Beyond Moore DOE Internal Lab Focused Workshop

ORNL Capabilities and Interests

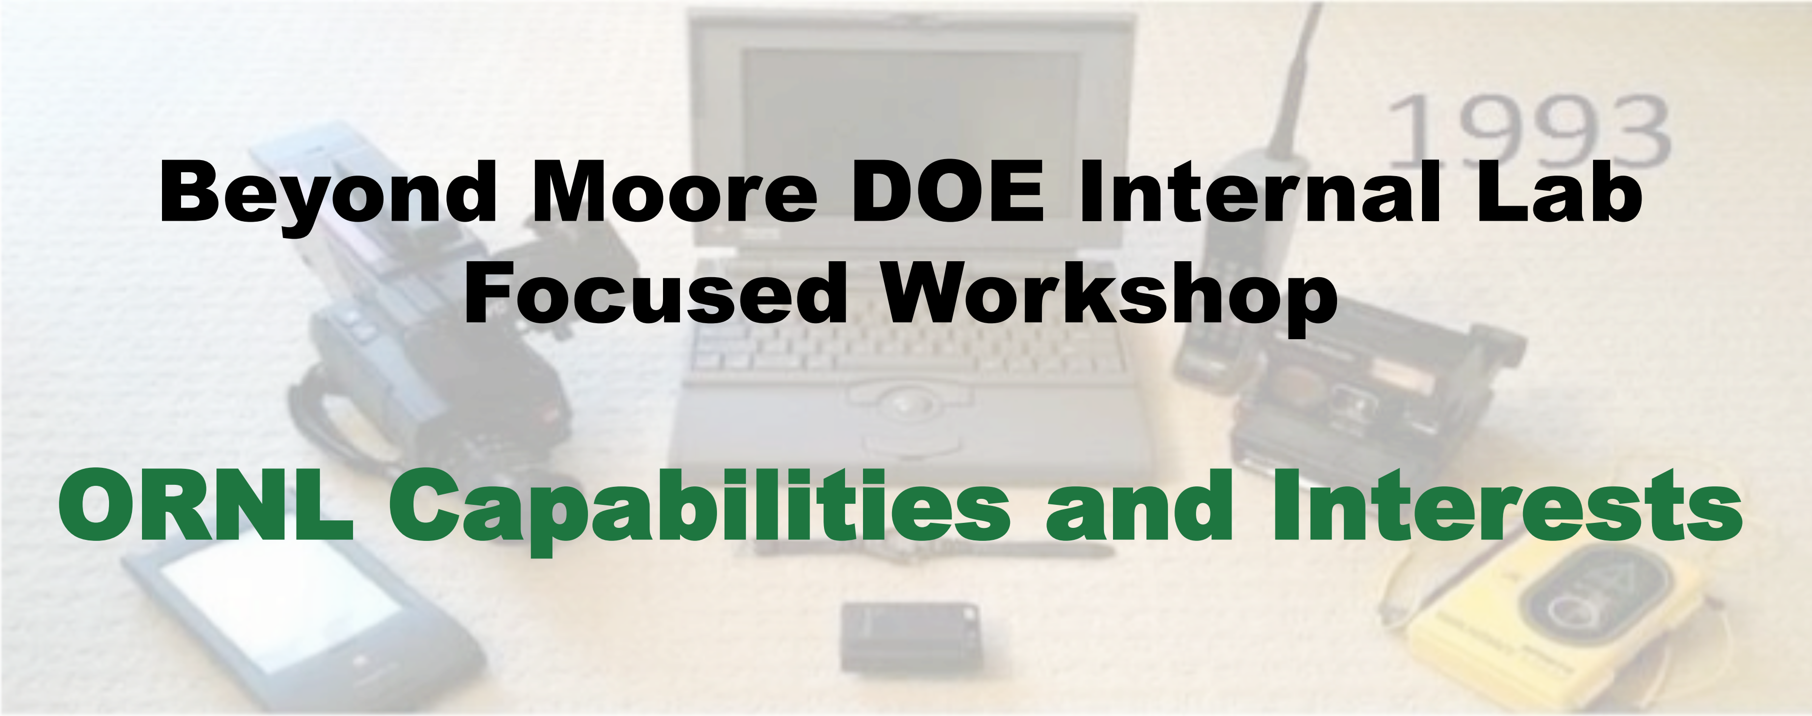

BERKELEY LAB

OAK RIDGE National Laboratory

LAWRENCE LIVERMORE NATIONAL LABORATORY

Sandia National Laboratories

Argonne NATIONAL LABORATORY

Los Alamos NATIONAL LABORATORY EST.1943

Pacific Northwest NATIONAL LABORATORY

Why go beyond exascale?

Sustained energy, national defense and economic security
US leadership in energy, defense and economic technology

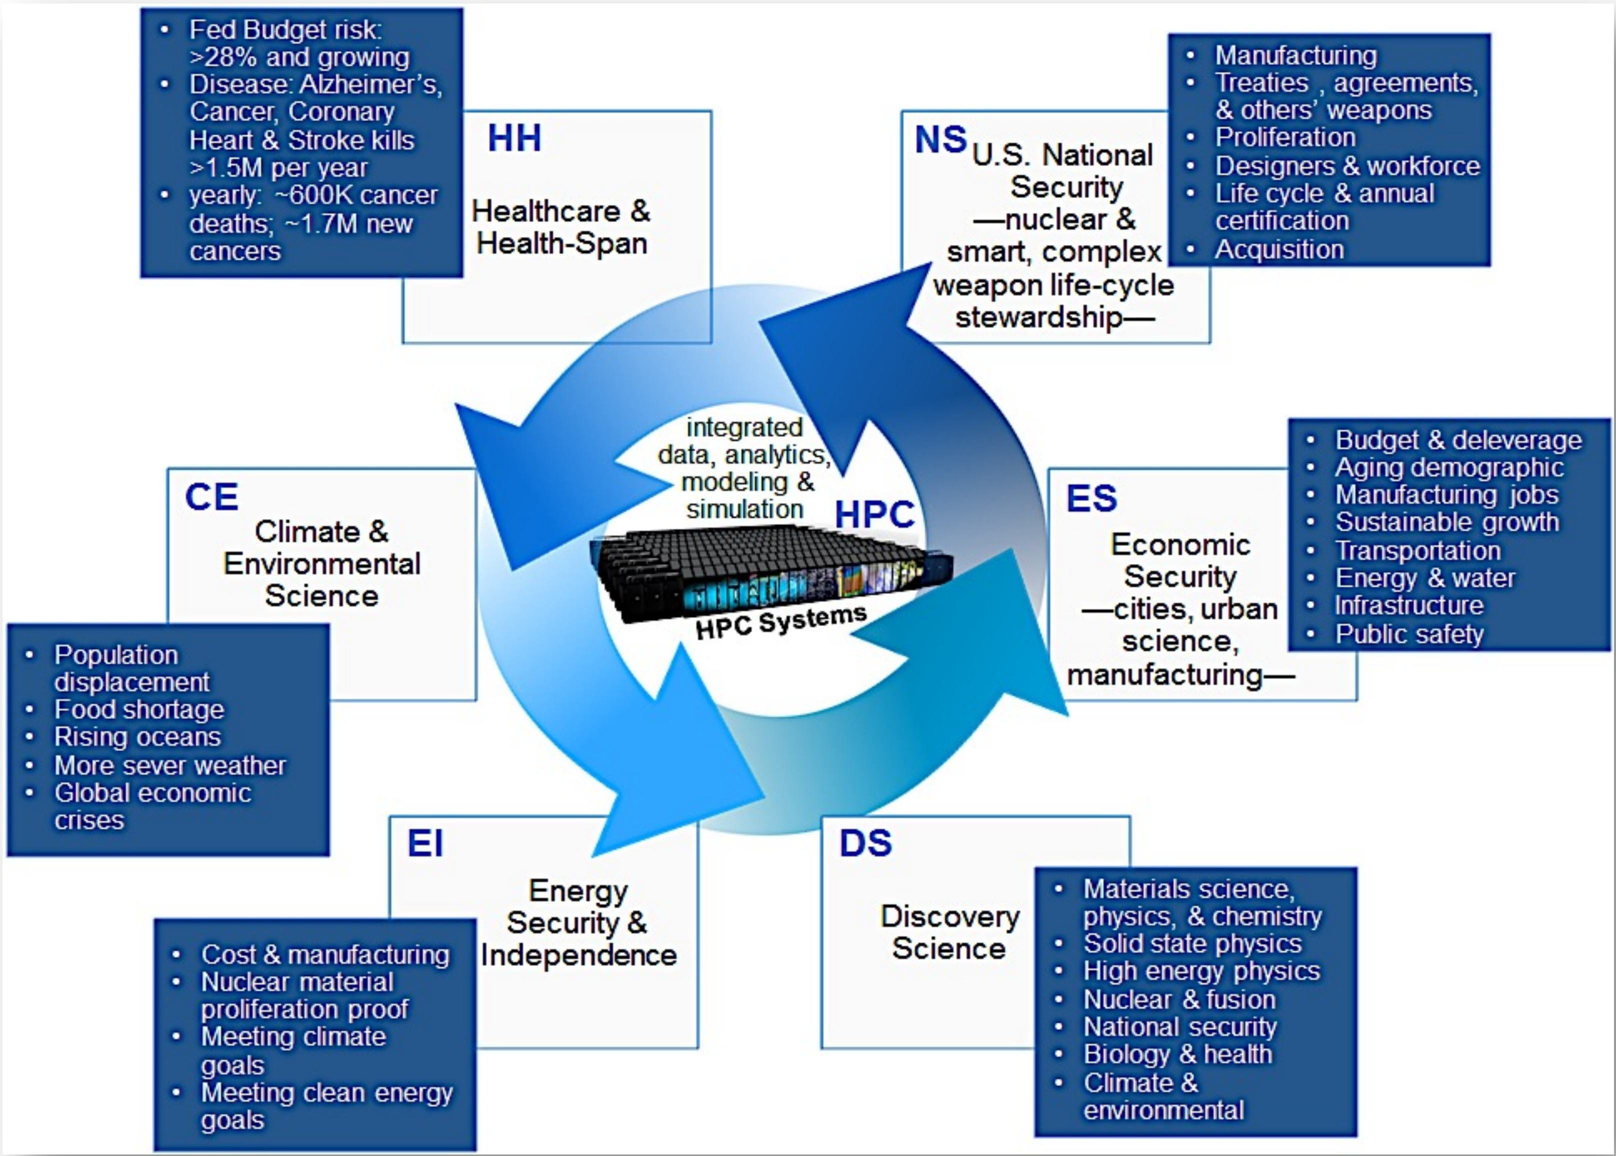

ES-mfg. -- The necessary synergy of computing and M&S to enable USA as "the" advanced manufacturing juggernaut

mfg. focus → Creating "Print" Mfg. Industry → Integrated Decision Support → Sustained U.S. Powerhouse Mfg. industry

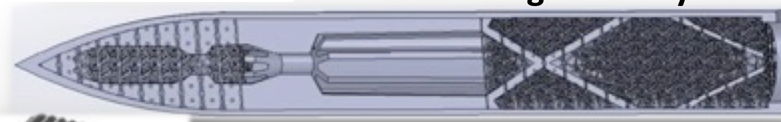

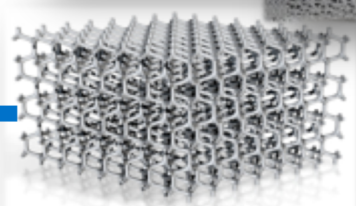

10^{24} yottascale

10^{21} zettascale

10^{18} exascale

10^{15} petascale

Predictive Part-Scale

- *initial local control of microstructure and properties*
- *data analytics and machine learning for process parameters*

Born Qualified Components

- *full local control of microstructure for desired properties*
- *physics-based shape and topology optimization of components*
- *integration of physics and data analytics for process optimization*

Qualified Integrated Devices

- *hybrid materials*
- *physics-based shape and topology optimization of integrated device*
- *in situ monitoring and real-time optimization of process parameters based on machine learning*
- *introduction & integration of robotics*

- *"Printing" Integrated Devices Industry-Wide*
- *Widespread Mfg. Automation*

Fully-automated Mission-Specific Systems & Devices

- *automated shape and topology optimization*
- *in situ monitoring and real-time optimization of process parameters based on integrated physics and machine learning*
- *fully hybrid materials*
- *Fully robotic & mobile mfg. systems*

- *"Printing" Systems & Devices Industry-Wide*
- *Jobs and Manufacturing Automation Juggernaut*

time

Beyond Moore Co-Design Framework

Modeling | **Experimental**

10,000x improvement:
20 fJ per instruction equivalent

Systems — Devices — Materials

Alg & SW Environments

Application Performance Modeling

Computer System Architecture Modeling
- Next generation of Structural Simulation Toolkit
- Heterogeneous systems HPC models

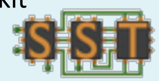

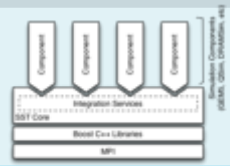

HW & Circuit Archs

Component Level Models
- Gem5, McPAT, HiHAT, CACTI, NVSIM

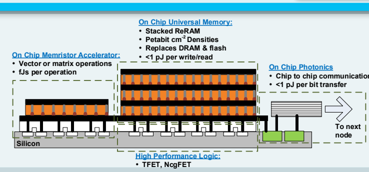

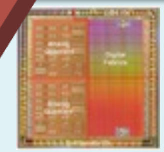

Component Fabrication
- Processors, ASICs
- Photonics
- Memory

Circuit/IP Block Design and Modeling
- SPICE/Xyce model

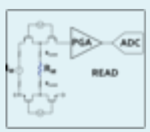

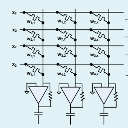

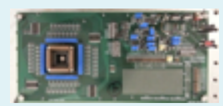

Test Circuit Fab and Measurement
- Sub-circuit measurement

Common & logic devices

Compact Device Models
- Single device electrical models
- Variability and corner models

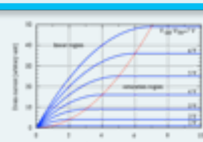

Device Measurements
- Single device electrical behavior
- Parametric variability

Device Physics Modeling
- Device physics modeling (TCAD)
- Electron transport, ion transport
- Magnetic properties

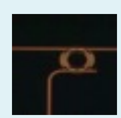

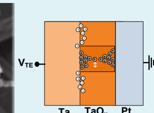

Device Structure Integration and Demonstration
- Novel device structure demonstration

Materials

Process Module Modeling
- Diffusion, etch, implant simulation
- EUV and novel lithography models

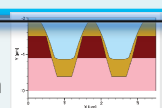

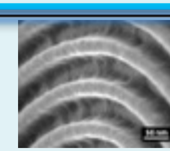

Process Module Demonstrations
- EUV and novel lithography
- Diffusion, etch, implant simulation

Atomistic and Ab-Initio Modeling
- DFT – VASP, etc.
- Many-body – QMCPACK
- MD – LAMMPS

Example modeling and characterization activities within an MSCD framework

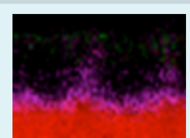

Fundamental Materials Science
- X-ray and neutron
- Microscopy: TEM, SEM, EELS
- Scanning probe technique

DOE User Facilities

CNMS provides capabilities & expertise for materials formation and function

About CNMS:

- **Unlike many user facilities, you don't need to have samples to apply for time**
- **Two calls per year for continuous access; anytime for short-term projects**
- **Simple 2-page proposal**
- **Free access to laboratories, equipment and expertise if you agree to publish**
- **Proposal deadlines: early May and mid-October**
- **Joint proposals with neutron sources (SNS, HFIR)**

Research areas:

- **Synthesis –** 2D, precision synthesis, selective deuteration
- **Nanofabrication –** direct-write, microfluidics, cleanroom
- **Advanced Microscopy –** AFM, STM, aberration-corrected TEM/STEM, atom-probe tomography
- **Functional Characterization –** laser spectroscopy, transport, magnetism, electromechanics
- **Theory and Modelling –** including gateway to leadership-class high performance computing

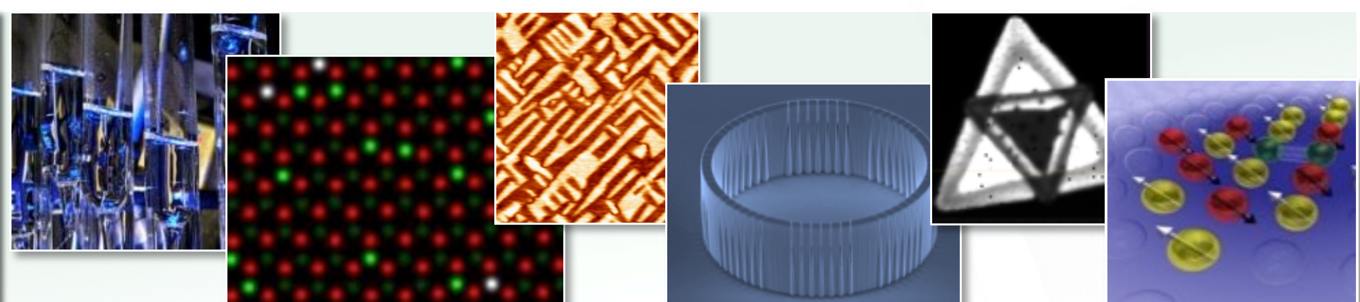

CNMS is a Nanoscale Science Research Center supported by the U.S. Department of Energy, Office of Science, Scientific User Facilities Division

nsrcportal.sandia.gov

cnms.ornl.gov

Oak Ridge National Laboratory

Synthesis and functional assembly

High purity materials are essential building blocks

Functional Hybrid Nanostructures

- Pulsed Laser Deposition – films
- CVD – 2D TMCs, graphene, 1D nanotubes, perovskites
- *Electrical characterization*
- *Optical characterization*

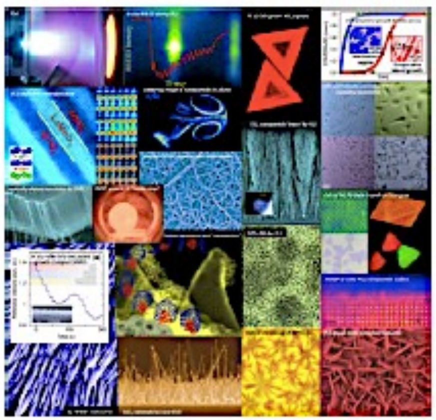

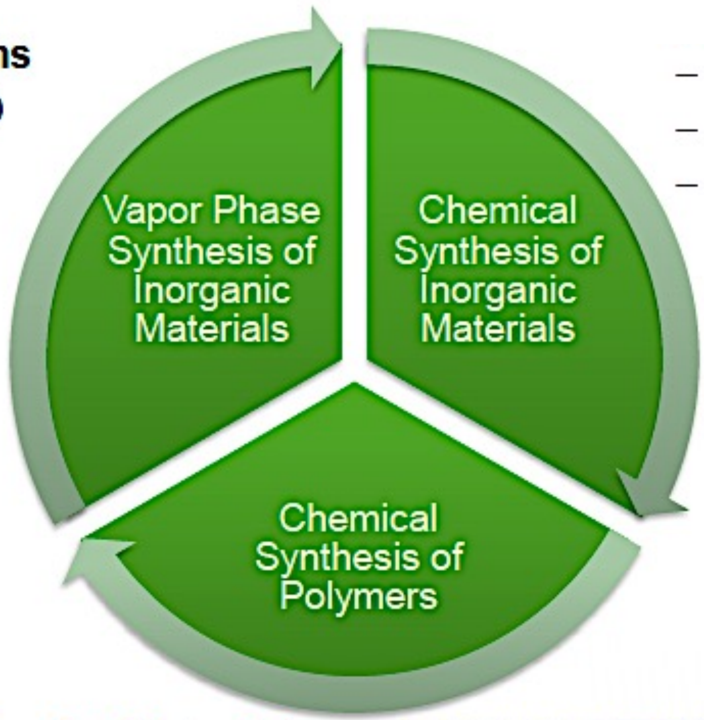

Chemical Functionality

- Hydrothermal & chemical
- Nanoparticles for catalysis
- *BET, MS, SEM, TEM, XRD, HIM*

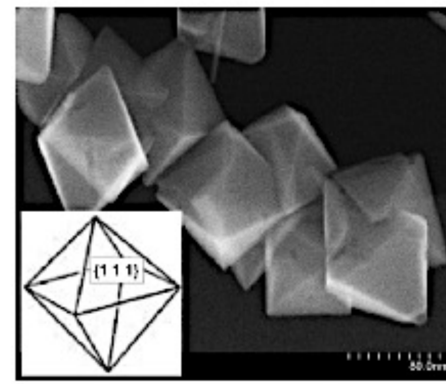

Macromolecular Materials

- Controlled polymerizations
- Deuteration for neutron scattering
- *NMR, Light scattering, MALDI*

Neutrons (SNS/HIFR) offer unique capabilities

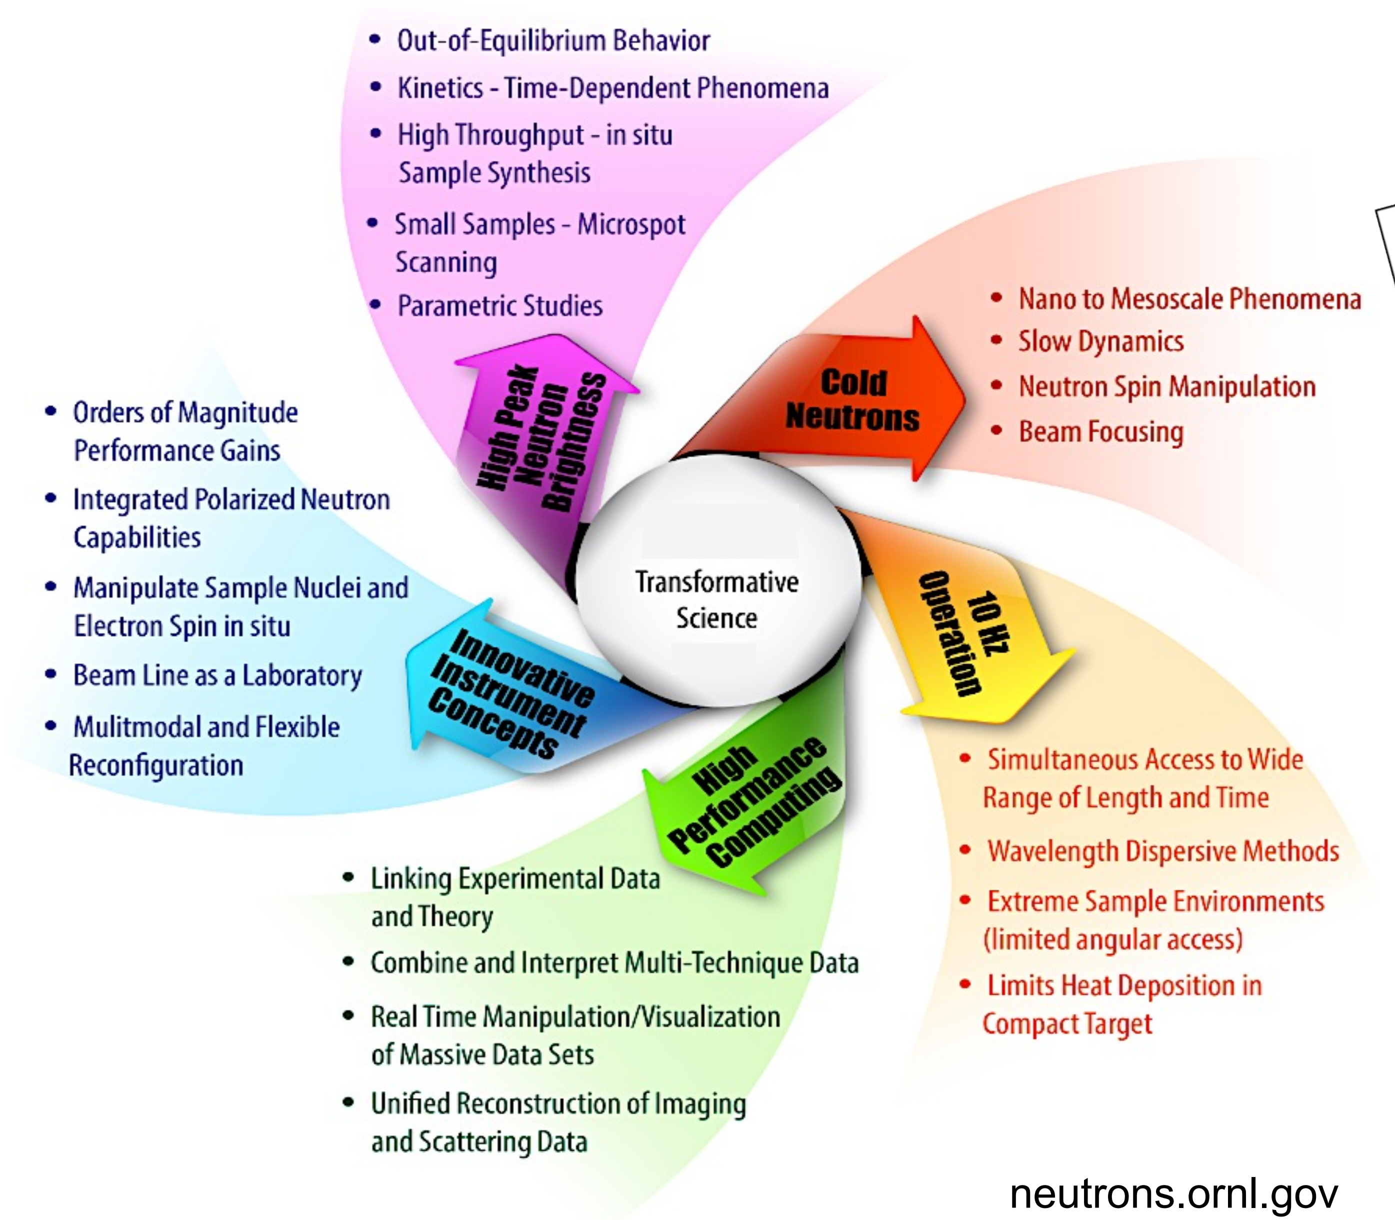

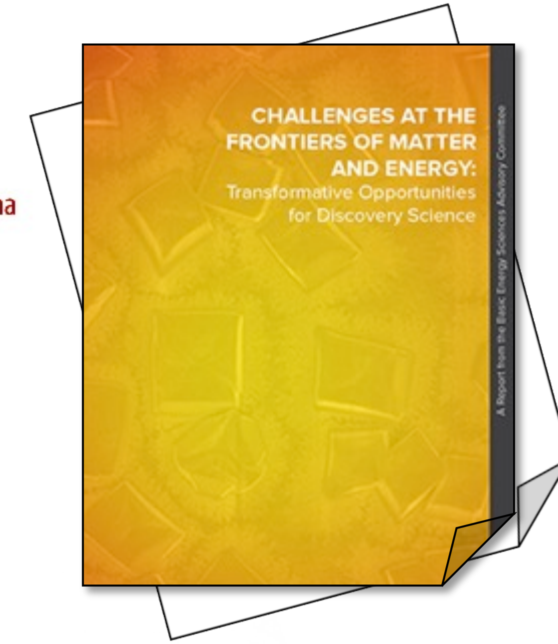

neutrons.ornl.gov

Neutrons (SNS/HIFR) offer unique capabilities

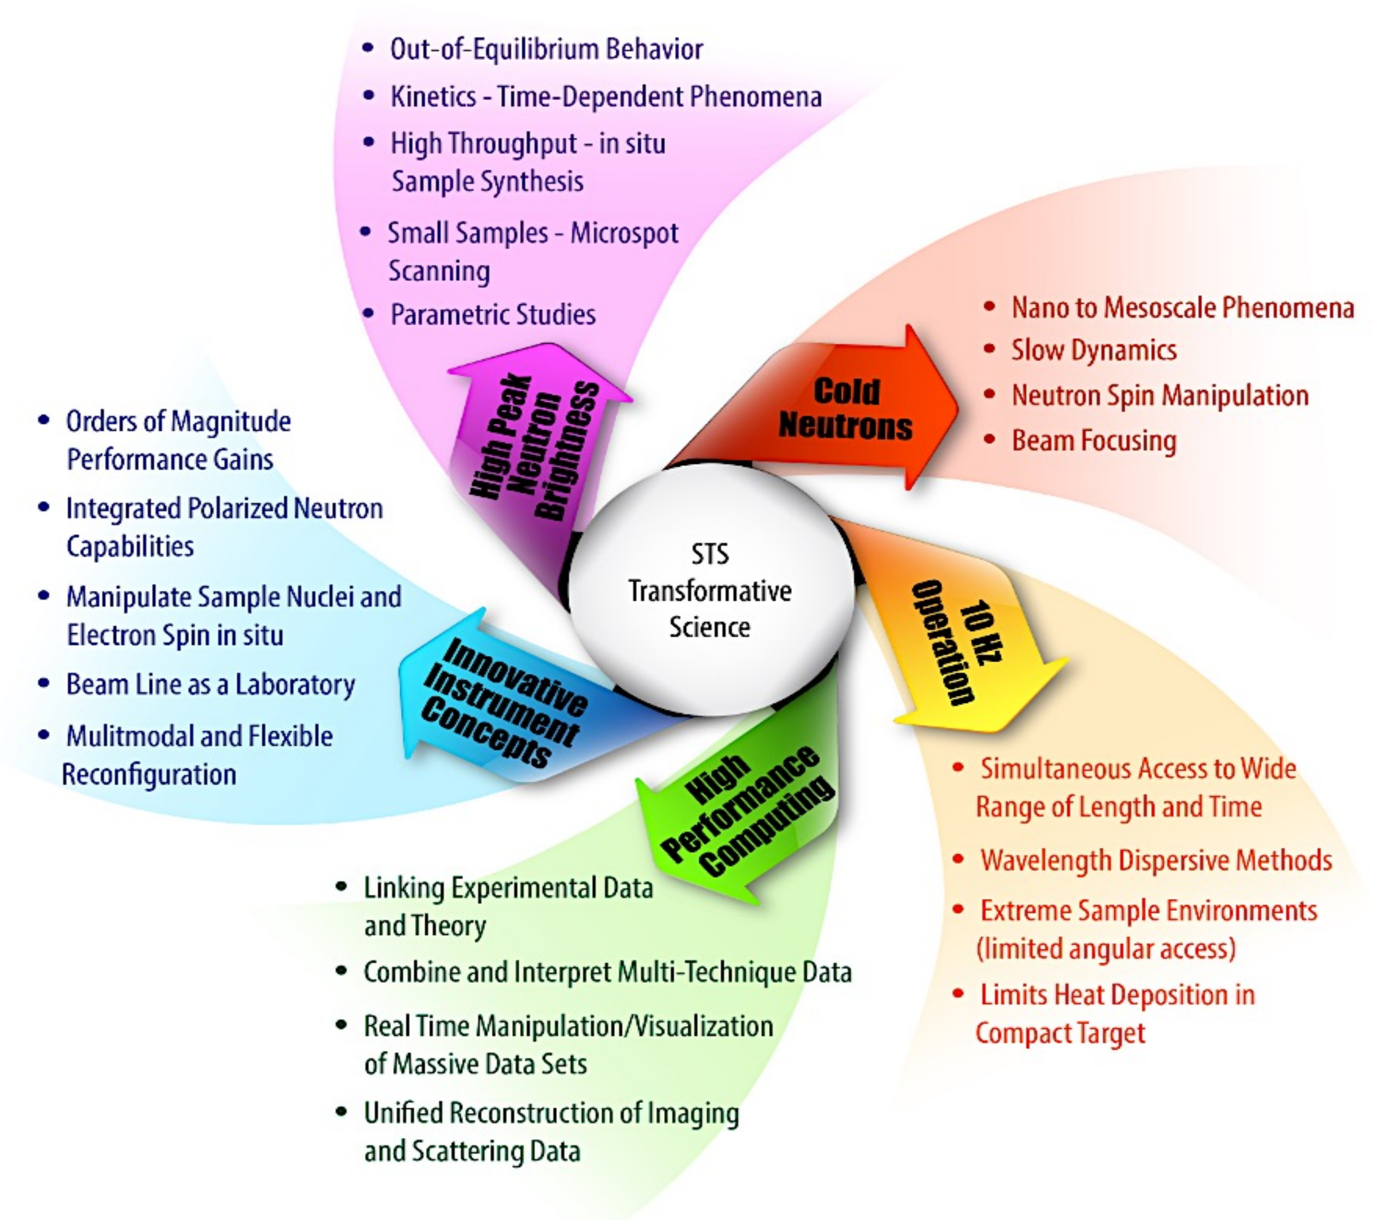

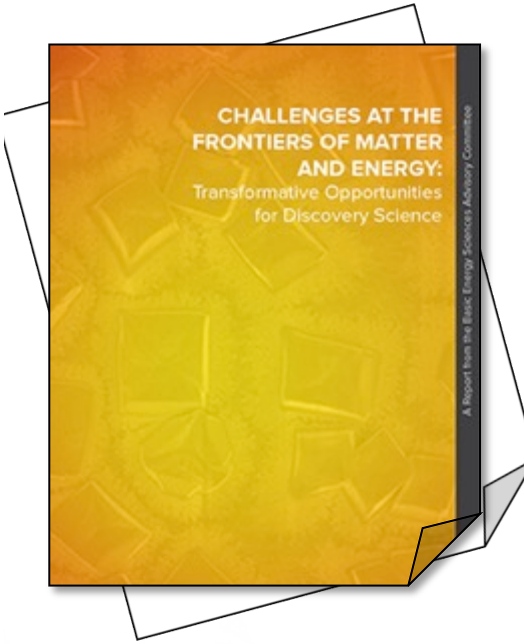

OAK RIDGE National Laboratory

DOE's Office of Science Computation User Facilities

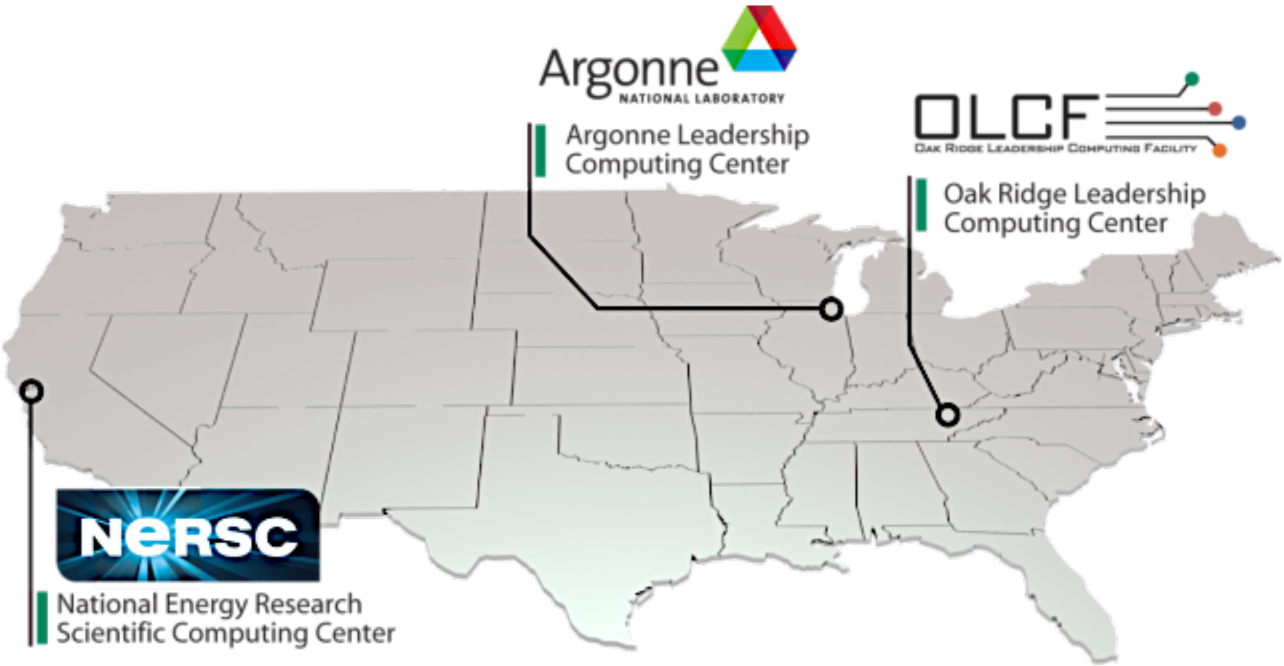

- DOE is leader in open High-Performance Computing
- Provide the world's most powerful computational tools for open science
- Access is free to researchers who publish
- Boost US competitiveness

NERSC
Edison is 2.57 PF

ALCF
Mira is 10 PF

OLCF
Titan is 27 PF

Multi-scale materials theory, modeling and simulation

Developing a fundamental understanding of physical and chemical properties to accelerate discovery and design of responsive materials

- Scalable methodologies for theoretical and computational materials and chemical science to establish new capabilities and to enhance links with experiment
- Understanding and rational tuning of transport (electron, phonon, spin, ion, molecule), reactivity and electronic structure
- Building the scientific foundation to study and design strongly correlated electronic materials (such as superconductors)
- Advancing theory and simulation for understanding multi-scale chemical and materials properties & processes to guide design of matter with atomic, molecular and mesoscale fidelity

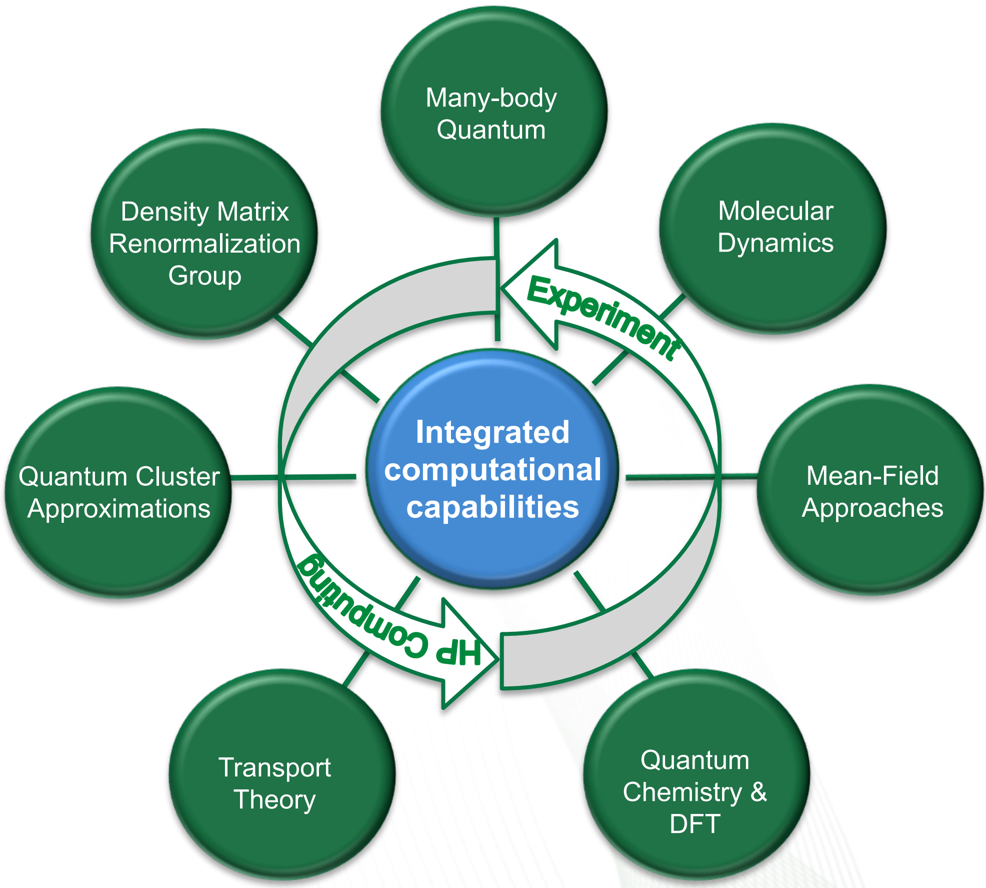

Atoms: what they do and how it gives rise to materials properties

Our scientific paradigm is shifting

Traditional approach

- Synthesis
- Characterization
- Theory
- Computation

Transitioning to include

- Data analysis
- Correlative functional imaging
- Local theory-experiment matching of multi-dimensional, multi-modal spatially and temporally resolved information

Trustworthy frameworks for multi-scale matter

- Bridge experimental imaging and scattering with theory/simulation via math, big data and data analytics to design new materials
- Leverage unique strengths in
  - Physics and chemistry on the atomic scale
  - Mesoscale structure and functional probing
  - **Big-deep-smart math & data integrated with predictive theories: Transformative with ultascale computation**

Concurrent neutron scattering experiments and first-principles modeling

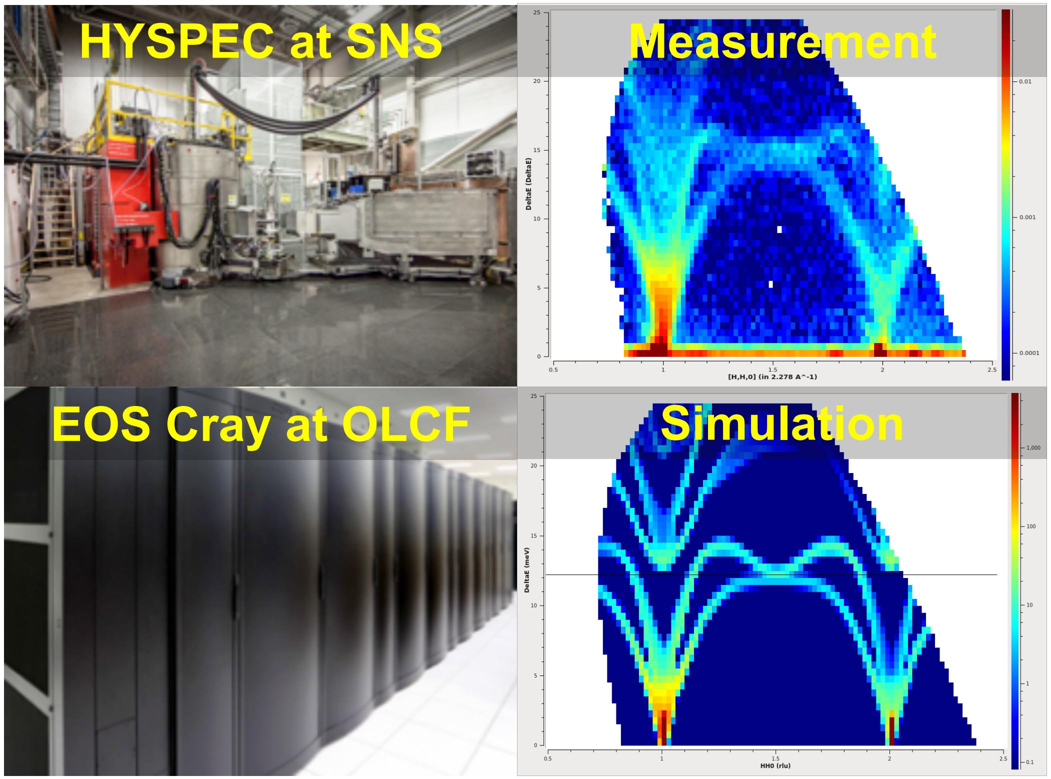

O. Delaire, J. Hong, H. Cao, A. Savici, B. Winn, L. Boatner, G. Shipman

MANTID

- Study of ferroelectric instabilities in $SrTiO_3$
- Measurements on HYSPEC using live data streaming (top)
- Dedicated access to Cray XC30 EOS cluster at Oak Ridge Leadership Computing Facility (11,000 cores)
- Full scale *ab-inito* molecular dynamics simulations on experiment timescale allowing real time decisions

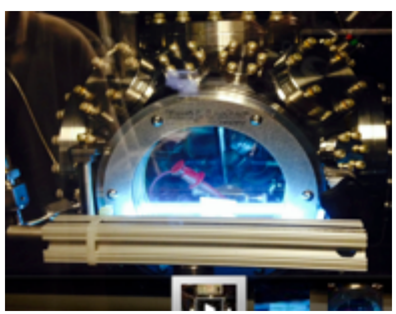

- Environmental chamber(s) for the *in-situ* study of structure and function of materials at different temperatures, pressures, humidity, composition of vapors and solvents

Key science questions link basic and applied materials research

Materials at extremes

Can materials be designed for superior performance at extremes of environments and properties?

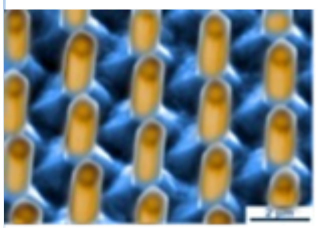

Interfacial Processes

Can molecular processes be predicted and controlled over broad length and time scales to provide optimized functionalities for separations, transport, chemical transportation?

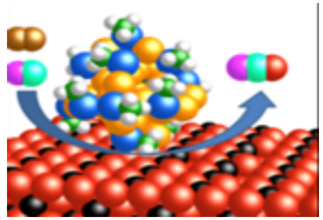

Multiscale phenomena

How do nanoscale effects emerge from the properties of individual atoms and electrons and how do macroscopic properties of assemblies and materials emerge from these nanoscale phenomena?

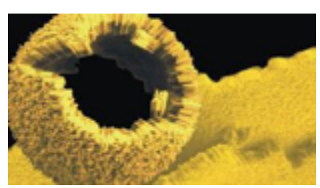

Soft matter

Can we develop a predictive understanding of the structure, dynamics and function of soft matter systems to enable the design and synthesis of new functional soft materials?

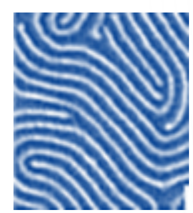

Directing matter

Can we direct and assemble matter in 3D with atomic precision?

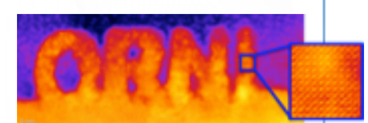

Directing Matter

Fundamental issues

- Atomic-scale control of materials assembly enable designer hierarchical structures
- As predictive materials sciences advance, approaches to create (desired) structures are often lacking
- Approaches at multiple length scales must work in synergy
- Fabrication and manipulation are intrinsically intertwined

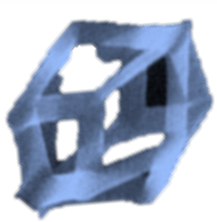

ORNL assets

- Expertise in materials synthesis, functional imaging (e.g., *in-situ* monitoring of synthesis/ assembly), multiscale modeling
- Capabilities to direct matter at multiple length scales:
  - Manipulation of individual atoms by electron beams and scanning probes
  - Precision synthesis of macromolecules
  - Fabrication at sub-10nm scales and above
  - Additive manufacturing of macroscopic structures

Application Areas

- Materials and structures for energy generation, storage, use
- Computing, including quantum and neuromorphic computing
- Multifunctional materials: smart materials combining mechanical and electronic properties for actuators; combined energy harvesting and storage, “beyond Moore” electronics

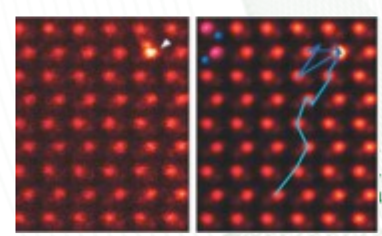

Nanoscale Additive Manufacturing

- Focused electron beam deposition has long been used to create comparatively simple structures
- The implementation of predictive calculations based on Monte Carlo and continuum simulations guides the formation of free-standing nanostructures
- He-ion beams offer the promise of further reducing the achievable feature sizes

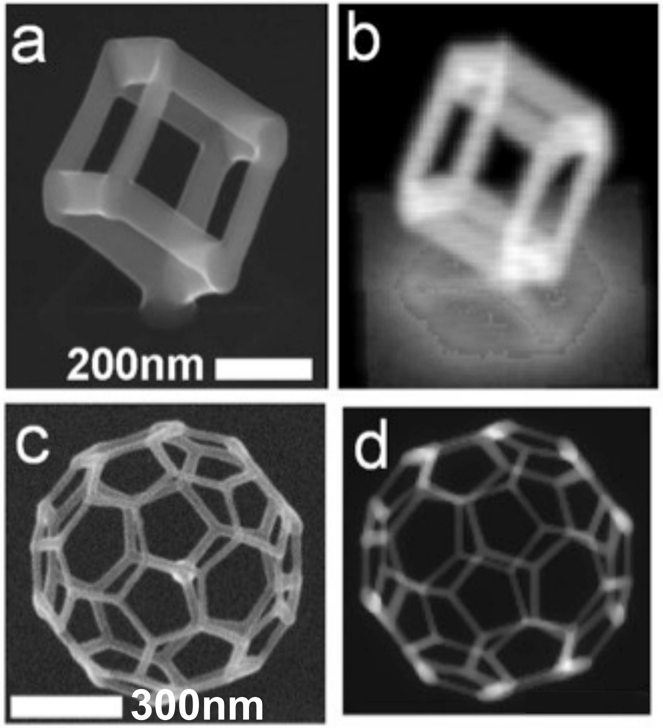

Actual (left) and predicted (right) 3D PtC structures created by focused electron beam deposition

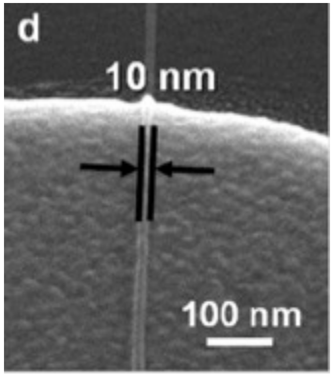

He ion-beam-induced single cobalt line

J. D. Fowlkes, et al. *ACS Nano* (2016). DOI: 10.1021/acsnano.6b02108

H. Wu et al., *J. Mater. Sci.: Mater. Electron* (2014)

Nanosculpting and atomistic repair

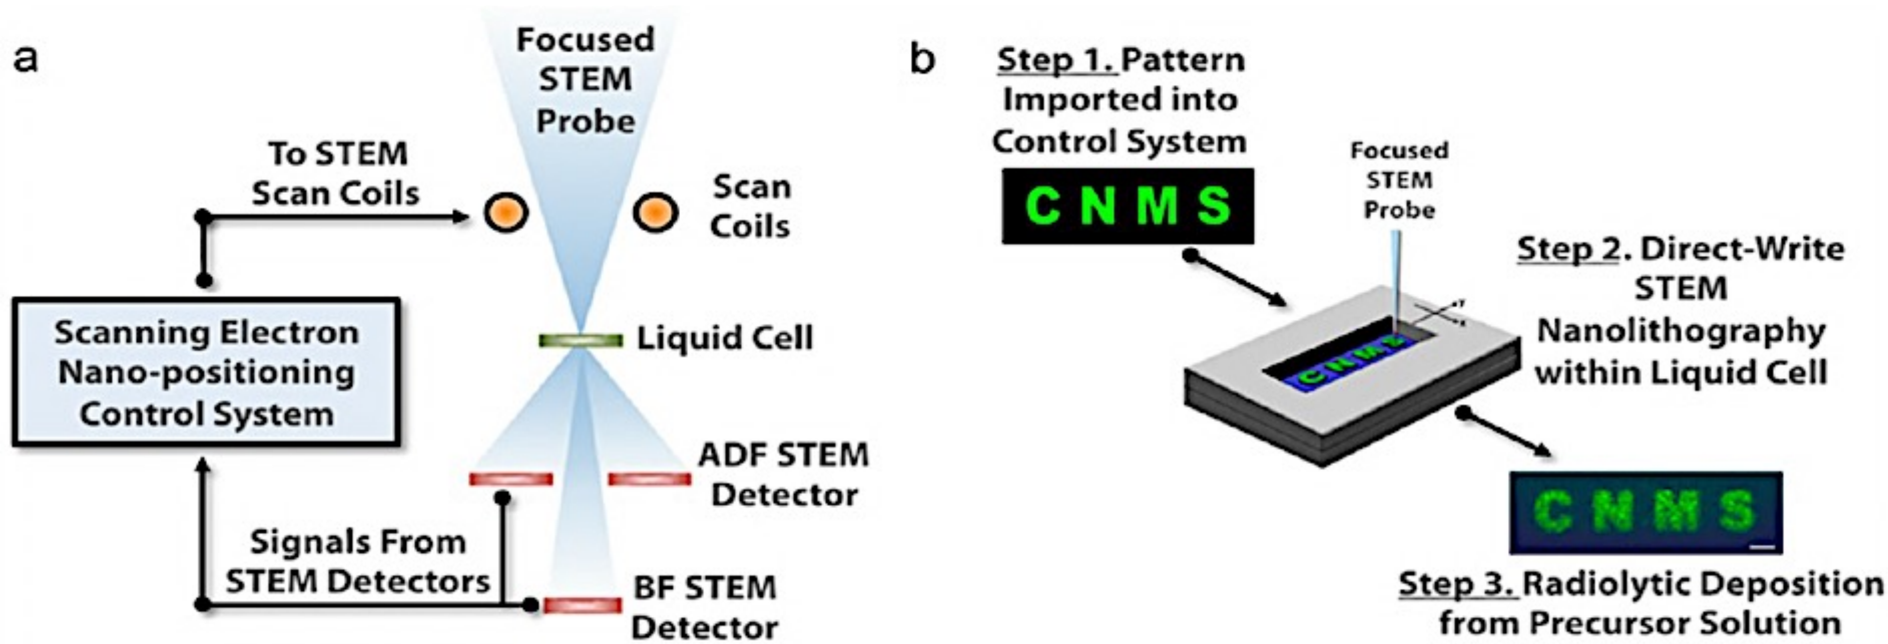

Beam directed transformation

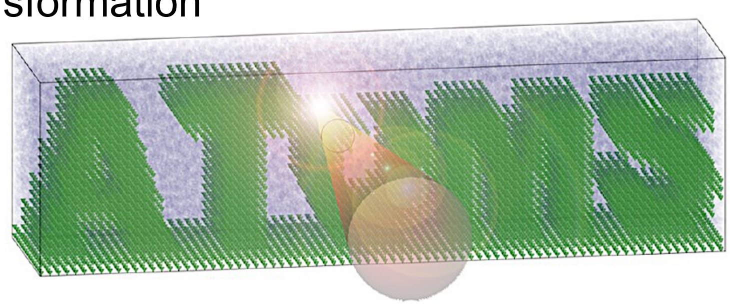

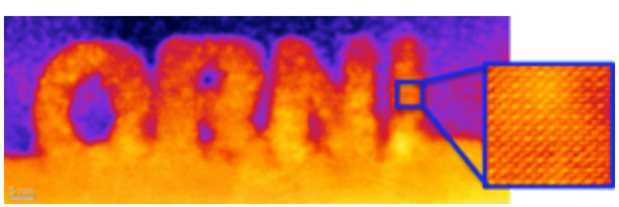

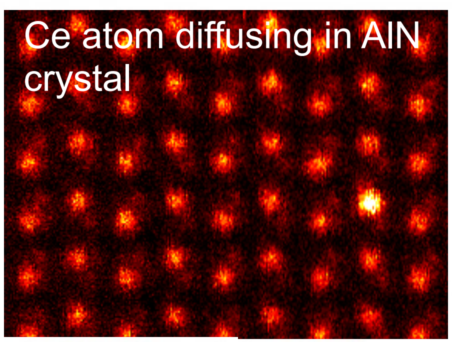

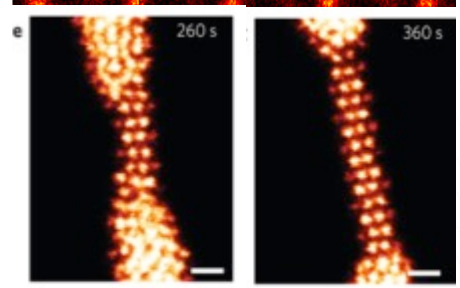

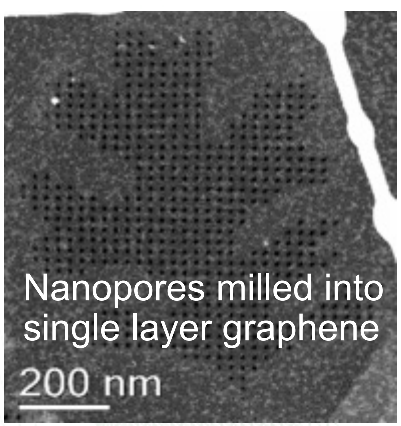

Jacob Swett, Lockheed Martin

Control beyond intrinsic vacancy levels

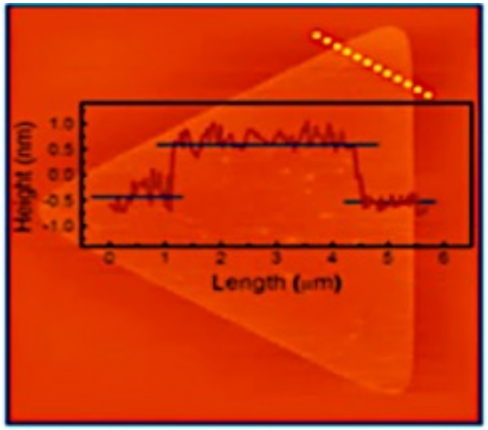

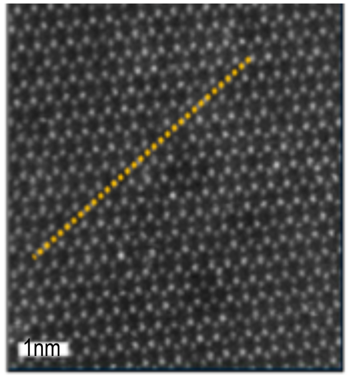

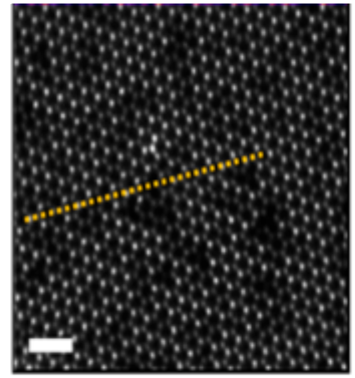

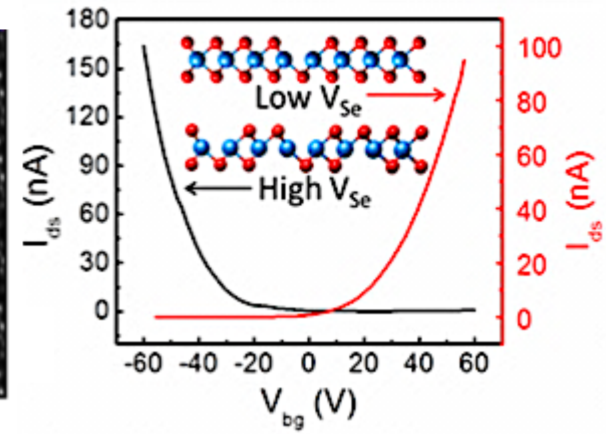

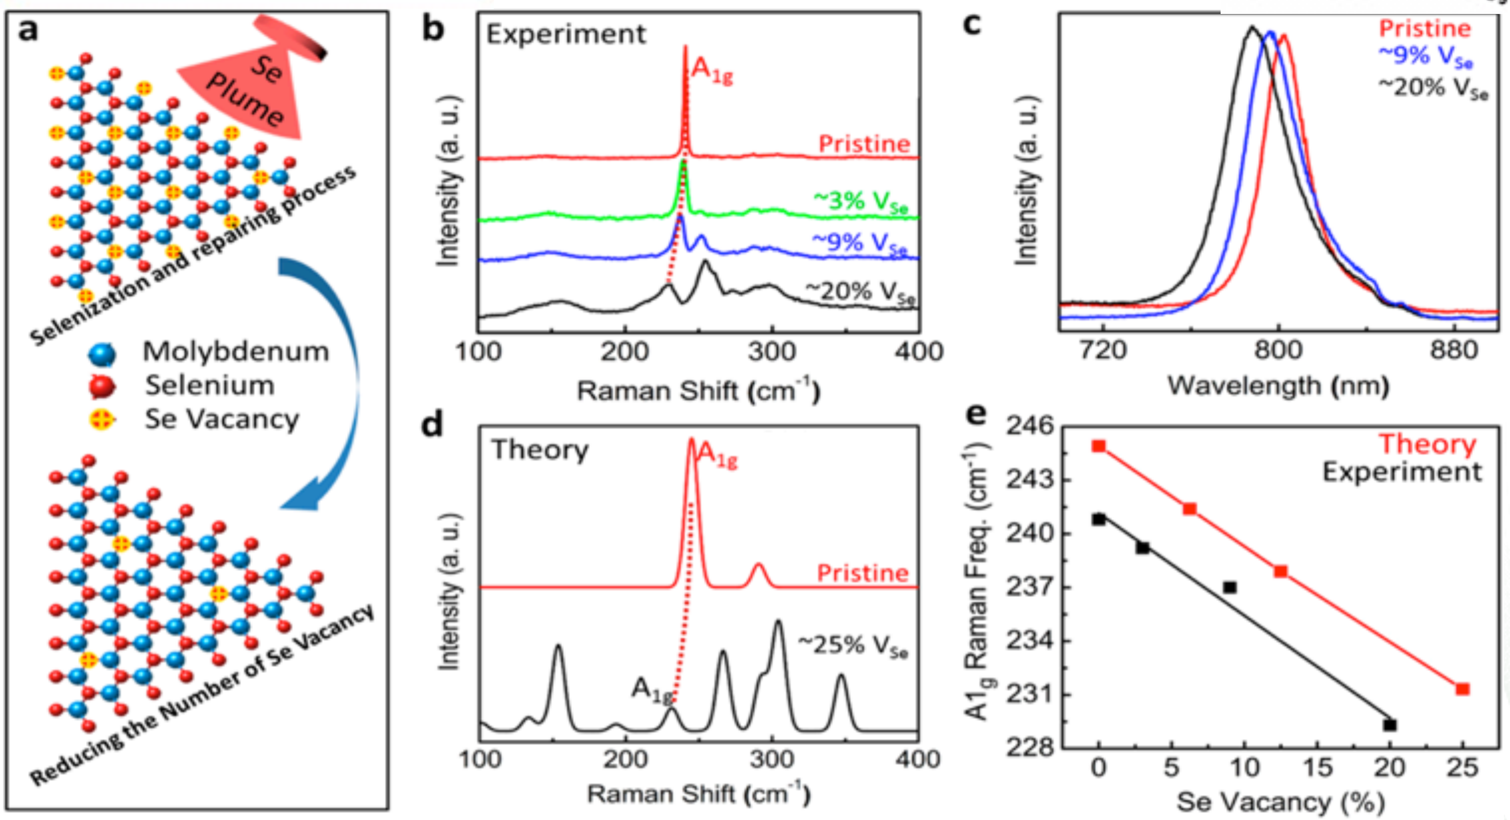

“Nano-sandblasting” of layered ferroelectric semiconductors

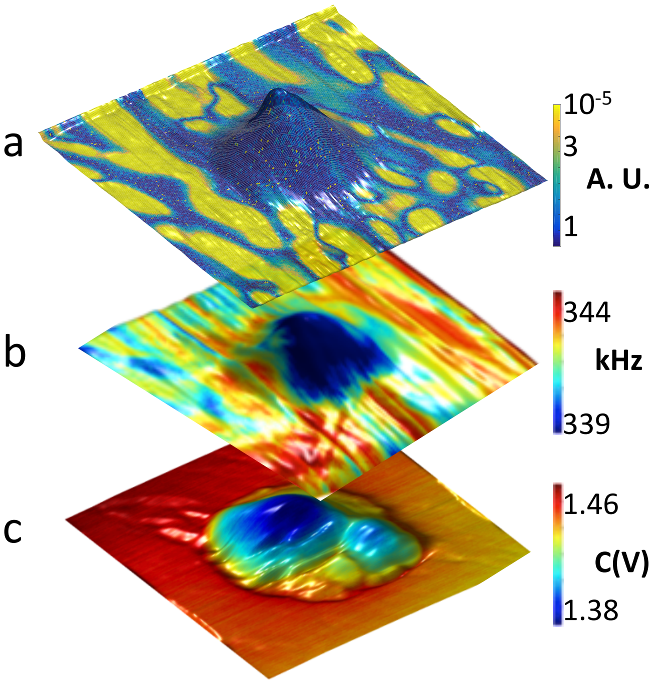

Effect of 8×10^{14} ions/cm^2 He ion beam dose on mechanical and electrical properties of $CuInP_2S_6$ painted over 3D topo image **(a)** Piezoresponce Force Microscopy. **(b)** Blue Drive contact Band Excitation measurement of resonant frequency. **(c)** Scanning Microwave Microscopy image.

- **Layered materials are broadly perceived as enabling component for the ultimate scaling down in electronics. However, critical is the structuring and functional tuning of these materials**
  - **Developed a pathway to direct-write circuitry in copper indium thiophosphate (CITP), a ferroelectric semiconductor; controlled conductivity, altered built-in potential, and modifying chemical composition**

Research Details

- He-ion microscopy was used on CITP to develop nanostructures, control the built-in material potential, and enhance conductivity.
- The size and functionality of the grown structures could be controlled by tuning He-ion exposure.

Work was performed at the Center for Nanophase Materials Sciences, ORNL.

A. Belianinov, et al. *ACS Appl. Mater. Interfaces* (2016). DOI: 10.1021/acsami.5b12056

Other nanofab capabilities
CNMS Nanofab Research Lab (NRL)

Advanced Lithography and Direct Write Nano
Electron Beam Lithography, Mask Writing, Electron and Ion Beam Processing, Two Photon Lithography

The NRL is housed in over 10,000 sq.ft. of clean room space and bio-affiliate labs to support the integration of nanoscale materials into multi-scale systems.

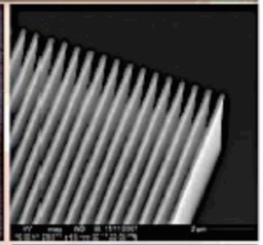

Thin Film Deposition and Characterization
High Temp Growth, LPCVD, PECVD, Atomic Layer Deposition, Rapid Thermal Processing, Profilometry, Stress Measurement, Atomic Force Microscopy

Material Etching
Fluorine and Chlorine chemistry RIE, Wet Chemical Processing, Beam Assisted Etching

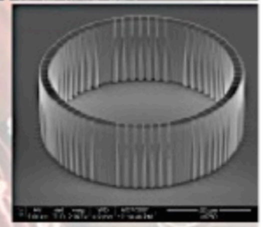

Soft Material Interfaces and Imaging
Bio-affiliate Labs, Microfluidics, Embossing, Stamping, and Soft Lithography

Nanostructure Synthesis
High aspect ratio silicon nanostructures, Carbon Nanofibers, Asymmetric Nanotextures, Nanoparticles from Dewetting

The NRL is supported by technical, operations, and scientific staff that drive process development, training, and maintenance to support in-house and user science

2. Nanofabrication Research Laboratory

OAK RIDGE National Laboratory | CENTER FOR NANOPHASE MATERIALS SCIENCES

Other unique capabilities for materials
Scanning tunneling microscopy(spectroscopy)

- Development of new STM(S) and 4-probe STM(S) based imaging and spectroscopy modes reveal local electronic, magnetic, and transport properties
  - tunneling thermopower microscopy
  - scanning tunneling potentiometry
  - spin-polarized STM, for spatial mapping of local crystallographic, electronic, magnetic, and transport properties

STM image and spectroscopy revealing confined electronic states at Gr/h-BN heterojunctions

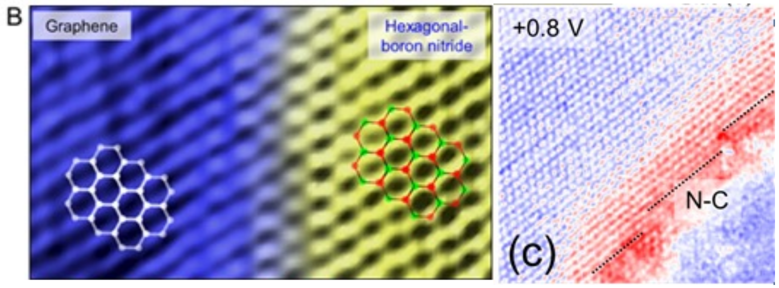

J. Park et al, *Science* (2014); *Nature Comm.*(2014)

Scanning tunneling potentiometry to map conductivity across grain boundaries

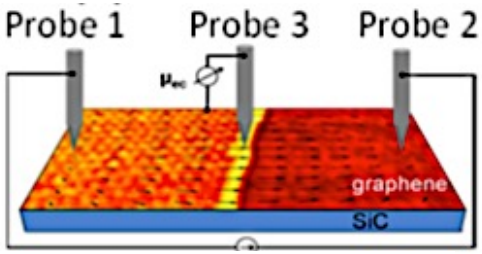

K.W. Clark et al, *Nano Lett*. (2013); *PRX* (2014)

SP-STM reveals spin polarization of Co/Cu

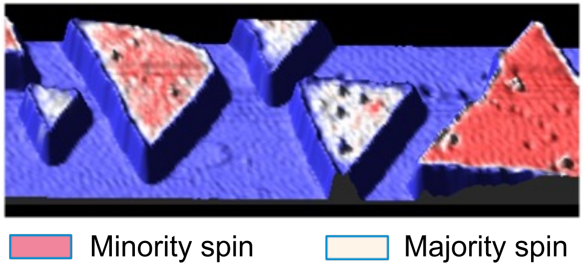

Force-based scanning probe microscopies

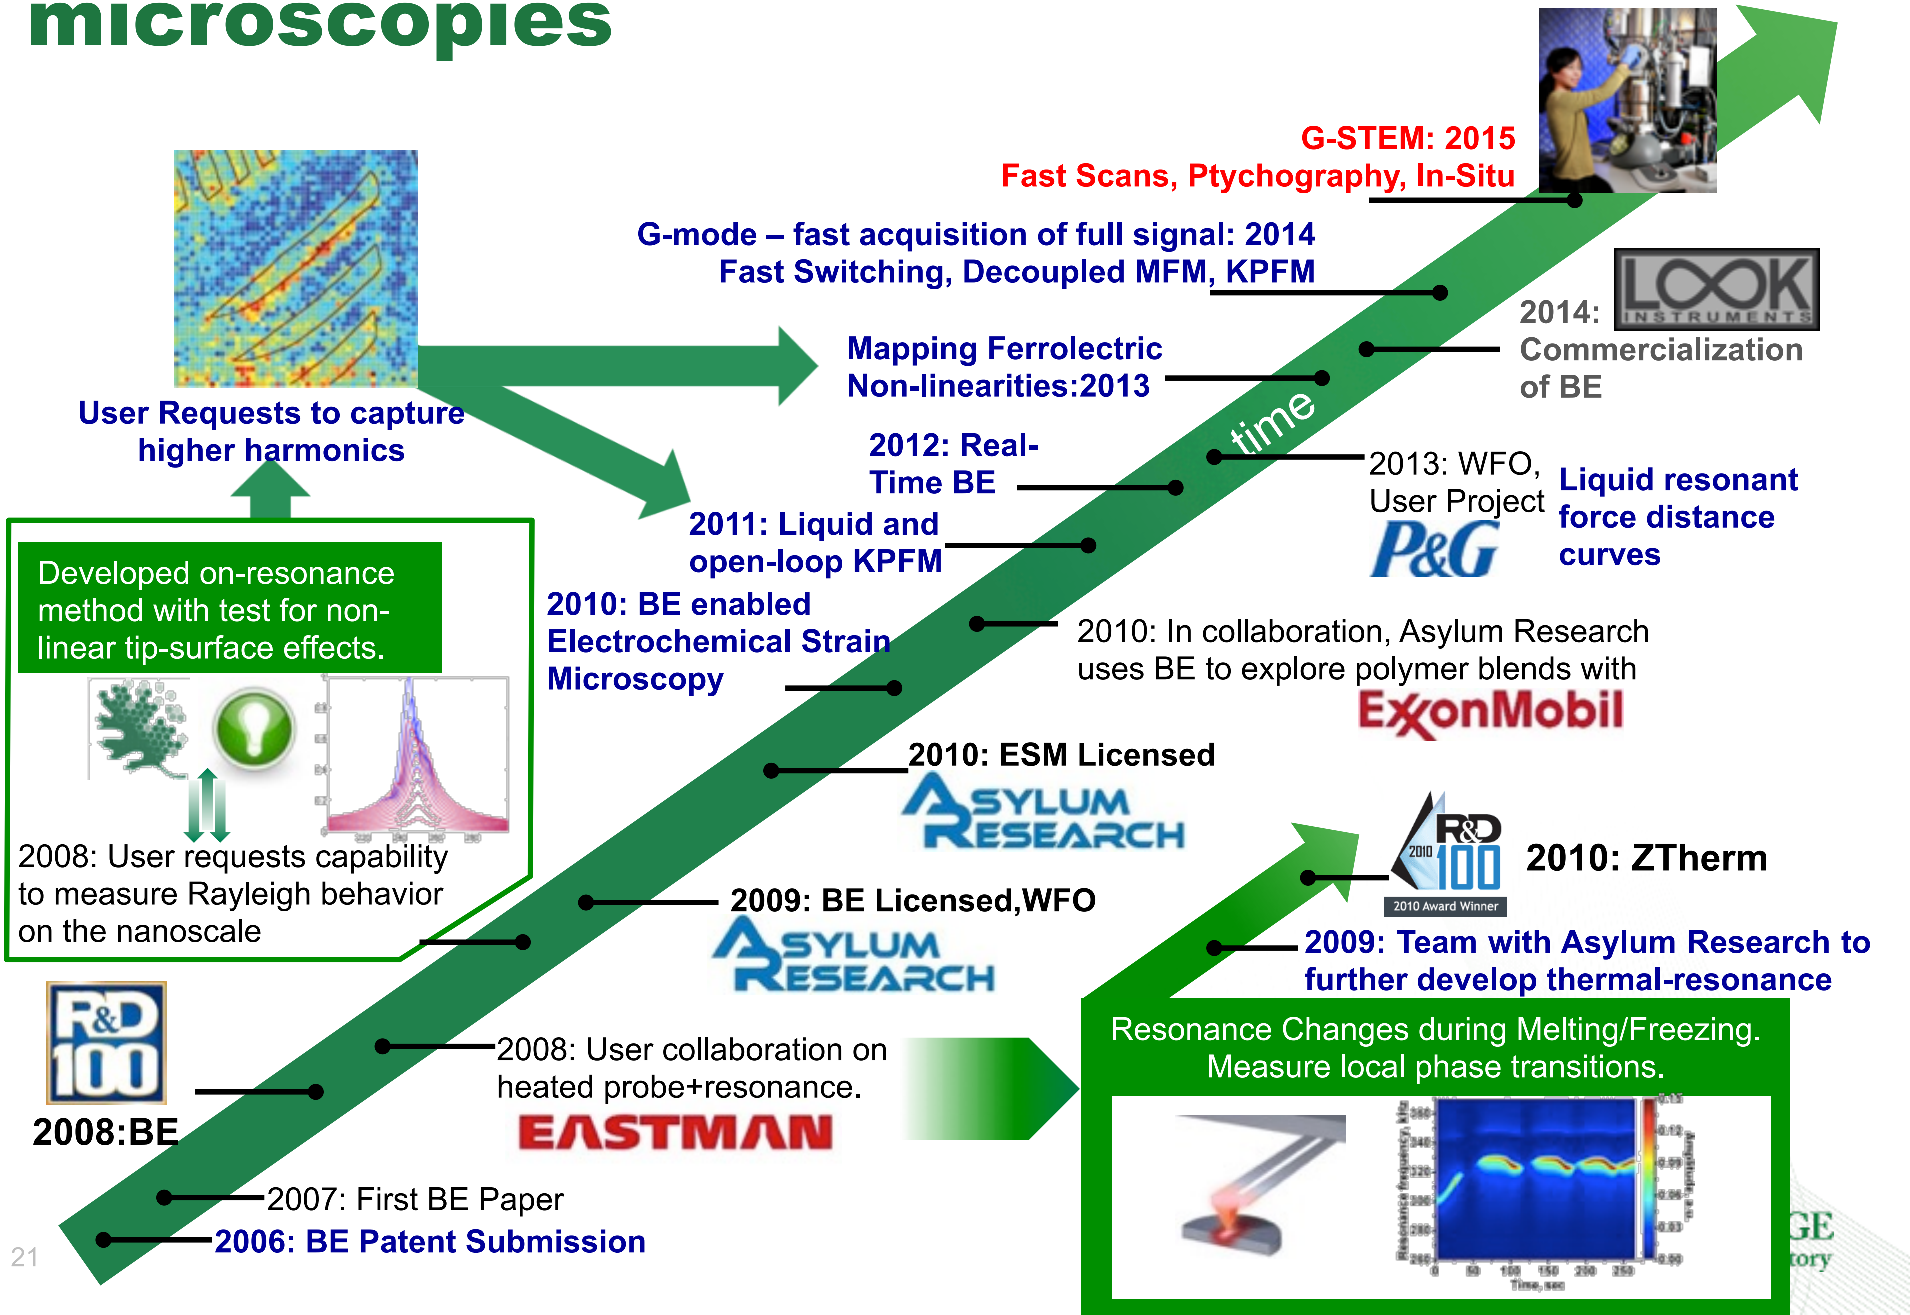

Imaging Capabilities

Former ShaRE facility (now CNMS)

– Operated as Electron Beam Microcharacterization Center since 1977 and was merged with CNMS in 2014
– Focus on Scanning Transmission Electron Microscopy and Atom Probe Tomography
– Moving toward aberrated electron probes

Example: developing high-fidelity interfaces through application of aberration-corrected Z-contrast STEM imaging, electron energy loss spectroscopy, and atom probe tomography to provide the single-atom-level understanding of defects, interfacial steps/terraces, chemistry, composition, and structural thermal stability

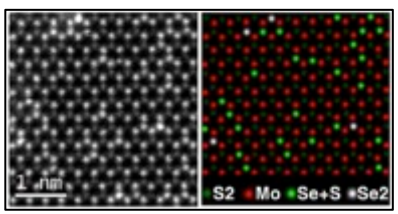

Nion UltraSTEM

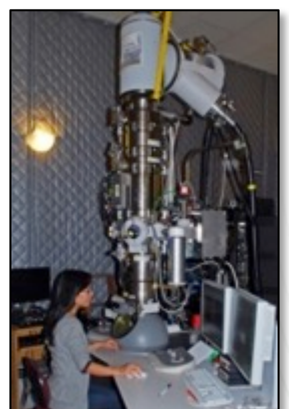

FEI Titan

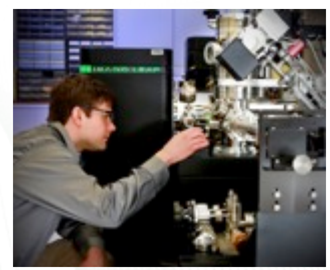

Cameca LEAP

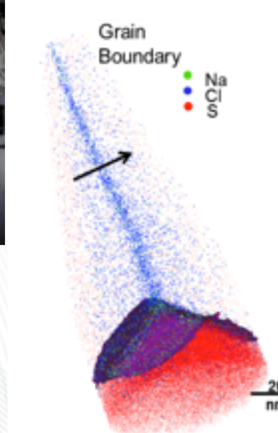

Aberration-corrected Scanning Transmission Electron Microscopy (AC-STEM) and Electron Energy Loss Spectroscopy (EELS)

Z-contrast STEM and image analysis provides an atomic-scale map of individual dopants

Y. Gong *et al.*, *Nano Letters* (2014)

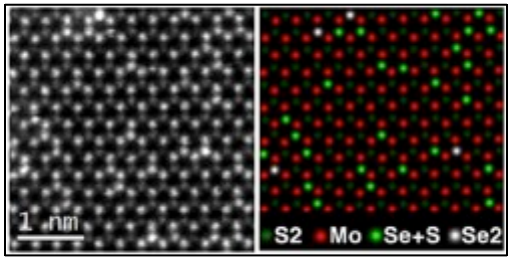

Quantitative EELS measurements are possible in a liquid cell

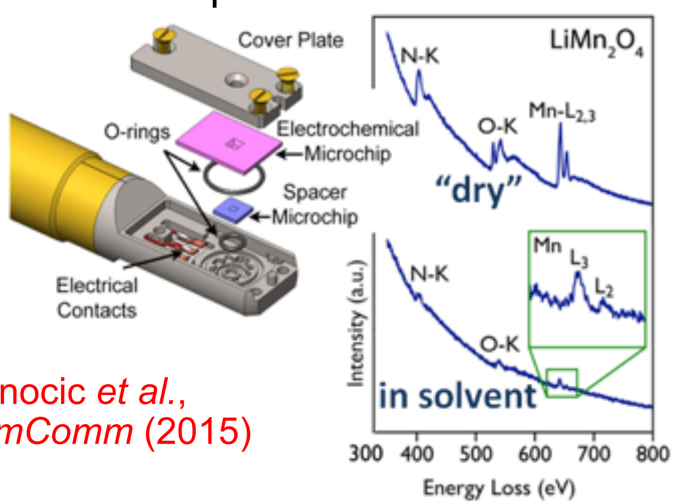

R. Unocic *et al.*, *ChemComm* (2015)

ORNL's new MAC-STEM (monochromated AC-STEM) will enable quantitative imaging of optical excitations and phonons at the nanoscale

With ~1Å probe and <10meV energy resolution, the MAC-STEM becomes the nanoscale/real-space counterpart to neutron scattering

Example: Imaging optical excitations in bilayer graphene

Non-Monochromated (350 meV)

Monochromated (60 meV)

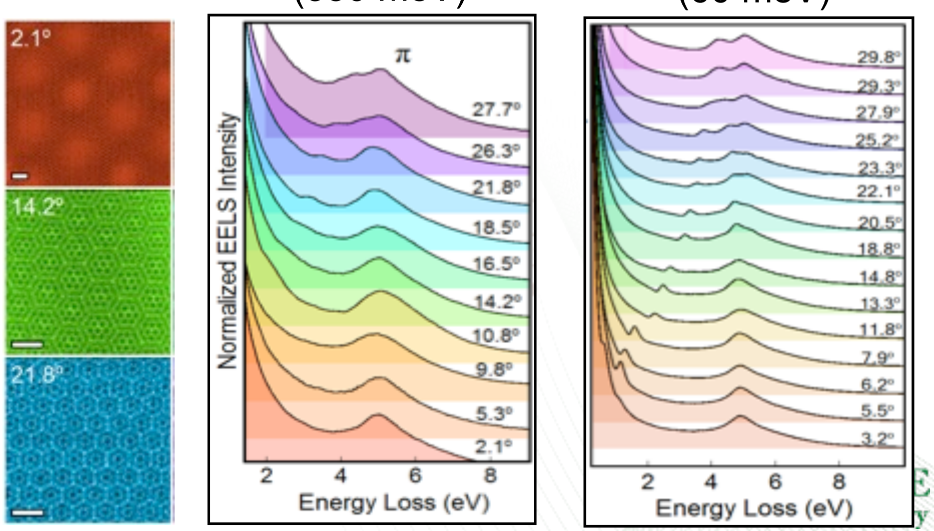

Capabilities in chemical imaging

A mass spectrometry upgrade to the HIM will provide chemical information at nanometer resolution

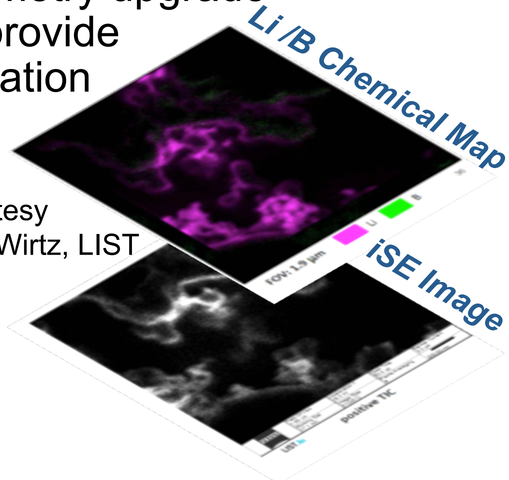

Courtesy
Tom Wirtz, LIST

The combination of scanning probes and optical spectroscopies provides a broad range of information

- AFM-Raman (NT-MDT Integra)
- AFM-Fluorescence (Asylum Bioscope)
- AFM-NanoIR (Anasys, FY2016)

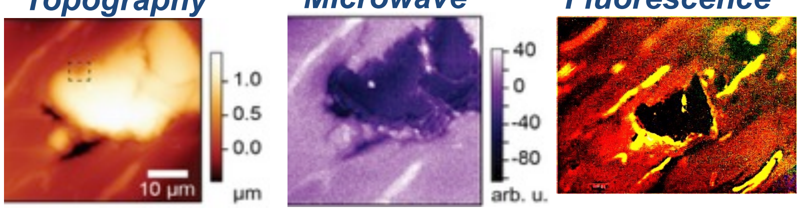

A. Tselev, *Fuel* (2014)

ORNL's new **AFM/TOF-SIMS** combines focused-ion-beam TOF-SIMS with scanning probe microscopy for co-registered chemical and physical analysis of materials

CNMS will lead the development of novel approaches, including data analytics.

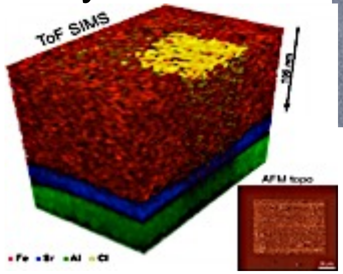

See Overview Document p. 229

$CH_3NH_3PbI_3$ (perovskite PV material)

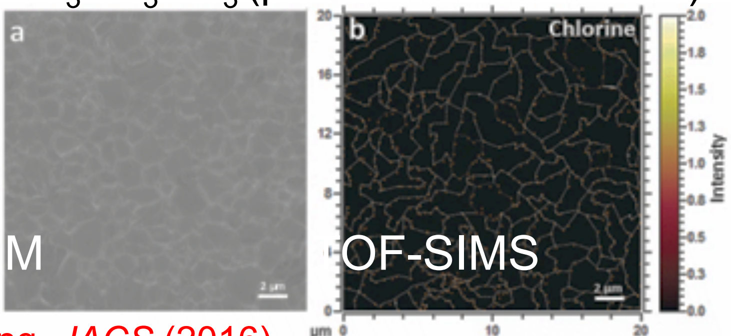

B. Yang, *JACS* (2016)

OAK RIDGE National Laboratory

ORNL Capabilities and Interests Summary Slide for Materials

Primary Expertise & Interest Areas

- Controlled/precision synthesis (molecules, crystals, thin films, bulk)
- Directing Matter: from atoms to function
- Multi-modal *in-situ/operando* experiments co-theory
- Quantum condensed matter
- Hybrid assemblies (low-D materials; soft matter composites)

Most Differentiating Factor

- Broad experience and knowledge base for materials, including synthesis (both structural and functional; basic to applied)
- Center for Nanophase Material Sciences and Spallation Neutron Source/High Flux Isotope Reactor – from materials to devices
- Aberrated probes, MAC-STEM, APT, direct-write nano (electrons and ions) + sculpting, AFM/TOF-SIMS, 4-Probe STM(S), He-ion, *in-situ* characterization
- Integrated modeling and simulation capability for materials

Main Contribution/Role

- Develop new materials and processing enabling specific properties/ performance for next generation HPC
- Provide gateway to innovation and enabling capabilities (CNMS, SNS/ HFIR, LCF)

There have been a variety of proposed compute architectures

Communication throughout hardware interface levels, including ultracold and superconducting layers

- Different potential technology choices have dictated different partitioning
- A general case would have three distinct temperature regions
  - Actual processing, Local Memory/Data transmission, aggregation/serialization occurs at 4K
  - Memory/Interface at 50-80K
  - System interface at 25C
- Technologies
  - Niobium Processing – 4K
  - CMOS memory – 4K-273K
  - RF data – SiGe – 50K-273K
  - Josephson Junctions

(Red denotes ORNL expertise)

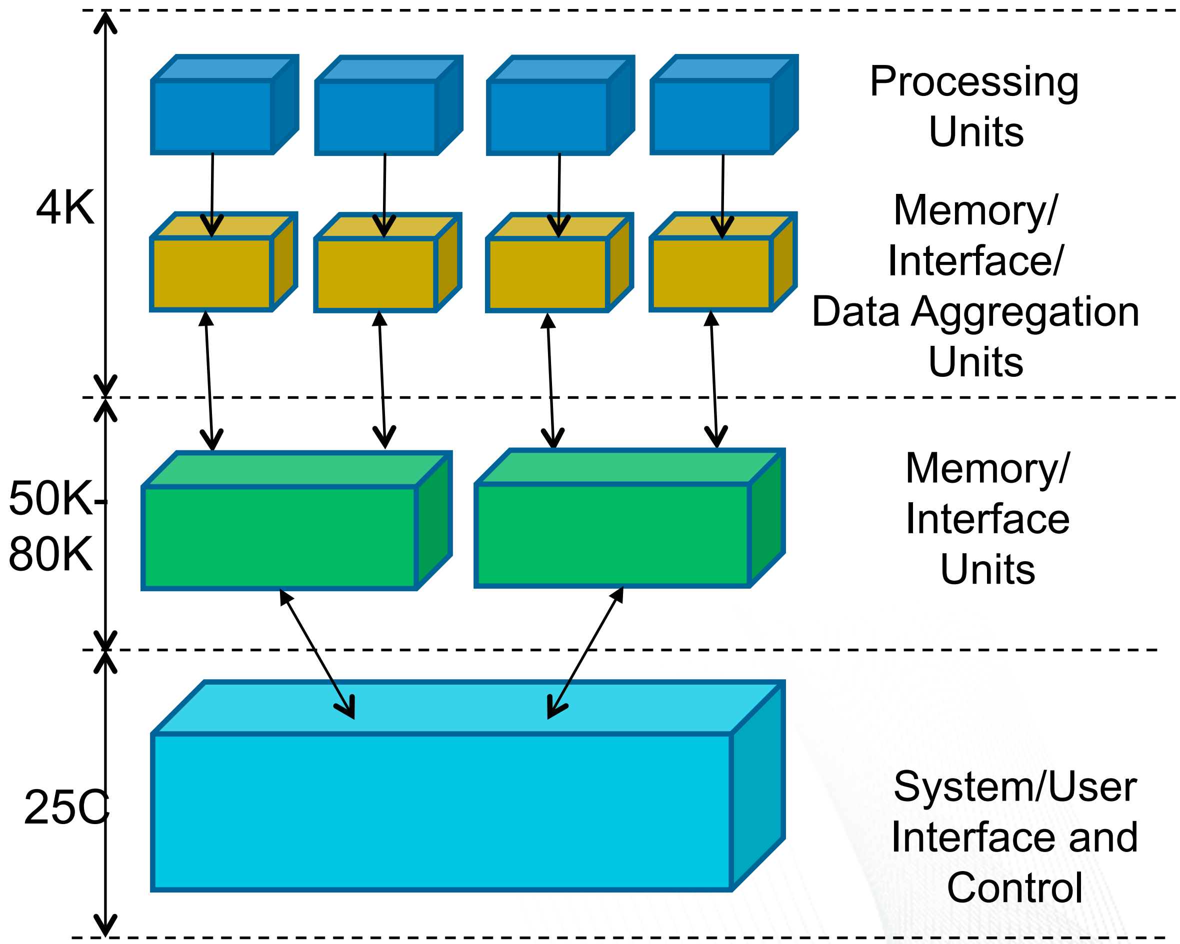

Photonic communication – from next generation architecture to macroscale

Coupling information from next gen computing unit to the rest of the world

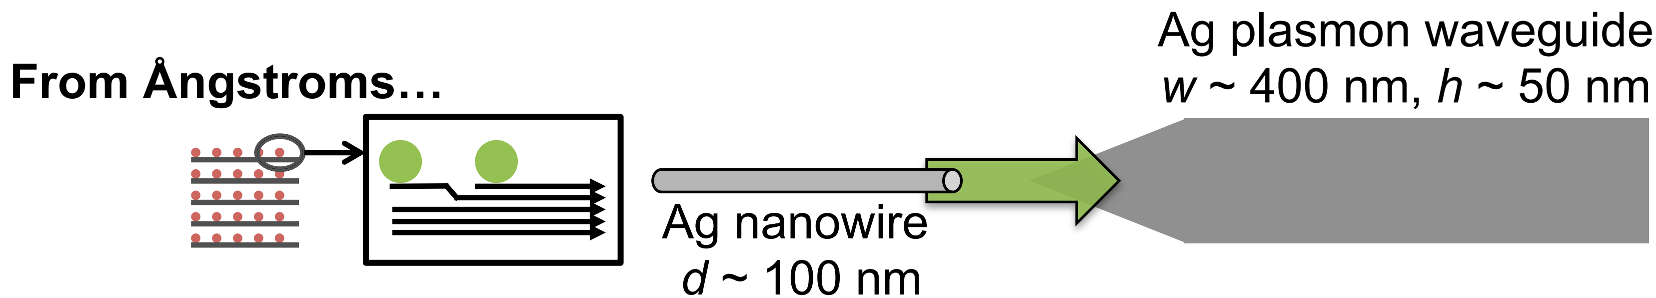

atom → nanowire ***Nanowire → strip plasmon***

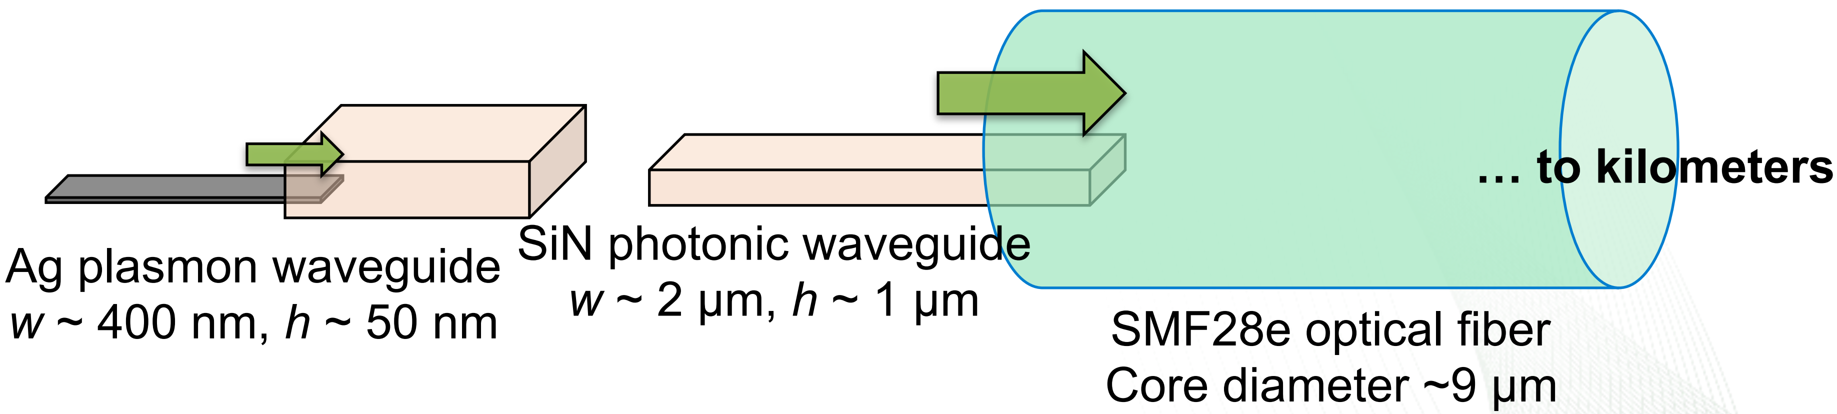

Strip plasmon → photon

OAK RIDGE National Laboratory

Nanofabrication for nanophotonic circuits and *ultralow energy* interconnects

Advanced photonic circuits for communications interconnects

DARPA InPho: new interconnects, modulation formats, *ultra low energy* interconnects

Si photonic circuits, photonic metamaterials, ultra efficient metamaterial optical antennae

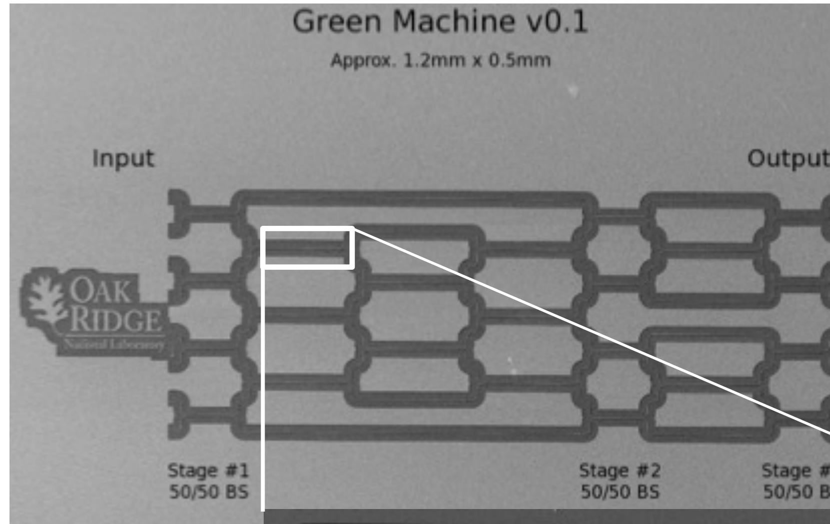

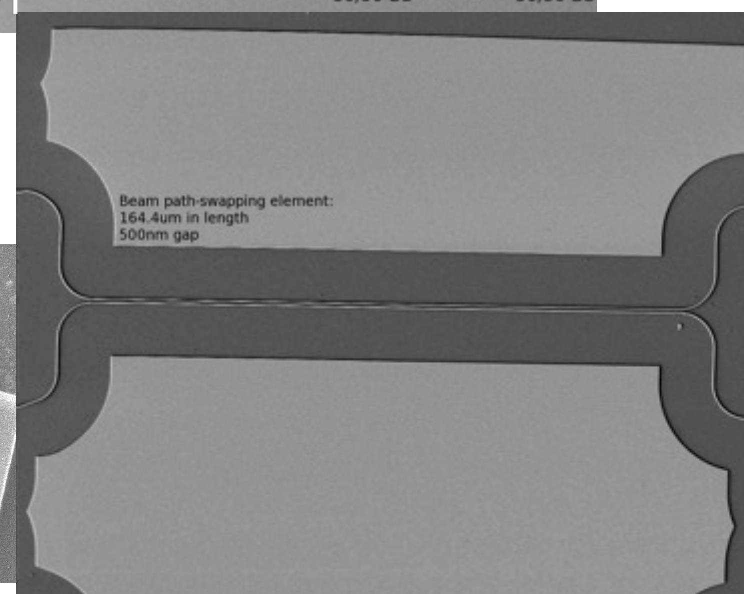

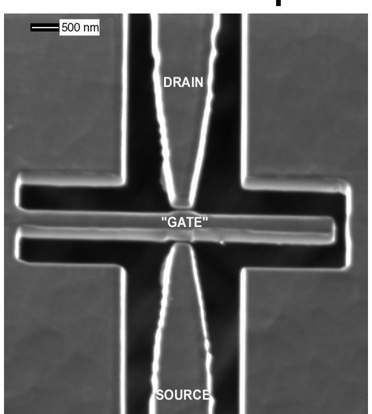

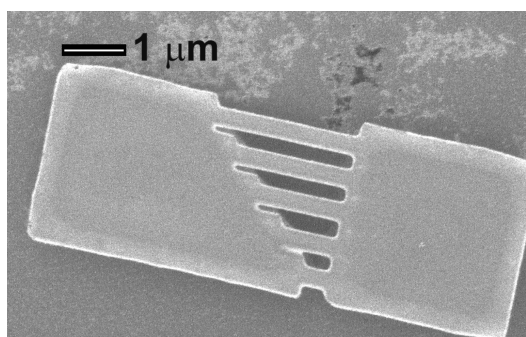

Very Low Power Analog Signal Processing (ASP)

- Continuous or periodic signal monitoring
- Utilizes commercial CMOS process (130nm, 1.2V)
- Tunable high order analog filters (8^{th}-16^{th} order)
- Building block for very low power sensing platforms
- Sample Application: Signal signature analysis

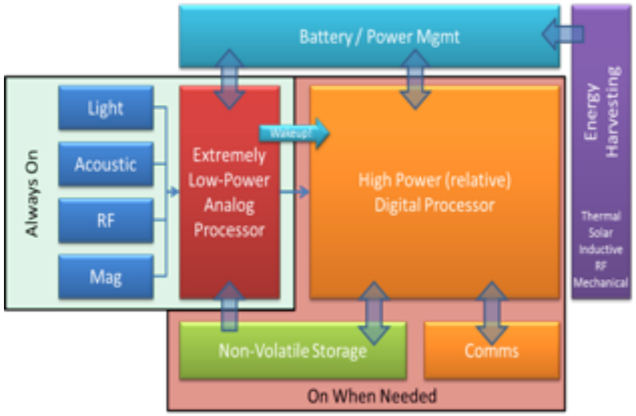

ASP –based Sensing Platform

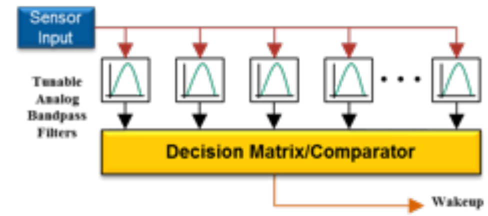

Signal Signature Analysis Example

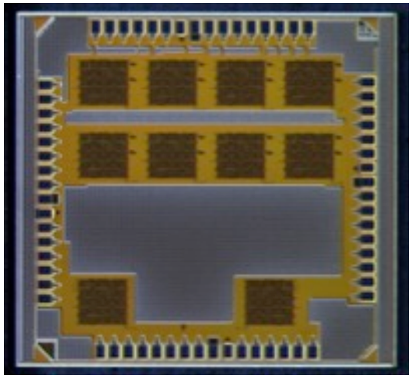

ASP Tunable Analog Filters – 130nm, 1.2V CMOS

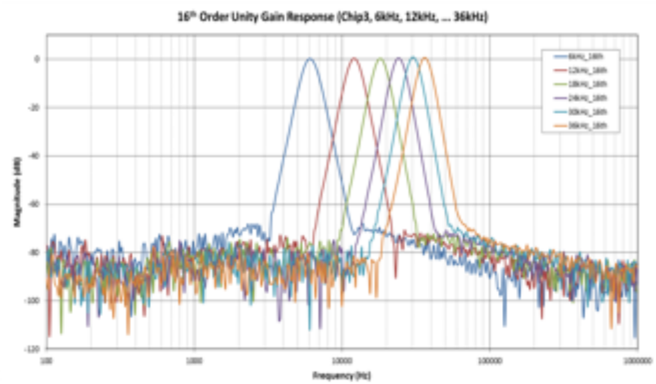

ASP Tunable Analog Filter Measured Spectral Response

DARPA UPSIDE: Low Power Machine Learning Hardware

- Goal: Implement very low power machine learning circuits
  - 1000 times reduction in energy consumption
- Analog computations using standard CMOS processes
  - Low precision, fewer components
- Graphene Field Effect Transistors
- Deep learning architectures

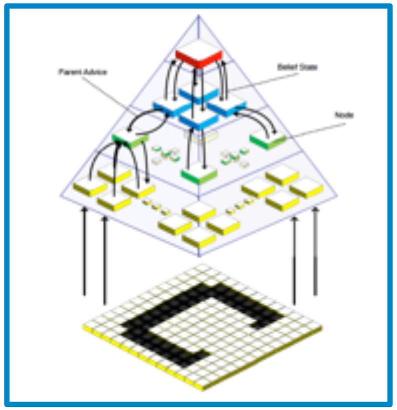

DeSTIN:
Deep Learning

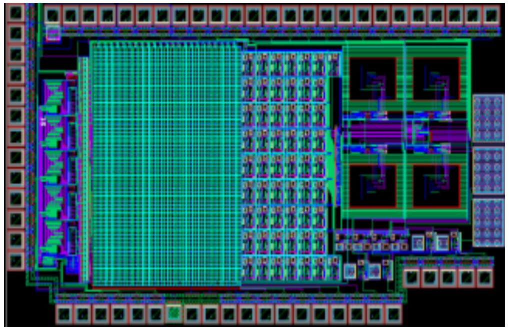

CMOS based analog computation

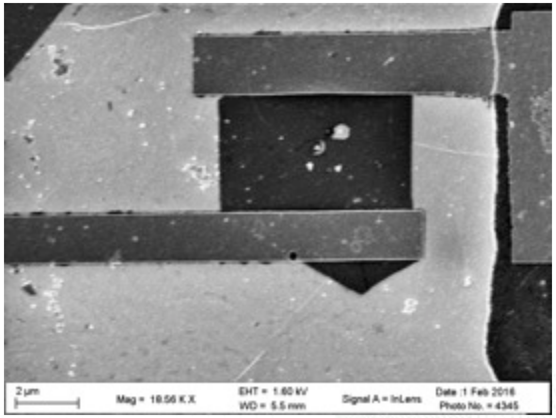

Graphene integrated on CMOS chip.

Contact: Dr. David Bolme, bolmeds@ornl.gov, 865-576-0300

Charge recycling and adiabatic approaches to both analog and digital circuitry

- Charge recycling allows charge to be scavenged from the leakage and dynamic switching of CMOS logic
- A potential of ~ 25% can be saved in total power required for operation
- Adiabatic techniques (low-power, reduced speed) can also be applied to both analog and digital functions

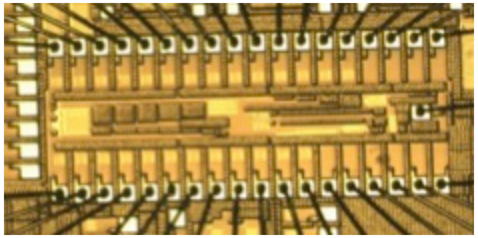

Die picture & layout of the CR gray-code counter in 90nm process

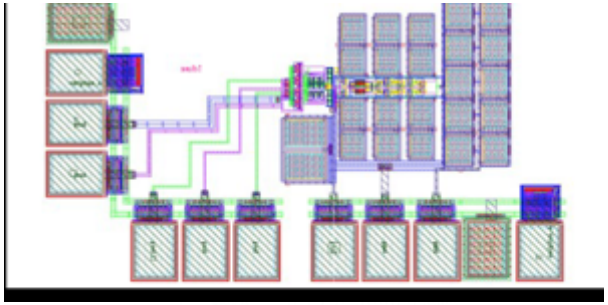

130-nm power-scavenging charge pump

C. Ulaganathan, B. J. Blalock, J. Holleman and C. L. Britton, *2012 IEEE 55th International Midwest Symposium on Circuits and Systems (MWSCAS)*, Boise, ID, 2012, pp. 206-209.

C. Ulaganathan, C. L. Britton, J. Holleman and B. J. Blalock, " *Mixed Design of Integrated Circuits and Systems (MIXDES), 2012 Proceedings of the 19th International Conference*, Warsaw, 2012, pp. 208-213.

ORNL Capabilities and Interests Summary Slide for Comms. & Logic

Primary Expertise & Interest Areas

- Low power electronics
- RF electronics
- Integrated optics / photonics
- Analog / digital hybrid signal processing
- Superconductors and electronics at very low power and temperature

Most Differentiating Factor

- Broad experience in all stages of integrated device development
- Embedded RF electronics, machine vision, photonics, and quantum optics groups all specializing in communications layers
- Rapid characterization of new logic devices
- Highly developed modeling and simulation capability for both full devices and higher level logic

Main Contribution/Role

- Develop new materials into true logic devices
- Provide interconnect designs and implementation
- Validate designs via modeling and simulation of low level devices and logic

ORNL Capabilities and Interests Summary Slide for Hardware & Circuits Architectures

Primary Expertise & Interest Areas

- HPC hardware architectural optimization
  - Nonvolatile memory organization
  - Novel memory modeling tools (e.g., Destiny)
- Power efficient memory architectures
- Low power mixed signal processing electronics & architectures
- Embedded systems
- Superconducting computing (cryogenic memory design based on Josephson Junctions)

Most Differentiating Factor

- Tools for software-hardware co-design and hardware organization
  - Focus on NVM, Heterogeneous processing
- Extensive simulation capability for multi-memory technologies
- Broad experience in all stages of architectures and circuits

Main Contribution/Role

- HPC architecture and organization with focus on novel memory systems
- Design of cryogenic memory with small footprint, very low energy dissipation, and ultra-fast access time.
- Low power analog processing circuits design and simulation
- Mixed signal circuits and architectures design, testing and characterization

Quantum & Neuromorphic

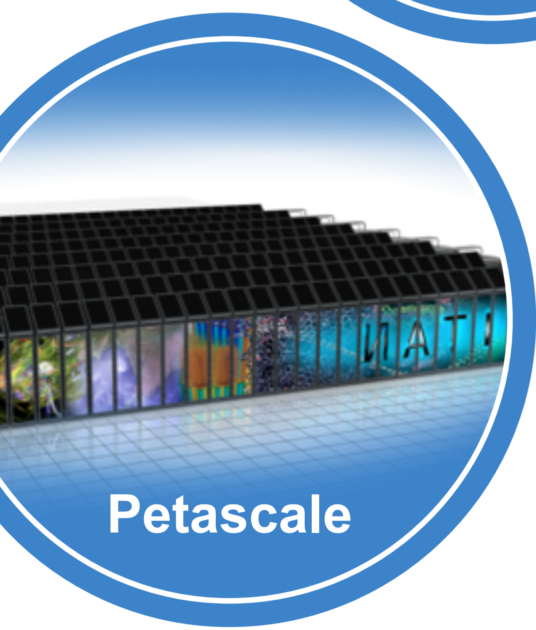

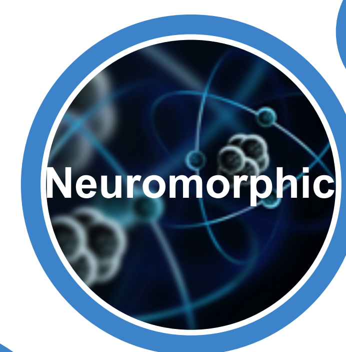

ORNL Cross-cutting initiative: Quantum computing materials and interfaces

Opportunity

- Integrate core competencies in materials, modeling, and isotopes to establish a broad R&D effort in quantum computing
- Create S&T base to drive computing beyond exascale

ORNL assets

- Expertise in quantum information science and quantum computing
- Unique resources for materials characterization
- Strengths in first principles theory, modeling, and simulation for quantum materials
- CNMS, SNS/HFIR

Strategy

- Develop tools necessary to characterize and design high-fidelity physical qubits
- Explore methods to interface qubits to traditional computers
- Develop a multi-qubit research test bed
- Research methods to program multi-qubit systems
- Foster multiagency ties to secure long-term funding

Neuromorphic Computing ("Top-down")

Models and Simulation

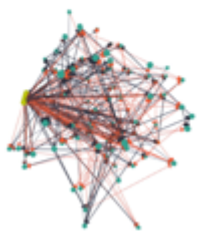

High-level Models:
- Spiking Neural Networks
- Event-driven simulation

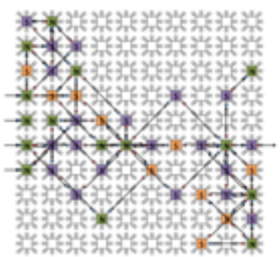

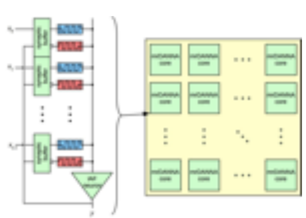

Low-level Models:
- Hardware-accurate simulations
- Digital (FPGA, VLSI)
- Memristor

Scalable for HPC

Algorithms

Evolutionary Algorithm Training

Bio-Inspired Learning

Supporting Software

"Compilers"

Visualization

Application Development Environment

Fundamental relation of reversibility to energy consumption for computing

- **Initial Question**
  What is the minimum energy needed/dissipated to "compute?"
  - **Initial thesis**
    Every ***bit operation*** dissipates a unit of energy ($kT\ln 2$)
  - **Next development**
    Not every ***bit operation***, but every ***bit erasure*** dissipates a unit of energy ($kT\ln 2$).
    Other bit operations can be implemented without energy dissipation
- **Follow-on Question**
  What is the minimum number of bit erasures needed to "compute?"
  - **Initial hypothesis**
    There would be a non-zero, computation-specific number
  - **Bennett's surprising solution:**
    Zero bit erasures! Bennett's "compute-copy-uncompute" algorithm avoids *all* bit erasures for *any* arbitrary (Turing) program
  - **Further refinements**
    Algorithmic complexity, tradeoffs
    ***Partial* reversibility**

Software for emerging computing systems

- **Software is needed to address multiple aspects of emerging computing systems**
  - Applications and programming
  - Execution and run-time
  - Device and architecture design
- **Modeling and simulation provides a useful proxy for hardware, predicting performance**
  - Separation of concerns isolates functional requirements, manages complexity
  - Simulation granularity offers varying scope, device to system scale

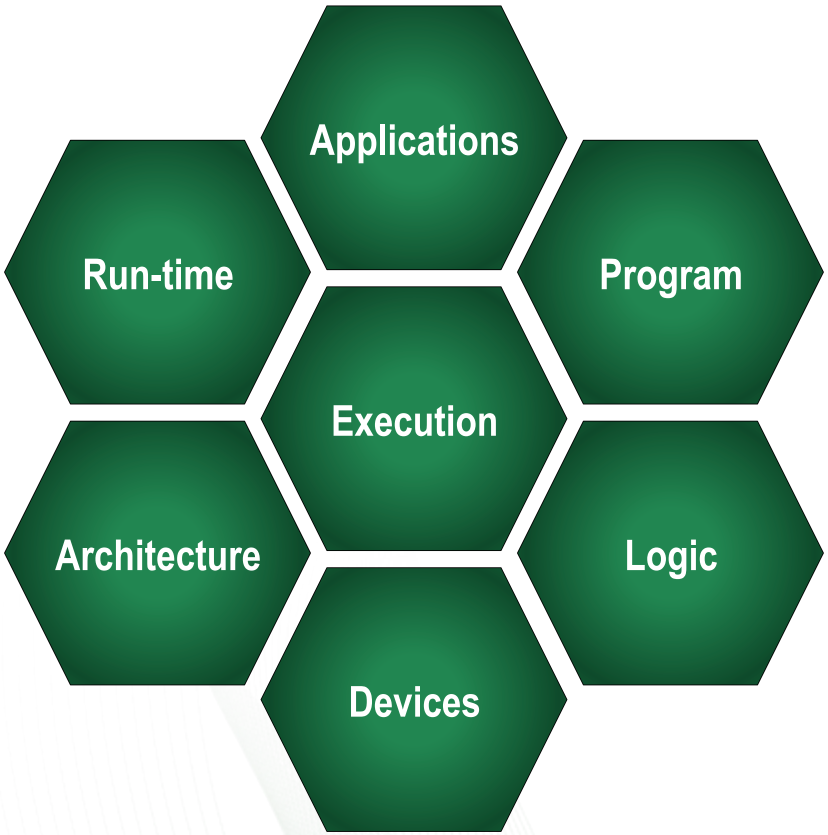

Performance Modeling and Simulation

- **Models of hybrid HPC architectures help stakeholders set expectations, planning**
  - Priority is performance models at the node-level
  - This works for classical parallelism, domain decomposition methods
  - Novel technologies may challenge node parallelism, more complicated scaling
- **Two problems need to be solved**
  - **Forward problem**: What is the performance of a fixed system design? Time-to-solution?
  - **Inverse problem**: What system meets given performance requirements?

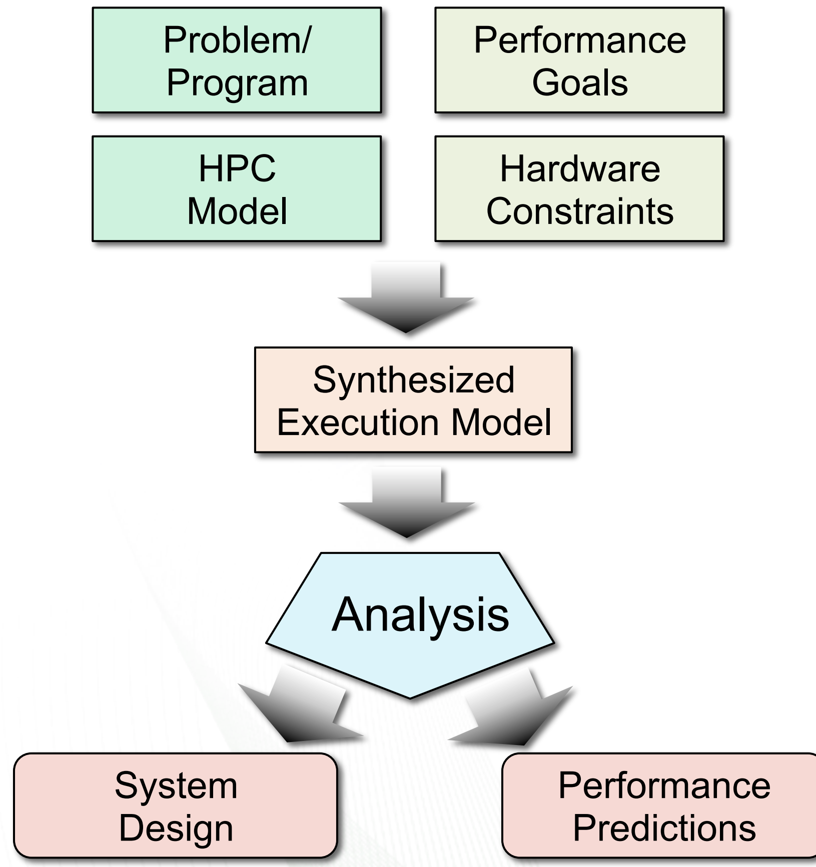

OpenARC: Open Accelerator Research Compiler

- Open-Sourced, High-Level Intermediate Representation (HIR)-Based, Extensible Compiler Framework.
  - Perform source-to-source translation from OpenACC C to target accelerator models.
    - Support full features of OpenACC V1.0 (+ array reductions and function calls)
    - Support both CUDA and OpenCL as target accelerator models
  - Provide common runtime APIs for various back-ends
  - Can be used as a research framework for various study on directive-based accelerator computing.
    - Built on top of Cetus compiler framework, equipped with various advanced analysis/ transformation passes and built-in tuning tools.
    - OpenARC's IR provides an AST-like syntactic view of the source program, easy to understand, access, and transform the input program.

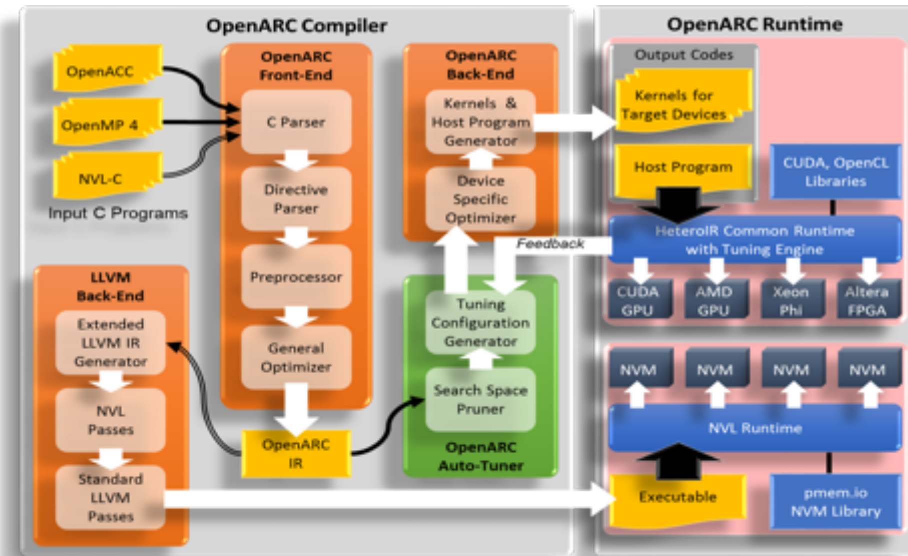

ACM Symposium on High-Performance Parallel and Distributed Computing (HPDC). Vancouver: ACM, 2014

Oak Ridge National Laboratory | Center for Nanophase Materials Sciences

OpenACC to FPGA: A Framework for Directive-Based High-Performance Reconfigurable Computing

- Problem
  - Reconfigurable computers, such as FPGAs, offer more performance and energy efficiency for specific workloads than other heterogeneous systems, but their programming complexities and low portability have limited their deployment in large scale HPC systems.
- Solution
  - Proposed an OpenACC-to-FPGA translation framework, which performs source-to-source translation of the input OpenACC program into an output OpenCL code, which is further compiled to an FPGA program by the underlying backend Altera OpenCL compiler.
- Recent Results
  - Proposed several FPGA-specific OpenACC compiler optimizations and pragma extensions to achieve higher throughput.
  - Evaluated the framework using eight OpenACC benchmarks, and measured performance variations on diverse architectures (Altera FPGA, NVIDIA/AMD GPUs, and Intel Xeon Phi).
- Impact
  - Proposed translation framework is the first work to use a standard and portable, directive-based, high-level programming system for FPGAs.
  - Preliminary evaluation of eight OpenACC benchmarks on an FPGA and comparison study on other accelerators identified that the unique capabilities of an FPGA offer new performance tuning opportunities different from other accelerators.

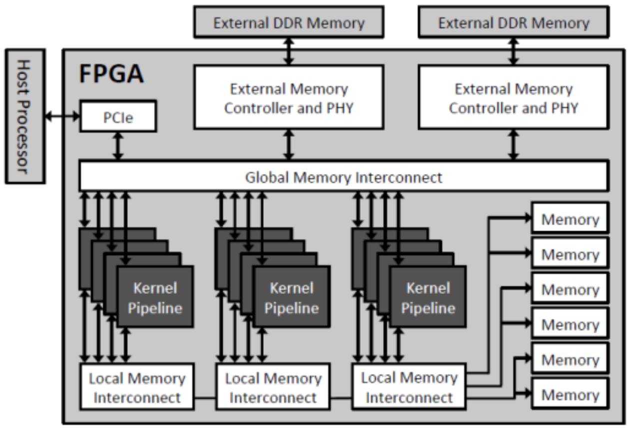

Figure 2: FPGA OpenCL Architecture

S. Lee, J. Kim, and J.S. Vetter, "OpenACC to FPGA: A Framework for Directive-based High-Performance Reconfigurable Computing," Proc. IEEE International Parallel & Distributed Processing Symposium (IPDPS), 2016.

ORNL Capabilities and Interests Summary Slide for Algs & Software

Primary Expertise & Interest Areas

- Software for emerging computers
  - Quantum, neuromorphic, optical, reversible computing
- Predicting system performance
  - Modeling and Simulation
  - ASPEN, xSIM, Colloquium
- Migrating apps to new platforms
  - Compilers, libraries, run-time systems,

- Software Ecosystems
  - ReveR-SES, reversible compiler
  - Annie Rose, quantum simulator
- Experimental Computing Laboratory
  - User access to new technology platforms, tools
- Career Awards in Emerging Technologies
  - Kalyan Perumalla, reversible, (ASCR)
  - Travis Humble, quantum, (ASCR)
  - Katie Schuman, neuromorphic, (ORNL)

Main Contribution/Role

- Develop algorithms to exploit emerging computing paradigms
- Develop predictive performance models for systems at scale
- Refine user applications/codes to leverage new libraries
- Develop software ecosystems to validate

Materials Innovation at ORNL

Integration of experiments, data and simulations across scales

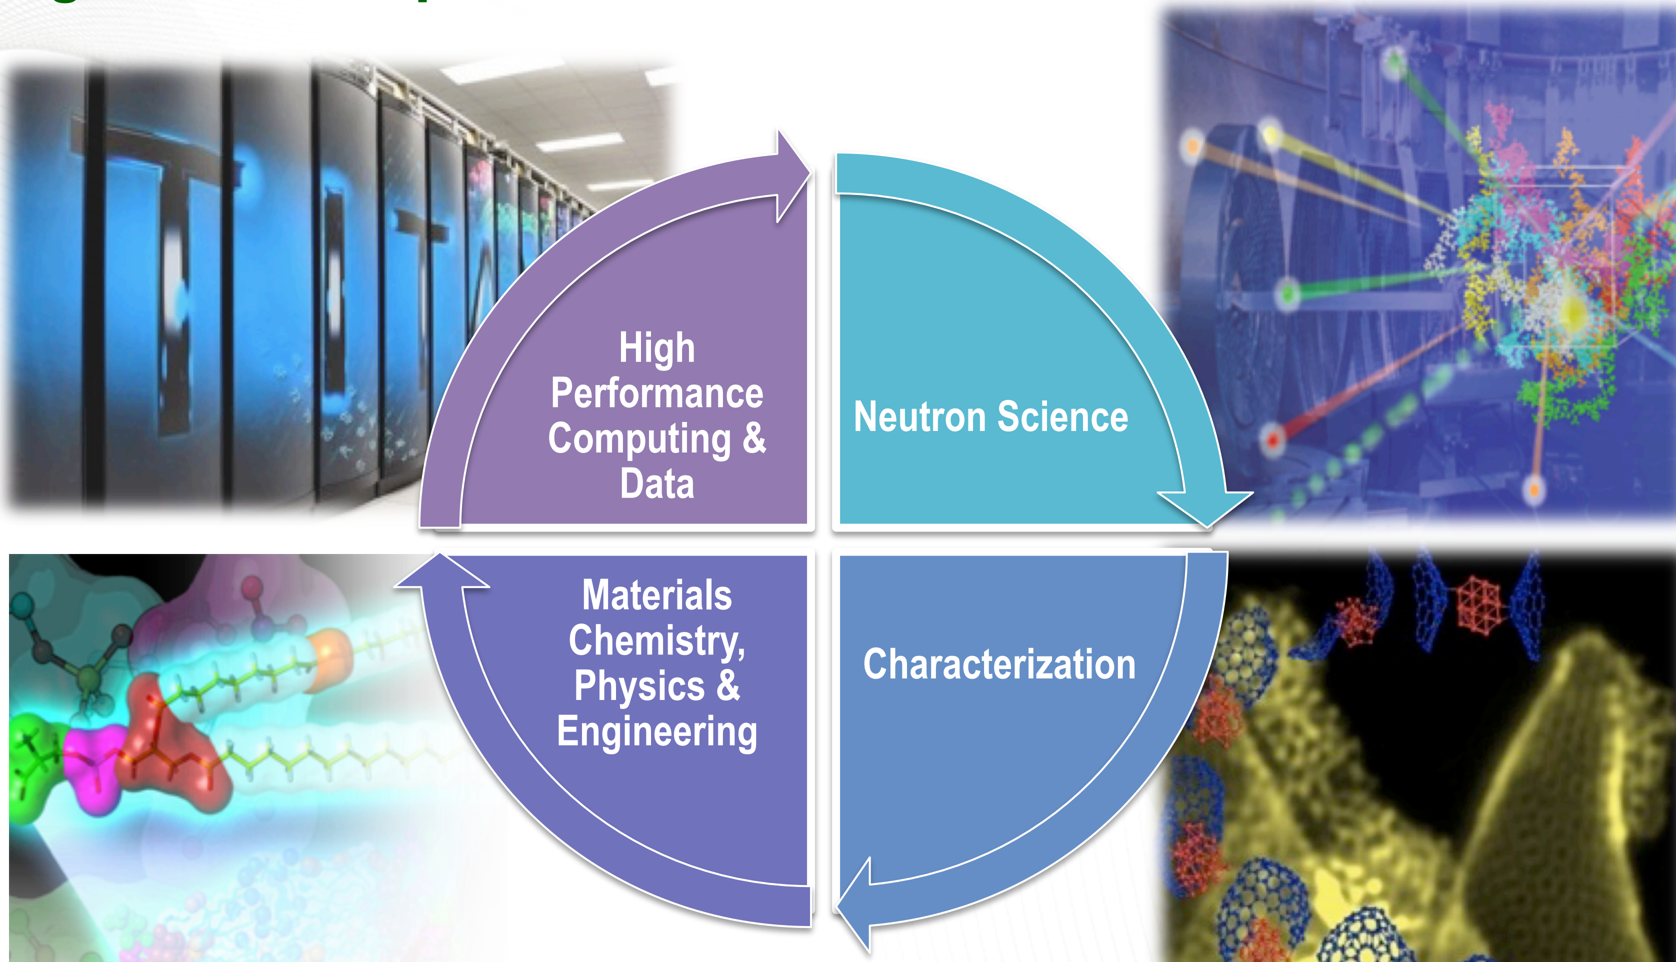

Revolutionary structural and functional materials
for energy generation, storage, & use