The Exascale Computing Project:
Update on Applications Relevant to the BER Mission

Douglas B. Kothe
Oak Ridge National Laboratory
Director, Exascale Computing Project

DOE Office of Science Biological and Environmental Research Program Advisory Committee (BERAC)
Gaithersburg, MD
Nov 3, 2017

Exascale applications are addressing national problems across six key strategic pillars

National security

Stockpile stewardship

Next-generation electromagnetics simulation of hostile environment and virtual flight testing for hypersonic re-entry vehicles

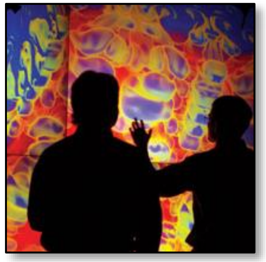

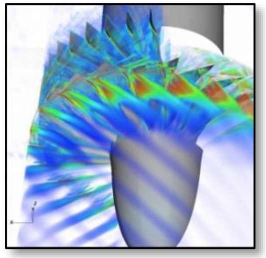

Energy security

Turbine wind plant efficiency

Design and commercialization of SMRs

Nuclear fission and fusion reactor materials design

Subsurface use for carbon capture, petroleum extraction, waste disposal

High-efficiency, low-emission combustion engine and gas turbine design

Clean fossil fuel combustion

Biofuel catalyst design

Economic security

Additive manufacturing of qualifiable metal parts

Urban planning

Reliable and efficient planning of the power grid

Seismic hazard risk assessment

Scientific discovery

Cosmological probe of the standard model of particle physics

Validate fundamental laws of nature

Plasma wakefield accelerator design

Light source-enabled analysis of protein and molecular structure and design

Find, predict, and control materials and properties

Magnetic fusion reactor stability and control

Demystify origin of chemical elements in the universe

Earth system

Accurate regional impact assessments in Earth system models

Stress-resistant crop analysis and catalytic conversion of biomass-derived alcohols

Metagenomics analysis of biogeochemical cycles

Health care

Accelerate and translate cancer research

ECP EXASCALE COMPUTING PROJECT

Exascale Bridges the Gap!

Simulation Gap	Needed to Address the Gap	Impact of Gap Remaining
Simplified or incomplete physics	• Higher fidelity models for all relevant physical phenomena • Compute memory and speed necessary to accommodate their numerical solutions and their non-linear coupling requirements	• Cleaner, more efficient combustion engine designs delayed, or not discovered • Inefficient agriculture & energy production sector planning - insufficient regional water cycle assessments • Fewer (likely more expensive) options for CO_2 sequestration, petroleum extraction, geothermal energy, due to lack of understanding long-term reservoir-scale behavior • Astrophysics discoveries (origin of elements in the universe, gravity waves) remain elusive
Simulation detail insufficient	• Larger simulation domain • Finer partitioning of the simulation domain • Computing speed, memory and I/O to accommodate larger simulation domains with finer partitioning • Workflows integrating an advanced technology and software toolset • Data streaming methods to steer simulation partition real time	• Conservative earthquake retrofits more costly and over-designed • Inability to predict and control material properties at the quantum level precludes advances in high temperature superconductivity • Higher power grid operating margins and lost cost savings potential • Wind plant efficiencies lag theoretical energy extraction potential by 20-30% • Delays in scale-up of new chemical looping reactors for clean fossil fuel combustion • Key cosmology and nuclear physics discoveries – dark matter/energy, standard model of particle physics, inflation of the universe – remain elusive
Can only simulate subset of scenarios of interest	• Robust and fast algorithms for the numerical solution of coupled multi-physics systems that expand limits of applicability • Workflow tools to analyze simulation ensembles	• Continued high rejection rates (80%) of additively manufactured metal alloy parts with tight specifications, increasing waste and cost • Tools for retrofitting & improving urban districts remain empirical • Limited ability to influence ITER design decisions and ultimate operations
Uncertainty insufficiently quantified	• High fidelity in situ data analytic techniques for reliable quantification of simulation uncertainties and sensitivities • Workflow tools to analyze simulation ensembles	• Protracted deployment of small & advanced nuclear reactors • Delays in design of small, low cost, and ultra high intensity plasma wakefield accelerators • Engr & materials design requires more expensive & time-consuming physical experiments
Unable to intersect to design cycle	• Improved computer throughput to shorten simulation turn-around times • Large ensembles of calculations to enable optimization thru rigorous exploration of design space	• Continue deploying materials for extreme environments with a cumbersome make-test cycle rather than by atomistic design • Efficient in silico design of new chemical catalysts not realized
Inadequate analysis and knowledge discovery in big data	• Scalable AI (deep learning) networks of large size and complexity for efficient training on big datasets. • Efficient computational workflows seamlessly integrating simulation, data analytics, and big datasets	• New cancer biology treatment options missed or delayed due to unrealized understanding of precision oncology • Analysis and reduction of data deluge from experimental science facilities requires weeks (instead of near real time) • Manufacture of new products and chemicals delayed or missed because microbiome DNA sequencing unable to keep up with available data • Inability to quantify uncertainties via a nexus of simulation and experimental facility data

ECP Mission Need Statement: Application Requirements

Mission Need Approval (CD-0): Jul 2016

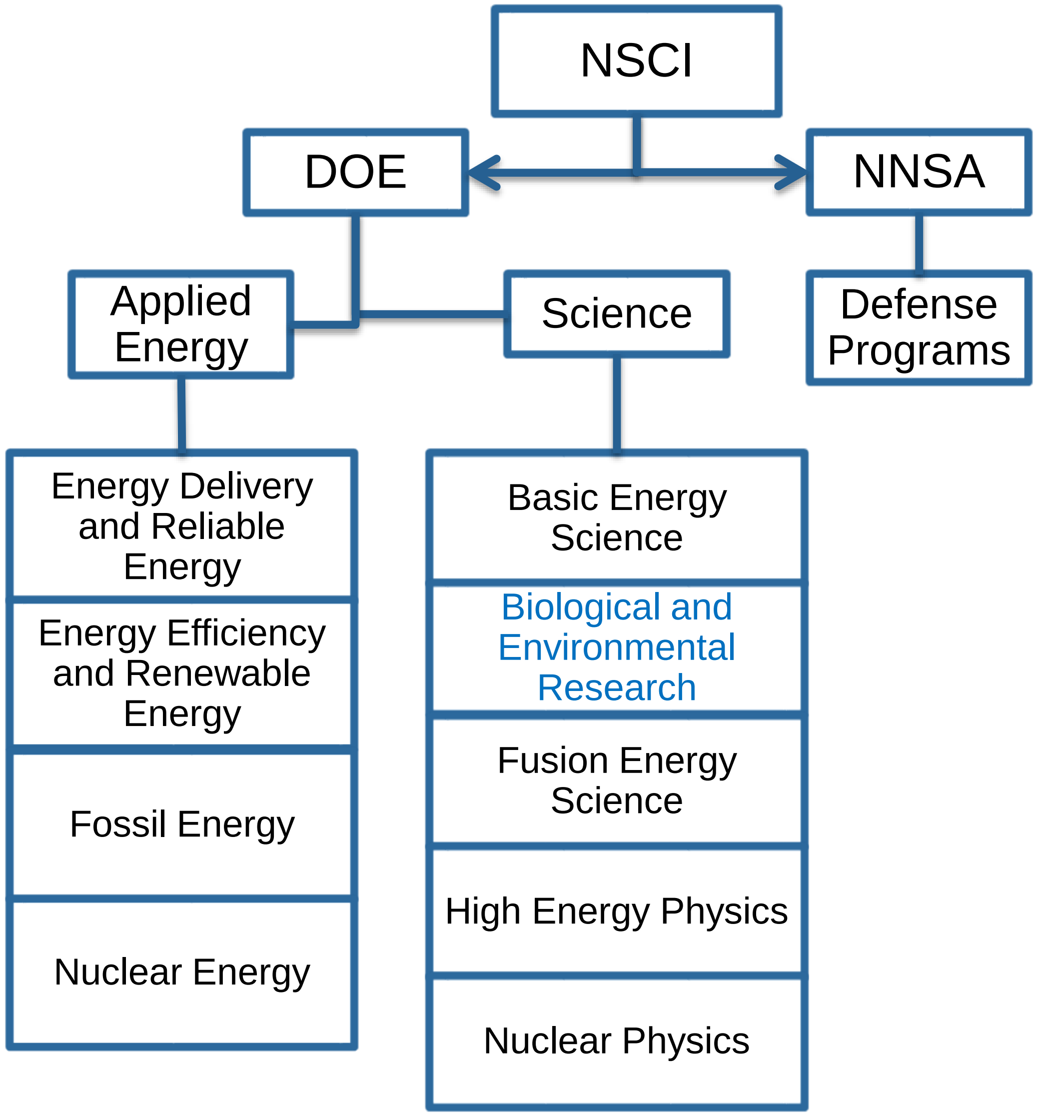

- ECP Mission Need must address a broad array of mission needs, including
  - ***DOE Science and Applied Energy Technology***
    - Materials Discovery and Design; Climate Science; Nuclear Energy; Combustion Science; Fusion Energy; Large Data Applications; Additive Manufacturing
  - ***DOE National Security***
    - Stockpile stewardship (Certification and Annual Assessment; Significant Finding Investigations; Stockpile Warhead Lifetime Extension)
  - ***NSCI*** mission support for NIH, NASA, NOAA, DHS, FBI

BER Requirements: We Hear You

- **BER Grand Challenges**
  1. Biological Systems
  2. Earth and Environmental Systems
  3. Microbial to Earth System Pathways
  4. Energy Sustainability
  5. Data Analytics and Computing

- **Cross-Cuts for Computational Climate Science**
  1. *Earth System Big Ideas*: Process-resolving capabilities; Integrated model credibility; Representing climate system complexity; Continuum model framework
  2. *Computation Cross Cuts*: Model/software complexity; Hardware; Performance; Data management

- **Exascale Requirements**
  1. Atmosphere
  2. Terrestrial / Subsurface
  3. Ocean / Cryosphere
  4. Earth System
  5. Integrated Assessments & Impacts-Adaptation-Vulnerability
  6. Model-Data Fusion
  7. Exascale Capabilities: Algorithms and Computation

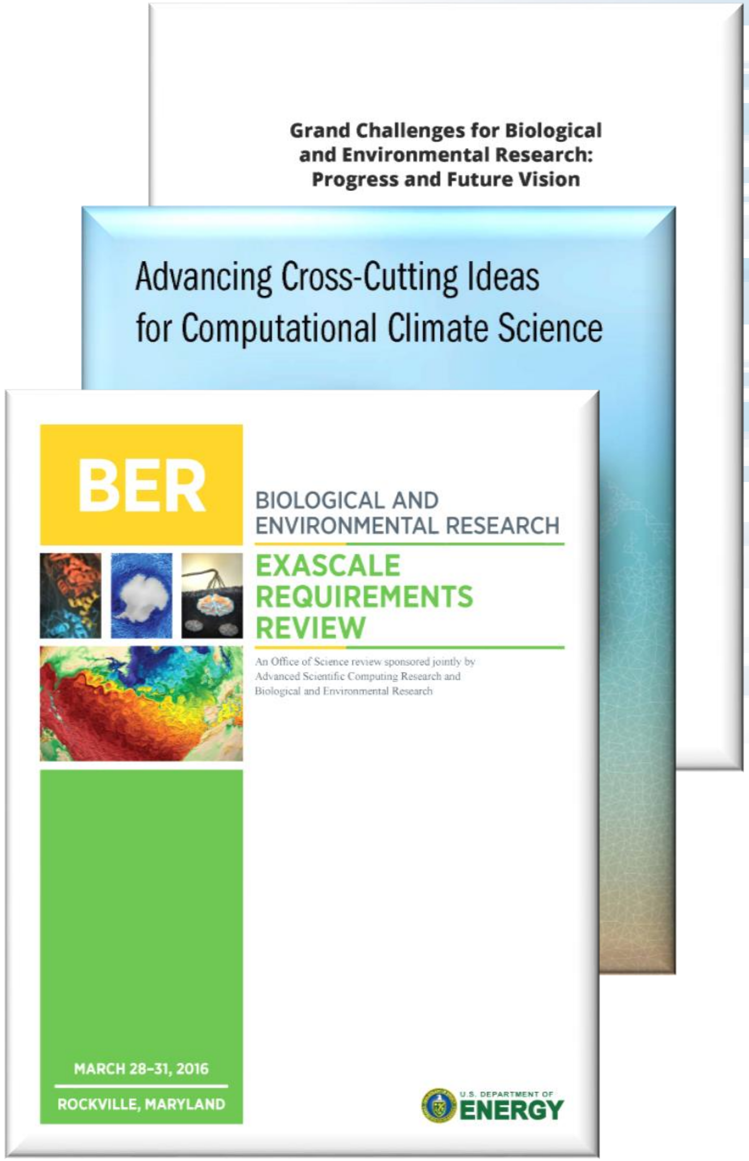

EXASCALE COMPUTING PROJECT

ECP Application Metrics

1. **Deliver improved and impactful science & engineering (*performance*)**
   - New or improved (ideally step change in) predictability on a problem of national importance (a “challenge problem”)
2. **As performance portable as possible and reasonable (*portability*)**
   - No “boutique” one-off applications able to only execute on one (and likely ephemeral) system
3. **Able to make effective use of a capable system (*readiness*)**
   - *Effective* is app specific (weak, strong, ensembles, single-node performance)
4. **Able to integrate latest relevant software technologies (*modern*)**
   - Needed to demonstrate agility, flexibility, modern architecture; overall app portfolio must apply pressure to all key attributes of the system design characteristics
5. **High priority (*strategic*)**
   - Some key stakeholder somewhere really cares about using application to make consequential decisions

Evolution of Subsurface Fractures and Wellbores

ECP Subsurface: WBS 1.2.1.05
PI: Steefel, Berkeley Lab
Members: LBNL, LLNL, NETL

In the subsurface, fractures and degraded wellbores are primary flow pathways.
They are subject to alteration by geomechanical stress and geochemical reactions.

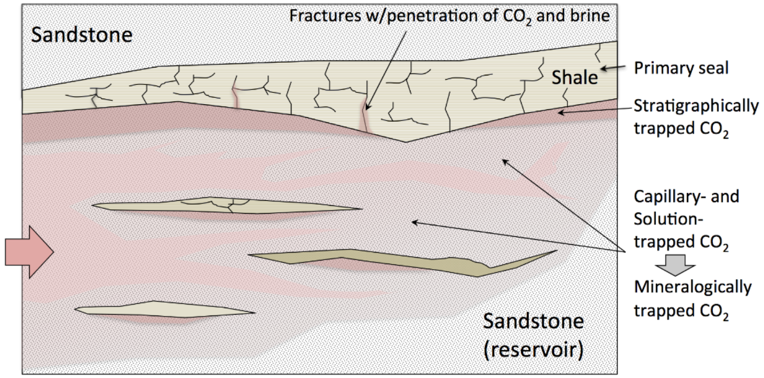

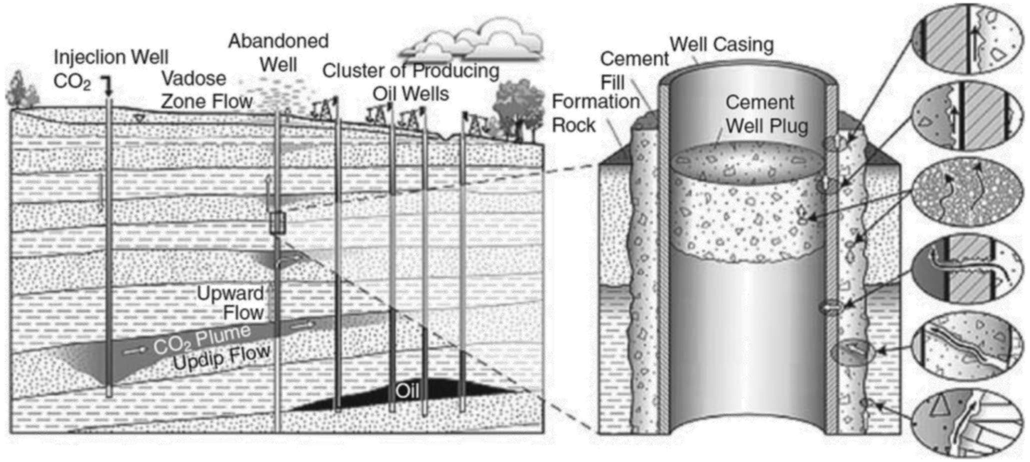

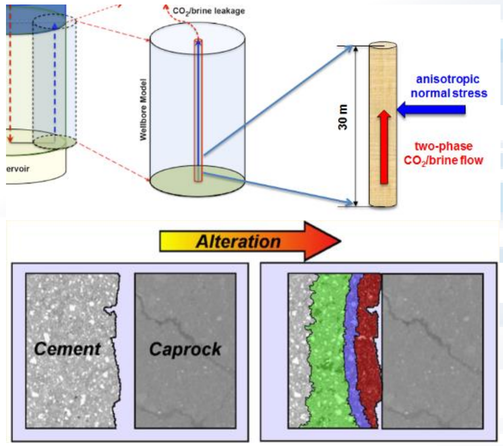

Better understanding of these processes are needed in order to predict potential leakage in the systems of Geologic Carbon Storage (shown above), and to evaluate the performance of other subsurface energy practices.

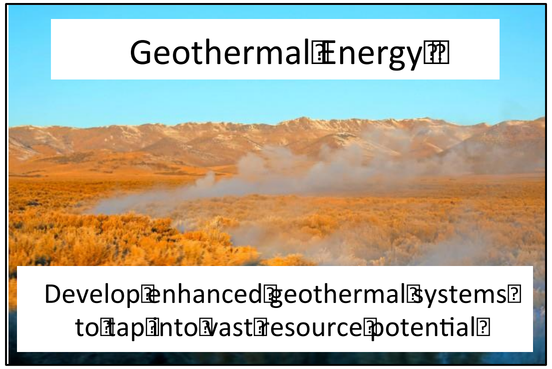

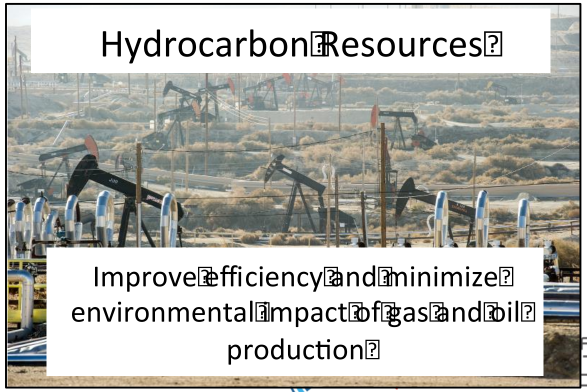

Coupling Approach

ECP Subsurface: WBS 1.2.1.05
PI: Steefel, Berkeley Lab
Members: LBNL, LLNL, NETL

Couple two separate software packages:

GEOS: Computes geomechanics and Darcy flow using finite elements

Chombo-Crunch: Computes flow and multicomponent reactive transport using integrated finite volumes on rectilinear grid, with adaptive mesh refinement (AMR) and embedded boundaries (EB)

Approach is to couple within a single time step, with GEOS passing deformed pore geometry based on Lagrangian finite element treatment to Chombo embedded boundaries, and vice versa

GEOS

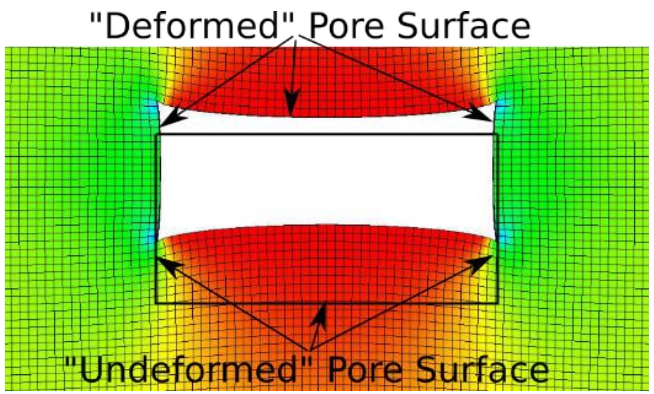

Pore translation/deformation resulting from change in stress loading in a Lagrangian mechanics treatment

Chombo-Crunch

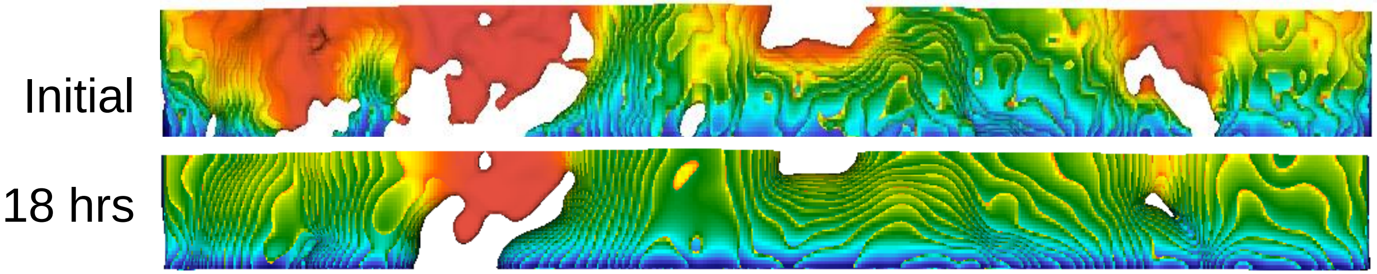

Flow driven chemical erosion of a fracture in $CaMg(CO_3)_2$

EXASCALE COMPUTING PROJECT

Optimizing Stochastic Grid Dynamics at Exascale

Application Domain

- **Application Area:** *Reliably and efficiently planning our nation's electric grid for societal drivers: rapidly increasing but variable renewable energy penetration; more active load management at the consumer level. Going beyond current steady-state constraints to incorporate complex dynamics. Potentially save billions annually by producing the best planning decisions.*
- **Challenge Problem**: *Transmission Planning with long time horizon (20 years) with integer variables, uncertainties, and transient constraints scales to 1M-10M-way parallelism*

Physical Models and Codes

- **Physical Models**: *Differential algebraic equation of generators interacting with other dynamic devices under stochastic and transient constraints from conservation laws and externally driven demand while undergoing non-linear optimization for cost and reliability.*
- **Codes**: *GridPACK, GridOPTICS™ Software System, StructJump, PIPS, PETSc, BLAS/PBLAS, Metis/ParMETIS, Global Arrayl*
- **Motifs:** *Monte Carlo, Sparse Linear Algebra, Dense Linear Algebra, Graph Transversal, Dynamical Program, Backtrack & Branch and Bound*

Partnerships

Co-Design Centers: *Co-Design Center for Online Data Analysis and Reduction (CODAR)*

Software Technology Centers: *Exascale MPI, Open MP Tool Interfaces, xGA, OMPI-X, ProTools, PerfTools, AID, Exa-PAPI, xSDK4ECP, SUNDIALS, PETSc, STRUMPACK, SLATE, HPCToolkit, VeloC, ExaHDF5, ADIOS*

Application Projects: *DOE Grid Modernization Initiative, ARPA-E GRID DATA (SDET & DR POWER) and HIPPO*

Multifaceted Mathematics for Complex Energy System (M2ACS)

Four Year Development Plans

- ***Year 1:*** *Linear programming problem for optimal power flow with uncertainties.*
- ***Year 2:*** *Expand to include integer variable for power grid unit commitment.*
- ***Year 3**: Expand to include dynamic constraints for high-fidelity power grid modeling.*
- ***Year 4**: Large-scale software testing, profiling, and extrapolation to exascale.*

EXASCALE COMPUTING PROJECT

PI: H. Huang (PNNL) with M. Anitescu (ANL) and W. Jones (NREL)

Power Grid Presents Complex Spatial-Temporal Multi-Scale Physics

- Multi-scale spatio-temporal modeling and simulation with stochasticity
  - From micro-second to decades
  - From 10^3 generators nodes to 10^9 end-use devices
- Large-scale data assimilation for state and parameter calibration
  - Petabyte data/year from high-speed sensors and smart meters.
- Modeling of multi-system dynamics and dependency
  - Grid, Communication, Gas-pipeline, Weather/Wind/Solar, Water

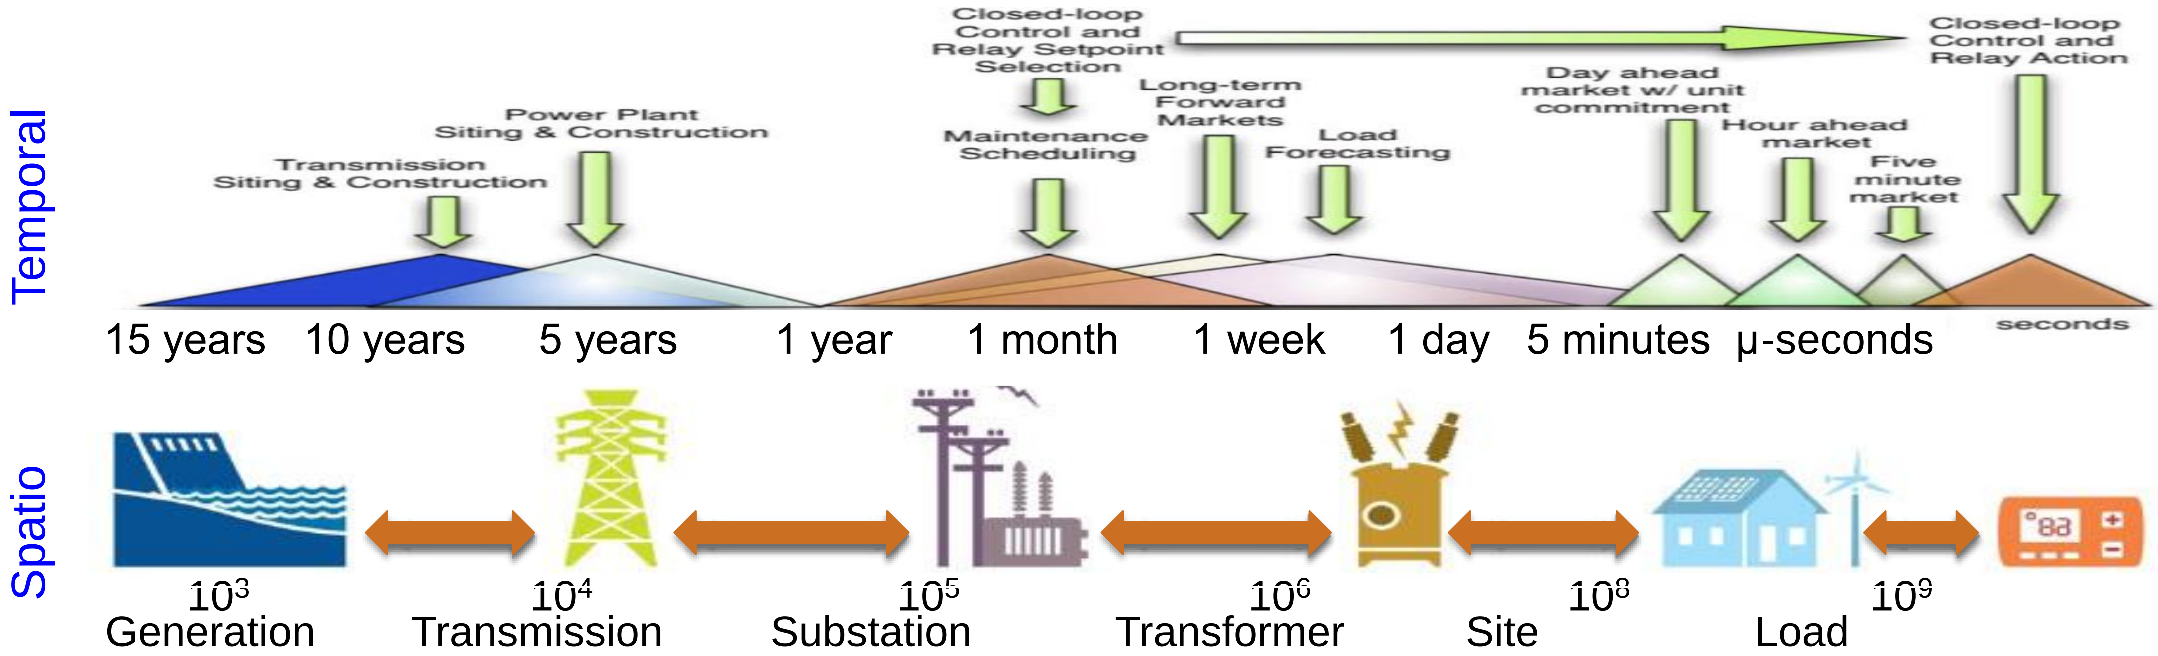

PI: H. Huang (PNNL) with M. Anitescu (ANL) and W. Jones (NREL)

Emerging dynamics and uncertainties poses exascale challenges in power system transmission planning

- Challenge Problem:
  - Base problem: Power grid transmission planning in long-time horizons (20 years) with integer variables, uncertainties, and transient constraints, scaling to 1M-10M-way parallelism:
    - A 5-year hourly-resolution interconnection planning with uncertainties and transients is a $O(10^{20})$ problem

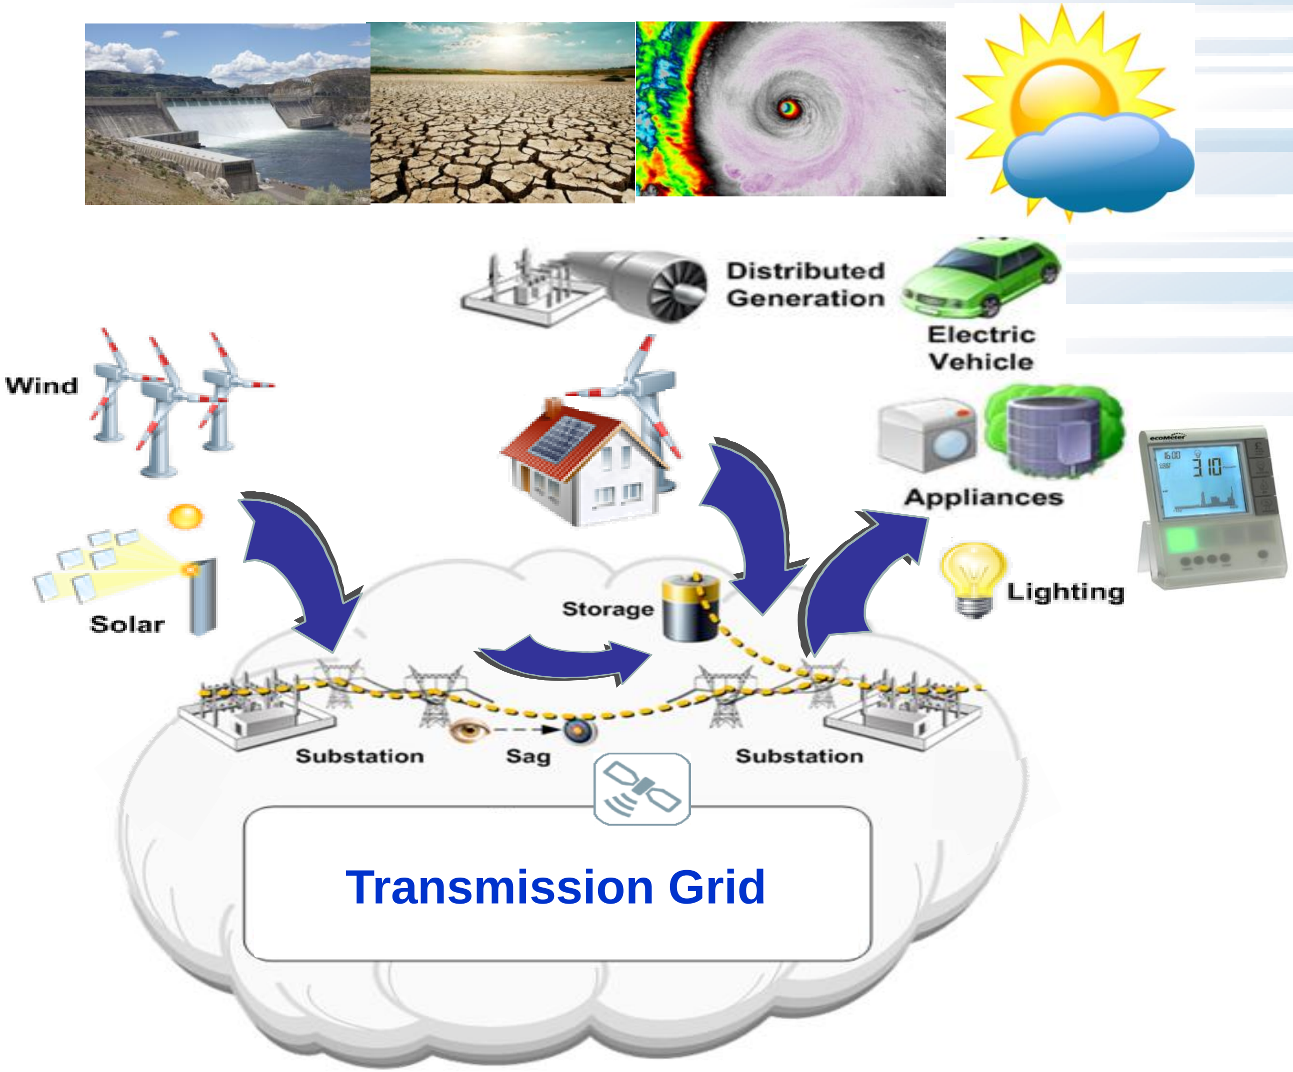

PI: H. Huang (PNNL) with M. Anitescu (ANL) and W. Jones (NREL)

ExaBiome: Exascale Solutions for Microbiome Analysis

- **Microbes:** single cell organisms, such as bacteria and viruses
- **Microbiomes:** communities of 1000s of microbial species, less than 1% individually culturable in a lab (and thus sequenced)
- **Metagenomics:** genome sequencing on these communities (growing exponentially)
- Understanding and controlling microbiomes has many applications

Environment

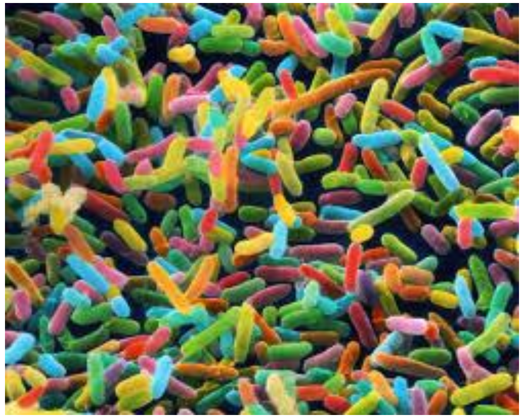

Health

Bio-Energy

Bio-Manufacturing

EXASCALE COMPUTING PROJECT

PI: Katherine Yelick (LBNL)

MetaHipMer: High Metagenome Assembly Quality & Scalability

ECP WBS: Exabiome
PI: Kathy Yelick, LBNL
Members: LBNL, JGI, LANL

Exabiome Scope & Objectives

- Metagenomic datasets are large and rapidly growing, and future assembly and analysis of these will require exascale resources.
- The Exabiome project will provide scalable tools for core metagenomics computations: assembly (MetaHipMer), protein clustering (HipMCL), and comparative analysis (GOTTCHA)
- Recently we developed MetaHipMer assembler and demonstrated quality comparable to other state-of-the-art assemblers.
- Current FY plan: demonstrate HipMCL clustering algorithm, demonstrate MetaHipMer performance, evaluate merAligner for GOTTCHA, release MetaHipMer software.

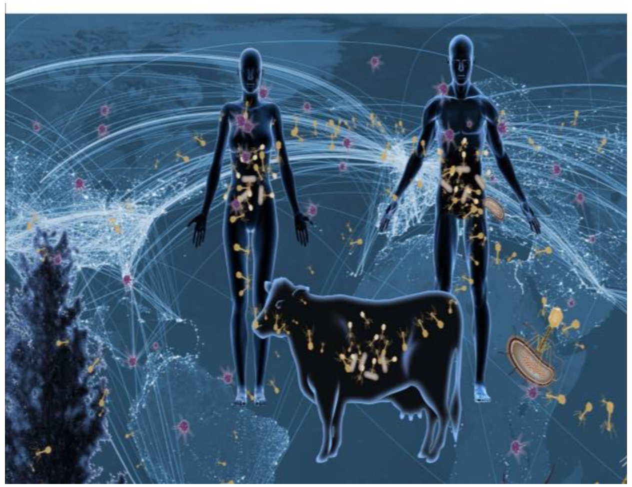

Metagenomics: genome sequencing communities of 1000s of microbial species. Key in environment, climate, agriculture, health and biomanufacturing.

Impact of MetaHipMer Quality Milestone

- MetaHipMer is of sufficient quality to be a viable alternative to current state-of-the art assemblers, and could be widely used by biologists for research into metagenomes.
- This is an important milestone because it demonstrates that the algorithms used are sufficiently accurate to form the basis for an assembler that can run on future exascale systems.

Project Accomplishments demonstrating Quality

- **MetaHipMer quality is as good as and often better than the best state-of-the-art assemblers**, including metaSPAdes and MEGAHIT
- For several commonly used evaluation datasets, MetaHipMer consistently provides the lowest number of mismatches with the highest error-free contiguity
- MetaHipMer can run efficiently on large-scale HPC systems, **reducing assembly times by orders of magnitude compared to other assemblers**. This enables more aggressive assembly approaches and accelerates scientific discovery.

Massively Parallel Protein Clustering

ECP WBS: ExaBiome
PI: Katherine Yelick, LBNL
Members: LBNL, JGI, LANL

Scope & Objectives

- Develop massively parallel Markov Cluster algorithm, called **HipMCL**, for clustering proteins.
- Major algorithmic components of HipMCL (a) sparse matrix-matrix multiplication, (b) connected component and (c) top-k selection. We develop communication- and memory-efficient parallel algorithms for them.
- The developed software will be used by biologist to cluster proteins from isolate-genome and metagenome datasets.
- The performance of HipMCL will be evaluated with novel metagenomic datasets on Cori, Titan and upcoming supercomputers.

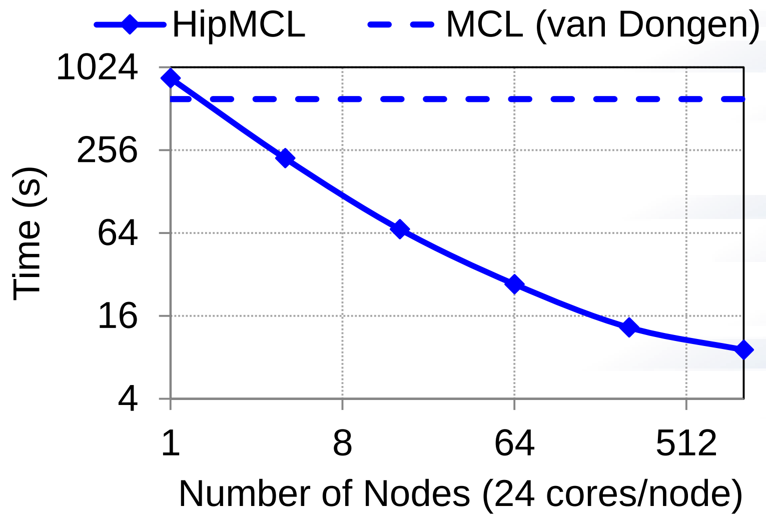

Scalability of HipMCL on NERSC/Edison and its performance relative to a state-of-the-art shared-memory code by van Dongen.

Impact

- **Science impact**: HipMCL enables unprecedented discovery in Biology. Current shared-memory implementation is very slow, requiring 45 days to cluster 50 million proteins. HipMCL can cluster the same data in less than an hour on 1K nodes of Edison.
- **ECP impact**: The computational need in biological clustering is reaching to exascale. HipMCL can address this need.

Project Accomplishment

- HipMCL clustered a network of 400 million proteins and 8 billion edges obtained from metagenomic data in 50 minutes on 2K nodes on Cori KNL partition. It is impossible to cluster this dataset with previous shared-memory implementation.
- Developed novel algorithms for thresholded sparse matrix-matrix multiply that adapts to the available memory.
- **Next step**: Develop communication avoiding algorithms for the sparse matrix kernels. Test with bigger datasets on higher concurrency.

Deliverables: Software for HipMCL and publications

EXASCALE COMPUTING PROJECT

ExaBiome: Demonstrate MetaHipMer at Scale

ECP WBS 1.2.1.20: ExaBiome
PI: Kathy Yelick, LBNL
Members: LBNL, LANL, JGI

Scope & Objectives

- Ensure MetaHipMer can run on at least 10K cores.
- Determine the maximum strong scaling for MetaHipMer on Edison, Titan or Cori.
- Produce microbenchmarks that capture key communication patterns in MetaHipMer.
- Run competing assemblers at the same set of concurrencies as MetaHipmer.
- Produce a report summarizing the performance of MetaHipMer and scalability, in the form of a draft of a technical paper.

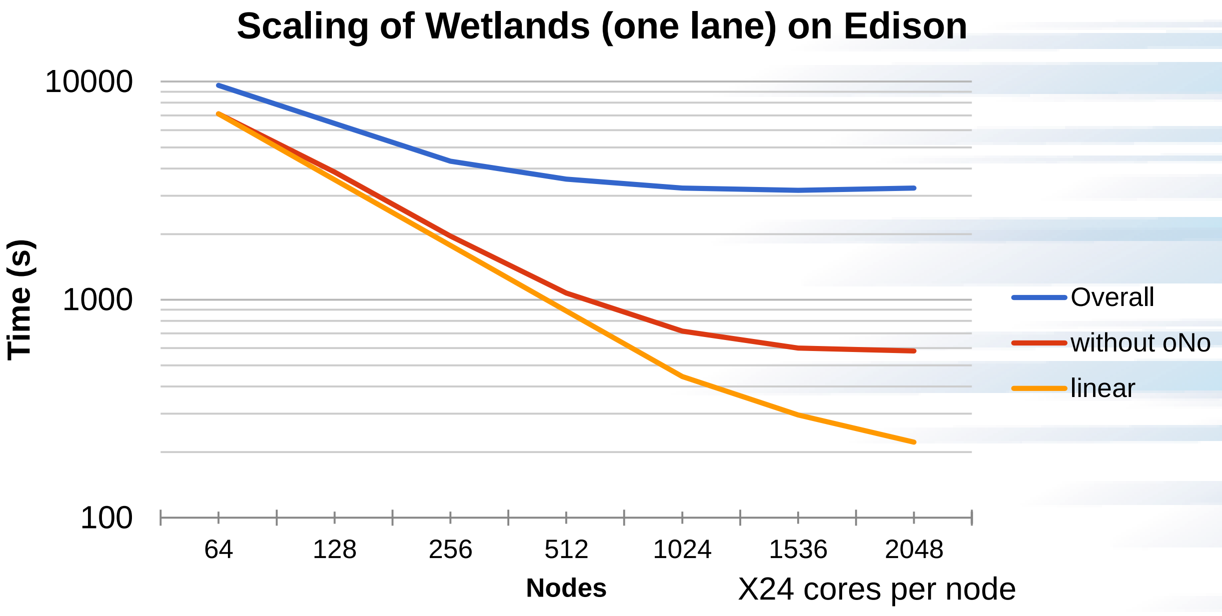

Impact

- New scalable metagenome assembler (MetaHipmer) reduces runtime by orders of magnitude at petascale
- Use processing speed to improve quality of metagenome assemblies
- Drive exascale architectures and programming systems to support irregular, graph-based analytics
-

Project Accomplishment and Next Steps

- Demonstrated MetaHipMer scalability on 1-lane Wetlands (above) and multiple synthetic metagenome data sets
- Measured impact of known serial heuristic (oNo) in scaffolding, to be replaced by parallel version
- Improved some memory scaling issue, but others still to be addressed
- Run multi-lane (full data) Wetlands and other metagenome data
- Improve quality of ribosomal DNA

EXASCALE COMPUTING PROJECT

ExaBiome: Release of MetaHipMer Software

Scope & Objectives

- Make the metagenome version of HipMer (aka MetaHipMer) publicly available
- This is an open source release available from a public website: https://sites.google.com/lbl.gov/exabiome/downloads
- The pipeline runs to completion without errors on several datasets and on multiple HPC systems.
- The quality is comparable to or better than in earlier milestones

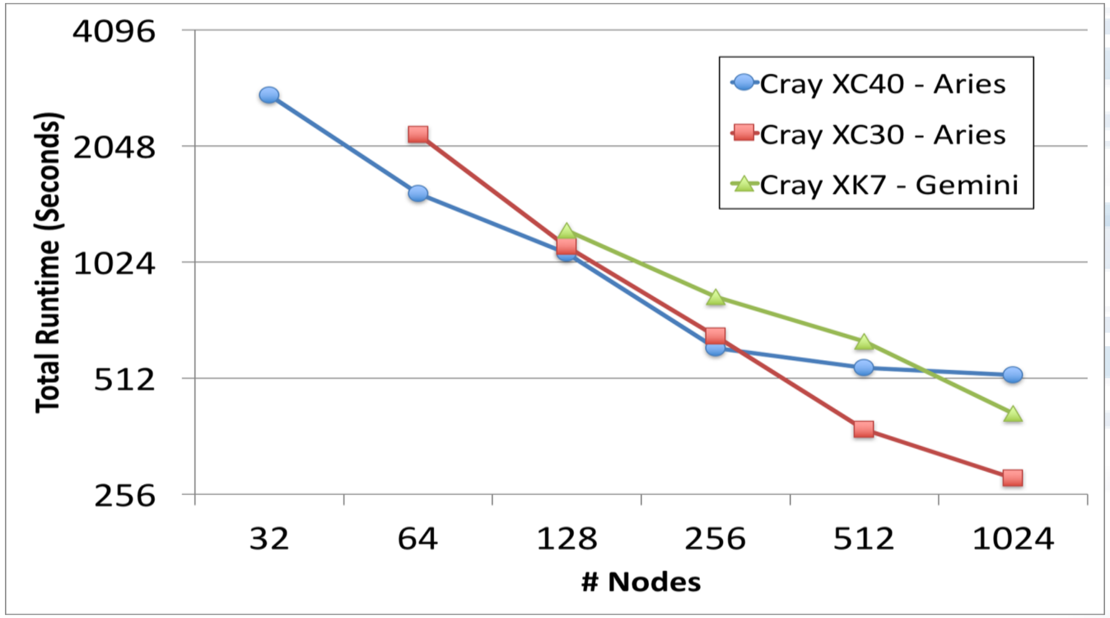

Impact

- Scalable metagenome assembler (MetaHipmer) reduces runtime by orders of magnitude
- Enables direct assembly (without error correction / filtering) of multi-terabyte data sets, which may be several times larger at runtime
- Drive exascale architectures and programming systems to support irregular, graph-based analytics

Project Accomplishment and Next Steps

- First release of metagenome version of HipMer
- Improved robustness, portability and usability of HipMer in addition to metagenome support
- HipMer next step: parallelize oNo with connected component computation; improve scaffolding quality, probably with backtracking graph walk
- GOTTCHA next step: merAligner analysis
- HipMCL next step: publication

EXASCALE COMPUTING PROJECT

Multiscale Coupled Urban Systems

- Urbanization is increasing demand for energy, water, transportation, healthcare, infrastructure, physical & cyber security and resilience, food, education—and deepening the interdependencies between these systems. New technologies, knowledge, tools needed to retrofit / improve urban districts, with capability to model and quantify interactions between urban systems
- Integrate modules for urban atmosphere and infrastructure heat exchange and air flow; building energy demand at district or city-scale, generation, and use; urban dynamics & activity based decision, behavioral, and socioeconomic models; population mobility and transportation; energy systems; water resources
- Chicago metro area as testbed for coupling agent-based social/economics model with transportation, regional climate, CFD microclimate, energy of (up to 800K) buildings
- Develop the Metropolitan Systems Energy and Economic Dynamics (Metro-SEED)

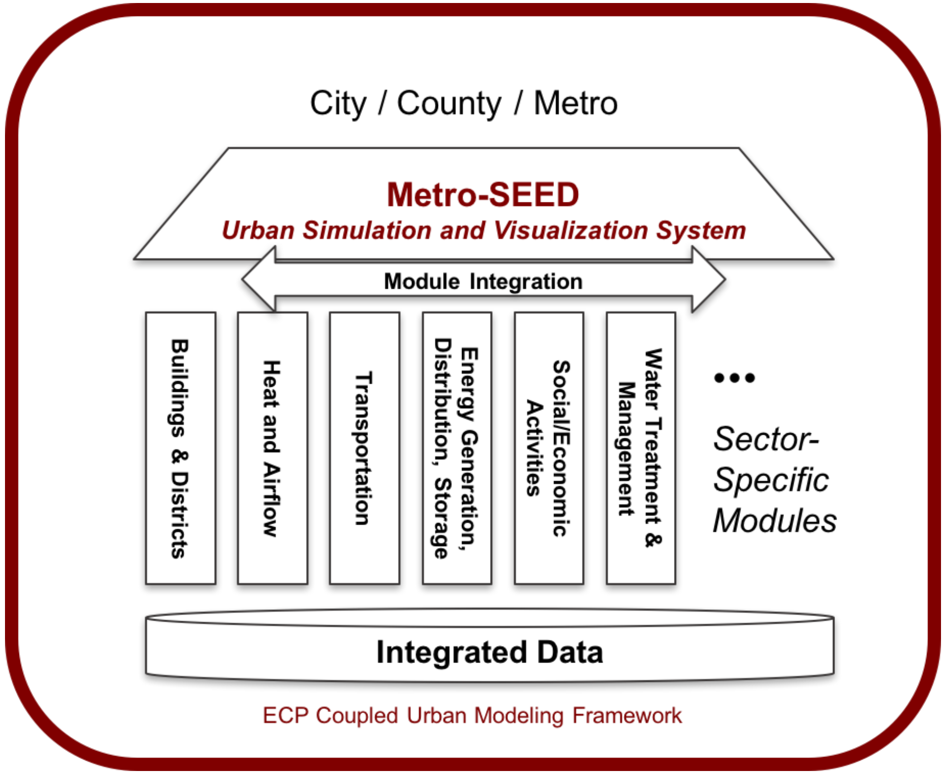

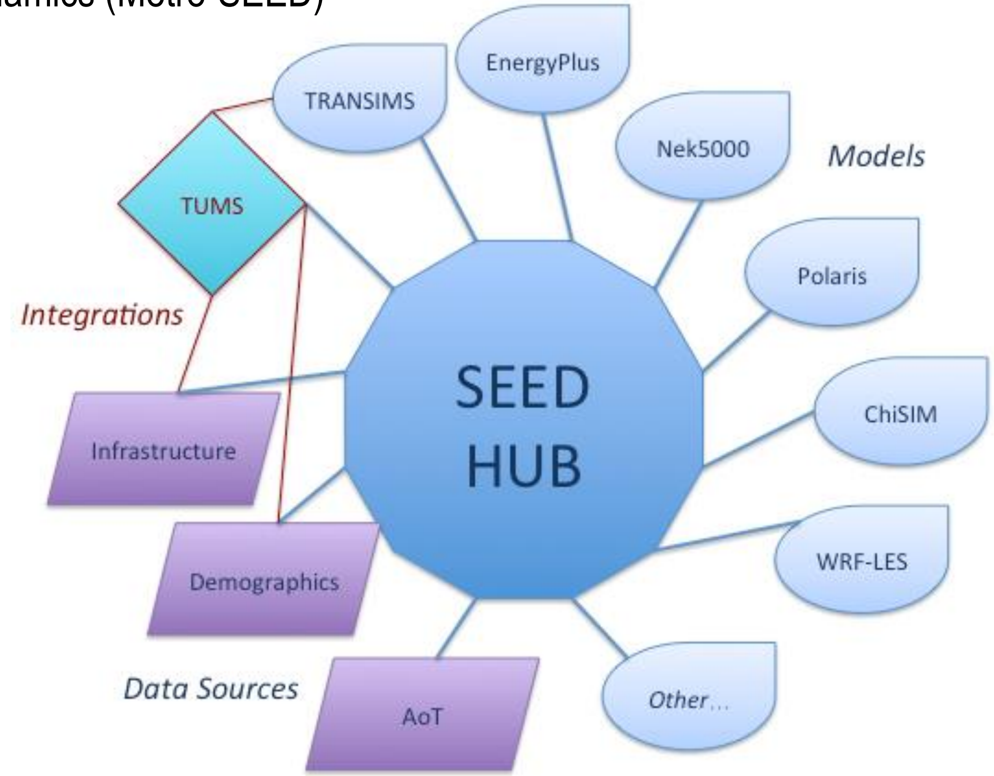

PI: Charles Catlett (ANL)

Coupled Urban Systems

Four Models to Understand Data Flow and Interconnections

- Transportation
- Socio-Economics
- Building Energy
- Urban Canyon Heat Flow

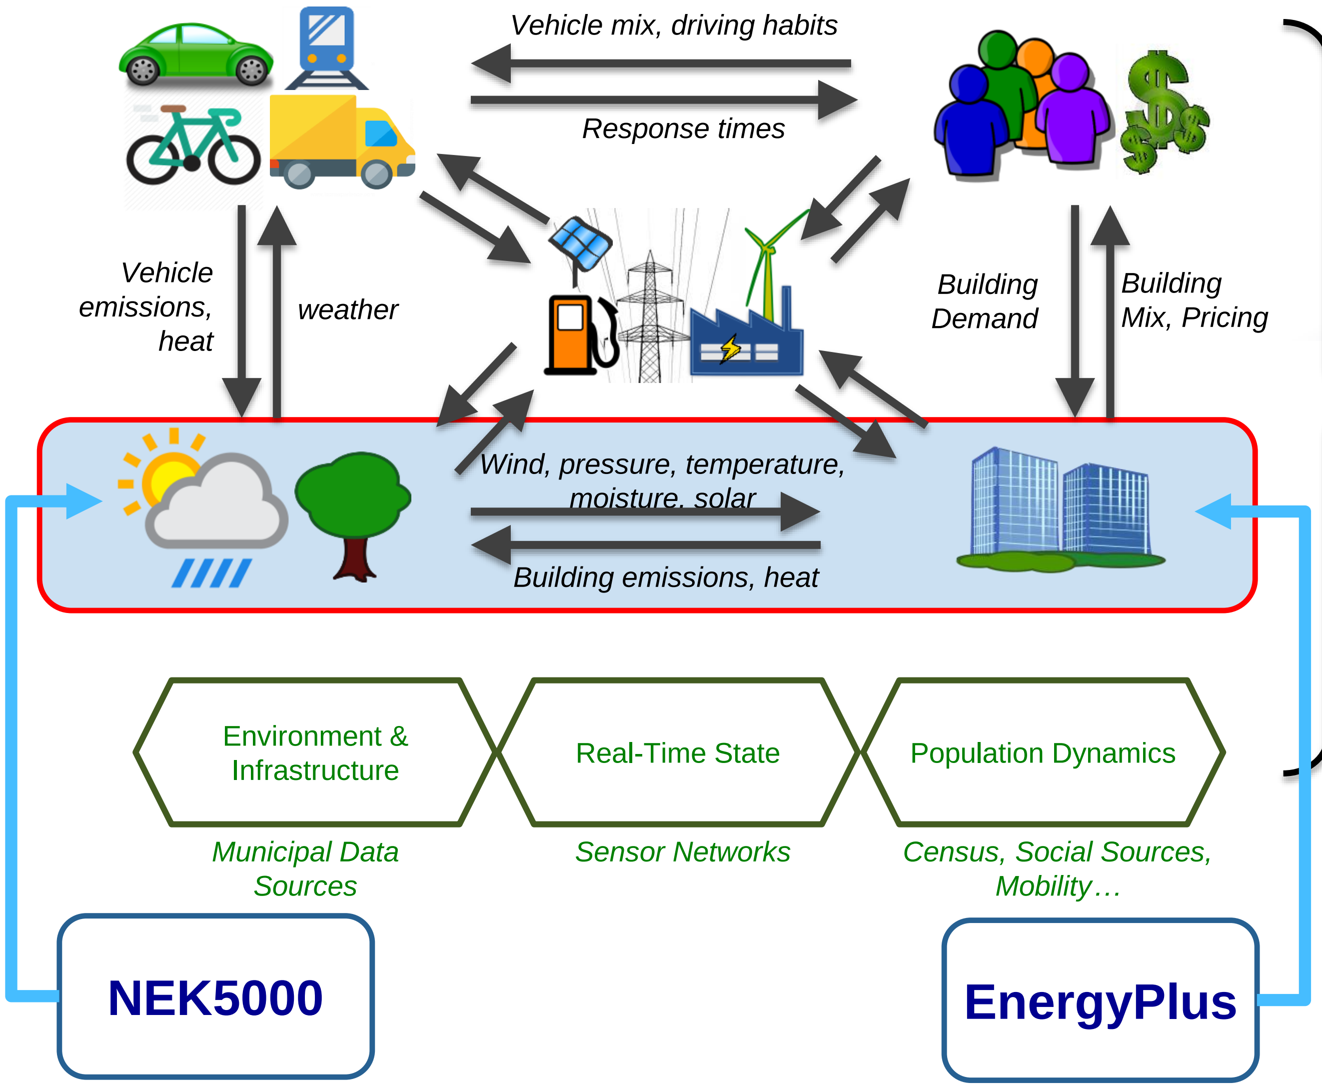

- Impacts of greenhouse gases (GHG) on local climate
- Resulting impacts on city function
- Incorporation of renewables into city energy portfolio
- Resilience of physical infrastructure
- Economic protection, resilience, and enhancement
- …

PI: Charles Catlett (ANL)

Coupling Strategy: Buildings and Urban Atmosphere

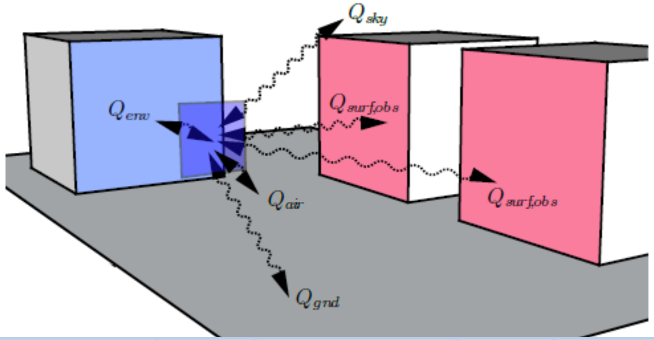

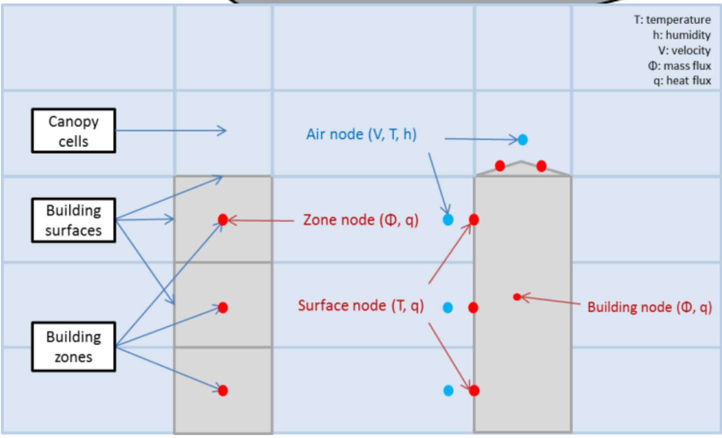


CAD model with atmosphere and building geometry

Integrating various urban data for building geometries

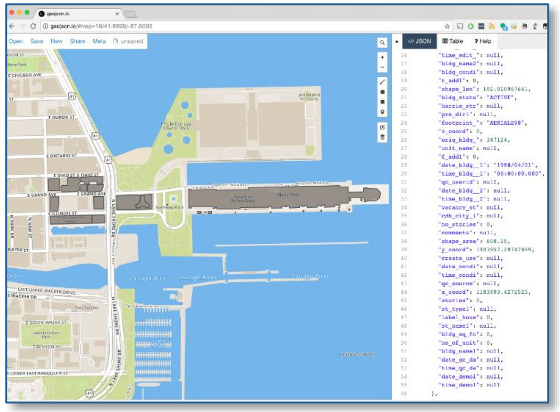
Building footprint.

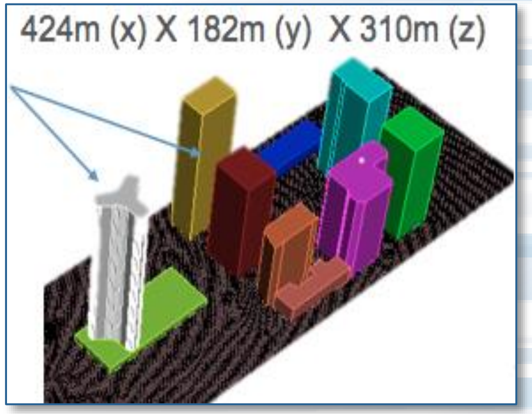

CAD model of the building geometry

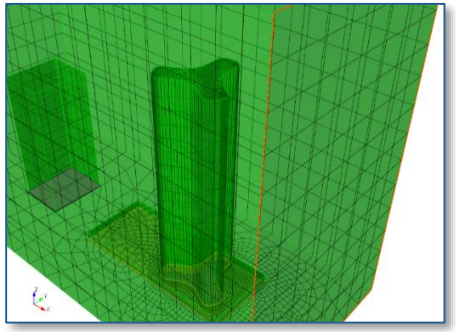
Isometric view of atmosphere and building (Lake Point Tower) mesh.

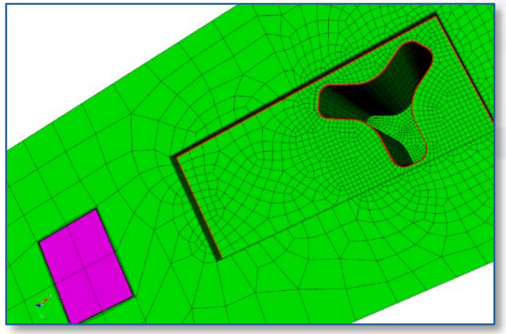
Bottom view showing boundary layer around buildings

Coupled Urban Systems

Prototype Data/Control Hub Framework

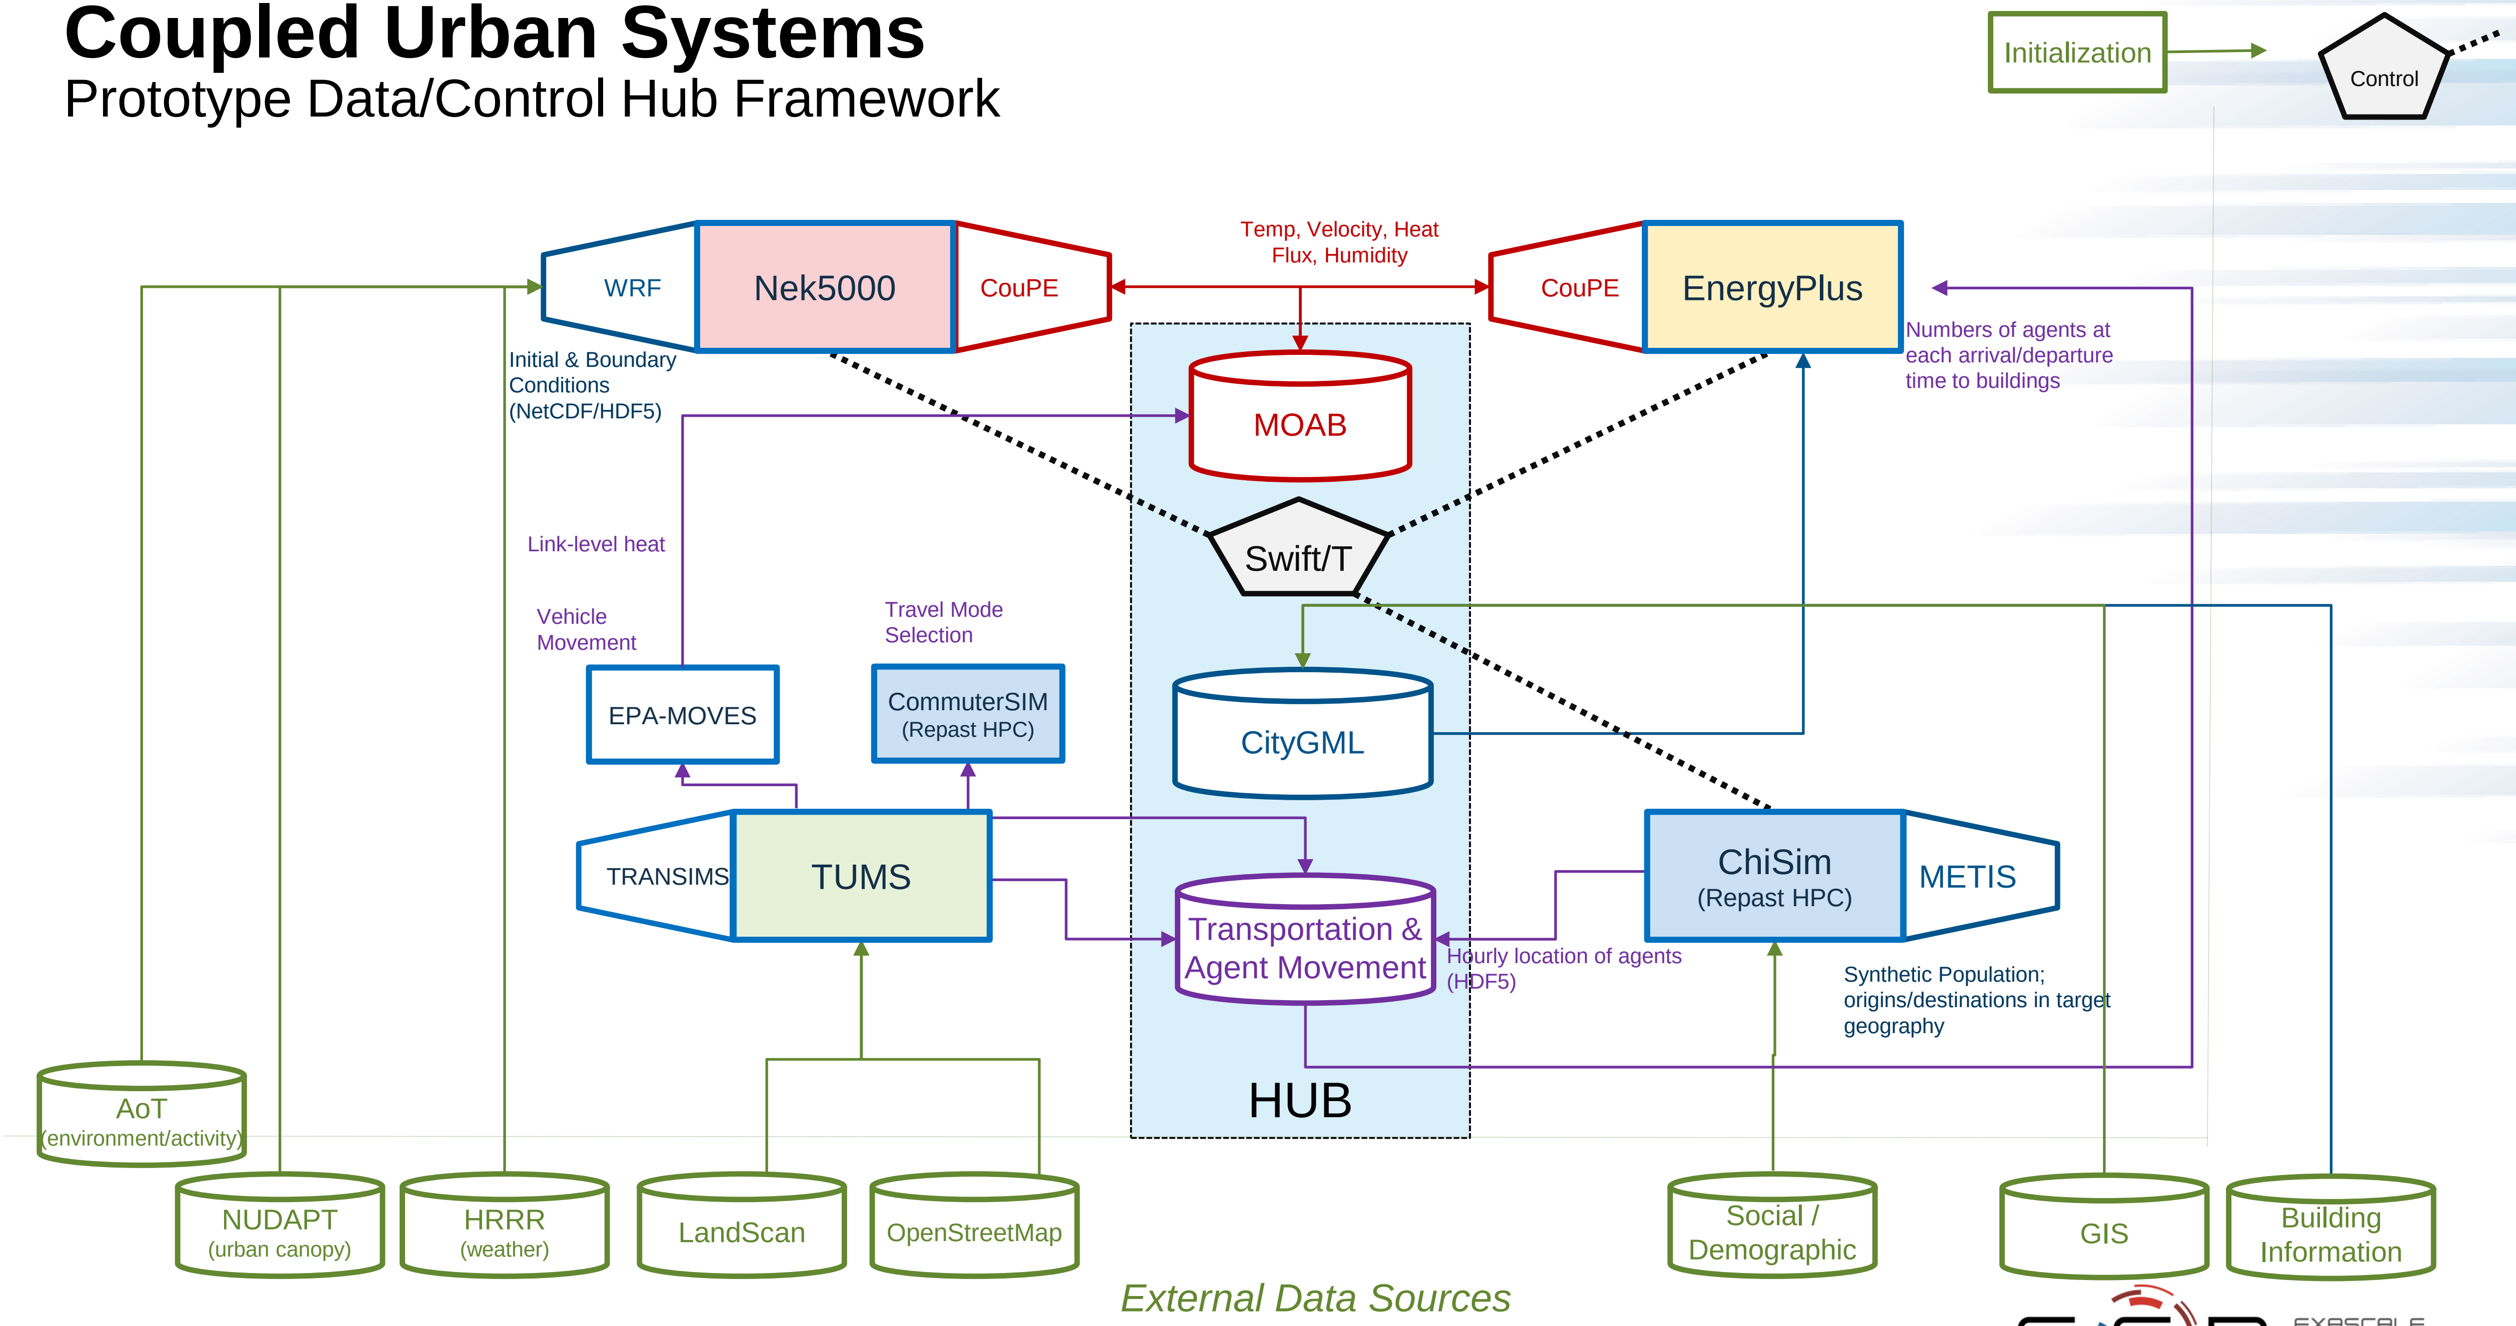

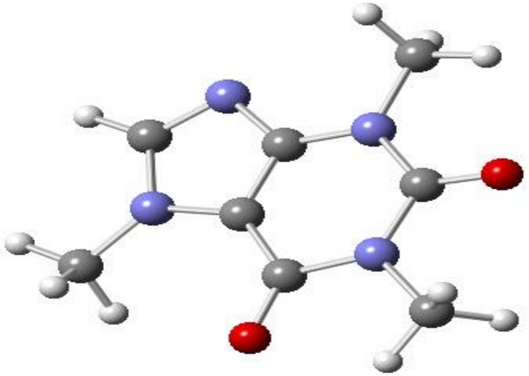

NWChemEx

Application Development in the Exascale Computing Project

NWChemEx Project
Moving Computational Chemistry to the Exascale

Thom H. Dunning, Jr., Northwest Institute for Advanced Computing
Theresa L. Windus, Ames Laboratory
Robert J. Harrison, Institute for Advanced Computational Science

AMES LABORATORY Creating Materials & Energy Solutions U.S. DEPARTMENT OF ENERGY
Argonne NATIONAL LABORATORY
BERKELEY LAB
BROOKHAVEN NATIONAL LABORATORY
OAK RIDGE National Laboratory
Pacific Northwest NATIONAL LABORATORY
1872 VirginiaTech Invent the Future®
XASCALE OMPUTING ROJECT

From NWChem to NWChemEx

- Redesign NWChem for Exascale Computing Technologies
  - Address limitations (new design, state-of-the-art algorithms, take advantage of massive levels of concurrency, …)
  - Update code base to C++ to enable more flexible design
- Creating New Computational Framework
  - Developing “Tensor Algebra for Many-Body Methods” (TAMM) to replace TCE
  - Working closely with XGA team on Global Arrays
  - Exploring other software technologies in the ECP
- Implementing Limited Set of Electronic Structure Methods
  - HF/DFT methods
  - Coupled cluster methods (*most challenging*)
  - Embedding method (DFET): CC-in-DFT
- Implementing Basic Set of Classical/Ab Initio Molecular Dynamics Methods

PI: Thom Dunning (PNNL)

Overview of NWChemEx 1.0

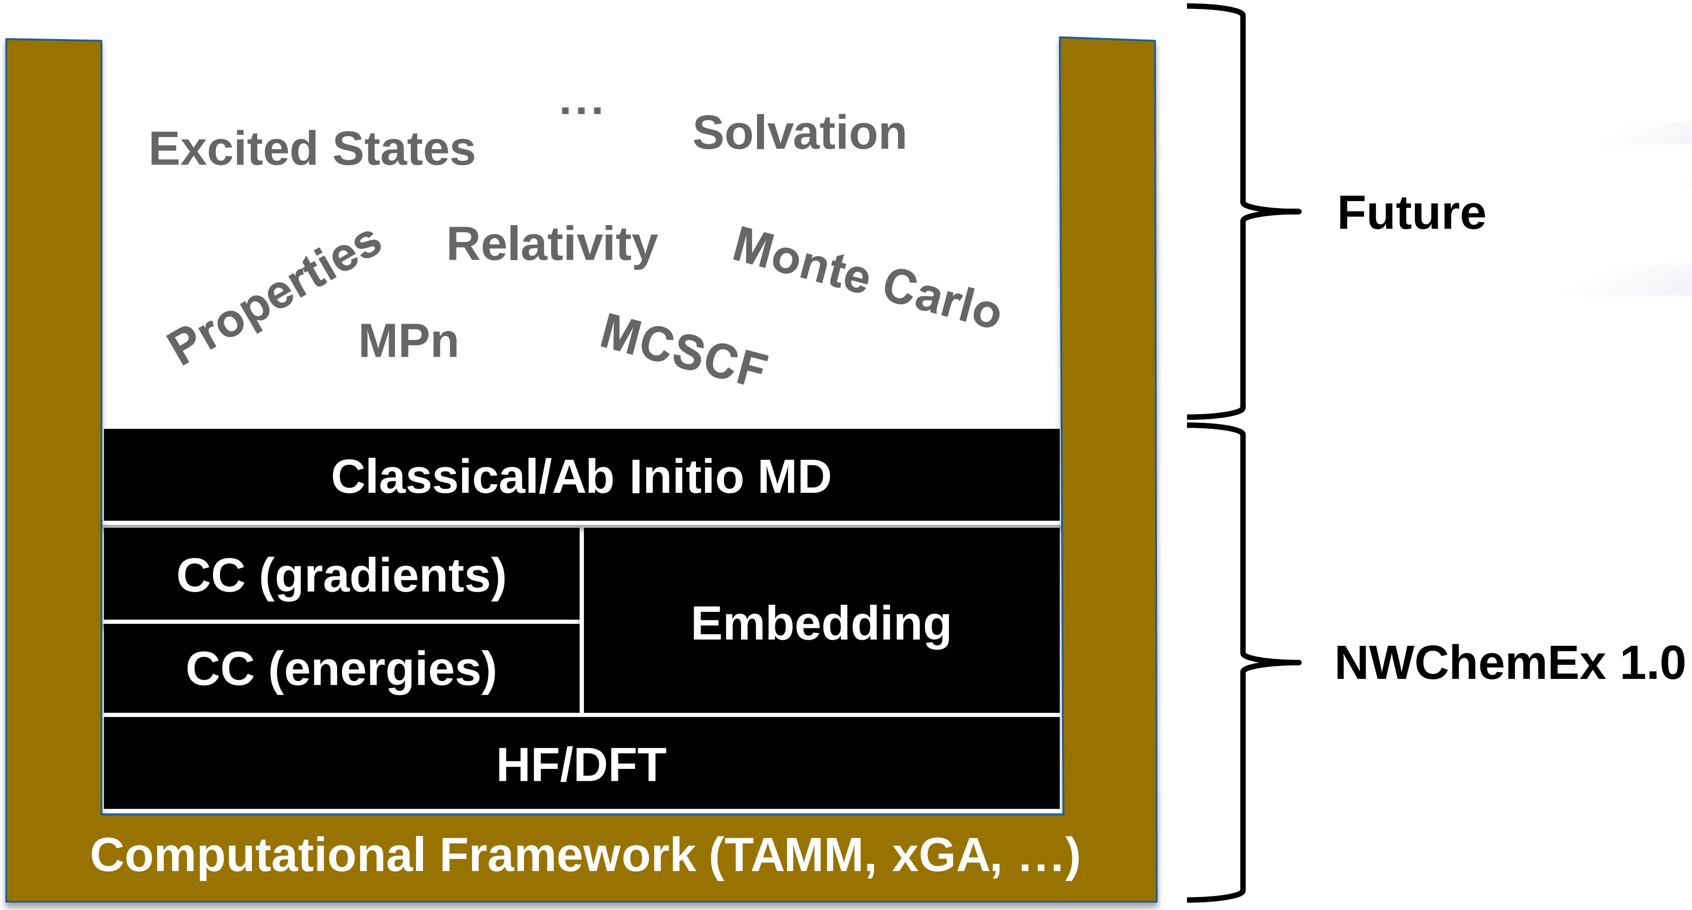

PI: Thom Dunning (PNNL)

23

EXASCALE COMPUTING PROJECT

Targeted Science Challenges for NWChemEx

- **DOE 2014-2018 Strategic Plan**
  - Development of high performance computing models that demonstrate that biomass can be a viable, sustainable feedstock for biofuels, hydrogen and other products
- **Two Inter-related Science Challenges**
  - ***Efficient Production of Biomass***: Development of detailed molecular model of transport processes across cellular membranes that control stress responses to aid in the development of stress-resistant crops

    A realistic model of proton-controlled membrane transport requires modeling hundreds of thousands of atoms to describe a suitable portion of the cellular membrane, a sufficient region of the immediate cytoplasmic environment, and the 3500-atom Bax inhibitor-1 protein containing two active sites.
  - ***Efficient Conversion of Biomass to Biofuels:*** Development of detailed molecular model of catalytic conversion of biomass-derived alcohols to biofuels to aid in the discovery of energy efficient conversion processes

    A realistic model of the catalytic dehydration of an alcohol such as 2-propanol requires modeling hundreds of thousands of atoms to describe a portion of the zeolite along with the alcohol and other species involved in the dehydration process. Truly predictive results require coupled cluster calculations to describe the active region, $O(10^2$–$10^3)$ atoms, embedded in a density functional theory description of thelarger environment, $O(10^4$–$10^5)$ atoms.

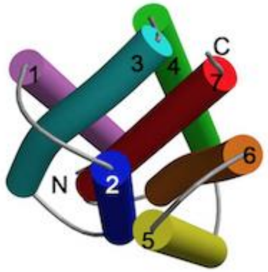

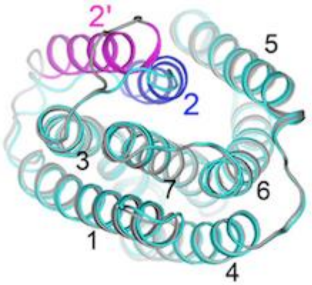

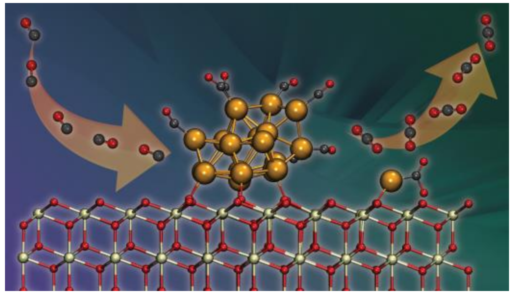

24

PI: Thom Dunning (PNNL)

NWChemEx's Tensor Algebra for Many-body Methods (TAMM)

SCOPE & OBJECTIVES

- Scalable, performance-portable implementations of many-body methods
- Tailored implementations for specific architectures
- Multiple implementations for different use cases
  - Communication-efficient for highly parallel calculations
  - Memory-efficient ensemble calculations
- Many-body methods are essential for robustly predictive science, validation of other models, and for high accuracy

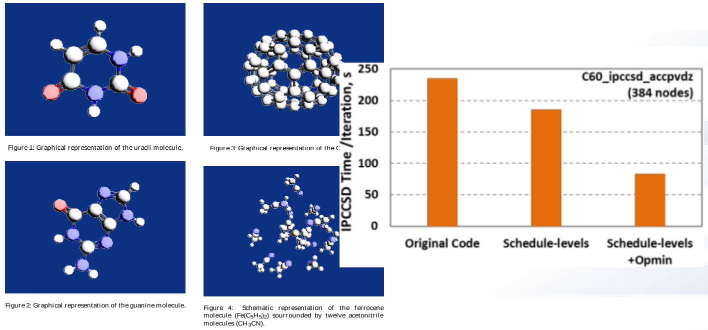

TAMM implementations can be used to calculate properties of large molecular systems. TAMM implementation is already 3X more efficient than existing TCE code.

IMPACT

- Taking "Gold" standard Coupled Cluster methods to exascale with tailored execution plans for specific architectures
- High-productivity and high-performance building blocks to enable new methods to be developed rapidly
- Library that can be used by the broader chemistry and material science communities.

PROJECT ACCOMPLISHMENTS

- Initial implementation of key Coupled Cluster (CC) methods in TAMM
  - More general, more scalable implementation in 100x fewer lines of code than NWChem's current TCE (tensor contraction engine)
  - Initial prototypes of new CC methods demonstrated
- Close engagement with multiple ECP software technology teams to prototype and determine best software technologies upon which to base NWChemEx

PI: Thom Dunning (PNNL)

EXASCALE COMPUTING PROJECT

Cloud-Resolving Climate Modeling of the Earth's Water Cycle
Energy Exascale Earth System Model (E3SM) Multiscale Modeling Framework

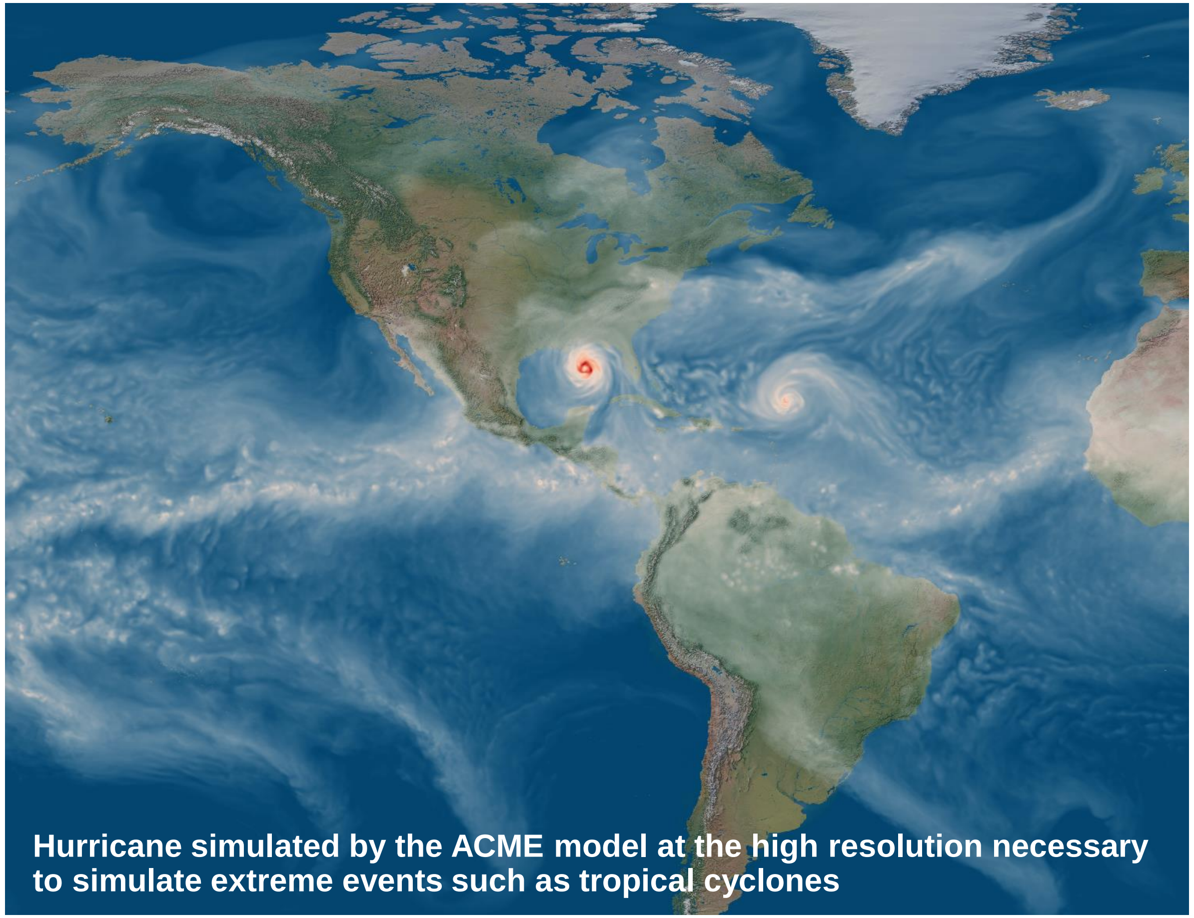

Hurricane simulated by the ACME model at the high resolution necessary to simulate extreme events such as tropical cyclones

- Cloud-resolving Earth system model with throughput necessary for multi-decade, coupled high resolution climate simulations
- Target substantial reduction of major systematic errors in precipitation with realistic / explicit convective storm treatment
- Improve ability to assess regional impacts of climate change on the water cycle that directly affect multiple sectors of the U.S. and global economies (agriculture & energy production)
- Implement advanced algorithms supporting a *super-parameterization* cloud-resolving model to advance climate simulation and prediction
- Design super-parameterization approach to make full use of GPU accelerated systems, using performance performance portable approaches, to ready the model for capable exascale

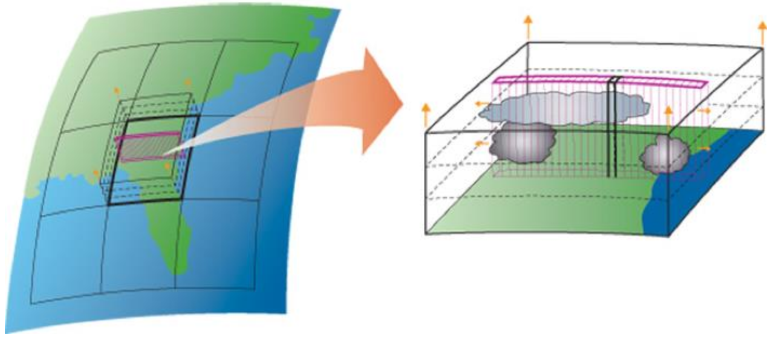

PI: Mark Taylor (SNL)

Cloud-Resolving Climate Modeling of Earth's Water Cycle

E3SM-MMF

Exascale challenge problem

- Earth system model (ESM) with throughput needed for multi-decadal coupled high-resolution (~1 km) climate simulations, reducing major systematic errors in precipitation models via explicit treatment of convective storms
- Improve regional impact assessments of climate change on water cycle, e.g., influencing agriculture/energy production
- Integrate cloud-resolving GPU-enabled convective parameterization into E3SM ESM using Multiscale Modeling Framework (MMF); refactor key E3SM model components for GPU systems
- E3SM-MMF goal: Fully weather-resolving atmosphere/cloud-resolving superparameterization, eddy-resolving ocean/ice components, throughput (5 SYPD) enabling 10–100 member ensembles of 100 year simulations

Applications

- E3SM Earth system model: E3SM-Atmosphere
- MPAS (Model Prediction Across Scales)-Ocean (ocean)
- MPAS-Seaice (sea ice); MPAS-Landice (land ice)
- SAM (System for Atmospheric Modeling)

Software technologies

- Fortran, C++, MPI, OpenMP, OpenACC
- Kokkos, Legion
- PIO, Trilinos, PETSc
- ESGF, Globus Online, AKUNA framework

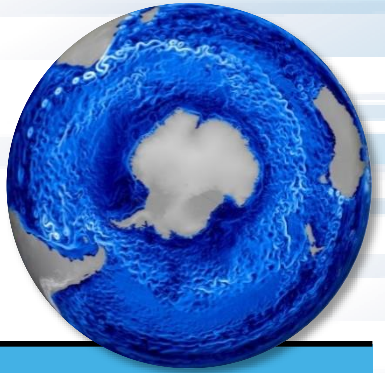

Risks and challenges

- Insufficient LCF allocations
- Obtaining necessary GPU throughput on the cloud-resolving model
- Cloud-resolving convective parameterization via multi-scale modeling framework does not provide expected improvements in water cycle simulation quality
- Global atmospheric model cannot obtain necessary throughput
- MPAS ocean/ice components not amenable to GPU acceleration

PI: Mark Taylor (SNL)

Development Plan

- **Y1:** Demonstrate E3SM-MMF model for Atmospheric Model Intercomparison Project configuration; complete 5 year E3SM-MMF simulation with active atmosphere and land components at low resolution and E3SM atmosphere diagnostics/ metrics
- **Y2:** Demonstrate E3SM-MMF model with active atmosphere, land, ocean and ice; complete 40 year simulation with E3SM coupled group water cycle diagnostics/metrics
- **Y3:** Document GPU speedup in performance-critical components: Atmosphere, Ocean and Ice; compare SYPD with and without using the GPU
- **Y4:** E3SM-MMF configuration integrated E3SM model; document highest resolution able to deliver 5 SYPD; complete 3 member ensemble of 40 year simulations with all active components (atmosphere, ocean, land, ice) with E3SM coupled group diagnostics/metrics

ECP WBS 1.2.1.15/ACME-MMF
Taylor, SNL
Members: ANL, LANL, LLNL, ORNL, PNNL, CSU, UCI

Milestone: E3SM-MMF AMIP capability

Scope & Objectives

- Develop capability to assess regional impacts of climate change on the water cycle that directly affect the US economy such as agriculture and energy production.
- ACME-MMF approach addresses structural uncertainty in cloud processes by replacing traditional parameterizations with cloud resolving “superparameterization” within each grid cell of global climate model. Super-parameterization dramatically increases arithmetic intensity.
- ACME-MMF: Use a multiscale approach ideal for new architectures to achieve cloud resolving convection on Exascale resources

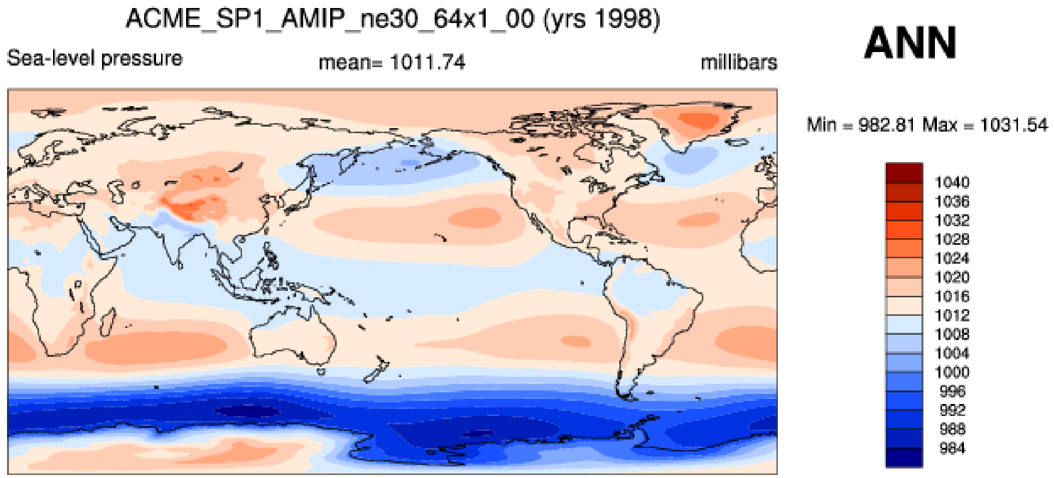

Sea level pressure (climatology) from the first long simulation of the ACME-MMF running in AMIP mode, using a 1 degree global grid and a 1km resolution super-parameterization.

ACME-MMF Impact

- A cloud resolving climate model is needed to reduce major systematic errors in climate simulations due to structural uncertainty in numerical treatments of convection – such as convective storm systems.

Project Accomplishments

- Further refine the E3SM-MMF modeling system so that it can run stably for long integrations in AMIP mode: active atmosphere, land and sea ice component models with prescribed sea surface temperature and sea ice extent.
- Evaluate the simulated climate with atmospheric diagnostics and verified the simulation quality is sufficient to serve as our baseline for ongoing model improvements.
- This code is also the basis of our ongoing work on GPU optimization/porting.

Deliverables: Milestone report at https://confluence.exascaleproject.org/display/ADSE15/Milestone+Reports

EXASCALE COMPUTING PROJECT

IDEAS productivity

www.ideas-productivity.org

Interoperable Design of Extreme-scale Application Software (IDEAS)

Motivation

Enable ***increased scientific productivity,*** realizing the potential of extreme- scale computing, through ***a new interdisciplinary and agile approach to the scientific software ecosystem***.

Objectives

Address confluence of trends in hardware and increasing demands for predictive multiscale, multiphysics simulations.

Respond to trend of continuous refactoring with efficient agile software engineering methodologies and improved software design.

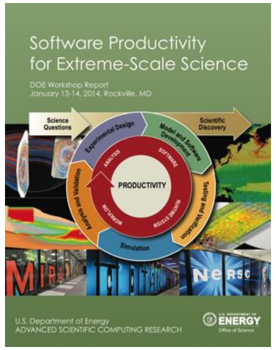

Impact on Applications & Programs

Terrestrial ecosystem ***use cases tie IDEAS to modeling and simulation goals*** in two Science Focus Area (SFA) programs and both Next Generation Ecosystem Experiment (NGEE) programs in DOE Biologic and Environmental Research (BER).

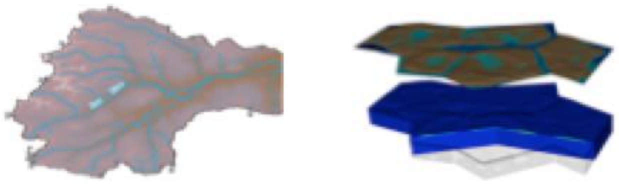

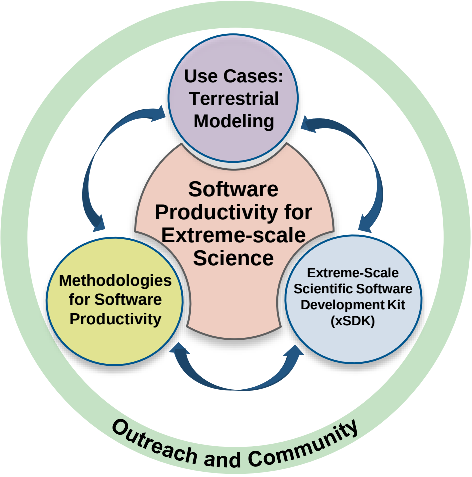

Approach

ASCR/BER partnership ensures delivery of both crosscutting methodologies and metrics with impact on real application and programs.

Interdisciplinary multi-lab team (ANL, LANL, LBNL, LLNL, ORNL, PNNL, SNL)

ASCR Co-Leads: Mike Heroux (SNL) and Lois Curfman McInnes (ANL)

BER Lead: David Moulton (LANL)

Integration and synergistic advances in three communities deliver scientific productivity; outreach establishes a new holistic perspective for the broader scientific community.

IDEAS history

DOE National Laboratory Announcement Number LAB 14-0001

ASCR/BER partnership began in Sept 2014

Program Managers:

- Paul Bayer, David Lesmes (BER)
- Thomas Ndousse-Fetter (ASCR)

First-of-a-kind project: qualitatively new approach based on making productivity and sustainability the explicit and primary principles for guiding our decisions and efforts.

U.S. DEPARTMENT OF ENERGY | Office of Science

PIs: Mike Heroux (SNL) and Lois Curfman McInnes (ANL)

Next Steps: xSDK4ECP & IDEAS-ECP

xSDK4ECP: Develop **community policies** and **interoperability layers** among numerical packages as needed by ECP scientific applications

- Coordinated use of on-node resources
- Integrated execution
- Coordinated and sustainable documentation, testing, packaging, and deployment

xSDK

Current xSDK packages:

- **Numerical libraries: hypre, PETSc, SuperLU, Trilinos**
- **Domain components: Alquimia, PFLOTRAN**

Packages working toward xSDK compatibility:

- **Chombo, MFEM, SUNDIALS, ALExa** (AMP, DTK, TASMANIAN)
- **Dense linear algebra packages:** MAGMA, PLASMA, DPLASMA, ScaLAPACK, LAPACK

xSDK4ECP Team:

Michael Heroux (SNL), Co-Lead PI
Lois Curfman McInnes (ANL), Co-Lead PI
Jim Demmel (UC Berkeley), Co-PI
Jack Dongarra (UTK), Co-PI
Xiaoye Sherry Li (LBNL), Co-PI
Carol Woodward (LLNL), Co-PI
Ulrike Meier Yang (LLNL), Co-PI
Satish Balay (ANL)
David Gardner (LLNL)
Piotr Luszczek (UTK)
Slaven Peles (LLNL)
Ben Recht (UC Berkeley)
Jacob Schroder (LLNL)
Barry Smith (ANL)
Keita Teranishi (SNL)
Jim Willenbring (SNL)

IDEAS-ECP: Collaborate with ECP app teams to understand productivity bottlenecks and improve practices

Productivity and Sustainability Improvement Planning Tools:

Helping a software team to increase software quality while decreasing the effort, time, and cost to develop, deploy, maintain, and extend software over its intended lifetime.

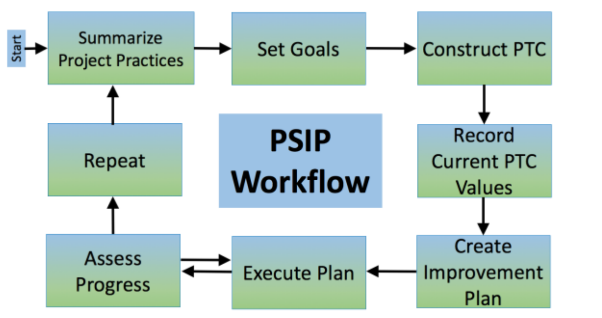

better scientific software

Coming soon: New web-based hub for scientific software improvement exchange

IDEAS-ECP Team:

Michael Heroux (SNL), Co-Lead PI
Lois Curfman McInnes (ANL), Co-Lead PI
David Bernholdt (ORNL), Co-PI, Outreach Lead
Todd Gamblin (LLNL), Co-PI
Osni Marques (LBNL), Co-PI
David Moulton (LANL), Co-PI
Boyana Norris (Univ of Oregon), Co-PI
Satish Balay (ANL)
Roscoe Bartlett (SNL)
Anshu Dubey (ANL)
Rinku Gupta (ANL)
Christoph Junghans (LANL)
Reed Milewicz (SNL)
Mark Miller (LLNL)
Jared O'Neal (ANL)
Elaine Raybourn (SNL)
Barry Smith (ANL)
Louis Vernon (LANL)
Greg Watson (ORNL)
James Willenbring (SNL)

Facilities liaisons:
Lisa Childers (ALCF)
Rebecca Hartman-Baker (NERSC)
Judy Hill (OLCF)
Hai Ah Nam (LANL)
Jean Shuler (LLNL)

ASCR/BER provided first of a kind funding to prove two concepts:

1. Creating the xSDK as a product and community provides a superior solution for app developers using libraries by enabling turnkey installation, compatible builds, and interoperability. This is especially important for multiscale and multiphysics projects that must use these libraries in combination.

2. Explicit investment in developing, adapting and adopting new and better software practices can improve developer productivity and software sustainability at a time when such improvements are essential for transforming capabilities for new platforms, coupling multiscale and multiphysics, and improving the effectiveness of our highly skilled computational scientists.

- *Our ECP efforts build on and expand what ASCR/BER started.*
- *BER Use Case collaboration with IDEAS and xSDK is continuing.*
- *Our ECP efforts would not be possible without the foresight of Paul Bayer, David Lesmes, and Thomas Ndousse-Fetter, who anticipated these needs.*

xSDK release 0.2.0: Packages can be readily used in combination by multiphysics, multiscale applications

Opportunities for more BER applications

Opportunities for more BER components

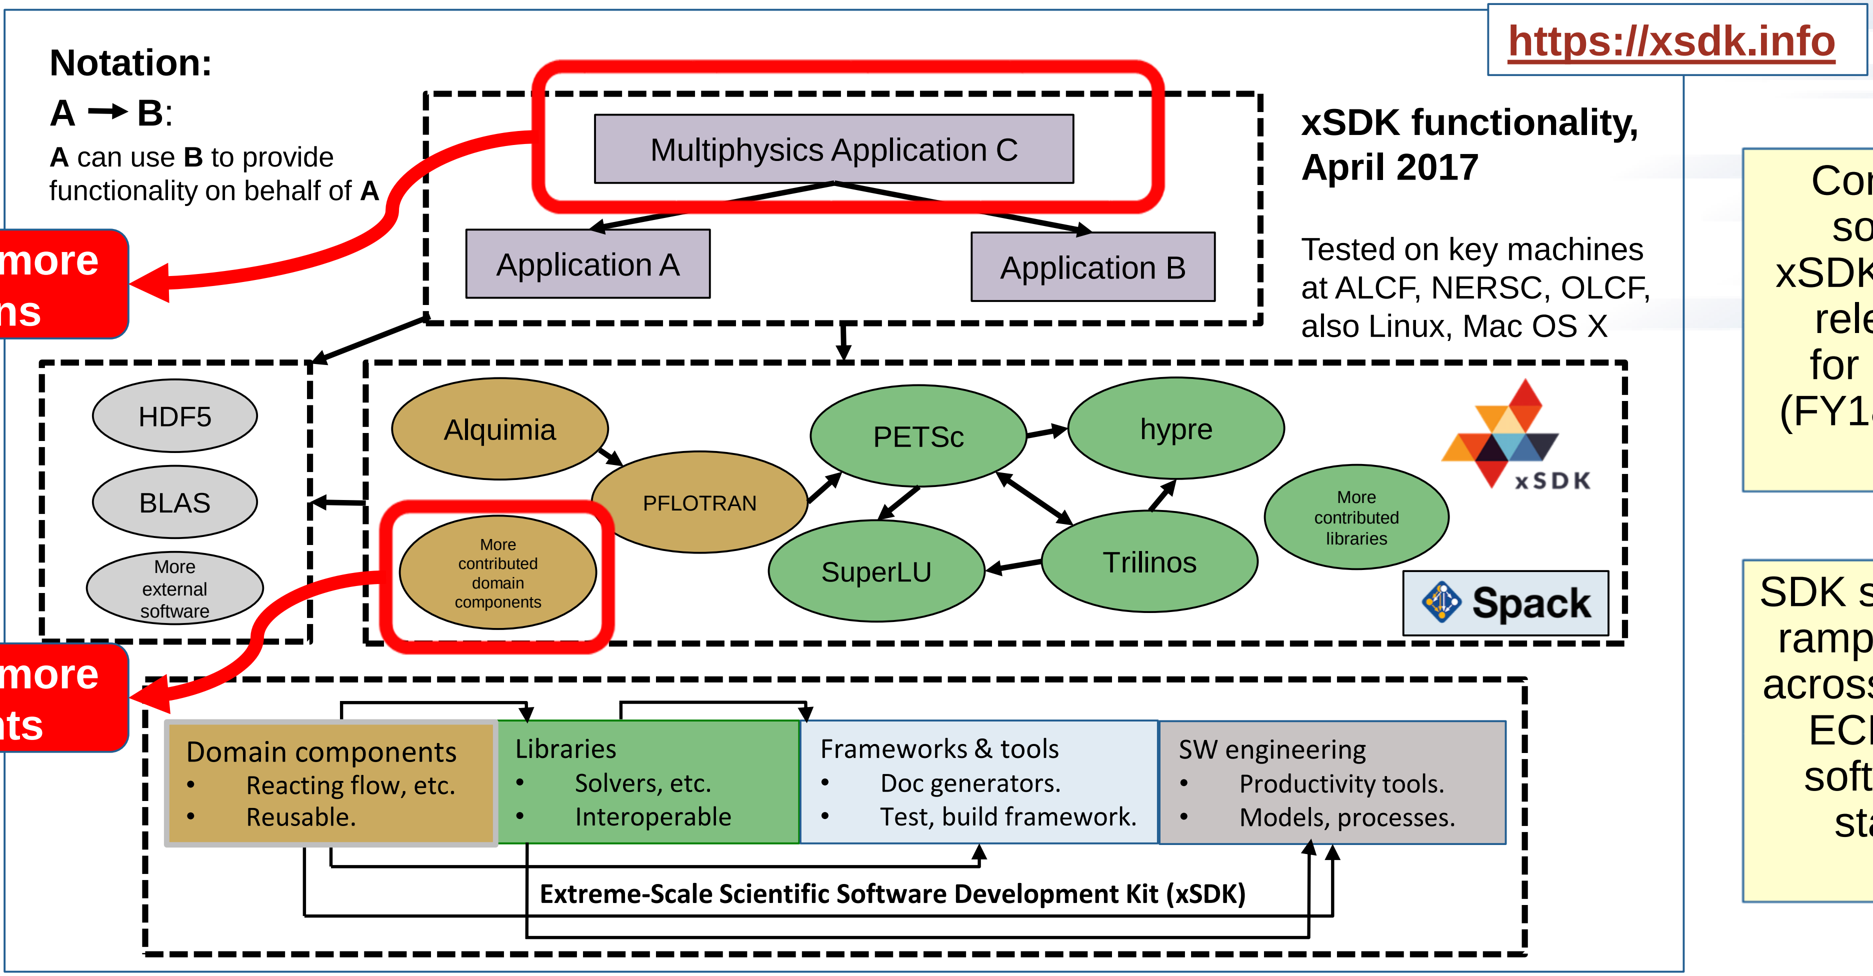

Coming soon: xSDK-0.3.0 release for ECP (FY18, Q1)

SDK strategy ramping up across entire ECP ST software stack

Ref: *xDSK Foundations: Toward an Extreme-scale Scientific Software Development Kit,* Bartlett et al, Feb 2017, https://arxiv.org/abs/1702.08425, to appear in **Supercomputing Frontiers and Innovations**.

U.S. DEPARTMENT OF ENERGY | Office of Science | IDEAS productivity

PIs: Mike Heroux (SNL) and Lois Curfman McInnes (ANL)

EXASCALE COMPUTING PROJECT

Opportunities for Broadened ECP Application Scope in Support of the BER Mission

Assess water security risk through hyper-resolution global terrestrial hydrology simulations

- Globally, “high levels of threat to water security” (e.g. Vörösmarty et al. 2012)
- Need integrated perspective coupling surface, shallow subsurface, and deep aquifers at high spatial resolution
- New computational strategies exploiting multiscale nature of watersheds are needed to reach exascale
- Leverage and extend strategies and software developed in the IDEAS project
- Enables hyper-resolution global models and supports detailed watershed-system simulations

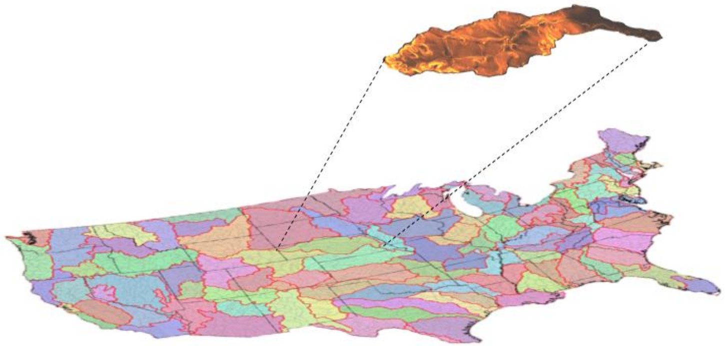

Resolve vegetation-soil ecosystem structure and function through individual-based model

- Knowledge exists to produce detailed individual-based models describing vegetation physiology and biochemistry.
- Integration of individual-based models over large spatial scales with explicit representation of plant-plant and plant-soil interactions has not been attempted.
- Global-scale parameterization of ecosystem process models can be improved by studying the emergent behavior of large and diverse populations of virtual plants interacting in/with a virtual soil ecosystem.
- Exascale computation required to achieve large-scale virtual ecosystem implementations.

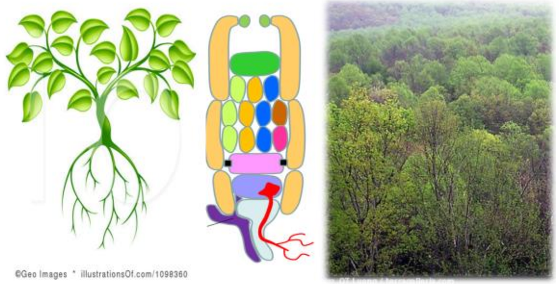

Quantify the Vulnerability of Critical Energy Infrastructures due to Intensified Flood Risks

- Enable physics-based, high-resolution real-time dynamical flood simulation
- Provide surface flood regime for more intuitive assessment of risk and vulnerability of critical infrastructures
- Large computational gain through GPU-acceleration demonstrated by multiple flooding studies and applications
- Can be coupled with either real-time streamflow forecasting models or climatic land surface models for different applications

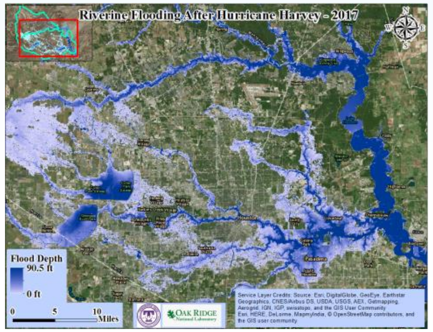

Hurricane Harvey flood regime simulated by TTU/ORNL GPU-accelerated 2D hydraulics flood model

EXASCALE COMPUTING PROJECT

Opportunities for Broadened ECP Application Scope in Support of the BER Mission

- **Biological complexity of plant and microbial metabolism and interfaces across scales spanning molecules to ecosystems.**
  - Systems biology modeling & mining of omics and microbiome/viriome phenotypes
  - Machine and deep learning algorithm development
    - Discovery of epistatic and pleiotropic interactions
    - Protein-protein interaction network prediction
  - Mining multiplexed neutron imaging for spatial and temporal resolution of microbiomes and biomarkers
- **Co-Evolutionary Networks: From Genome to 3D Proteome**
  - Exhaustive 3D characterization of an organism's protein structures and their molecular machines and assemblies.
    - Integration of evolutionary (ortholog) data and biophysical simulation
    - Binding site prediction for drugs, metabolites and nucleic acids
  - Pharmacogenomics, precision medicine and high-throughput compound screening

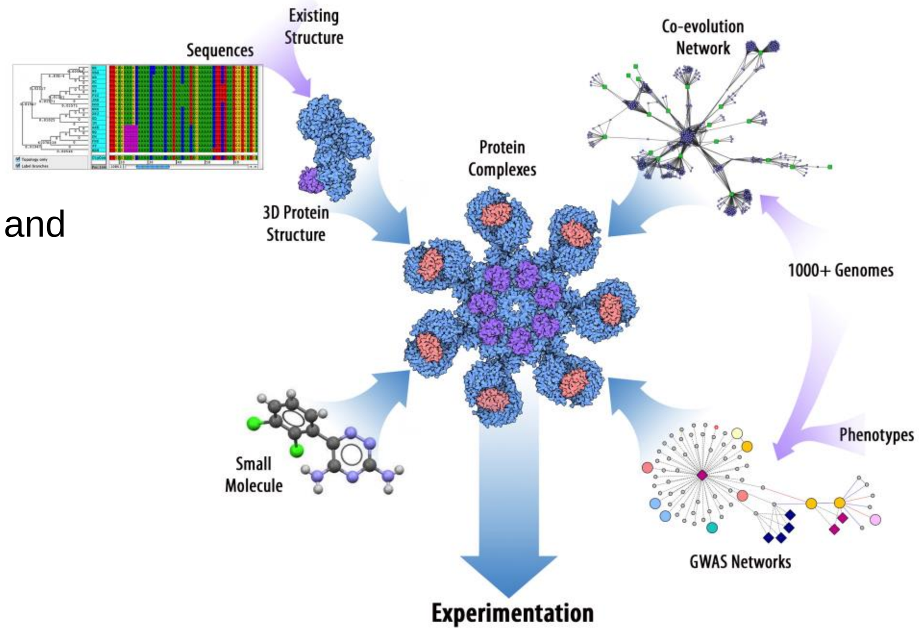

ECP Strategic Goals and Outcomes

- ***Goal: Applications***
  - Foundational element of ECP and vehicle for delivery of results from the exascale systems enabled by ECP. Each addresses an exascale challenge problem—a problem of strategic importance and national interest that is intractable without at least 50 times the computational power of today's systems
- ***Goal: Software Technologies***
  - Underlying technologies that applications are built and relied on: essential for application performance, portability, integrity, and resilience. Spans low-level system software to high-level application development environments, including infrastructure for large-scale data science and an expanded and vertically integrated software stack with advanced mathematical libs and frameworks, extreme-scale programming environments, tools, and visualization libs
- ***Goal: Hardware and Integration***
  - Partnerships between U.S. vendors and DOE's application and software developers to develop a new generation of commodity computing components. Assure at least two diverse and viable exascale computing technology pathways for the Nation to meet identified mission needs and proactively engage and integrate with DOE HPC facilities in the process.
- ***Outcome: Accelerated delivery of a capable exascale computing ecosystem*** to provide breakthrough solutions addressing our most critical challenges in scientific discovery, energy assurance, economic competitiveness, and national security
  - *Capable*: Wide range of applications can effectively use the systems developed through ECP, ensuring that both science and security needs will be addressed (affordable, usable, useful)
  - *Exascale*: Ability to perform $>10^{18}$ operations per second
  - *Ecosystem*: Not just more powerful systems, but all methods and tools needed for effective use of ECP-enabled exascale systems to be acquired by DOE labs

ECP Focus Areas

Application Development

The Application Development effort develops and enhances the predictive capability of applications critical to the DOE, including the science, energy, and national security mission space. The scope of the AD focus area includes

- targeted development of requirements-based models, algorithms, and methods,
- integration of appropriate software and hardware via co-design methodologies,
- systematic improvement of exascale system readiness and utilization, and
- demonstration and assessment of effective software integration.

Software Technology

Software Technology spans low-level operational software to high-level applications software development environments, including the software infrastructure to support large data management and data science for the DOE SC and NNSA computational science and national security activities at exascale. Projects will have:

- line of sight to application’s efforts
- inclusion of a Software Development Kit to enhance the drive for collaboration, and
- delivery of specific software products across this focus area.

Hardware and Integration

This focus area is centered on the integrated delivery of specific outcomes (ECP Key Performance Parameters, or KPPs) and products (e.g., science as enabled by applications, software, and hardware innovations) on targeted systems at leading DOE computing facilities. Areas include:

- PathForward
- Hardware Evaluation
- Application Integration at Facilities
- Software Deployment at Facilities
- Facility Resource Utilization
- Training and Productivity

ECP Work Breakdown Structure

- **Exascale Computing Project** 2.
  - **Project Management** 2.1
    - Project Planning and Management 2.1.1
    - Project Controls & Risk Management 2.1.2
    - Business Management 2.1.3
    - Procurement Management 2.1.4
    - Information Technology and Quality Management 2.1.5
    - Communications & Outreach 2.1.6
  - **Application Development** 2.2
    - Chemistry and Materials Applications 2.2.1
    - Energy Applications 2.2.2
    - Earth and Space Science Applications 2.2.3
    - Data Analytics and Optimization Applications 2.2.4
    - National Security Applications 2.2.5
    - Co-Design 2.2.6
  - **Software Technology** 2.3
    - Programming Models and Runtimes 2.3.1
    - Development Tools 2.3.2
    - Math Libraries 2.3.3
    - Data and Visualization 2.3.4
    - Software Ecosystem and Delivery 2.3.5
  - **Hardware and Integration** 2.4
    - PathForward 2.4.1
    - Hardware Evaluation 2.4.2
    - Application Integration at Facilities 2.4.3
    - Software Deployment at Facilities 2.4.4
    - Facility Resource Utilization 2.4.5
    - Training and Productivity 2.4.6

Thank you!

kothe@ornl.gov