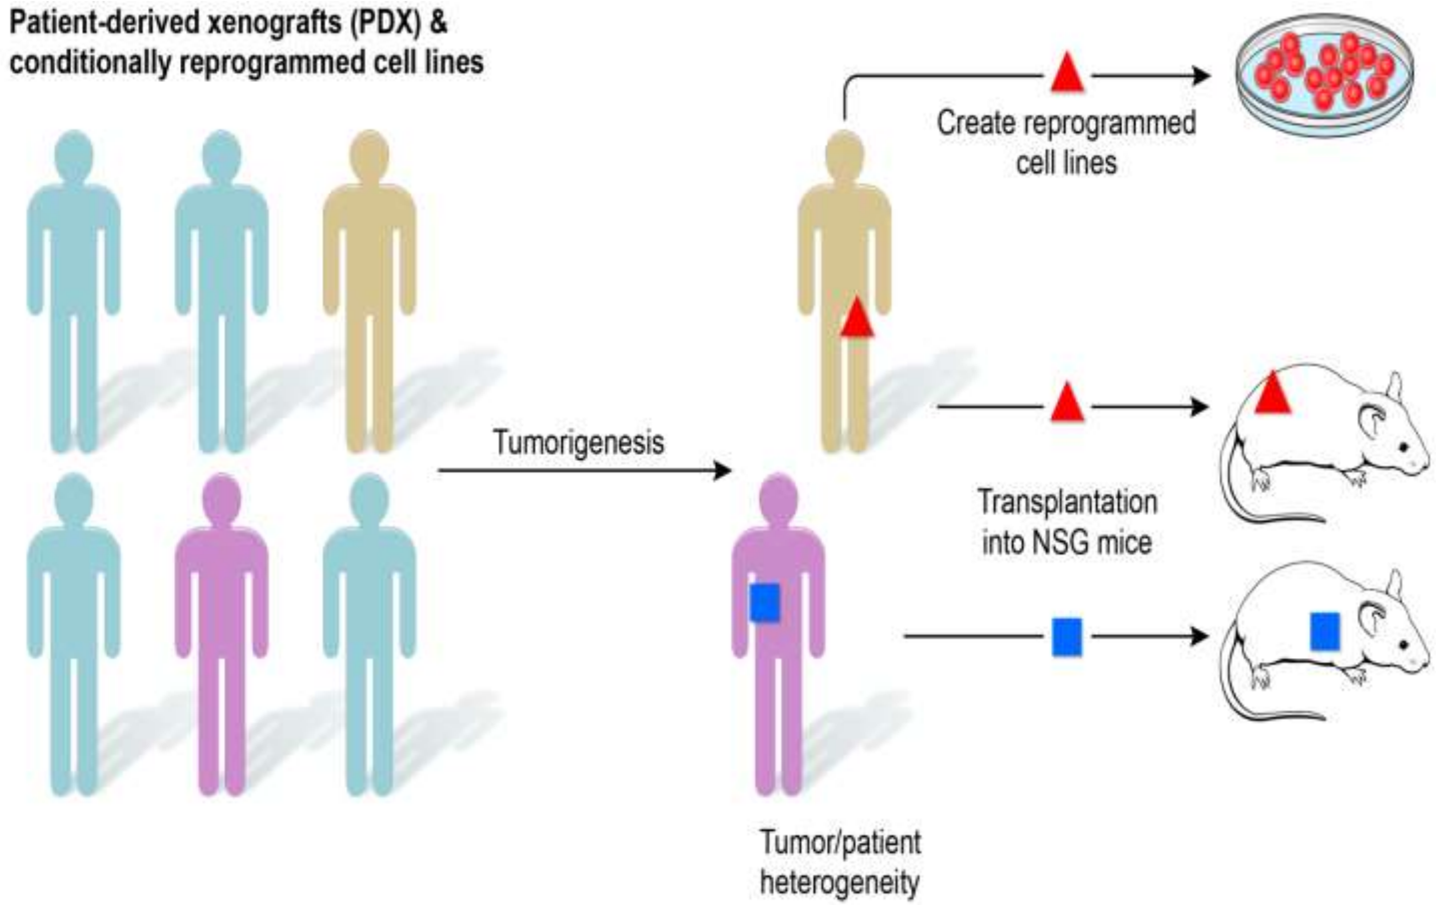

Nature Rev. Clin. Oncol. 11: 649-662, 2014.

Cancer Cell Lines

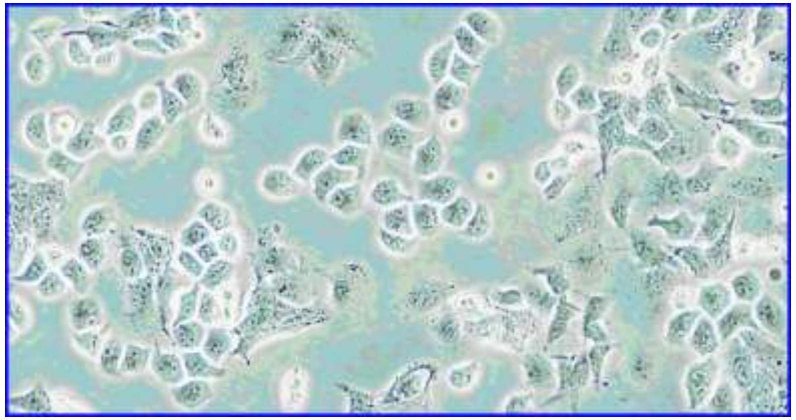

CL and PD Xenografts

Co-leads: Yvonne Evrard (NCI), Rick Stevens (ANL)

ANL: **Rick Stevens**, Frank Alexander (BNL), Cori Anderson, Jillian Aurisano, Prasanna Balaprakash, Tom Brettin, Jim Davis, Emily Dietrich, Nicoli Dryden, Veronika Dubinkina, Monisha Ghosh, Ushma Kriplani, Ravi Madduri, Sergei Maslov, Bob Olson, Dan Olson, Mike Papka, Lorenzo Pesce, Alex Partin, John Santerre, Maulik Shukla, Venkat Vishwanath, Fangfang Xia, Harry Yoo, Mikhail Fonstein, Austin Clyde, Jigar Shah; **LANL**: **Marian Anghel**, Tanmoy Bhattacharya, Judith Cohn, Paul Dotson, Will Fischer, Kumkum Ganguly, Jason Gans, Cristina Garcia-Cardona, Nick Hengartner, William Hlavacek, John Hogden, Patrick Kelly, Miranda Lynch, Ben Mcmahon ; **LLNL: Jonathan Allen**, Ya Ju Fan, Stewart He; **ORNL**: **Fernanda Foertter**, Arvind Ramanathan; **NCI: James Doroshow, Yvonne Evrard, Susan Holbeck**, **Eric Stalhberg**, George Zaki, Melinda Hollingshead, Rose Aurigemma, Chris Karlovich,

DOE-NCI Pilot 1: Predictive Models for Preclinical Screening

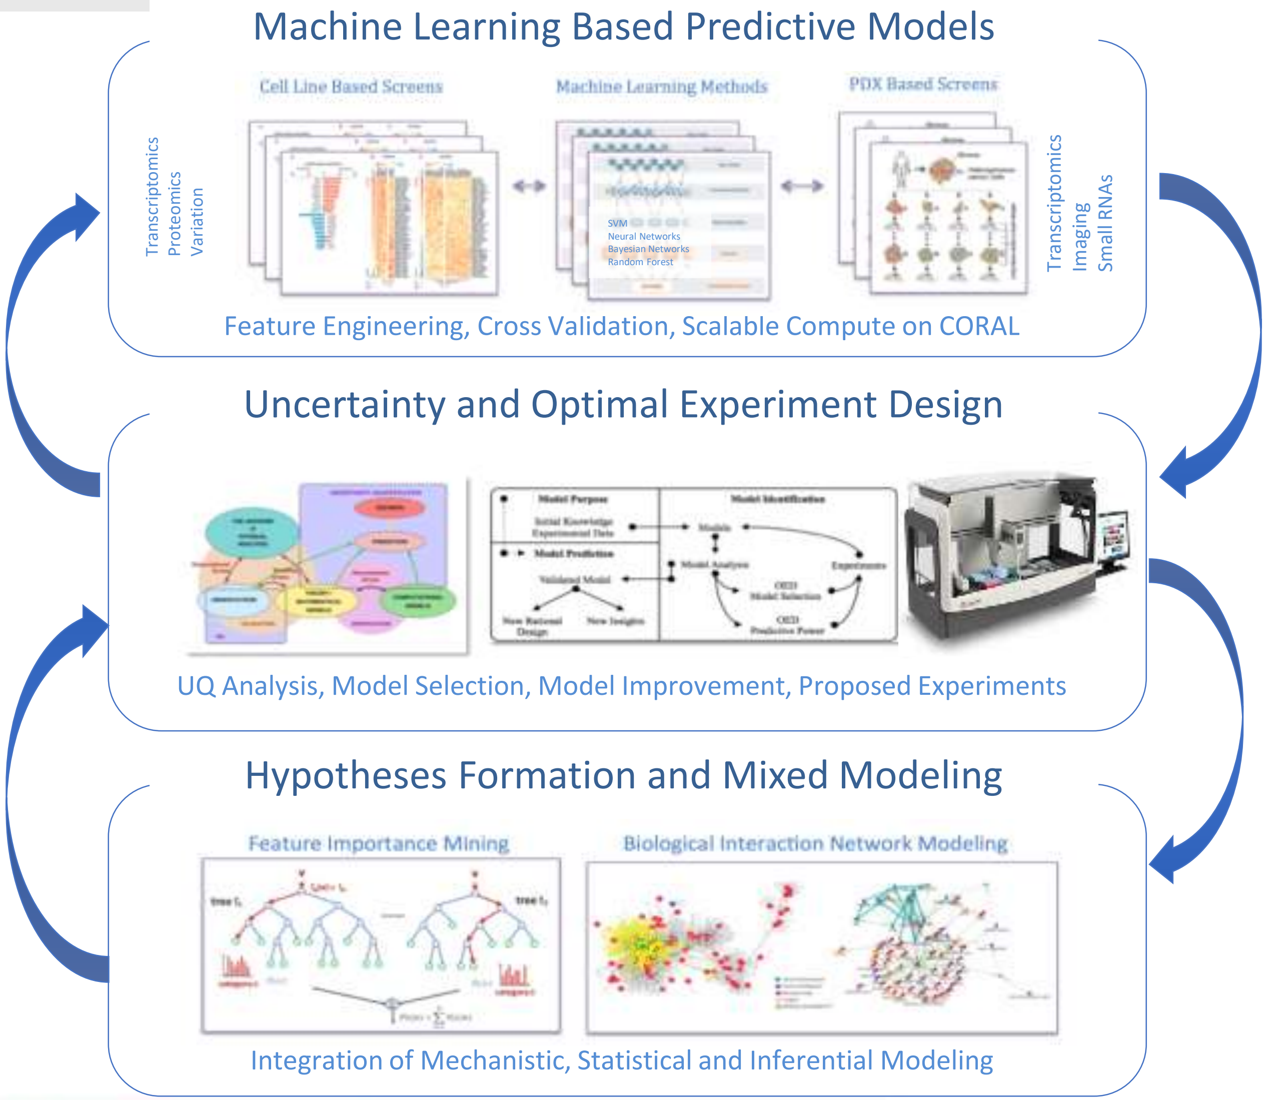

Specific Aims of Pilot1

Reliable machine learning based predictive models of drug response that enable the projection of screening results from and between cell-lines and PDX models

Uncertainty quantification and optimal experimental design to assert quantitative limits on predictions and to recommend experiments that will improve predictions

Improved modeling paradigms that support the *graded introduction of mechanistic models* into the machine learning framework and to rigorously assess the potential modeling improvements obtained thereof and to generate hypotheses from learned models

DOE-NCI Pilot 3 – The Surveillance Perspective

Improve the effectiveness of cancer treatment in the “real world” through computing

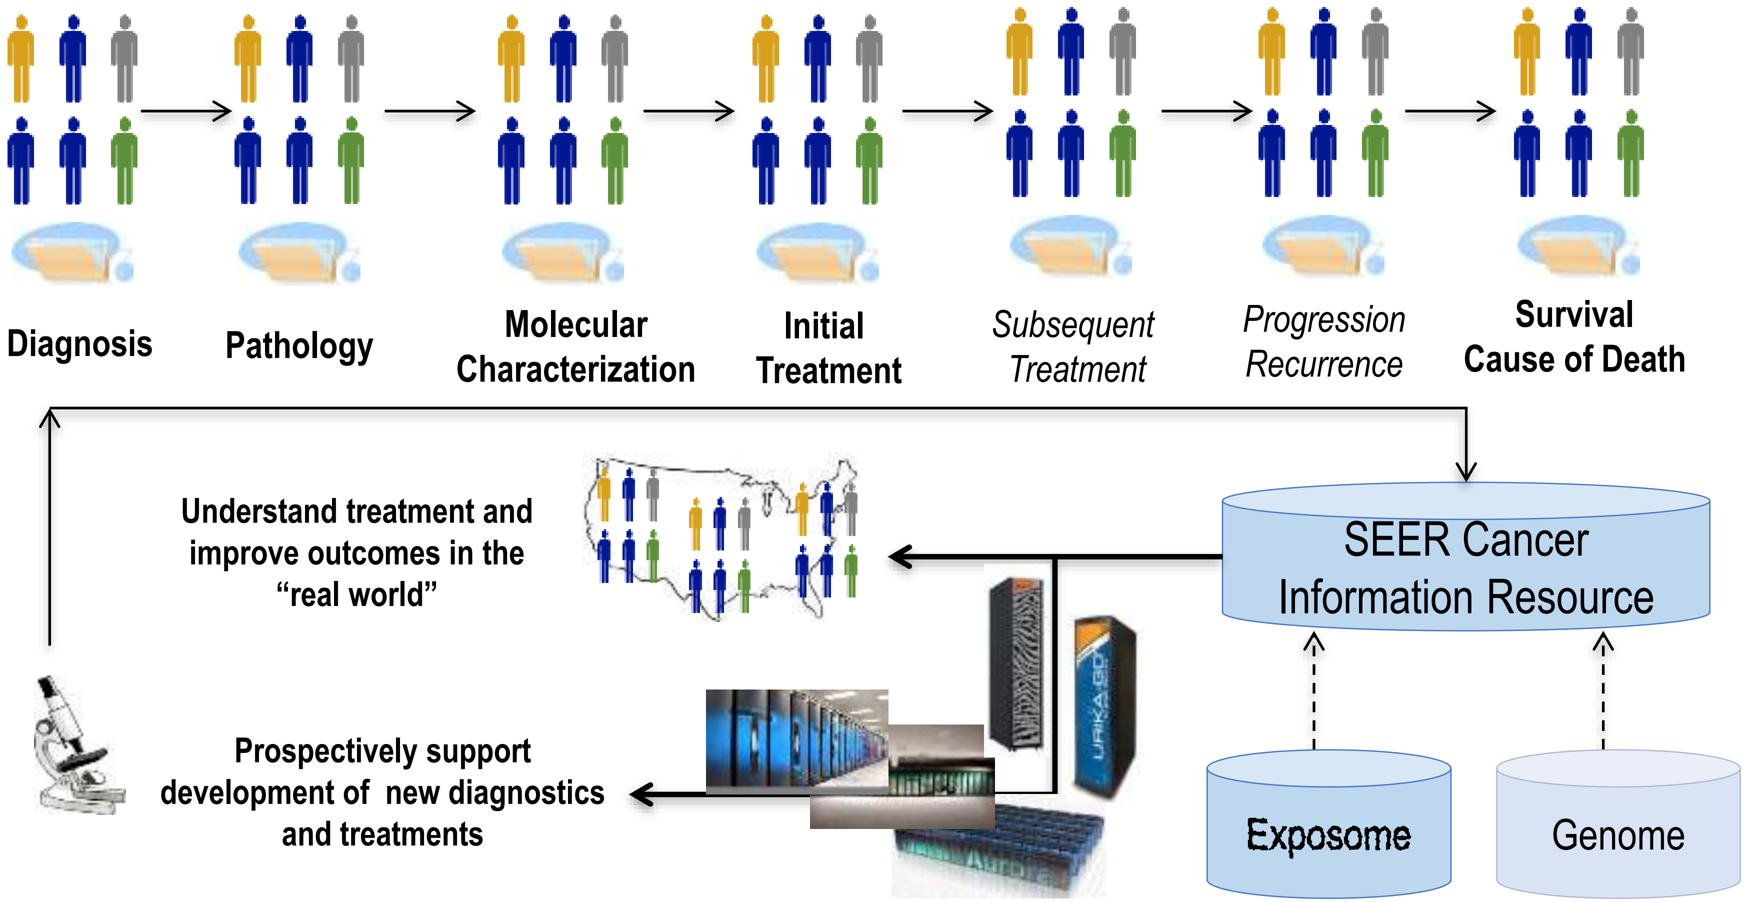

Co-leads: Lynne Penberthy ,Paul Fearn (NCI) Georgia Tourassi & Blair Christian (ORNL)

Pilot 3 - Aims

Deep NLP for information capture

Advanced machine learning for scalable patient Information capture from unstructured clinical reports to semi-automate the SEER program

Novel data analytic techniques for patient information integration

Scalable graph and visual analytics to understand the association between patient trajectories and patient outcomes

Data-driven integrated modeling and simulation for precision oncology

Precision modeling of patient trajectories

In silico clinical trials

RAS oncogenic protein is responsible for 1/3 of all human cancers

95% of all pancreatic;
45% of all colorectal;
35% of all lung cancers

>1 million deaths/year

No effective inhibitors

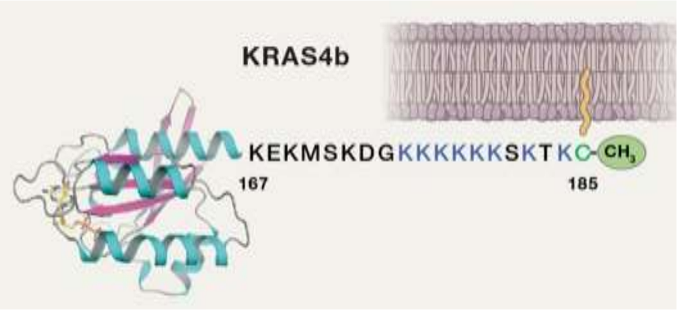

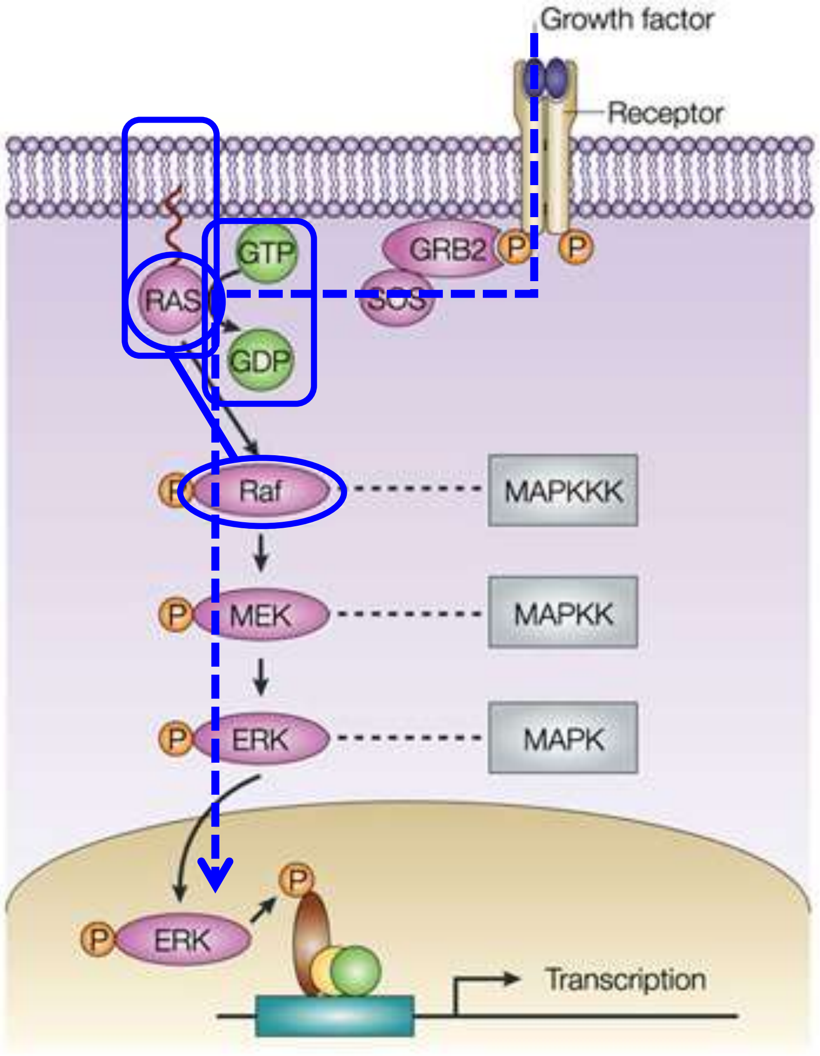

Nature Reviews | Molecular Cell Biology

transmits signals

RAS is a switch

RAS is localized on the membrane

RAS binds effectors to activate growth

Co-leads: Dwight Nissley (NCI), Fred Streitz, Jim Glosli (LLNL) Arvind Ramathan (ORNL)

DOE-NCI Pilot 2: RAS biology on membranes

RAS activation experiments

Experiments on nanodisc

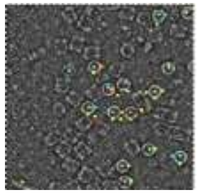

CryoEM imaging

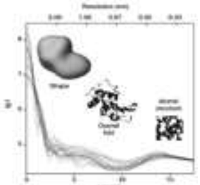

X-ray/neutron scattering

Multi-modal experimental data, image reconstruction, analytics

Protein structure databases

Predictive simulation and analysis of RAS

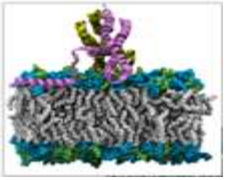

Phase Field model of lipid membrane

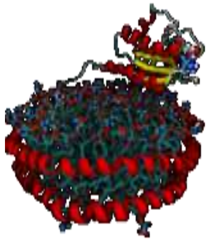

Granular RAS membrane interaction simulations

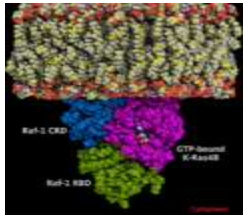

Atomic resolution RAS-RAF interaction

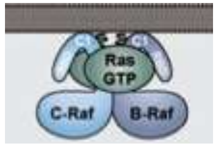

RAS Activation

Adaptive sampling molecular dynamics simulation codes

Adaptive time stepping

Adaptive spatial resolution

Phase field model

Coarse-grain MD

Classical MD

High-fidelity subgrid modeling

Machine learning guided dynamic validation

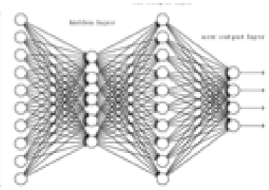

Unsupervised deep feature learning

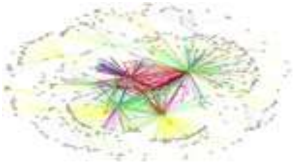

Mechanistic network models

Uncertainty quantification

Macro to micro scale coupling has been achieved

Marco Scale --Continuum Phase Field Model

Micro Scale-–Coarse Grain Bead Martini Model

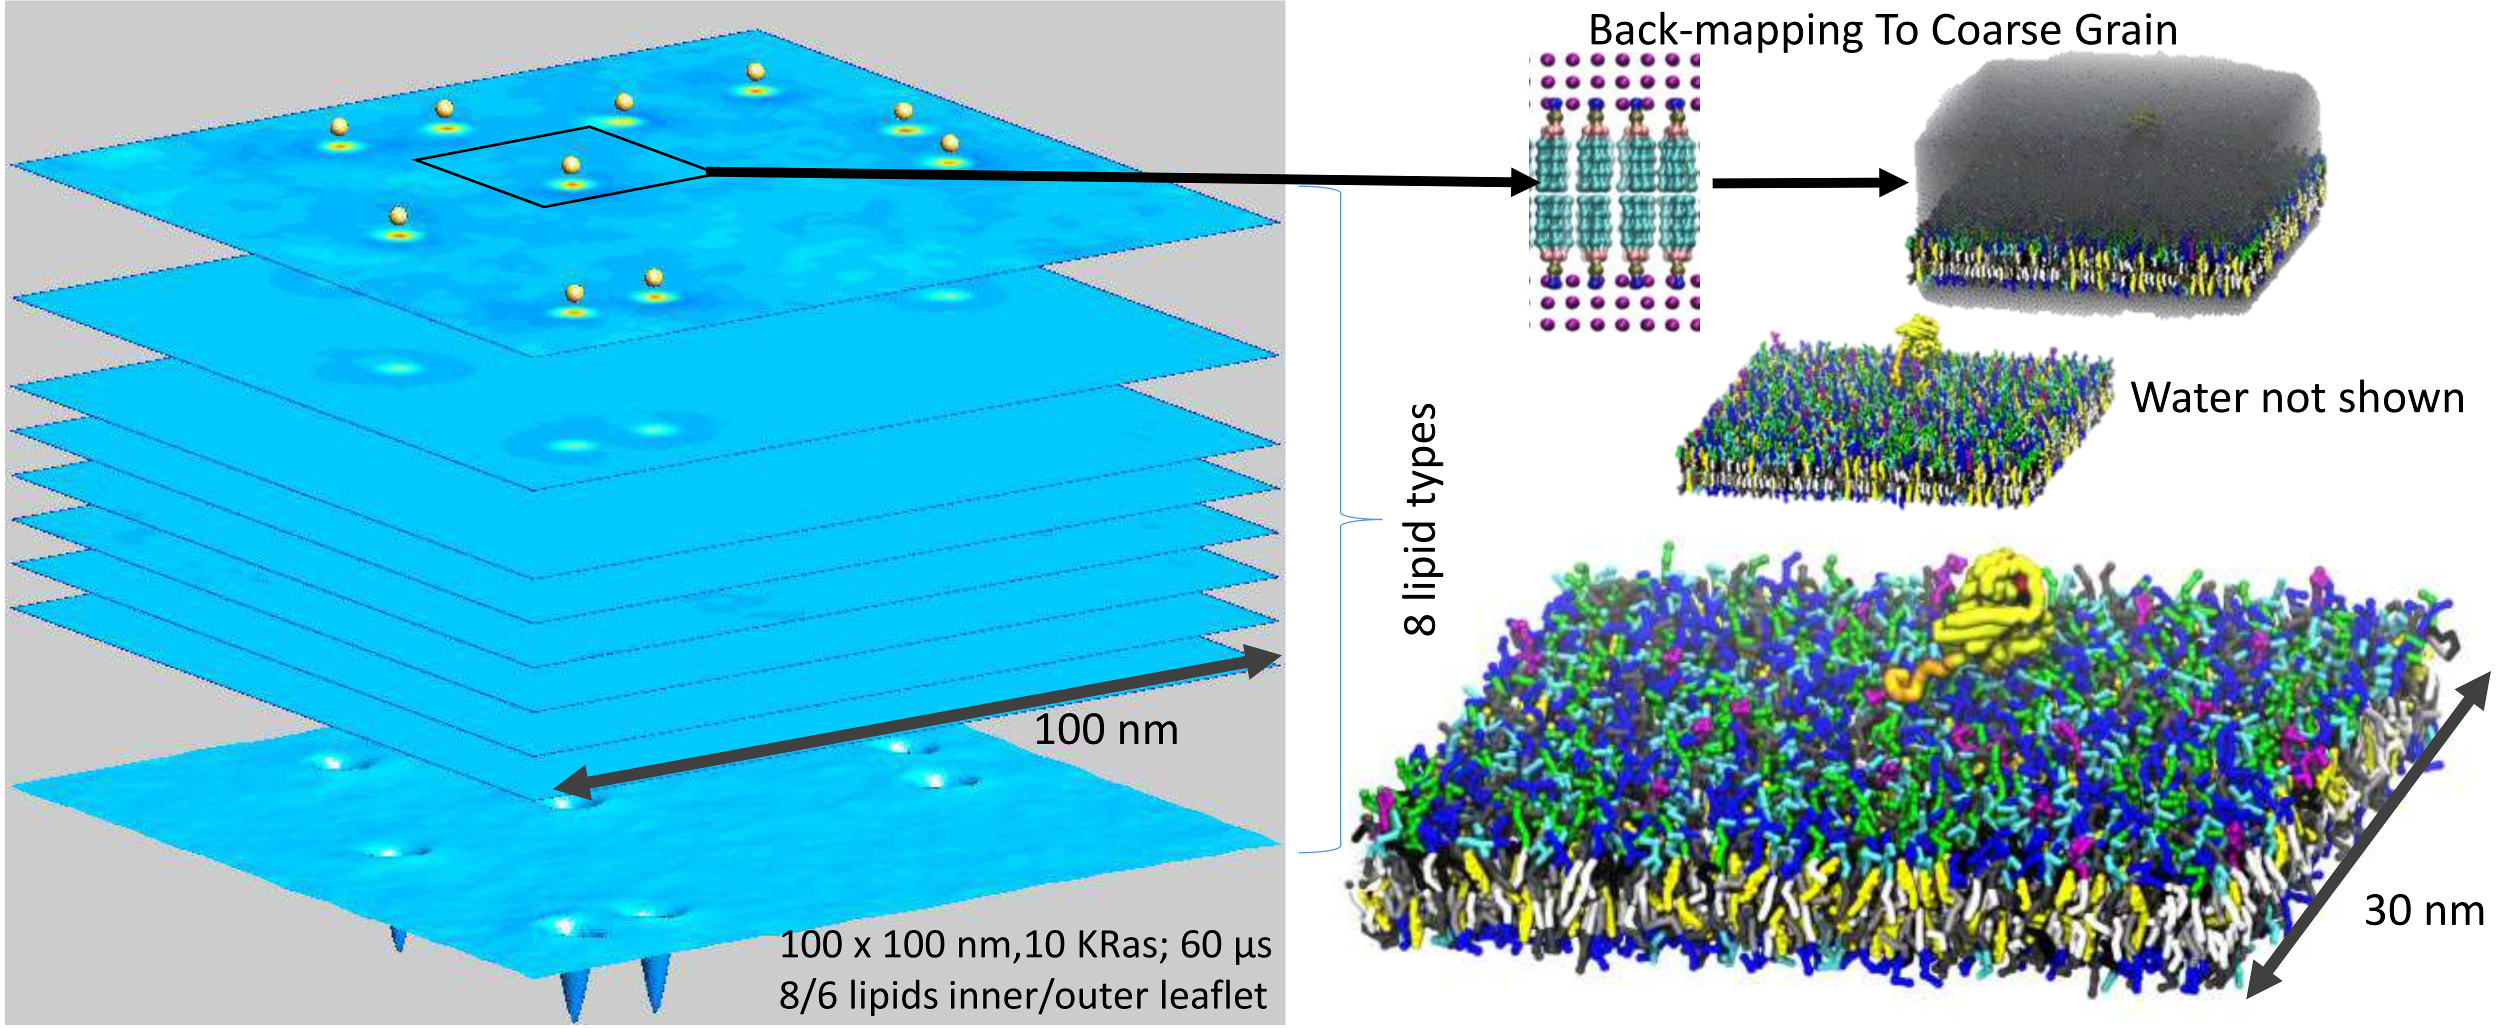

1 µs into production simulation

Entire Pilot 2 Team

Veterans effort focuses on outcomes:

- **Pilot programs** to address three top priority challenges that could deliver early impacts for veteran's health have been established to:
  - develop patient-specific analysis for **suicide prevention**
  - help doctors discern which **prostate cancers** require treatment and which patients can forgo invasive and deleterious treatments
  - deliver new predictive tools to inform individualized drug therapies for **cardiovascular disease**
- + polypharmacy,…

CDW Phenotype Data Examples

- Major biobanks in the world
  - UK Biobank (500K)
  - Kadoori Biobank (512K)
  - Kaiser Permanante (>200K)
  - **MVP (1 million)- current >550,000**
  - NIH *All of Us* Cohort- aiming for 1 million
- **Stable and altruistic patient population**
- **Rich electronic health record**
- **Research program embedded within a healthcare system**
- **Privacy and security of Veterans' data is the highest priority**
  - MVP data stored in a secure central repository
  - Researchers come to the data instead of data moving to the researchers
  - Knowledge repurposed

ATOM Consortium

Accelerating Therapeutics for Opportunities in Medicine

Preclinical issues:

- Average time: 5.5 years
- 33% of total cost of medicine development
- Millions of molecules tested, 1000s made, and most fail
- (Recall from India HPC panel last night: 10^{60} small molecules)
- Clinical success rates still only 12%, indicating poor translation in patients

3 year goal

From target to patient-ready in <1 year

Mission

To accelerate the development of more effective therapies for patients

Vision

A new starting point:
Transform drug discovery from a slow, sequential, and high-failure process into a rapid, integrated, and patient-centric model

How?

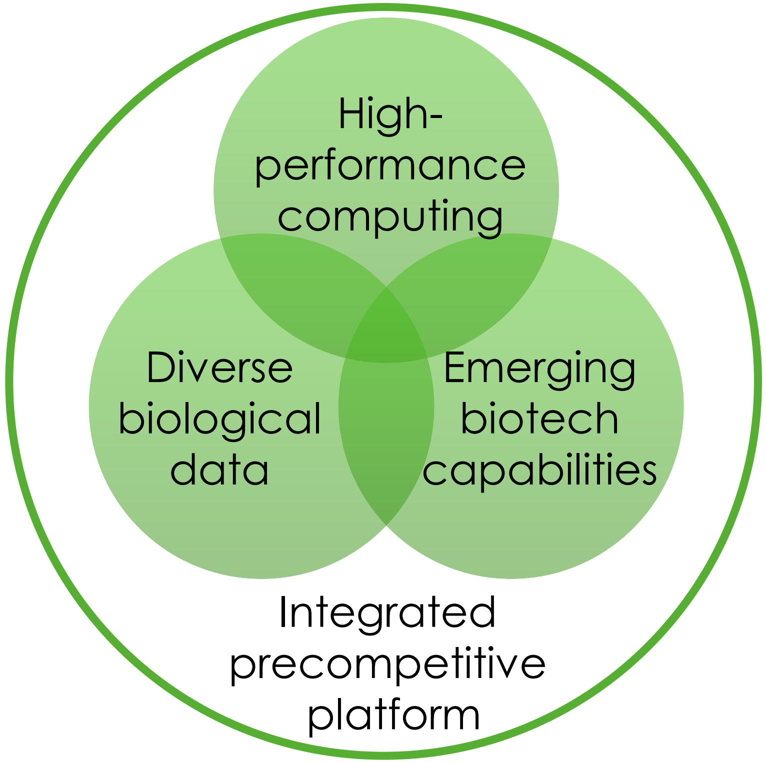

The ATOM partnership for rapid drug development is now in place

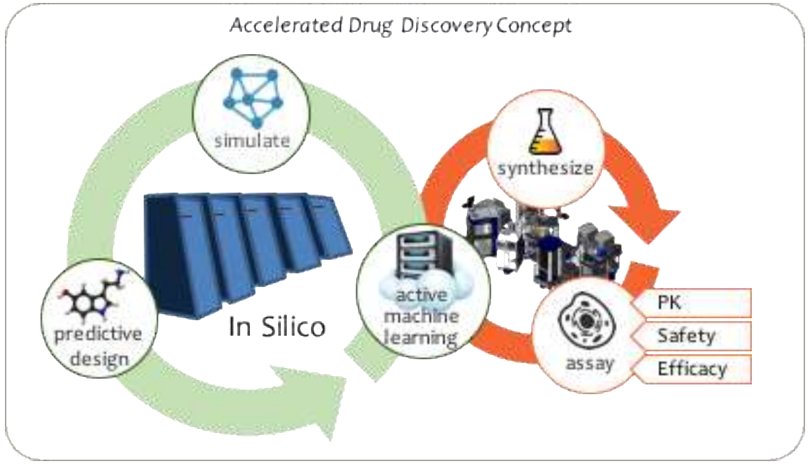

- Partnership between GlaxoSmithKline, National Cancer Institute, UCSF, and LLNL
- CRADA executed on Oct 10, 2017
- Research team working at UCSF's Mission Bay campus
  - 12 GSK, 8 LLNL
  - 3 NCI/FNLCR, UCSF TBD

2M compounds and associated data from terminated projects and screening collections

http://atom.cancer.gov

Partnership with Cancer Registry of Norway provides an opportunity for unique data access

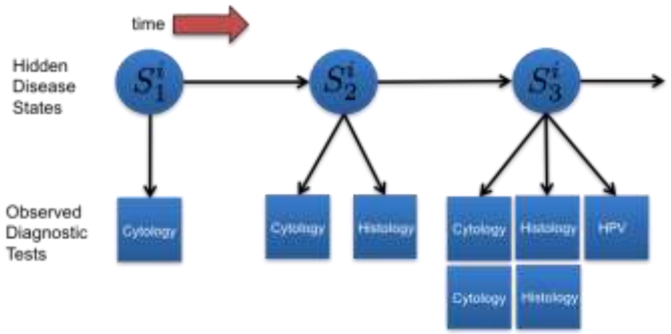

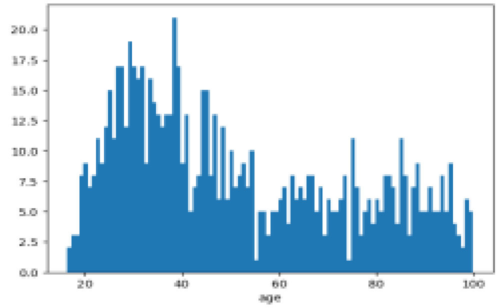

First data set:

National database of 1,728,336 unique Norwegian women's cervical cancer screening results

Multiple screening results are recorded over time for each patient

Records cover a 25-year period: 1991–2015

Screening times are unique to each patient

10,753,752 individual records

- Comprehensive data collection in a single-payer health-care system
- Well-understood population
- Enables direct application of systematic statistical modeling
- Pilot project: Predictive models of cervical cancer screening sequences will enable optimization of testing frequency.
    - Patient questionnaires as covariates
    - HPV status as covariates.
    - Going from 100K to >1M cases
- Extension to other Nordic countries being discussed

Summary

Precision Medicine

- We are piecing together perhaps to world's richest (precision) medicine database.
- Every step is new; every step is hard, but much progress since mid-2015. For DOE scientists – all data is available.
- We are building new architectures, technologies and tools based on the complexity of the data and the questions we are asking.
- We have an opportunity to push the frontier of precision medicine and next generation high-performance computing.
- For technology and simulation: new models for UQ, new paradigm for predictive simulation, novel architectures, and a step into the future.
- Ideas and broadening the conversations and partners is an important part of making progress.

Thank you for your time today