DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof. Reference herein to any social initiative (including but not limited to Diversity, Equity, and Inclusion (DEI); Community Benefits Plans (CBP); Justice 40; etc.) is made by the Author independent of any current requirement by the United States Government and does not constitute or imply endorsement, recommendation, or support by the United States Government or any agency thereof.

ORNL/TM-2026/4474

Laboratory Directed Research and Development Program

FY 2025 COMPLETED PROJECTS

OAK RIDGE National Laboratory

ORNL/TM-2026/4474

Oak Ridge National Laboratory

LABORATORY DIRECTED RESEARCH AND DEVELOPMENT PROGRAM

FY 2025 COMPLETED PROJECTS REPORT

March 31, 2026

Prepared by
OAK RIDGE NATIONAL LABORATORY
Oak Ridge, TN 37831
managed by
UT-BATTELLE LLC
for the
U.S. DEPARTMENT OF ENERGY
under contract DE-AC05-00OR22725

CONTENTS

ABBREVIATED TERMS

AEM	alkaline exchange membrane
AFQ	accelerated fuel qualification
AI	artificial intelligence
ALCF	Argonne Leadership Computing Facility
AM	additive manufacturing
APPL	Advanced Plant Phenotyping Laboratory
ARM	Accelerating Radiotherapeutics through Advanced Molecular Constructs
ASGarD	Adaptive Sparse Grid Discretization
CDR	carbon dioxide removal
CFD	computational fluid dynamics
CFG	coal flue gas
CLAM	Curvilinear Aperture Monopulse
CMOS	complementary metal–oxide–semiconductor
CNF	cellulose nanofiber
CONUS	contiguous United States
COR	catalytic oxygen reduction
CRI	Convergent Research Initiative
cryoEM	cryogenic electron microscopy
CSP	cryogenic shell pellet
DEAE	diethylaminoethanol
DEPA	di(4-ethyl-2-pyridyl)amide
DETA	diethylenetriamine
DIAL	Distributed INTERSECT Active Learning
DMAE	dimethylaminoethanol
DOE	US Department of Energy
DOTA	1,4,7,10-tetraazacyclododecane-1,4,7,10-tetraacetic acid
DRI	direct reduced iron
DRIFT	Design of Resilient Internet of Things/Edge with Federated Tiny Machine Learning
DTPA	diethylenetriaminepentaacetic acid
EC	edge computing
Eco-NetI	Edge-to-ecosystem ConNect Integration
EGR	exhaust gas recirculation
ELM	edge-localized mode
ELM	Land Model of the DOE Energy Exascale Earth System Model
EMF	ectomycorrhizal fungi
ERW	enhanced rock weathering
FL	federated learning
FLiBe	$2LiF–BeF_2$
FY	fiscal year
GAN	generative adversarial network
GC	gas chromatography
GITR	Global Impurity Transport Code

GRAVITAS	Graphics Processing Unit–Based Resolution And Visualization Interface for Triple-Axis Spectrometers
HHF	high–heat flux
HPC	high-performance computing
HPMC	hydroxypropyl methyl cellulose
HS	Hamiltonian simulation
ICE	instrument-computing ecosystem
ICP-MS	inductively coupled plasma mass spectrometry
IFF	interatomic force field
INTERSECT	Interconnected Science Ecosystem
IoT	Internet of Things
LA	laser ablation
LCA	life cycle assessment
LDRD	Laboratory Directed Research and Development
LIBS	laser-induced breakdown spectroscopy
LiNGAM	Linear Non-Gaussian Acyclic Model
LLM	large language model
MC	methylcellulose
MCNP	Monte Carlo N-Particle
MEA	monoethanolamine
MHD	magnetohydrodynamic
MHT	methyl halide transferase
ML	machine learning
MICE	Mapping of Isotopic, Chemical, and Elemental
MPEX	Material Plasma Exposure eXperiment
MP-LDGM	Multi-Physics Lattice Discrete Grain Model
MRI	magnetic resonance imaging
MS	mass spectrometry
MSG	magic squares game
m/z	mass-to-charge ratio
NDIP	Neutron Data Interpretation Platform
NEAT	Nonequilibrium and Emergent Transients in Advanced and Soft Materials
NGFG	natural gas flue gas
NIR	near-infrared
NLG	nonlocal game
NOVA	Neutrons Advanced Visualization and Analysis
NSI	Nonproliferation Science Initiative
NTRC	National Transportation Research Center
OLCF	Oak Ridge Leadership Computing Facility
ORNL	Oak Ridge National Laboratory
ORQCS	Oak Ridge Quasi-Clifford Simulator
PETINA	Privacy Preservation Algorithms
PFC	plasma-facing component
PFM	plasma-facing material
PGA	polyguanidinium
PoL	pattern-of-life
PRESTO	Privacy Recommendation and Security Optimization
PSP	plasma separation process
PVDF	polyvinylidene fluoride
Py4PA	pyridine-based tetrapicolinic acid
QC	quantum computing

QCCD	quantum charge-coupled device
R&D	research and development
re-ID	re-identification
ROMP	ring-opening metathesis polymerization
RPM	radiation portal monitor
RPT	radiopharmaceutical therapy
RT	radiation transport
SANS	small-angle neutron scattering
SAR	synthetic aperture radar
SDL	scientific data layer
SG	sparse grid
SNS	Spallation Neutron Source
SPP	Strategic Partnership Projects
TAT	targeted alpha therapy
TBC	terpyridine biscarboxylate
TEA	techno-economic analysis
TEG-DME	triethylene glycol dimethyl ether
TEM	transmission electron microscopy
TETA	triethylenetetramine
tinyML	tiny machine learning
TISCC	Trapped-Ion Surface Code Compiler
TOR	Train of Reactors
TRISO	tristructural-isotropic
UAS	uncrewed autonomous systems
UT-ORII	The University of Tennessee–Oak Ridge Innovation Institute
V44	V–4Cr–4Ti
V44C	V–4Cr–4.4Ti–0.1C–0.2Y
WF	wood flour
WPT	wireless power transfer

INTRODUCTION

Oak Ridge National Laboratory (ORNL) is the US Department of Energy's (DOE's) largest multiprogram science, technology, and energy laboratory. It possesses distinctive capabilities in a variety of fields, such as neutron science, computing, advanced materials, and nuclear science and technology. Using these capabilities, ORNL conducts basic and applied research and development (R&D) to support DOE's overarching mission "to ensure America's security and prosperity by addressing its energy, environmental and nuclear challenges through transformative science and technology solutions."[1] As a national resource, ORNL also applies its capabilities and skills to the specific needs of other federal agencies and customers through the DOE Strategic Partnership Projects (SPP) Program. Information about the laboratory and its programs is available on the ORNL website.[2]

The Laboratory Directed Research and Development (LDRD) Program at ORNL operates under the authority of DOE Order 413.2C, *Laboratory Directed Research and Development*,[3] which establishes DOE's requirements for the program while providing the laboratory director broad flexibility for program implementation. The LDRD Program funds are obtained through a charge to all laboratory programs. Although it represents a relatively small portion of the overall research budget, the LDRD Program plays an essential role in maintaining the laboratory's ability to respond to national needs. The program allows ORNL to improve its distinctive capabilities and enhance its ability to conduct cutting-edge R&D. In accordance with the DOE order, R&D projects funded through the LDRD Program at ORNL support the goals of

- maintaining the scientific and technical vitality of the laboratory,
- enhancing the laboratory's ability to address future DOE missions,
- fostering creativity and stimulating exploration of forefront areas of science and technology,
- serving as a proving ground for new concepts in R&D, and
- supporting high-risk, potentially high-value R&D.

This report provides an overview of the LDRD Program at ORNL in FY 2025 and contains summaries of all the LDRD research projects that concluded between October 1, 2024, and September 30, 2025.

[1] DOE, "Mission," https://www.energy.gov/mission.
[2] ORNL, "Home page," https://www.ornl.gov/.
[3] DOE, *Laboratory Directed Research and Development*, DOE Order 413.2C (US Department of Energy, 2018), https://www.directives.doe.gov/directives-documents/400-series/0413.2-BOrder-c-chg1-minchg/@@images/file.

LDRD Program Summary

ORNL has established a program with six complementary subprograms to meet its LDRD Program objectives and meet the particular needs of the laboratory. A provision for multiple routes of access to ORNL LDRD Program funds maximizes the likelihood that novel ideas with scientific and technological merit will be recognized and supported.

The six subprograms are as follows:

- The **Director's R&D Program** supports projects that advance research frontiers, capabilities, and expertise at ORNL in key strategic areas.
- The **Seed Program** supports innovative high-risk/high-reward research to the proof-of-principle stage.
- The **Strategic Hire Program** supports the research of key new staff whose expertise and capabilities address a critical strategic need for the laboratory.
- The **Distinguished Staff Fellowships** assist the laboratory in bringing in exceptional early career scientists to refresh and expand the lab's scientific and technical expertise.
- The **Early Career Competition** invests in early career scientists and engineers to cultivate new research leadership.
- The **University of Tennessee–Oak Ridge Innovation Institute (UT-ORII) Fellowships** support research relating to key topic areas that bring together the expertise and infrastructure of ORNL and the University of Tennessee System.

The total ORNL LDRD Program expenditure of $65.5 million was approximately 3.0% of the laboratory's total budget of $2,192,543,853 for operating and capital expenses, which is well below the maximum of 6% allowed by DOE Order 413.2C and is in accordance with Section 309 of Division D of the Consolidated Appropriations Act.[4]

Table 1. ORNL FY 2025 LDRD costs

	Director's R&D Program	Seed Program	Strategic Hire Program	Distinguished Staff Fellowships	Early Career Competition	UT-ORII Fellowships	Total
Cost ($)	47,097,269	7,038,705	1,706,790	3,581,299	2,329,498	3,707,496	65,461,057
Number of active projects	105	52	6	16	22	16	217

Director's R&D Program

The development of research initiatives is the primary tool by which the laboratory develops, builds, and enhances capabilities in particularly promising areas of science and technology to meet anticipated national needs. The Director's R&D Program plays an important role in helping ORNL realize successful outcomes for its initiatives by providing resources to conduct R&D in key areas.

In FY 2025, the Director's R&D Program supported research under the following initiatives, which aligned with the ORNL Lab Plan and emerging sponsor needs:

- Accelerating Radiotherapeutics through Advanced Molecular Constructs (ARM)
- Artificial Intelligence (AI)
- Enzyme Engineering

[4] Consolidated Appropriations Act, Public Law 113-76 (US Congress, 2014), https://www.congress.gov/113/plaws/publ76/PLAW-113publ76.htm.

- Heterogeneous Quantum Systems
- Integrated Studies of Complex Biological and Environmental Systems
- Nonequilibrium and Emergent Transients in Advanced and Soft Materials (NEAT)
- Neutron Data Interpretation Platform Ecosystem
- Nonproliferation Science Initiative (NSI)
- Self-Driven Experiments for Science/Interconnected Science Ecosystem (INTERSECT)
- Transform Initiative
- Validated Design and Evaluation of Fusion Wall Components
- Vertex: Advanced Multiphysics Simulations for Core Applications

Additionally, the program supported a small number of projects outside of the focused initiatives. These projects, selected to address emerging strategic needs, are called Deputy Director Discretionary projects.

The research priorities for each of the initiatives are described in the following sections.

Accelerating Radiotherapeutics through Advanced Molecular Constructs (ARM)

The goal of the ARM Initiative is to advance technology in designing, synthesizing, and characterizing molecular cages to enable targeted transport of therapeutic radionuclides to cancer sites anywhere within the body. This goal will be achieved by conducting fundamental research into the coordination chemistry of actinide, lanthanide, and alkaline earth radioisotopes to guide in silico design and synthesis of novel molecular constructs for optimal capture and retention of these radionuclides and their decay daughters. The goal includes elucidation of the mechanism of action of these targeted radionuclides and builds a foundation for advances in precision medicine for cancer and drug-resistant diseases. This initiative is built on four thrust areas: (1) designing ultrachelators for targeted radiotherapy, (2) developing targeted nanoconstructs for theranostics and combination therapies, (3) advancing novel targeting vectors, and (4) developing precision dosimetry at scale to improve the understanding of radiation doses at target and off-target sites within the body.

Artificial Intelligence (AI)

The AI Initiative focuses on secure, trustworthy, and energy-efficient AI to advance and transform the scientific, engineering, and national security missions of the lab. The initiative has three major application areas:

- **AI for Scientific Discovery and Complex Systems,** which seeks to use translational AI to accelerate scientific discoveries and to model, optimize, and control complex engineered systems;
- **AI for Experimental Facilities,** which aims to automate science workflows at experimental facilities using AI-driven methods and includes self-directing labs; and
- **AI for National Security,** which aims to support national security efforts by developing foundational AI technologies required for advanced threat detection, decision-making, and planning.

Enzyme Engineering

The Enzyme Engineering Initiative seeks to rapidly engineer enzymes that catalyze reactions at material interfaces. Enzymes can perform challenging chemical transformations with high specificity under benign conditions. Natural enzymes can be further improved through directed evolution to improve activity, specificity, and stability. However, current enzyme engineering methods are slow and, largely because of analytical limitations, difficult to use with enzymes that act at material interfaces. Addressing these limitations will enable novel biochemical solutions to challenges in materials science. This initiative pursues computational enzyme design using machine learning (ML), high-throughput enzyme expression, rapid analytical techniques, and atomistic simulations focused on polymer hydrolysis and protein binding of platinum-group elements.

Heterogeneous Quantum Systems

The purpose of the Heterogenous Quantum Systems Initiative is to enable scalable networks of heterogeneous quantum computing and sensing platforms through coherent transduction of quantum information. This initiative aims to enable new applications and workflows by using the power of in-house experimental quantum computers and sensors, commercial quantum computers, and unique DOE user facilities available at the Spallation Neutron Source, the High Flux Isotope Reactor, and the Center for Nanophase Materials Sciences. Research across quantum materials, quantum sensing, quantum networks, and quantum computing will be used to advance DOE mission objectives in quantum information science. The initiative supports crosscutting research in three thematic areas: quantum transduction across scalable quantum networks; quantum sensing, spectroscopy, and imaging; and workflows for hybrid quantum–classical computing systems.

Integrated Studies of Complex Biological and Environmental Systems

A large proportion of microbial and plant genes have unknown functions, and characterized genes often cannot be linked to functional traits. This lack of knowledge limits the ability to develop mechanistic models that accurately simulate an organism's activities and response to change in environmental conditions. Furthermore, although traits have become fundamental to understanding ecology and evolution, predicting how these traits affect large-scale systems, such as ecosystem, watershed, and Earth-system levels, remains a grand challenge. Global change may elicit unusually large feedback from some of these ecosystems, for which vulnerability and response are uncertain. Addressing this challenge will accelerate research for critical needs in bioenergy and bioproducts, biosystems design, and the predictability of ecological and Earth systems. The Integrated Studies of Complex Biological and Environmental Systems Initiative focuses on advancing the understanding, design, and engineering of the fundamental processes influencing complex systems' response to, resilience against, and recovery under extreme environmental and climatic events. To accomplish this objective, the initiative explicitly targets multiple scales: cells, organisms, communities, and ecosystems. The focus on resiliency acknowledges that biological and ecological systems possess an incredible potential to adjust and adapt to wide fluctuations in their surrounding environments. This initiative supports the following key areas of research:

- Mechanistic insights into the interfaces among microbiology, geochemistry, and root–soil processes in coupling nutrient and trace metal cycling and how biotic–abiotic interactions translate from the molecular level to genetic and ecosystem scales
- High-throughput genomic characterization, functional validation, and phenotyping that supports work at the intersection of trophic scales and kingdoms
- Coupling of terrestrial and aquatic systems to numerically describe how climate, land use and cover, and trace element fate and transformation influence stream biogeochemistry
- Modeling studies to identify threshold to ecosystem resiliency (e.g., when ecosystems reach a tipping point to thermal perturbations) and determine what aspects of ecosystems are most and least resilient to extreme perturbation

The initiative was completed in FY 2025.

Neutron Data Interpretation Platform Ecosystem

The High Flux Isotope Reactor and Spallation Neutron Source provide world-class experimental capabilities to a wide range of scientific users. Interpretation of neutron data is necessary to convert measurements into scientific knowledge. As problems under investigation increase in complexity and neutron instrumentation provides larger volumes of data from more sophisticated experiments, simple approaches to data analysis are often insufficient, and advanced modeling and analysis methods are required to generate the most effective science from these neutron sources. The purpose of the Neutron Data Interpretation Platform Ecosystem Initiative is to push the boundary of advanced neutron data

analysis and modeling and lay the foundation for future developments, including autonomous or guided experiments. Projects in this initiative focus on three areas:

- **Real-time data analysis**—Develop tools and workflows to enable automated data analysis. This analysis will enable scientific content to be extracted rapidly from neutron data and will lay the foundation for automated or guided experiments.
- **High-performance computing materials modeling**—Incorporate computationally intensive materials modeling approaches into data interpretation workflows. Modular workflow implementation will enable the development of digital twins of neutron instruments, which will allow for virtual experiments to be performed and synthetic data to be generated.
- **AI/ML applications**—ML approaches have shown potential for advancing the understanding of complex datasets, such as those generated at the Spallation Neutron Source. These methods can enable capabilities such as feature identification, background and artifact suppression, materials and signal classification, and property prediction. New approaches and workflows will be developed to enable rapid, automated interpretation of neutron scattering data with the potential for experiment optimization or guiding.

The initiative was completed in FY 2025.

Nonequilibrium and Emergent Transients in Advanced and Soft Materials (NEAT)

Almost all systems in nature are influenced by nonequilibrium processes that span broad scales in space and time. However, most research has focused on investigating the equilibrium properties of matter. Developing an understanding of nonequilibrium processes presents unique opportunities for directing, manipulating, and controlling matter into states not accessible at equilibrium that are essential for the discovery of new technologies. At the same time, such processes pose deep questions for basic science that apply across disciplines, including synthesis science, biology, soft matter, quantum matter and advanced materials, and quantum information sciences. The overarching objective of the NEAT Initiative is to develop an understanding of nonequilibrium processes and ultimately to direct, manipulate, and control matter into states that are essential for discovering new technologies but are not accessible at equilibrium.

Nonproliferation Science Initiative (NSI)

The NSI advances fundamental scientific understanding and develops applications needed for the security of the nation. This initiative is structured around two crosscutting, foundational thrust areas in nuclear nonproliferation research: (1) pattern-of-life (PoL) and (2) signal detection in data with a low signal-to-noise ratio. The PoL thrust is focused on demonstrations of significant improvements to existing PoL analysis techniques and new types of PoL analyses of data obtained using emerging sensing techniques or algorithms, as well as new visualization and validation techniques for PoL analysis. The signal detection thrust focuses on the ability to detect or verify operational processes related to the nuclear fuel cycle. Research under this thrust focuses on advanced, ab initio computational models of large, complex, solid-state actinide systems; novel methods for extracting chemical and physical information connecting material properties to formation; and novel data analyses designed to search for new signatures.

The initiative was completed in FY 2025.

Self-Driven Experiments for Science/Interconnected Science Ecosystem (INTERSECT)

The objective of the Self-Driven Experiments for Science/INTERSECT Initiative is to develop the capabilities needed to automate entire scientific workflows and control them with AI/ML to create revolutionary efficiencies that will allow researchers to explore high-dimensional problems previously considered impossible and discover new, subtle correlations. INTERSECT focuses on two key areas of development: automation (incorporates tools that perform well-defined, systematic processes with limited

human intervention) and autonomy (smart decision-making techniques, such as AI and ML) of processes, experiments, and laboratories, including the following:

- Creating new automated methods and tools that incorporate on-the-fly data analysis to provide autonomous and rapid feedback that steers the course of a synthesis workflow
- Integrating automation with experimental synthesis techniques
- Developing multidomain interconnected laboratory capabilities
- Integrating AI-based approaches to designing and steering experiments for scientific discovery or for designing new systems, parts, materials, or processes
- Integrating reinforcement learning techniques to demonstrate autonomous control of instruments, facilities, or smart laboratories

Transform Initiative

The Transform Initiative seeks to advance the fundamental science of carbon dioxide capture or conversion processes and aims to translate that knowledge into scalable demonstration and the deployment of energy-efficient and cost-effective capture and conversion systems. The initiative is concentrated on advancing scalable technologies in critical sectors of the economy, such as power generation and chemicals manufacturing, that have large direct greenhouse gas emissions. The initiative also focuses on negative-emission technologies that, if deployed at scale, could offset emissions from sectors in which no plausible low-emissions approaches exist. Techno-economic and life cycle analyses to demonstrate a viable path to scalable deployment are integral to the initiative.

Validated Design and Evaluation of Fusion Wall Components

The Validated Design and Evaluation of Fusion Wall Components Initiative combines fusion and fission expertise to accelerate the deployment of fusion energy. The initiative focuses on integrated fusion component modeling that addresses the question of how scientists can predict and control the phenomena around a fusion wall component's plasma–material interface. Work in this initiative is directed toward the need for well-verified, well-validated modeling and simulation, which includes uncertainty quantification for fusion components. Its outcomes are intended to be the foundation for eventual deployment of a pilot plant as well as to facilitate the intermediate step of validating the critical technology for the pilot plant. The ultimate goal of this initiative is to develop comprehensive models and sensor technology integrated into efficient workflows, which would allow validation testing at facilities such as the Material Plasma Exposure eXperiment—a fusion prototypic neutron source and a component high–heat flux facility. Projects in this initiative fit within four thrust areas focused on answering the fundamental science question of how to predict and control the phenomena around a fusion wall component's plasma–material interface, including the following:

- **Thermal management at extremes**—Investigate and develop new methods for thermal management that incorporate real-time feedback control of a fusion component.
- **Plasma–material interactions**—Develop a further understanding of material erosion, migration, and plasma–material interactions.
- **Fusion neutron effects and damage**—Build a foundational basis for blanket technology and new plasma-facing and structural materials that can sustain 14 MeV neutrons.
- **Computation**—Introduce computational approaches for integrated performance and response of the plasma–material interface that simultaneously incorporate all three of these science areas.

The initiative was completed in FY 2025.

Vertex: Advanced Multiphysics Simulations for Core Applications

The Vertex Initiative aims to establish a new computational and data science–ready framework that simulates coupled multiscale and multiphysics phenomena in core mission spaces. An advanced

predictive capability is being created in which new individual and coupled physics are developed to meet science-driven needs and are demonstrated on problems residing in the areas of (1) next-generation enrichment science for applications, including the plasma separation process, gas centrifuge, and electromagnetic isotope separation; (2) fusion science, including blanket design and the evaluation of novel fusion concepts; and (3) advanced scientific computing workflows integrating data-based models.

Seed Program

The Seed Program supports innovative ideas that have the potential to enhance the laboratory's core scientific and technical capabilities. It also provides a path for funding new approaches that fall within the distinctive capabilities of ORNL but outside the more focused research priorities of the major laboratory initiatives.

Strategic Hire Program

The Strategic Hire Program was formed to add critical skills to the laboratory by supporting individuals whose research is aligned with ORNL's strategic needs and have the potential to advance key national priorities. Candidates for strategic hires are expected to be established investigators who are well-qualified for leading research programs, capable of developing substantial programs, and able to take organizational leadership roles. The program also supports research of promising individuals at earlier stages of their careers who bring expertise critical to major ORNL strategic initiatives.

Distinguished Staff Fellowships

The Distinguished Staff Fellowship programs were formed to provide research opportunities for exceptional early career scientists in honor of Dr. Eugene Wigner, the first director of R&D at ORNL; Dr. Alvin Weinberg, a former ORNL director; and Dr. Liane Russell, ORNL's groundbreaking geneticist. The appointment of fellows at ORNL provides an opportunity for outstanding scientists and engineers in life, physical, computer, computational, and social sciences to pursue research in areas related to national energy, environment, and security problems and interests. Awardees of Distinguished Staff Fellowships receive funding for their research projects through the LDRD Program.

Early Career Competition

The LDRD Early Career Competition cultivates the research and professional development of high-potential early career research staff. The competition is open annually, with typically one awardee selected from each of ORNL's eight research directorates.

UT-ORII Fellowships

Researchers hired as UT-ORII Fellows receive LDRD Program support to conduct research related to the objectives of one of the UT-ORII Convergent Research Initiatives (CRIs). The CRIs bring together the expertise and infrastructure of ORNL and the University of Tennessee System to advance key national objectives in areas such as radiopharmaceutical therapies, efficient transportation, and affordable building construction.

Report Organization

This report, which provides a summary of all LDRD Program projects that concluded in FY 2025, is divided by subprogram into six sections:

- Director's R&D
- Seed
- Strategic Hire
- Distinguished Staff Fellowships

- Early Career Competition
- UT-ORII Fellowships

The Director's R&D Program projects are categorized according to the programs' initiatives. Seed projects, Strategic Hire projects, and Early Career Competition projects are categorized by the research division of the principal investigator. The Distinguished Staff Fellowship projects are grouped by fellowship. UT-ORII Fellowship projects are grouped together in a single section. The summaries are arranged within each section by project number. Each completed project summary contains

- the project description,
- a brief description of the project's relevance to DOE and ORNL missions, and
- results and accomplishments.

SUMMARIES OF PROJECTS SUPPORTED BY THE DIRECTOR'S R&D PROGRAM

ACCELERATING RADIOTHERAPEUTICS THROUGH ADVANCED MOLECULAR CONSTRUCTS (ARM)

11410: Multimodal Analysis of Nuclear Medicine by Cryogenic Electron Microscopy

A. Williams, A. Borisevich, D. Mukherjee

Project Description

This project focused on developing and applying advanced multimodal cryogenic electron microscopy (cryoEM) techniques to understand the structure, chemistry, and targeting behavior of radiolabeled nanoparticle systems used in nuclear medicine. Radiopharmaceutical therapies depend on efficient delivery of radioactive payloads to specific cellular targets. However, limitations in payload loading, off-target effects, and an incomplete understanding of nanoparticle–receptor interactions hinder therapeutic efficacy and dose optimization. This project addressed these challenges by providing molecular and nanoscale insight into radioisotope packaging, nanoparticle structural dynamics, and receptor-specific targeting mechanisms.

This project integrated cryogenic transmission electron microscopy, cryogenic scanning transmission electron microscopy, and cryogenic electron energy loss spectroscopy to enable correlated structural and chemical characterization of hydrated, beam-sensitive nanoparticle systems. These techniques were applied to lipid, polymeric, and inorganic nanoparticle delivery vehicles to quantify morphology, heterogeneity, membrane structure, and radioisotope localization. Multimodal imaging and spectroscopy were used to determine how nanoparticle composition and architecture influence isotope loading efficiency and stability. In parallel, the project investigated targeted delivery strategies by attaching synthetic nanobodies to nanoparticle surfaces and resolving nanoparticle morphology and isotope loading capacity in these functionalized systems.

Overall, this project established new experimental workflows, data integration strategies, and cryoEM methodologies that bridge materials science and structural biology. The project also produced automated high-performance computing data analysis software, CryoBlob, which enables autonomous, statistically robust, reproducible, large-scale analysis of heterogeneous nanoparticle morphology. Together, these capabilities support comprehensive characterization of complex radiopharmaceutical delivery systems.

Mission Relevance

This project advanced DOE's mission by applying transformative science and technology to challenges in human health, nuclear medicine, and national scientific infrastructure. Specifically, the project developed advanced multimodal cryoEM methods to improve the understanding of radiopharmaceutical delivery systems—an area limited by dose efficiency, off-target effects, and incomplete mechanistic knowledge. By elucidating how radioisotopes are packaged, transported, and targeted at the nanoscale, the project contributed to improving the safety and effectiveness of radiotherapeutics, which supports national health security and US leadership in nuclear technologies.

The project also aligned with ORNL's mission by delivering new scientific capabilities and technical breakthroughs in multimodal electron microscopy, spectroscopy, large-scale data synthesis, and automated analysis. By leveraging ORNL's unique strengths in these areas, the project established new experimental pipelines that span materials science and biological systems. These developments accelerate the deployment of advanced characterization tools for isotope science relevant to energy, biomedical applications, and global security.

Results and Accomplishments

The project successfully defined and implemented experimental workflows that integrate low-dose cryoEM, cryogenic scanning transmission electron microscopy, and cryogenic electron energy loss spectroscopy for functionalized nanoparticles. These workflows enabled correlated structural and chemical characterization of hydrated, beam-sensitive nanoparticle systems across length scales ranging from subnanometer to micrometer resolution.

The project developed and applied cryoEM-based methods to characterize nanoparticle morphology, heterogeneity, and isotope packaging strategies, along with the development of the automated CryoBlob analysis software. Using low-dose cryogenic scanning transmission electron microscopy, the project demonstrated the ability to measure nanoparticle size distributions, membrane thickness, internal cargo localization, and structural variability in lipid, polymeric, and inorganic delivery vehicles. Complementary cryogenic scanning transmission electron microscopy and cryogenic electron energy loss spectroscopy measurements provided chemical sensitivity to radioisotope localization (interior vs. surface-associated) and relative loading capacity within nanoparticle systems. Analysis of nanoparticles with functionalized nanobodies bound to the surface allowed for an increase in understanding specific cellular targeting abilities. These results established an experimental basis for evaluating how nanoparticle composition and structure influence isotope loading efficiency, delivery, and stability.

In addition to scientific results, the project produced technical and methodological outcomes, including the generation of large, multimodal microscopy datasets; development of data handling and analysis workflows suitable for terabyte-scale cryoEM experiments; development of software; and integration of structural and spectroscopic information into unified interpretations of nanoparticle behavior. The resulting protocols, low-dose imaging strategies, and multimodal analysis approaches are broadly applicable to other biological and materials systems at ORNL and establish a foundation for future cryogenic characterization of complex, heterogeneous nanostructures relevant to nuclear medicine and beyond.

11500: Advancing Biomineralization-Inspired Chelators for Targeted Radionuclide Therapy

C. Parker, I. Popovs

Project Description

Targeted alpha therapy (TAT) is one method proposed to fight certain cancers. TAT involves an alpha-emitter (usually an actinide) chelated by some biological molecule that makes the alpha-emitter deliverable to biological systems. By complexing rare f-block elements with biologically inspired chelators to better characterize overall f-electron behaviors, researchers can more richly understand the behavior of TAT-relevant elements. Chelating the actinides is a vital first step in understanding the overall chemistry of a TAT system and modulating the chemistry to become compatible with biological delivery mechanisms without disrupting alpha-emitter effectiveness. This study examined periodic trends across the lanthanides, including promethium, and actinides, including americium, curium, and californium, with a terpyridine biscarboxylate (TBC) ligand. This is the second ligand system to include a promethium structure and the first to include one that is isotypic with actinides beyond curium. Changes in bonding behavior across f-elements resulted in three different primary structures, which highlighted subtle differences between central metal behaviors. Primary findings included subtle coordination differences between europium, curium, and americium; changes in metal positioning within the chelating cavity across the series; and competition between auxiliary ligands for the remaining four coordination sites on the metal center after chelation. This work offered insights into periodic behavior of trivalent f-elements and augmented baseline data for crystallographic structural data for these rare elements. Molecular structures of TBC with these elements provided insights into chelation processes in lanthanides and actinides.

Mission Relevance

This work addressed the DOE and ORNL missions of supporting energy and national security, as well as economic growth, by obtaining data related to a promising cancer treatment via probing key fundamental properties of f-elements in the solid state. This work was also relevant because of its applications to nuclear energy. Specifically, the work supported by this funding provided insights into the entire lanthanide series as the second published study to include promethium crystallographic structures, as well as comparisons of promethium to americium, curium, and californium. This study created the framework for richer molecular modeling that will allow researchers to access TAT via biologically inspired chelators and the f-elements.

Results and Accomplishments

This work included ligand synthesis by Ilja Popovs (former principal investigator, now at the University of Tennessee, Knoxville) and crystallographic experiments (lanthanide and actinide) by Connor Parker.

The major findings of this work include three primary focal points:

- Coordination differences occurred between europium, americium, and curium—an important trivalent f-block separation.
- Metal centers were offset within the series, and later metals exhibited a better fit into the chelating cavity.
- Heteroleptic coordination demonstrated a competitive ligand environment around the metal center.

Coordination differences between europium, americium, and curium are important because subtle differences between these elements apply to fundamental f-block studies and nuclear fuel cycle separations. These bonding differences also provide information for studies on f-orbital covalency. Differences in the bonding indicate less covalency in curium, whose half-filled f^7 configuration provides stability against enhancements of covalent interactions. Structurally, the metal centers exhibited a better fit into the chelating cavity later in the series, which indicated a better fit within the terpyridine ligand and helped inform models that apply biological molecules to TAT-relevant metal centers. A heteroleptic system is one in which a metal center has multiple different ligand types bonded to it; in the case of this work, this environment included the pentadentate TBC ligand synthesized during this project, water, nitrate, and bridged carboxylates from neighboring TBC molecules. The heteroleptic environment coordinating each metal in the series was an analogue for the competition that could be experienced by TAT molecules in biologic systems and provided reference points to ground computational studies about the competitive nature of complex biochemical systems with respect to f-elements.

Major operational successes stemming from this project included successful crystallization of the TBC ligand with highly radioactive materials (promethium-147, americium-243, curium-248, and californium-249). This work solidified ORNL's capability to study these radioisotopes in the solid state safely and effectively using an established workflow and researcher expertise to make these key measurements. This project also demonstrated successful chelation of actinides with biologically inspired molecules, which is a crucial step in designing TAT systems that can deliver treatment using biological transport and efficiently and safely maintain strong chelation with f-block radioisotopes.

ARTIFICIAL INTELLIGENCE (AI)

11049: Automating Cyber Defense Using Reinforcement Learning

T. Oesch, A. Chaulagain, P. Austria

Project Description

Current cyber defense technologies insufficiently protect against future automated adversaries. Given the increase in attacks targeting operational technology that controls critical infrastructure, it is important to realize cognitive cyber capabilities to stay ahead of attackers and ensure the security of the nation. Reinforcement learning develops the optimal policy for attack/defense in situations in which multiple actions must be taken before a reward signal is received. By using reinforcement learning, researchers can develop an agent that anticipates an adversary's future actions and develop an automated responses to an adversarial campaign. This project advanced the state of the art in automating cyber defense using reinforcement learning agents trained in a custom-built, realistic network test bed environment called Cyberwheel.

Mission Relevance

This work advanced the DOE mission of national security by advancing the state of the art in autonomous cyber defense to protect critical infrastructure, including energy infrastructure, as attacks increase in speed and scale.

Results and Accomplishments

This project developed a high-fidelity network simulation and emulation environment to train autonomous cyber defense agents. Using this agent training environment, this project showed that AI agents can autonomously deploy decoys in unseen network topologies that delay or deny adversarial actors. The project also showed, via the emulation environment, that these agents can be transferred from simulation environments to more realistic emulation environments.

Cyberwheel is built on robust network definition code and configuration files that enable rapid experimentation across topology sizes and configurations. The emulation environment uses a modified version of the Firewheel test bed for emulation, which is an experimental platform built on QEMU that enables experimental reproducibility and instantiation of networks with over 100,000 virtual nodes. The goal of this training environment is to aid in developing autonomous defense agents deployable in live networks that are adaptable to a variety of (1) network sizes and topologies and (2) adversary strategies.

In this study's environment, a scenario configuration file defined the network in which the experiment took place, and that scenario was used to generate a simulated and an emulated network. The reinforcement learning agent then used observations of the environment produced by the observation converter to take actions in the environment, such as deploying a decoy to help detect the adversary and delay or deny their attack.

In this study, the runtime of this simulator increased linearly with the number of hosts, whereas many existing simulators performed poorly with more than 100 hosts in this testing. The red agent attacked a decoy deployed by the blue reinforcement learning agent instead of the actual server, effectively denying the attacker from exploiting the network.

11053: Next-Generation Security for Interconnected Systems

P. Moriano Salazar, R. Bridges, S. Hollifield, M. Li, S. Hespeler

Project Description

This project aimed to build a more effective way to defend interconnected systems from evolving cyberattacks. The goal was to move beyond static intrusion detection and develop algorithms that can

learn and adapt as threats change. Working with ORNL's Vehicle Security Laboratory, this project created a framework that processes control- and data-plane signals in real time. By combining time-series analysis, clustering, and graph mining, the system can detect abnormal behaviors early and adjust to new attack patterns as they emerge. Over the course of the project, the team built and tested a modular data pipeline, developed adaptive learning methods that reduce false alarms, and generated new datasets by injecting controlled faults and attacks. The final system showed that adaptive, self-learning defenses are feasible and effective in realistic test beds.

Mission Relevance

This project supported DOE's mission to ensure US security and prosperity through transformative science and technology. As cyber threats grow more sophisticated, the safety and stability of the nation's energy and transportation systems depend on intelligent defenses that can adapt in real time. This project developed algorithms that learn from data and respond dynamically to evolving attacks. These capabilities enhance the resilience of cyber-physical infrastructures that are central to the nation's energy and scientific enterprise. By improving the reliability and trustworthiness of these interconnected systems, the project contributed to DOE's broader goal of safeguarding critical assets that underpin economic growth and national security.

Results and Accomplishments

This project created a new computational framework for adaptive cybersecurity in interconnected systems. The team designed algorithms that detect and respond to malicious activity in real time by combining multivariate time-series analysis, clustering, and graph-based learning. The approach allowed the system to adjust automatically as new threats appeared, which addressed the long-standing challenge of static, nonadaptive intrusion detection. A scientific advance made possible by this project was the ability to distinguish system faults from genuine cyberattacks, which improved the reliability of safety-critical systems such as vehicles and industrial control networks.

This project implemented a modular data acquisition and analytics pipeline that streams control- and data-plane telemetry from ORNL's Vehicle Security Laboratory. The project benchmarked leading algorithms for near–real-time detection and developed adaptive methods that integrate ensemble learning, signal clustering, and temporal cross validation. These algorithms reduced false alarm rates and shortened detection latency relative to standard approaches. The project team also executed a comprehensive analysis of how temporal validation techniques influence anomaly detection outcomes in multivariate time-series data.

The project demonstrated that adaptive, data-driven methods can make cyber-physical systems more resilient by detecting attacks earlier and reducing false positives. The combination of real-time analytics, unsupervised learning, and test bed–driven experimentation created a foundation for future self-learning defense systems that evolve alongside modern interconnected technologies.

11450: Design of Resilient Internet of Things/Edge with Federated Tiny Machine Learning

M. Li, P. Laiu, J. Nichols, M. Huettel, I. Sikkema, M. Mather

Project Description

Internet of Things (IoT) and edge computing (EC) create new cybersecurity threats. IoT/EC device connectivity means their vulnerabilities put the wider system at risk, and these devices provide avenues for major cyberattacks. There is increasing IoT/EC usage in US energy systems (e.g., smart grid, distributed energy resources, smart home, and radioactive/nuclear material transport). Current IoT/EC devices lack cognitive capabilities to analyze data, recognize impending attacks, and adaptively respond. The inability to share and enhance knowledge across system-of-systems also leads to ineffective

cybersecurity postures. This project applied AI/machine learning (ML) to equip IoT/EC devices with the necessary knowledge to intelligently detect cybersecurity attacks and adaptively respond. The project also explored federated learning (FL) at the cloud level for cross-system model sharing and enhancement. FL has the potential to share local detection knowledge across system-of-systems while preserving IoT/EC edge data privacy. This project built physical and virtual IoT test beds by expanding the ORNL Cyber Operation Research Range emulator to include IoT/EC simulation. These test beds facilitated data collection under simulated IoT attacks and enabled Design of Resilient IoT/Edge with Federated TinyML (DRIFT) prototype validation and performance evaluation.

Mission Relevance

IoT cybersecurity is part of the DOE mission to secure IoT/EC-capable critical infrastructure operations such as the smart grid and distributed energy resources. The DRIFT framework provides greater cognitive cyber detection and better adaptive response. This advancement supports DOE missions by driving transformative science and technology solutions that safeguard critical infrastructure.

Results and Accomplishments

This project developed and validated DRIFT architecture, which integrated tiny machine learning (tinyML) and FL technologies to address cybersecurity challenges in resource-constrained IoT/EC environments. The DRIFT framework targeted modern IoT/EC ecosystems with growing attack surfaces and vulnerability threats because these technologies have become increasingly integrated into critical infrastructure systems such as smart grids, healthcare, and smart homes.

The DRIFT architecture introduced a novel and scalable three-layer design to enhance the resilience of IoT/EC systems against cyberattacks. The architecture combined cognitive capabilities at the IoT layer, collaborative learning at the edge layer, and global intelligence aggregation at the cloud layer. IoT devices at the IoT/microcontroller unit layer deployed tinyML models to monitor and analyze network traffic locally to detect anomalies and intrusions. These lightweight models ensured on-device intelligence for threat recognition and minimized detection latency. Edge devices acted as FL clients by aggregating data from multiple IoT devices within their local environments and sending refined models up to the cloud for global aggregation. TinyML models were trained in edge nodes to offload computationally intensive tasks from resource-constrained IoT/microcontroller units. The edge layer can also implement adaptive countermeasures. In the cloud/enterprise layer, the FL server aggregated knowledge and insights gained from edge training across distributed IoT environments to create enhanced global models. These models were redistributed back to IoT devices, which allowed collaborative learning across heterogeneous, geographically varied setups without exposing raw data or violating privacy constraints.

Compared with conventional two-layer FL architectures, DRIFT introduced hierarchical coordination between IoT, edge, and cloud layers. This structure alleviated IoT computation bottlenecks, minimized communication overhead, accelerated model convergence, and enhanced system-wide knowledge sharing, detection accuracy, and scalability.

This project created optimized tinyML models capable of running efficiently on resource-constrained IoT devices (e.g., ESP32 microcontrollers) while maintaining high detection accuracy for sophisticated cyberattacks such as Mirai and Bashlite botnets. The DRIFT framework demonstrated that FL improved detection accuracy and preserved local data privacy across heterogeneous IoT ecosystems. The project developed a federated preprocessing pipeline to aggregate data distributions from local edge devices into a global quantile transformer. This global preprocessor mitigated the risk of overfitting into local distributions. Feature selection across edge nodes produced intersected feature sets, which harmonized data for consistent FL model training. This process ensured robust learning even under strong data heterogeneity between clients. Additionally, FL-trained global models consistently outperformed locally trained models in terms of attack detection accuracy. Superior detection was observed across multiple traffic classes, including unseen attacks, which validated FL's ability to aggregate and share knowledge.

The federated global models also maintained high accuracy regardless of preprocessing strategies (e.g., feature selection methods) or post-training techniques (e.g., quantization). Model quantization reduced memory usage without degrading classification performance, which made FL models practical for deployment in low-energy IoT systems.

Two test environments—physical and virtual test beds—were constructed to validate the DRIFT framework under realistic and controlled settings. The physical test bed consisted of three isolated networks: a home network for normal IoT operations, an attack network simulating botnet propagation (including Mirai, slow-and-low denial-of-service, and Command and Control attacks), and a monitoring network for capturing traffic through port mirroring. Legacy malware-prone devices, such as cameras and hubs, were used to generate datasets. The test bed supported iterative refinement of intrusion detection pipelines, which provided real-world validation for DRIFT's capability in enabling iterative refinement of intrusion detection systems. The virtual test bed featured replicas of Linux and Windows IoT devices to test scalability and heterogeneity in larger setups. Traffic data from the physical test bed were replayed to emulate attack scenarios on enterprise-scale networks within ORNL's Cyber Operations Research Range. Challenges such as firmware emulation gaps were overcome by reusing collected traffic data to effectively approximate real-world attacks, which supported further validation of DRIFT's workflows in a variety of environments.

11481: Assurance of Artificial Intelligence for Science Applications

D. Lu, X. Wang, A. Tsaris, J. Choi, M. Fan, D. Wang, X. Shi

Project Description

This project advanced AI-driven Earth system modeling to improve predictive accuracy, interpretability, and computational efficiency, as well as advance Earth sciences and support DOE's energy mission. The team developed efficient and reliable explainable AI and uncertainty quantification methods and applied them to enhance hydropower inflow forecasting, improve atmospheric vapor pressure prediction, analyze hydrometeorological influences on reservoir releases to inform hydropower operations, and advance Earth system model calibration. Additionally, the project developed an AI foundation model with demonstrated capabilities in weather forecasting, downscaling, and land model acceleration. The model accurately predicted atmospheric conditions up to 30 days in advance, generated high-resolution global weather data, and accelerated land model simulations by up to tenfold.

Mission Relevance

This project advanced DOE and ORNL missions by developing AI technologies that accelerate scientific discovery and improve the nation's ability to address complex energy and environmental challenges. Through explainable, scalable, and foundation AI models, the project enhanced predictive understanding of the Earth system, which enables more accurate and efficient forecasts of weather, water availability, and land processes that are important to the reliability and resilience of the US energy infrastructure. These innovations can reduce computational costs, improve model interpretability, and strengthen decision-making for Earth, human, and energy systems, which support DOE's mission to ensure national security and prosperity through transformative science and ORNL's mission to deliver breakthroughs in energy abundance and global security.

Results and Accomplishments

This project advanced AI-driven Earth system modeling to support DOE's mission of ensuring abundant, reliable, and secure energy. The team developed explainable AI and uncertainty quantification methods that improved hydropower inflow forecasting, flood prediction, reservoir operations analysis, and Earth system modeling. An explainable probabilistic framework for reservoir inflow forecasting was developed to demonstrate robust performance across multiple US reservoirs. An encoder–decoder long short-term memory model achieved high predictive skill for inflow forecasts with lead times of up to 30 days, and

the uncertainty quantification framework reliably characterized forecast uncertainty consistent with the project's confidence intervals and predictive errors. Explainability analysis identified key hydrometeorological drivers, which enhanced trust and operational relevance for hydropower decision-making.

The project also delivered an explainable long short-term memory–based framework for predicting fluvial floods, storm surges, and low-frequency surge events and captured seasonal basin-scale hydrologic dynamics. The model demonstrated strong streamflow prediction skill and accurately classified flood types, achieving 92% accuracy for fluvial floods, 93% for storm surges, and 80% for low-frequency surges. It also produced realistic flood inundation maps for 10-, 20-, 50-, and 100-year return periods and had spatial patterns consistent with observations. Analysis revealed precipitation and snowmelt runoff as dominant flood drivers in the Susquehanna River Basin and Delaware River Basin in the mid-Atlantic.

To enable large-scale hydrologic prediction, this project developed a hybrid deep learning framework that integrated Vision Transformer–derived spatial features from remote sensing with geomorphological and meteorological inputs in a long short-term memory architecture. This approach enabled accurate streamflow forecasting across the United States with lead times of up to 30 days and consistently outperformed the process-based Variable Infiltration Capacity model. This result represents a significant advance in scalable, data-driven hydrologic forecasting.

Additionally, the project developed an AI foundation model with capabilities in extended-range weather forecasting, hyper-resolution data generation, and accelerated land surface modeling. The model predicts atmospheric conditions up to 30 days ahead, generates global high-resolution weather fields, and accelerates land model simulations by up to 10 times, which enable nationwide, 30-day streamflow forecasting. The model was successfully scaled on the Frontier supercomputer at the Oak Ridge Leadership Computing Facility and achieved up to 98% strong scaling efficiency across 65,536 GPUs.

Finally, the foundation model was adapted to address region-specific energy and water challenges, including groundwater mapping in the Pacific West to support drought resilience and bioenergy planning, improved flood prediction in the mid-Atlantic to enhance infrastructure reliability, and reservoir inflow forecasting in the Southeast to optimize hydropower operations and support grid stability.

11740: Artificial Intelligence Models Vulnerability Research at Scale

E. Begoli, P. Balaprakash, A. Sadovnik, J. Hutchins, J. Arnold

[The results of this project have been determined to be Controlled Unclassified Information. Therefore, additional information protection and distribution restrictions have been applied.]

11960: Credible, Trustworthy, Automated Materials Discovery Beyond Inorganics

S. Roy, A. Sedova, P. Kent, M. Coletti, C. A. Bridges, M. Ryder, I. Popovs, C. Engelmann, S. Bagchi

Project Description

This project aimed to build a connected generative deep learning computational foundation for automated materials discovery. The project built a framework to explore denoising diffusion generative deep learning methods for structure prediction and used high-performance computing to couple the generative methods with simulations of dynamics and disorder using deep neural network interatomic potential foundation models. The framework also focused on refining deep neural network interatomic potentials and structures via transfer learning approaches on high-accuracy data to accurately treat noncovalent interactions in organics. The project team built the platform for validation pipelines for vibrational dynamics simulations with data from experiments at ORNL's Spallation Neutron Source. This work contributed to supporting a full range of material types, including polymers and metal–organic and other

types of frameworks, as well as creating connected information pipelines between the Spallation Neutron Source, the Autonomous Chemistry Laboratory, and ORNL's leadership computing facilities.

Mission Relevance

This research addressed DOE's mission to deliver scientific and technical breakthroughs in energy security and energy abundance by combining basic energy science with deep learning and automation of energy-efficient material discovery.

Results and Accomplishments

Using ORNL's supercomputers, this project explored the utility of MatterGen, a denoising diffusion probabilistic model for crystal structure prediction developed by Microsoft Research. The model operates by diffusing fractional coordinates, lattice vectors, and atomic labels and learns to reconstruct crystal structures through denoising. Pretrained models are available that can generate stable or metastable structures with conditional properties, such as magnetism. Developing new models, however, requires training on datasets such as MP20 (a smaller dataset containing nearly 40,000 structures) and Alex-MP20 (a larger dataset containing nearly 600,000 structures). Initial testing indicated that when using eight H100 GPUs with 80 GB of RAM each, training the smaller dataset takes approximately 2.5 h, and training the larger dataset requires approximately 1 week. The project also found that tuning the learning rate can help strike a balance between accuracy and training time.

Along with exploration of MatterGen, this project developed a workflow for (1) running DeePMD-kit/MACE with machine learning potentials with fine-tuning capability, (2) performing disorder and vibrational simulations, and (3) fine-tuning data generation, testing, and validation.

One of the specific systems studied in this project was beryllium fluoride. In simulations, the project team introduced thermal disorder by melting a spherical region around an arbitrarily chosen atom and slowly cooling the system down to 0 K using molecular dynamics simulations. This process was carried out by varying the radius of the spherical region to access smaller and larger disordered regions. Notably, for smaller radii—namely, 5 Å or less—simulations suggest that the fully crystalline beryllium fluoride structure can be reconstructed through this heating-and-cooling approach. However, for larger radii, a mixture of crystalline and amorphous phases was formed. By computing pair correlation functions for different extents of disorder, this project developed a Bayesian ridge regression model capable of predicting the pair correlation function for a given structure.

DEPUTY DIRECTOR DISCRETIONARY

10900: Ion Trap Artificial Intelligence/Machine Learning Quantum Performance and Routine Optimization

C. Seck, G. Buchs, J. Cui, T. Naughton III, I.-S. Suh, A. Shehata, C. Winstead, A. Van De Brook

Project Description

With the onset of quantum computers that can run small algorithms, domain scientists have begun testing the efficacy of trapped ion quantum devices to deliver useful scientific results. This increase in research has driven strong demand for quantum computers and simulation devices to support quantum computer programming stack development, algorithm development, benchmarking, simulation, and general quantum computer science. However, existing on-site quantum computing/simulation capability is limited in several aspects. This project addressed the limitation in control electronics by identifying, evaluating, and procuring commercially available experimental control hardware from Riverlane for ORNL's on-site trapped ion quantum simulator/computer systems. This project collaborated and contracted with the Georgia Tech Research Institute for vacuum system modifications and installation of a Sandia National Laboratories Phoenix ion trap chip to provide a working experimental system to aid in standing up the Riverlane experimental control hardware and assist with developing the necessary associated software environment. Further work funded by this project resulted in exploring hybrid computational workflows between classical high-performance computing and quantum computing hardware.

Mission Relevance

Ion trap quantum computing has the potential to benefit areas of research ranging from protein folding calculations and vaccine development (biology applications) to macromolecular properties calculations (materials science applications) and environmental research applications. This device also provides an analog simulation platform for quantum spin liquids.

Results and Accomplishments

This project enabled a pathway to expand on-site quantum resource capability. This project explored hybrid computational workflows between classical high-performance computing and quantum computing hardware.

11393: Next-Generation Microelectronics for Advanced Sensing and Computing

S. Randolph, I. Vlassiouk, D. Briggs

Project Description

Microelectronics comprise fundamental electronic components, including transistors, diodes, capacitors, and integrated circuits, which form the core of modern computing architectures such as microprocessors, graphics processors, and memory devices. The global disaggregation of microelectronics and integrated circuit manufacturing has elevated their availability and trustworthiness to critical issues for supply chain security. Concerns regarding the availability and veracity of components sourced from foreign suppliers continue to grow. At the same time, the escalating energy demands of conventional computing, coupled with exponential growth in data processing, storage, and transmission, have intensified the need for fundamentally new approaches to energy-efficient computation, including quantum and neuromorphic architectures. Addressing these challenges requires a multifaceted strategy that rethinks materials development, integration, packaging, and testing across the entire microelectronics life cycle. Realizing next-generation sensing and computing technologies demands workflows that are capable of harnessing emerging materials beyond conventional silicon-based and III–V semiconductors. These materials,

including soft materials, ultrathin films, and low-dimensional systems, offer transformative functionality but present substantial challenges related to scalability, process compatibility, and environmental stability.

This project strengthened microelectronics research by establishing a comprehensive, end-to-end workflow enabled by targeted infrastructure investments. This workflow supports the synthesis, integration, packaging, and analysis of microelectronics components based on integrated hybrid systems and their heterostructures. This project developed a clearer understanding of material transformations that govern environmental sensitivity, degradation, and process incompatibilities to enable scalable growth, patterning, property optimization, and reliable connectivity. This project focused on the integrated instrumentation suite that culminated in wafer-scale synthesis and study of hexagonal boron nitride, which is a transformative wide-bandgap dielectric with applications spanning corrosion protection, photovoltaic passivation, gate dielectrics, and optically active quantum devices.

Mission Relevance

This project supported DOE missions in energy security and economic growth. The development of electronic devices that enable low-energy computing beyond conventional von Neumann architectures, including quantum and neuromorphic systems, requires an advanced understanding of how to integrate a broader range of materials into devices and integrated circuits that exhibit predictable and stable performance over extended lifetimes and across varying environmental conditions. This project enables codesigning next-generation sensing and computing architectures based on advanced thin-film and 2D materials. The wafer-scale metal–organic chemical vapor deposition synthesis platform represents a cornerstone capability that enables scalable materials growth compatible with microelectronics manufacturing.

Results and Accomplishments

This project examined instrumentation that could be used to synthesize and analyze soft materials, ultrathin films, and low-dimensional systems. This project established a comprehensive, end-to-end workflow based on this instrumentation. This workflow supports the synthesis, integration, packaging, and analysis of microelectronics components based on integrated hybrid systems and their heterostructures. A 4 in. wafer scale hexagonal boron nitride process has been demonstrated using a nickel/iron metal catalyst foil.

11563: Design and Synthesis of Covalent Radiopharmaceutical Therapies

B. Sanders, S. Davern, J. Parks, C. Brantley, B. Manard, S. Szakas

Project Description

This project developed and demonstrated a new class of small-molecule targeted radiopharmaceutical therapy (RPT) constructs designed for precise intracellular delivery of therapeutic and diagnostic radionuclides. The central concept was to repurpose known covalent small-molecule inhibitors as targeting vectors to transport chelated radionuclides directly to intracellular protein targets, thereby decoupling cell killing from pharmacological pathway inhibition and overcoming limitations of conventional targeted radionuclide therapy. Using KRAS G12C as a model system, this project established a modular construct architecture consisting of a covalent inhibitor, linker, click chemistry functionality, metal chelator, and metal ion or radionuclide. Multiple constructs were synthesized using copper-catalyzed and copper-free click strategies and were successfully metalated with surrogate metals and the alpha-emitting radionuclide actinium-225. Radiolabeling efficiencies were high, and the chemistry was shown to be compatible with rapid, late-stage radionuclide incorporation. Biochemical and cellular studies demonstrated selective covalent binding of lead constructs to mutant KRAS G12C over wild-type protein, which supported intracellular target engagement. In vitro cell assays showed synergistic cancer cell killing for radiolabeled constructs compared with the inhibitor or radionuclide

alone, which suggested effective delivery and retention of the radionuclide payload within target cells. Analytical methods, including intact protein mass spectrometry, gel-based assays, and metal-specific detection techniques, were developed or adapted to verify construct binding and cellular uptake. Overall, the project delivered a proof of concept for molecular-precision internalizing RPTs and established a versatile platform that can be extended to other covalent inhibitors, intracellular targets, and therapeutic or diagnostic radionuclides.

Mission Relevance

This project advanced DOE's mission by strengthening the nation's capabilities in nuclear science, radiochemistry, and advanced molecular design for applications in human health and national security. This work leveraged investments in isotope production, radiochemistry infrastructure, and analytical science to develop new approaches for the controlled use of radionuclides in broad applications, including biotechnology.

Results and Accomplishments

This project successfully conceived, executed, and validated a new molecular-precision strategy for intracellularly targeted RPTs. Work progressed from conceptual design to chemical synthesis, radiochemistry, biochemical validation, cellular evaluation, and analytical method development, which resulted in a cohesive and experimentally supported platform. A series of novel RPT constructs was designed and synthesized using a modular architecture that integrated covalent small-molecule inhibitors, chemically tunable linkers, click chemistry motifs, and metal-chelating groups compatible with diagnostic and therapeutic radionuclides. Multiple linker lengths, chelators, and metal ions were evaluated, and copper-catalyzed and copper-free click strategies were successfully implemented.

The project achieved proof-of-principle milestones at protein and cellular levels. Biochemical experiments confirmed selective covalent binding of lead constructs to the target mutant protein targets over wild-type controls. This result established that the added linker–chelator–metal components did not prevent target engagement. Cellular studies further demonstrated that radiolabeled constructs produced synergistic cancer cell killing relative to either the covalent inhibitor or radionuclide alone. These results validated the central project hypothesis that intracellular delivery and retention of a radionuclide payload can drive cell lethality independent of classical pharmacological inhibition.

In parallel, the project delivered advances in analytical capability. Novel, metal-specific analytical workflows were developed and validated to directly measure covalent protein engagement and cellular uptake of radiopharmaceutical constructs in complex biological samples. These methods fulfilled key analytical milestones outlined in the project and provided enabling tools for future radionuclide and molecular construct research.

11681: Radioisotope Production at the Spallation Neutron Source

J. Griswold, Y. Lee, E. Asano, S. Kim, S. Cousineau, D. Rotsch, J. Duran, D. Stracener, R. Hoerres, P. Kankanamalage, C. Rasco, F. Pilat

Project Description

A unique opportunity exists to investigate alternative radioisotope production technologies using the high-energy proton beams available at the Spallation Neutron Source (SNS) at ORNL. At SNS, the Second Target Station is being built to address emerging science challenges in energy, security, and transportation. Recently, the SNS accelerator successfully completed the Proton Power Upgrade project to enhance the beam power capability to 2.8 MW and beam energy to 1.3 GeV. The accelerator is expected to ramp up to a full beam power of 2.8–3 MW within a few years. Because the First Target Station has a power limit of 2 MW, an additional 0.8–1 MW of beam power will be available until the Second Target Station becomes operational in the 2030s. After completion of the Second Target Station, 100–300 kW will remain accessible and can be used for alternative applications, such as radioisotope production,

without compromising the DOE Basic Energy Sciences neutron production mission. This project was conducted to investigate the possibility of radioisotope production at SNS through four main goals: (1) identify isotopes of interest through modeling and simulation of prospective irradiation parameters and target compositions, (2) develop a target design concept that can receive high-energy beam pulses from the SNS accelerator, (3) identify enhanced isotope separation methods for SNS-produced radionuclides, and (4) develop a design concept for an experimental or demonstration test stand.

Mission Relevance

Radioisotope production at SNS has the potential to leverage surplus high-energy proton capacity at SNS to establish a domestic, scalable source of critical medical radioisotopes, which would enhance national security and prosperity and advance nuclear science and technology. This program maximizes the use of DOE infrastructure without affecting the neutron mission and drives innovations in high-power targetry, modeling, and separations. This work strengthens global security by stabilizing isotope supply chains and creates economic opportunities through US leadership in radiopharmaceutical production.

Results and Accomplishments

The feasibility of producing high-demand radioisotopes has been investigated with modeling and simulation techniques. Using high-energy protons, several Monte Carlo particle transport codes—such as Monte Carlo N-Particle (MCNP), Geant4, FLUKA, and Particle and Heavy Ion Transport code System (PHITS)—were employed to calculate the yield of radioisotopes resulting from interactions with a thorium target. Thorium was selected as the prospective target material because of its ability to produce medically relevant isotopes such as actinium-225, which is a critical radioisotope for targeted alpha therapy treatments. Simulations revealed a broad range of potential isotopes that can be generated, although the predicted yields varied significantly depending on the nuclear model that was used.

A concept for a robust target assembly was designed to handle the intense energy deposition from high-power proton irradiation. The design centers on thin thorium discs encapsulated in stainless steel, with graphite layers included to absorb shock waves and mitigate material stress. The assembly was engineered to dissipate heat efficiently using a system of narrow cooling channels that ensure optimal water flow rates to prevent overheating and extend the target's lifespan. Simulations confirmed that this design can withstand the harsh thermal and mechanical conditions imposed by the proton beam, including prolonged irradiation cycles.

A wide assortment of radioisotopes will be generated through 1.3 GeV proton irradiation of thorium radioisotope production targets. Following irradiation, the targets must undergo chemical processing to separate the desired elements from the by-products. ORNL's existing facilities, such as the Radiochemical Engineering Development Center, and the planned Radioisotope Processing Facility are well-equipped to support these separation processes. Using innovative methods such as slurry resin techniques and peroxide precipitation, these processes demonstrated scalability for larger target masses and improved waste management efficiency in other projects. In addition to chemical separation of various elements, electromagnetic separation would be extremely useful in isolating radioisotopes, especially those deployed in medical applications that require high specific activity. Implementation of an offline radioactive electromagnetic isotope separator was identified as a critical technique to significantly increase the number of radioisotopes that could be viably produced and purified.

Developing an experimental test stand is vital to validating yield predictions and testing isotopic production techniques. A linac dump area at SNS was identified as a candidate location for a small-scale test stand. A test stand capable of handling a 1 kW proton beam irradiating thorium targets for approximately 48 hours was designed. Analyses were performed and concluded that the target can withstand this process without the need for active cooling. Shielding and particle transport simulations verified that the stand's operation is safe and has manageable radiation exposure levels during and after irradiation.

11780: Bringing a Digital Underground to the Advanced Plant Phenotyping Laboratory

L. York, J. Warren, J. Lagergren

Project Description

The Advanced Plant Phenotyping Laboratory (APPL) is a unique asset at ORNL that uses conveyors to move plants from a greenhouse through multimodal imaging stations to measure form and function. The system measures chlorophyll fluorescence, 3D plant structure with a laser scanner, spectral properties using hyperspectral cameras, plant morphology with color cameras, shoot water content with near-infrared (NIR) light, and water use using thermal imaging of leaves. However, the root system which has remained the so-called hidden half even though it supports plant water and nutrient uptake, serves as the interface with soil microorganisms, and is the main conduit for carbon flow from the atmosphere to soil carbon fractions. This project revealed this hidden half by developing novel belowground imaging capabilities for APPL and using these new capabilities to understand how root traits contribute to plant resilience. New rhizobox trays featured a thin, planar container with soil and a clear window covered to block light. These trays were used to grow plants such that roots grew along the window. Two new imaging stations were installed; the first removed the window cover and used an RGB camera with a white light source to image the roots along the window. The second used an NIR camera with two NIR light sources at 950 and 1,450 nm to image soil water content. In addition to these imaging upgrades, the APPL IT and computational infrastructure was upgraded to connect the system to the ORNL network for the first time and to allow real-time, AI-enabled analysis. Efficient AI algorithms (e.g., neural networks) were developed to identify roots, and the segmented images were analyzed in an updated version of RhizoVision Explorer[1] with headless mode and other enhancements to measure root traits. Experiments with bioenergy crops—sorghum and poplar—were conducted with well-watered and drought conditions to demonstrate the effectiveness of the methods while also developing novel analysis algorithms.

Mission Relevance

Ecosystem and feedstock system resilience is a major concern because of unpredictable weather and the decreasing availability of productive land caused by encroaching human habitation. Resilience is defined as the ability of plants to recover from disturbances, but the resilience of belowground organs (e.g., the root system) is rarely observed. This lack of research is largely because of a lack of applicable methodologies, as well as a general lack of knowledge of root dynamics that limits the ability to form and test relevant hypotheses. Therefore, technological advancements in nondestructive, time-series root imaging and analytics will greatly increase the understanding of ecosystem resilience. This understanding can be used to better understand root–microbe relationships in soil, which ecosystems will be most affected by drought, and how to breed for better bioenergy feedstocks using root traits, thus advancing DOE's energy and security missions.

Results and Accomplishments

This project worked with the APPL system vendor, Photon Systems Instruments, to design a unique imaging system meant to offer high-resolution imaging of the roots and lower-resolution imaging of soil water content. This system was procured and installed, new safety equipment was added, and human–machine interfaces were added to facilitate system use, such as when loading and unloading plants into the system.

New rhizobox trays were designed and installed that had a 30 × 40 cm rhizobox with 2 cm depth and a glass window on the front and a metal sliding cover. The window was pointed down at user-defined angles of 30° or 45°. Because roots are gravitropic, they tended to grow down along the glass while still maintaining general patterns of growth direction and branching patterns that could be imaged over time.

[1] *RhizoVision Explorer*, GitHub, https://github.com/predictivephenomics/RhizoVisionExplorer.

The rhizobox had three water tanks with tubing going to three different depths in the rhizobox to user-defined positioning of water based on programmatic control by the APPL scheduling software. When the rhizobox entered the new bay 6 belowground imaging chamber, a linear actuator slid the cover to reveal the window. A white light was turned on, and the RGB camera took a 20 megapixel image at approximately 10 pixels per millimeter. The camera was moved, and an NIR camera was placed in position to acquire one image with an LED light at 940 nm and a second image at 1,450 nm. All three images were stored in the APPL database for subsequent analysis. The rhizobox tray was closed and moved to the water chamber where nozzles are automatically positioned over the inlet tanks for each depth level of the rhizobox. Water is added to a defined target weight. This project leveraged this watering and weighing system to estimate changes in plant mass.

For analysis, RhizoVision Explorer was transitioned to the *cmake* compiling method for C++, and a headless or command line interface mode was created. Support was added for Linux OS, and a new memory allocation tool was used that made the software 40% faster on average. The length measurement was improved using a smoothing algorithm to account for the small pixel-level steps that inflate length, and additional traits for rhizoboxes were added. The tool can now be used in scripts on the APPL analysis server for high-throughput image analysis. However, RhizoVision Explorer expects high-contrast images and thus struggles natively with root-on-soil images such as those in rhizoboxes. Therefore, the project team began training AI models, starting with public data, to create a U-Net for segmentation.[2]

Two experiments have initial results. Contrasting species are being used: sorghum as a grass and poplar as a tree. Each study contains two types, two water levels (well-watered and drought) and five replicates. Initial results showed consistent differences in rooting patterns, differences between types and species, and depletion of water during drought that reduces root growth. Although the NIR camera was expected to highlight soil moisture, the data also suggested it can be used to quantify root water content directly, which may be a novel trait for plant breeding. The results indicated the system performed as expected, can resolve statistical differences in root and soil parameters, and has utility for validation of plant form and function.

11797: Pioneering Multiscale Biological and Environmental Solutions for a Sustainable Earth

M. Mayes, K. Carter, F. Santos, M. Shu, L. York

Project Description

The overarching goal of this project was to develop capabilities for a multiscale observation–experiment–modeling integration framework. This framework can be used to facilitate data transfer, analysis, simulation, and experimental steering. Across three platforms (the SoilCosm laboratory-scale microcosm system, the Advanced Plant Phenotyping Laboratory (APPL) greenhouse, and field-scale ecosystem warming), the team used this project as a test bed to standardize and integrate data–model simulation synthesis capabilities within a computational infrastructure to provide high-frequency observation and comparison with simulation. The project implemented a pilot-scale experimental approach quantifying genome-controlled carbon uptake in plants, and the resulting stability of that plant carbon within soils was assessed across laboratory and field studies. The approach integrated currently independent capabilities in cross-scale experimental systems and modeling to achieve results across a range of temporal and spatial scales in the laboratory and greenhouse, as well as in natural ecosystems.

Mission Relevance

This project sought to advance capabilities in plant–microbe interfaces, plant genetics, microbial engineering, and biodesign by building a flexible and multiscale observation–experiment–modeling

[2] J. Lager, *PRMI Root Segmentation*, GitHub, https://github.com/jlager/prmi-root-segmentation.

integration framework for integrated lab, greenhouse, and field experiments. The key outcome from this pilot study was the ability to link independent experimental systems into a common visual data ingest, model output infrastructure on a website, and thereby facilitate increased scientific understanding through near–real-time experimental steering, graphical representation, and agentic AI interpretation. These outcomes are key for designing the next generation of ecosystem, greenhouse, and lab experiments to address DOE's mission.

Results and Accomplishments

This project developed an experimental dashboard suitable for collecting and visualizing experimental results in near–real time, as well as supporting documentation of the architecture of the data flows. The dashboard automatically incorporates data streams from two major instruments—a LI-COR Environmental analyzer for carbon dioxide and METER Group Inc. sensors for soil temperature, moisture, and water potentials—and provides graphical depictions of experimental progress, downloadable files from the incoming data, and an AI agent for analyzing results. The dashboard is applicable to the lab and the field and is sufficiently flexible for use in other experimental platforms, such as the APPL greenhouse, and for slide imaging of microbial activities. The architecture underlying the data flows was documented throughout the project.

An experiment was performed to provide testing and application data. The experiment used leaves selected for their genotype prescription of carbon-to-nitrogen ratios. The leaves were applied to laboratory-scale soil incubation experiments using the SoilCosm system and to a field-scale soil warming experiment, both of which used LI-COR instruments and Meter sensors. In both lab and field experimental platforms, adding leaf litter increased carbon dioxide measured by the LI-COR, but the extent of the increase was affected by the soil water potential as measured by the Meter sensors. Transient changes in soil moisture strongly affected the response of the microbial community and resultant carbon dioxide signals, which underscored the need for near–real-time data ingestion, visualization, and AI-enabled interpretation because of the complexity of observations and the time frames over which changes occurred. Additionally, a soil carbon cycling model (DAYCENT) was incorporated into the cloud environment to facilitate analysis of lab and field experiments to enhance the model–experiment (Mod-Ex) capabilities of the dashboard.

This project investigated a key indicator of photosynthesis efficacy: the emission of volatile organic compounds such as isoprene. Isoprene and other volatile organic compounds are emitted by plants because of a variety of stressors and accompanying key processes, such as photosynthesis. These compounds are measured by proton transfer reaction time-of-flight mass spectrometry, and the massive datasets of mass-to-charge ratios are then modeled with a proprietary software to predict gas chemistry. The manual photosynthesis measurements, the data-intensive aspect of the volatile organic compound data streams, the proprietary software, and the need for connection to plant genotypes all represent significant bottlenecks in data analysis, modeling, and interpretation from the flagship APPL facility. This project developed a method to automate transfer of manual photosynthesis measurements to the cloud, link photosynthesis and volatile organic compound measurements, automate large portions of the data analysis of volatile organic compound measurements, and model how photosynthesis responds to isoprene production. Including measurements supplementary to APPL have enhanced the analysis of the role of plant genotypes on key photosynthetic parameters and processes, thereby enabling a greater understanding of plant functions and sensitivities.

12038: Development of High-Performance Low-Activation Nanoprecipitate-Strengthened Vanadium Alloys for Fusion First-Wall and Blanket Applications

Y. Yang, Q. Q. Ren, Y. Wang, X. Chen, Y. Katoh, Y. F. Su, K. Hanson, R. Dalton, C. Eakens, S. Zinkle

Project Description

Vanadium–chromium–titanium alloys have been identified as promising candidates for structural materials used in fusion reactor blanket systems because these alloys are capable of operating at temperatures higher than those achievable with reduced-activation ferritic-martensitic steels. However, the current leading alloy, V–4Cr–4Ti (V44), exhibits an upper operating temperature below 700°C because of its limited creep resistance, which significantly constrains its design space. Moreover, V44 alloys are prone to irradiation-induced hardening and embrittlement when exposed to irradiation below 400°C, primarily because of the absence of stable microstructural features that could suppress the formation of high-density defect clusters. This project aimed to address both limitations by developing nanoprecipitate-strengthened vanadium alloys with enhanced high-temperature strength and improved theoretical irradiation tolerance using integrated computational–experimental design. The newly developed alloy, V–4Cr–4.4Ti–0.1C–0.2Y (called V44C), demonstrated promising mechanical performance and microstructural stability before postirradiation evaluation. These advancements have the potential to expand the operational window of vanadium-based structural materials to support the deployment of high-performance liquid metal blanket concepts, such as lithium-cooled systems. Furthermore, the improved strength of the V44C alloy may enable the design of integral tungsten armors supported by vanadium structural substrates; these materials would benefit from the chemical compatibility and minimal hardness and modulus mismatch between the two materials, which are important factors for robust fusion component integration.

Mission Relevance

This project advanced the technical foundation for deploying vanadium alloys in a self-cooled lithium blanket system, which is a key enabling technology for future fusion energy systems. By addressing the scientific and engineering challenges associated with these materials, this effort aided in accelerating the commercialization of fusion energy through targeted research and technology development. These goals align with the DOE mission to ensure America's security and prosperity related to energy and nuclear challenges. This project also supported ORNL's mission to deliver scientific discoveries and technical innovations that speed the development and deployment of fusion energy to create new economic opportunities for the nation.

Results and Accomplishments

A new low-activation vanadium alloy, V44C, was successfully developed through an integrated computational–experimental approach. Using the in-house CALPHAD database, high-throughput thermodynamic calculations were conducted to explore the compositional space of vanadium-based alloys. The thermodynamic screening guided the identification of a promising composition, leading to the design of V–4Cr–4.4Ti–0.1C–0.2Y (wt %), designated as V44C. To optimize strengthening through nanoprecipitation, a two-step heat treatment route using solutionizing and precipitation aging was designed based on thermodynamic simulations. The predicted precipitation behavior was experimentally validated via lab-scale alloy fabrication and characterization.

A lab-scale prototype of the V44C alloy was fabricated using ingot metallurgy—specifically, arc melting followed by drop casting. The as-cast ingot underwent hot forging and rolling to achieve the desired microstructure, after which it was subjected to the designed heat-treatment sequence. Transmission electron microscopy analysis confirmed the formation of a high number density of nanoprecipitates, which verified the effectiveness of the thermodynamic design approach.

Mechanical testing revealed substantial improvements in high-temperature strength. Tensile tests at 700°C under argon cover gas demonstrated a yield strength of 439 MPa for V44C, compared with 241 MPa for the reference V44 alloy. Room-temperature testing also showed significant enhancement, with yield strengths of 675 MPa for V44C versus 526 MPa for V44.

Microstructural characterization using scanning electron microscopy, transmission electron microscopy, and electron backscatter diffraction indicated a homogeneous, equiaxed grain structure with an average grain size of approximately 71 μm, which reflected the success of the thermomechanical processing route. Charpy impact testing established a ductile-to-brittle transition temperature near room temperature for V44C that is significantly higher than the ductile-to-brittle transition temperature of V44 (−200°C). However, this newly optimized heat treatment suggested the ductile-to-brittle transition temperature can be improved, which makes this alloy suitable for targeted structural applications in fusion environments.

Further advanced characterization provided deeper insight into the strengthening mechanisms. Atom probe tomography revealed that the V44C alloy contains a high density of TiC nanoprecipitates, which contrasts with the oxygen-rich, low-density titanium–X precipitates observed in V44. Because of the neutron transparency of vanadium alloys, synchrotron x-ray diffraction was employed in place of small-angle neutron scattering to identify and quantify precipitate phases. This method supplied critical data for refining the thermodynamic database and predictive models.

Overall, this work demonstrated the successful development and validation of a novel vanadium alloy, V44C, as a candidate material for fusion blanket structural applications. The combination of computational alloy design, scalable ingot metallurgy processing, and detailed microstructural–mechanical evaluation established the feasibility of introducing strengthening nanoprecipitates into vanadium alloys. The alloy exhibited a marked improvement in high-temperature strength and a moderate reduction in ductility, and ongoing optimization efforts indicate that a balanced strength–toughness combination is achievable.

12040: Trapped Ion G-Factor Research of the Calcium Ion

C. Seck, B. McMahon, V. Sandhu, J. Gray, C. Herold, K. Brown, B. Sawyer

Project Description

Experimental precision measurements of the Landé g-factor g_J in atomic, molecular, and optical physics serve multiple purposes: testing and bounding relativistic quantum mechanics and higher-order quantum electrodynamics effects, assisting in improving multielectron correlation effect calculations, and aiding searches for time variation of the fundamental constants of physics. This project collaborated with the Georgia Tech Research Institute to measure the Landé g-factor ratio of the $^{40}Ca^{+}$ 3d $^{2}D_{5/2}$ and 4s $^{2}S_{1/2}$ electronic states through low- and high-magnetic field (approximately 10 and >9,000 gauss, respectively) spectroscopy of the transitions in these manifolds. This project improved upon state-of-the-art precision by roughly three orders of magnitude by using microwave and radio frequency spectroscopy.

Mission Relevance

This project applied to DOE missions of advancing state-of-the-art experimental science to advance fundamental and quantum physics. Improving g-factor measurements is relevant to the fields of atomic, molecular, and optical physics to test and bound relativistic quantum mechanics and higher-order quantum electrodynamics effects, assist in improving multielectron correlation effect calculations, and aid searches for time variation of the fundamental constants of physics.

Results and Accomplishments

This project performed precision measurements of the ratio of Landé g-factors between the 3d $^{2}D_{5/2}$ and the 4s $^{2}S_{1/2}$ electronic states of a single trapped $^{40}Ca^{+}$ ion using novel atomic spectroscopy schemes. The measurements were performed in two distinct ion trap apparatuses: a low–magnetic field cryogenic

surface electrode radio frequency Paul trap and a high–magnetic field, room-temperature permanent magnet Penning trap. The two results agreed within the measurement errors.

This project reduced the uncertainty of this g-factor ratio of the $^{40}Ca^{+}$ 3d $^{2}D_{5/2}$ and the 4s $^{2}S_{1/2}$ electronic states by roughly three orders of magnitude compared with previous work, ultimately achieving a fractional imprecision of 340 parts per trillion in the Penning trap system. Because the ground-state g-factor is known with higher precision, this also improves the metastable 3d $^{2}D_{5/2}$-state g-factor precision. The project also evaluated systematic effects for both systems; many of these effects were common, but some were unique to one platform or another. This precision measurement is expected to test and bound relativistic quantum mechanics and higher-order quantum electrodynamics effects to assist in improving multielectron correlation effect calculations and aid searches for time variation of the fundamental constants of physics, as well as nuclear physics models.

12043: Thermal Gradient Capsules for Materials Irradiation

A. Schrell, E. Martinez, R. Howard, M. Zach, N. Russell

Project Description

Novel, cutting-edge materials need to be developed and tested to achieve the goal of deploying advanced commercial fusion and fission reactors. Understanding structural material damage caused by neutron flux is important for robustly and safely deploying these systems. Neutron irradiation in existing nuclear reactors is a way to investigate these degradation pathways in materials of interest. The High Flux Isotope Reactor has one of the highest neutron fluxes of any reactor in the world and, as a result, offers the fastest mechanism to accumulate significant damage in relevant nuclear materials. Typically, the capsules used as irradiation vehicles for these experiments are relatively small and quite complex. As a result, significant effort is required for the design and safety of the capsules. These two issues, in combination with available space in the reactor, contribute to lower throughput than desired for irradiated material specimens used to design and select next-generation nuclear materials. In this project, a new irradiation capsule design was developed to broaden the spectrum of thermal data points achievable within a single irradiation capsule via a temperature gradient, which additionally targeted higher-than-traditional temperatures (2,000°C). This design produced the benefit of a wider range of specimen temperatures (1,000°C–2,000°C) in a single capsule. This concept is useful for materials that will be used in areas of nuclear structures that are expected to see large thermal gradients. The first two capsules for this design were built.

Mission Relevance

This project aimed to provide critical data about the material science needed to develop next-generation nuclear energy in the fission and fusion space, which serves DOE's mission of addressing America's energy, environmental, and nuclear challenges to provide security and economic growth for the nation.

Results and Accomplishments

A high-temperature thermal gradient irradiation vehicle was developed and built. This project included High Flux Isotope Reactor–approved drawings, safety documentation, and a fundamental understanding of machining and assembly limitations associated with this design. The parts were assembled and, along with associated documentation, delivered to High Flux Isotope Reactor for irradiation of two capsule containing silicon carbide specimens.

This capsule vehicle enables irradiation of 2,000°C ceramic specimens and 1,500°C refractory metal specimens, with a broad range of distributed temperatures. These extremely high target temperatures exceed the capability of typical materials irradiation. The initial capsules containing silicon carbide will offer insights into the performance of this material under extreme temperature and neutron irradiation. Additionally, this work eased the way for irradiation of other refractory materials expected to experience large thermal gradients and extreme temperatures.

12052: Fusion's Extreme Thermal Management Demonstration Platform

M. Gehrig, P. Humrickhouse, T. Gray, K. Pieper, V. Badalassi, A. Sircar, B. Jordan, C. Wiggins, C. Hauck, S. Ghosh, M. DeVinney, D. Ottinger, T. Xu, M. Kirka, T. Beck, T. Graening

Project Description

One of the challenges that fusion energy faces is managing the extreme thermal loads to the reactor structures from the plasma and fusion processes. This project explored the feasibility of a novel approach to address this issue by integrating three strands of research. The first strand was designing a highly configurable physical test bed that accesses some of the conditions relevant for the first wall and blanket of a fusion pilot plant. One requirement is to support developing and testing first wall and/or blanket cooling structures: plasma-facing tile designs, tile support concepts, and heat extraction options. This work included a demonstration of advanced manufacturing of some of the key complex components. The second strand was a high-fidelity digital twin of the test bed: a rigorous, physics-based, virtual representation of the test bed that exploits high-performance computing to provide accurate simulations of its performance. The final strand was a mathematical methodology for developing a surrogate model for the first wall structure. This methodology provided a strategy to embody the model into a fast but accurate module suitable for multiple parameter scans and/or integration into a fusion pilot plant design tool.

Mission Relevance

The scientific and technical contribution of this project was a new methodology that combines data from physical test bed experiments, advanced engineering simulations, and mathematical modeling to develop and test components. By forwarding the development of a fusion pilot plant, this project advanced DOE's mission in energy security and abundance.

Results and Accomplishments

This project formed a design-to-manufacture-to-validate pipeline. Through the different aims of this project, design and analysis were performed to design a flow loop, 3D print a sample test article in steel, and perform baseline simulations exploring friction factor on rough walls.

For the first aim, this project was able to design a blanket sector-scale helium flow loop capable of testing under high heat flux and working with flow rates sufficient to cool blanket-component-size test articles. The upgraded helium flow loop design was produced with costed-out components to allow for trade-offs between components and an easily upgradable system to enable this experimental apparatus to grow as capabilities are needed. Additionally, this project successfully demonstrated 3D printing using 316L stainless steel with a component occupying a box with a volume of approximately 0.33 m. The component shoes poloidally oriented first wall cooling channels and a bulk breeding duct. This print was performed using directed energy deposition to demonstrate the utility of this technology for this purpose and provided the opportunity to allow for possible coatings using this same method. To refine the cooling channel and further push the bounds of capability, a cooling channel with a chevron rib for the first wall channel was designed, and print simulations were performed to test the ability to print the complex geometry.

The second aim focused on high-fidelity simulation. This activity's primary outcomes were technical and logistical. The technical outcomes involved producing Reynolds-Averaged Navier–Stokes simulations in STAR-CCM+ and OpenFOAM to examine the friction factor in the channels. Simulations examining higher-fidelity analysis using large eddy simulation were also performed for comparison. The logistical outcomes of the second aim communicated with the third aim of producing surrogate models.

The third aim focused on surrogate modeling and uncertainty quantification efforts as well as how they could be implemented in a demonstration of this project. The surrogate modeling was determined to be

most useful in optimizing parameters necessary for heat transfer enhancement and tuning turbulence model parameters in the validation process. The major technical outcome of this was a strategy that can be employed for incorporating the surrogate and reduced-order modeling in a heat transfer validation and optimization experiment.

This project created the foundation for best practices regarding designing an experiment loop and components for manufacturing, identifying the limitations of manufacturing to target for future development, and then identifying the process for which high-fidelity simulation can inform surrogate modeling and vice versa.

12060: Agile Manufacturing at the Edge for Drones

M. Larson, D. L. Page, A. Duncan, B. Post, J. Baldwin, N. Burchfield, D. Cotten

[The results of this project have been determined to be Controlled Unclassified Information. Therefore, additional information protection and distribution restrictions have been applied.]

12064: Classified High-Performance Computing at the Edge

J. Cohen, S. Young

[The results of this project have been determined to be Controlled Unclassified Information. Therefore, additional information protection and distribution restrictions have been applied.]

ENZYME ENGINEERING

No projects in this initiative were completed in FY 2025.

HETEROGENOUS QUANTUM SYSTEMS

No projects in this initiative were completed in FY 2025.

INTEGRATED STUDIES OF COMPLEX BIOLOGICAL AND ENVIRONMENTAL SYSTEMS

11578: Investigating Ecosystem-Scale Resilience to Thermal Extremes

N. Griffiths, F. Santos, V. Salmon, A. Webb, B. Wang

Project Description

The goal of this project was to identify the scale-dependent mechanisms that may confer resilience of a soil–plant–microbial ecosystem to environmental disturbances. An important component of this goal was to disentangle the interactions between biotic processes and the abiotic environment as well as the mechanisms by which these interactions scale up to affect ecosystem- and biogeochemical-process resilience. This project used a coupled model–experimental approach to (1) examine ecosystem responses to press (warming) and pulse (extreme heat events) disturbances using a field experiment, (2) identify mechanisms of ecosystem resiliency responses to thermal perturbations using laboratory experiments, and (3) examine how thermal extremes influence leaf litter microbial physiological and ecological dynamics with data-informed modeling.

Mission Relevance

This project improved the understanding of how ecosystems respond to and recover from extreme perturbations. Improving the understanding of ecosystem resiliency to environmental change and disturbance supports DOE's missions related to energy security and economic prosperity.

Results and Accomplishments

Infrastructure for field manipulations was installed in four plots in the Walker Branch Watershed. These plots received extreme warming treatments in FY 2024. Multiple tests were conducted to evaluate the infrastructure's ability to create an extreme heated environment. Monthly soil respiration and nutrient availability measurements were conducted in all plots at Walker Branch (four control, four warmed [+2°C to +3°C] and extreme, four extreme). Seasonal measurements of soil and litter chemistry were conducted in all plots. This project found that warming increased soil respiration, but the elevated respiration response to heat waves was lower when soils were previously heated, possibly because of compounded water limitation. Five overlapping decomposition experiments were completed to examine the interactive effects of warming and heat waves on mass loss rates. This project found that decomposition was slower under warming because of moisture limitation, and heat waves did not have a clear effect on decomposition. Laboratory incubation systems were completed in FY 2024 to allow for separate manipulation of temperature and moisture levels and assessment of carbon and nitrogen cycle responses. This project found that soil respiration was higher with warmer and wetter soils and that soil–water content had a stronger effect than temperature. The microbial trait-based modeling framework was built in DEMENTpy. This model simulated the effects of thermal and compound extremes. Model predictions showed elevated (warming) and suppressed (drought) decomposition responses. Similar to field and lab results, the combined disturbance effect was slower decomposition caused by the dominant effect of soil moisture.

NEUTRON DATA INTERPRETATION PLATFORM ECOSYSTEM

11438: Smart System for Neutron Crystallography

X. Wang, Z. Morgan, Z. Xiao, G. Zhang, V. Reshniak, T. Proffen

Project Description

This project delivered an AI- and machine learning–enabled single-crystal neutron data reduction and analysis application (i.e., single-crystal graphical user interface) that combined low-latency edge inference with exascale training on measured datasets. Implemented as NeuXtalViz, the single-crystal graphical user interface provided automated **UB** (i.e., standard orientation matrix) determination, peak indexing and integration, volumetric slicing, experiment planning, and seamless handoffs to crystallographic programs. At the Spallation Neutron Source's TOPAZ beamline, integration with the Experimental Physics and Industrial Control System and External Instrument Control enabled ingesting live neutron time-of-flight event data and supported real-time guidance. On ORNL's Frontier, scalable sequence models such as the Temporal Fusion Transformer were trained and validated on archived and measured data to inform steering. A Trame (Python-based web application framework) web interface, CrystalPilot, was packaged for the Neutron Data Interpretation Platform using the Neutron Open Visualization and Analysis framework, which simplified building applications that interact with Neutron Data Interpretation Platform to enable accessible, agentic workflows. The delivered stack is operational at the Spallation Neutron Source; it improves time to decision and raises data fidelity for complex materials studies.

Mission Relevance

The project outcomes advanced the DOE mission to address the nation's energy, environmental, and security challenges through transformative science and technology solutions by enabling data-driven neutron science that increases user facility efficiency and scientific yield through earlier, defensible decisions and reduced beam time waste. The project supported ORNL's mission to leverage unique capabilities in neutron sciences and high-performance computing and translate fundamental science into real-world applications by providing AI-assisted experimentation and analysis accelerated by high-performance computing to boost throughput, enhance data quality, and broaden access via reproducible, web-delivered workflows across ORNL neutron user facilities.

Results and Accomplishments

NeuXtalViz, the single-crystal graphical user interface for 3D reciprocal-space exploration, experiment planning, and structural interpretation, was deployed on the Spallation Neutron Source analysis cluster for user access. TOPAZ Single-Crystal Data Reduction and JANA2020 data and structure analysis tools are available to users on the Neutron Data Interpretation Platform using the Neutron Open Visualization and Analysis framework.

At the Spallation Neutron Source's TOPAZ beamline, Experimental Physics and Industrial Control System and External Instrument Control integration enabled live event ingestion, low-latency data quality assessment, and in-experiment guidance. The Trame-based CrystalPilot interface was packaged and ready for incorporation into the Neutron Data Interpretation Platform to provide a web-based, guided workflow for experiment steering.

The project implemented 4D Gaussian mixture feature extraction to characterize event space peak and a hierarchical Bayesian multiscale Poisson integrator with wavelength-aware regularization. These methods improved robustness for weak reflections and early data and stabilized integrated intensities compared with traditional approaches.

Sequence models such as the Temporal Fusion Transformer were trained and validated on Frontier using measured datasets as part of an integrated edge-to-exascale workflow. Inference runs at the edge for responsiveness that is compatible with steering needs at the beamline.

Overall, the integrated single-crystal graphical user interface, real-time steering capability, and modeling backed by high-performance computing were validated on Spallation Neutron Source single-crystal datasets and demonstrated reductions in analyst effort and time to decision, along with improved consistency and quality of structure analysis results.

11466: Graphics Processing Unit-Based Resolution and Visualization Interface for Triple-Axis Spectrometers

A. Savici, S. Hahn, B. Li, P. Fackler, H. Mankad, J. Chae, D. Dahlbom, A. Ayres, S. Chi

Project Description

The aim of the Graphics Processing Unit–Based Resolution and Visualization Interface for Triple-Axis Spectrometers (GRAVITAS) project was to develop software to visualize data from triple-axis spectrometers, perform rapid GPU-based resolution function calculations, and fit the data to resolution-convoluted models of the scattering function. The original goal was to create a graphic user interface that plots, convolutes, and fits input user data in real time. This work addressed two major obstacles in the experiment-to-publication pipeline. First, the developed software allows users to perform preliminary analysis of their data while being collected to make informed decisions about further data collection to ensure that beam time is used most efficiently. The software also reduces the need for follow-up experiments or data that cannot be used because of poor quality. Second, the software facilitates rapid fitting to obtain physical parameters of the material during an experiment.

Mission Relevance

The project supported scientific discovery and accelerated material development, which contribute to national economic opportunities and DOE priorities. GPU-based resolution calculations enhance experiment planning, optimize neutron beamtime use, and improve the understanding of measured samples.

Results and Accomplishments

Resolution-Corrected Diffraction Measurements on Triple-Axis Spectrometers

This project developed a calculator for the resolution ellipsoids using the Cooper–Nathans method. It was programmed in Python. Experimental measurements of Bragg peaks were used to benchmark the results.

The HB-1 triple-axis spectrometer at the High Flux Isotope Reactor is often used for diffraction experiments to perform parametric studies. Previously, measurements of the intensity of Bragg peaks were found to depend on the position in the reciprocal space and the measurement scan direction. As part of the research for this calculator, this project showed how the resolution matrix contributes to the Lorentz factor (a measure of the reciprocal space coverage of the measurements). When the resolution corrected Lorentz factor was removed, the transverse and longitudinal scans gave consistent integrated intensities. This procedure is now used at HB-1.

Web-Based Resolution Calculator for Triple-Axis Spectrometers at the High Flux Isotope Reactor

This project developed a graphical user interface that allows users to view and edit the instrument and sample parameters, calculate the resolution function at a given location in the reciprocal (Q, E) space, and view the resulting resolution ellipsoid. The software can also visualize triple-axis measurements and overlay the instrumental resolution function. This helps users understand resolution effects in experimental data and better plan their experiments.

Prototype Resolution Convolution

A main part of the GRAVITAS project was understanding the multidimensional convolution process. This work found that, in the case of 4D integrals, evaluating the function is one of the bottlenecks. Additionally, generating an integration grid that is aligned with the physics of the material (i.e., aligned with the reciprocal space) is often inefficient for calculations because of the orientation of the resolution ellipsoid. A hyperparallelepiped might contain most of the integration points outside the resolution ellipsoid, meaning that this project evaluated the model when the weight was negligible. For analytical models, this project proposed to check the overlap of the dispersion with the resolution function and only evaluated the intensity if the distance was small. The Python code is available for use.[1]

This project developed a Julia package, SunnyAnalysisTools, to provide a simple and efficient method of applying instrument resolution effects to linear spin–wave and Landau–Lifshitz simulations of magnetic dynamics. The package was built as a supplement to the popular Sunny package by adding tools to perform instrument convolution on top of Sunny's existing calculators. In particular, it provides a straightforward interface for specifying instrument and sample configurations. This information is then used to query the existing neutron software ecosystem to generate convolution kernels. Python interoperability was achieved via the PythonCall package, which interfaces with TAVI for triple-axis instrument information and PyChop for time-of-flight instruments. The resulting resolution data are passed to extensions of Sunny's functions for calculating scattering intensities, which automatically perform the appropriate convolutions. The calculations are fast enough to complete in seconds on an ordinary laptop, which facilitates direct comparison between experimental data and theoretical models. This capability aids data analysis and experiment planning. The package is freely available as part of the SunnySuite project on GitHub.[2]

Accelerated Phonon Calculations

Inelastic neutron scattering is an indispensable tool for experimentally probing phonons. Triple-axis spectrometers contain one or a small number of detectors, so measurements are usually made at hundreds of points in reciprocal space, typically along high-symmetry directions. The resolution calculator and prototype convolution previously described found the need to evaluate scattering intensity and many more points. To investigate performance optimization opportunities, this project developed a proxy application based on Phonopy to capture relevant computational aspects. The project team used the Numba library to write GPU kernels. These kernels included calculation of the dynamical matrix, solving for eigenvalues and eigenvectors, calculating the Debye–Waller factor, and using the results of these kernels to calculate neutron or x-ray scattering intensities. The project obtained speedups of more than one order of magnitude. The project also found that these fast evaluations of phonon intensities are also useful to train various AI models.

Accelerated Magnetic Excitation Modeling

Magnetic excitation modeling is more challenging than phonon calculations. The project team decided to prototype accelerated calculations using the Sunny.jl package, which is written in the Julia programming language. It simulates experimental observables from a candidate magnetic Hamiltonian. Forward solvers include spin wave theory solving a generalized eigenvalue problem and semiclassical simulations calculating a time-dependent solution. As a first step, this project showed that this package can be integrated it with the Neutron Data Interpretation Platform by deploying it in a docker container. User documentation is provided.

This project developed a road map for a full GPU port of Sunny.jl. The project created a document that describes how to move linear spin wave theory and SU(N) spin wave calculations in Sunny.jl from CPU

[1] *TAVI* (Github, 2025), https://github.com/neutrons/TAVI/tree/main/scripts/convolution.
[2] *SunnySuite* (Github, 2025), https://github.com/SunnySuite/SunnyAnalysisTools.

to GPU accelerators. This plan involves using a generalized eigenvalue solver from a combination of batched BLAS and LAPACK functions and includes potential variations to explore.

11471: Experiment Steering for Powder Diffraction

M. McDonnell, B. Vacaliuc, R. Gregory, K. Gofron, J. Liu, L. Drane, G. Cage, A. Ayres

Project Description

This project developed new capabilities for automation and smart data pipelining for neutron powder diffraction experiments. New automated powder diffraction data analysis tools were developed for the Neutron Data Interpretation Platform and INTERSECT software. These new tools were used to steer two neutron powder diffraction experiments and create reusable tools required for others to perform the same developments. Neutron powder diffraction experiments are routinely used to understand the evolution of materials' processes, such as structure changes for battery materials during charge/discharge cycles or temperature-dependent phase changes for carbon capture materials. Neutron powder diffractometers have the measurement capabilities to track these structural changes within seconds. In this study, hematite samples were used to perform commissioning experiments at two instruments at the Spallation Neutron Source: the Nanoscale Ordered Materials Diffractometer (NOMAD) in May 2025 and the Powder Diffractometer (POWGEN) in September 2025. Two active learning algorithms were used from the Distributed INTERSECT Active Learning (DIAL) project in the INTERSECT Initiative. From this project, new experiments have been confirmed for beamlines at the High Flux Isotope Reactor as well as the x-ray diffraction synchrotron experiment at the Cornell High-Energy Synchrotron Source for additive manufacturing samples.

Mission Relevance

The experimental steering capabilities and further work that will come from this project directly affect the speed and efficiency of autonomous experiments at large-scale experimental user facilities. These types of scattering experiments are used across a wide range of domain science and techniques, including materials science, protein crystallography, and engineering diffraction, and can extend to inelastic scattering for domains such as quantum material studies. This project helped complete some of the first autonomous neutron scattering experiments at a DOE neutron source user facility. These new capabilities and autonomous techniques help reduce the time to new scientific knowledge, more efficiently use awarded beam time, measure more samples during given times, and discover hard-to-find phenomena for new emerging behaviors of materials such as transient state studies, hidden phase transformations, and metastable state studies.

For upcoming x-ray scattering experiments for additive manufacturing, many areas of research and industry can benefit from the data generated. Direct energy deposition wire arc additive manufacturing has proven to be a cost-effective method for large components such as airplane repairs and turbines in hydroelectric dams. The insights from this work may also create economic competitiveness for additively manufacturing these types of parts. The autonomous science workflows can help reduce the time-to-product of additive manufacturing technologies for industry.

Results and Accomplishments

This project performed the first autonomous neutron powder diffraction experiments on NOMAD and POWGEN beamlines at the Spallation Neutron Source. Autonomous powder diffraction experiments allowed for quickly and efficiently measuring the phase transition of iron(III) oxide (i.e., hematite) using different samples (bulk and nanoparticles) as well as different algorithms provided by the DIAL service. This project also installed the External Instrument Control software for secure, out-of-beamline network requests for steering experiments on NOMAD and POWGEN. The INTERSECT External Instrument Control service was developed to connect External Instrument Control capabilities into INTERSECT workflows. INTERSECT Neutron Data Interpretation Platform service was developed for launching

Neutron Data Interpretation Platform tools and securely retrieving outcome data to integrate into INTERSECT autonomous workflows.

For Bayesian optimization and Gaussian process algorithms, active learning algorithms were codeveloped for phase exploration in the DIAL service software, including the NIST-developed Autonomous Neutron Diffraction Explorer algorithm for phase exploitation. DIAL is open-source software available on GitHub.

New INTERSECT services for neutron file parsing and peak analysis were used in the autonomous workflows. Virtual instrument server pipelines were created to host on the ORNL Research Cloud for testing NOMAD beamline controls.

Finally, this project worked with the ORNL Neutron Sciences team in charge of developing and maintaining the web monitor data pipeline of neutron beamlines. The project worked to include the POWGEN data from that web monitor into the INTERSECT infrastructure for autonomous workflows.

11474: Machine Learning Assisted Small-Angle Neutron Scattering Data Analysis Platform

C. Do, W.-R. Chen, Y. Wang, J.-M. Carrillo, L. Ding

Project Description

This project developed and demonstrated machine learning (ML)- and AI-assisted methods for advancing the interpretation and analysis of small-angle neutron scattering (SANS) data to address the physical inversion problem and long-standing usability barriers in neutron data analysis. By combining molecular simulations, Gaussian process regression, deep learning, and large language model (LLM)-based agents, the project established new data-driven pathways for extracting interaction parameters, structural descriptors, and physical insight directly from scattering data beyond the scope of traditional analytic models. This effort expanded into a broader AI-enabled neutron data interpretation ecosystem. This included LLM-driven multiagent systems for automated SANS data analysis (SasAgent), natural language–controlled molecular simulation workflows (ToPolyAgent), and LLM-based decision support tools for proposed evaluation at large user facilities. These outcomes demonstrated that AI agents can reliably automate complex scientific workflows, integrate domain-specific simulation and analysis tools, and enhance consistency, efficiency, and accessibility in neutron science. Collectively, the project established a scalable, extensible foundation for AI-assisted neutron data interpretation and scientific decision-making at ORNL.

Mission Relevance

This project advanced the DOE mission to deliver transformative science and technology solutions by strengthening the nation's capability to interpret complex neutron scattering data, which are critical to materials and soft matter research relevant to energy, environmental, and advanced manufacturing challenges. By developing ML- and AI-assisted tools that overcome long-standing limitations of traditional scattering analysis methods, the project enabled more accurate, efficient, and reproducible extraction of physical interactions and structural information from experimental data. By improving the accessibility and reliability of neutron data interpretation for a broad user community, this work enhances the scientific impact of DOE user facilities and accelerates discovery across materials systems central to national energy and security priorities.

Results and Accomplishments

This project achieved its central objective of establishing a new, data-driven paradigm to interpret SANS data with the goal of successfully overcoming long-standing limitations of traditional analytic and integral equation–based approaches. A central scientific outcome of this effort was developing and validating ML-based inversion frameworks that extracted effective interaction parameters, structural descriptors, and physical insight directly from scattering data. By combining large-scale molecular simulations with

Gaussian process regression and deep learning models, the project demonstrated that scattering inversion can be performed accurately and efficiently for systems that are strongly interacting, polydisperse, anisotropic, or driven far-from-equilibrium regimes in which conventional models often fail or become unreliable. These advances were validated across a wide range of soft matter systems, including charged colloids, polymers, lamellar phases, and mechanically driven materials, using simulated and experimental SANS data.

A key technical accomplishment of this work was the creation of extensive simulation-based libraries of structure factors and scattering observables generated via molecular dynamics and Monte Carlo simulations. These libraries enabled probabilistic learning of the mapping between physical interaction parameters and scattering signatures, which allowed uncertainty-aware inference rather than single-point parameter estimation. The resulting models demonstrated improved robustness, speed, and generality compared with traditional fitting methods, particularly in cases in which analytic structure factors were unavailable or relied on restrictive assumptions. A major conclusion from this project was that simulation-informed ML provided a physically grounded and extensible solution to the inverse scattering problem, which enabled quantitative interpretation of complex soft matter systems without sacrificing physical fidelity.

The project also expanded the scope of ML-assisted scattering analysis by addressing structurally complex and experimentally relevant systems. This included successful application of ML-based inversion to polymeric systems with internal constraints, polydispersity, and mechanical deformation, as well as lamellar and mesophase systems in which the conventional form factor–structure factor separation broke down. These studies demonstrated that scattering data encoded richer structural information than what could be accessed by traditional models and that modern ML methods could reliably extract this information when trained on physically meaningful simulation data. Collectively, these results advanced the state of the art in quantitative neutron scattering analysis and established a general framework applicable across multiple classes of materials.

A second accomplishment was developing AI-enabled software systems that translated these methodological advances into practical, user-accessible tools. The project delivered multiple LLM-driven, multiagent AI systems that automated complex neutron scattering workflows that previously required extensive expert intervention. SasAgent was developed as an AI system that integrated LLMs with established SasView analysis tools, which enabled natural language–guided SANS data analysis, scattering length density calculation, synthetic data generation, and experimental data fitting. SasAgent demonstrated that domain-specific scientific tools could be reliably orchestrated by AI agents while preserving the rigor and reproducibility of expert-driven analysis.

In parallel, the project delivered ToPolyAgent, a multiagent AI framework that enabled natural language–controlled molecular simulations of coarse-grained polymer systems using established simulation engines. ToPolyAgent supported interactive and autonomous workflows, generated polymer configurations, ran simulations, performed conformation analysis, and compiled structured reports. These capabilities were demonstrated across a wide range of polymer architectures and solvent conditions to show that AI agents can serve as effective intermediaries between users and high-performance simulation tools. Together, these systems help reduce the barriers to advanced simulation-based neutron data interpretation while maintaining scientific rigor.

12070: Neutrons Open Visualization and Analysis Framework

G. R. Watson, S. Yakubov, A. Ayres, G. Cage, J. Duggan, K. Maroun, D. Pugmire, K. Moreland

Project Description

Creating advanced visualization tools required for analyzing neutron scattering experiments is a complex and challenging task. Expertise in user interface design and advanced visualization techniques must be coupled with deep domain knowledge to make the tools effective. The sheer number of instruments

supported by ORNL's Spallation Neutron Source and High Flux Isotope Reactor, combined with the need to efficiently deliver tools to facility users, makes this a particularly intractable problem. The Neutron Data Interpretation Platform (NDIP)—an LDRD project that was completed in FY 2024—provided the basis for creating modular scientific workflows for analyzing and interpreting neutron scattering data. The platform currently supports over 60 different tools, as well as six workflows, ranging from tools for automated structure refinement, such as single peak fitting, to pattern matching and GSAS2 refinement, Monte Carlo ray tracing for simulating incident beams, and atomistic modeling and machine learning, such as DeePMD and LAMMPS. Additionally, several data reduction tools, including GP-SANS, USANS, and TOPAZ, are available in the platform, as well as user interfaces for managing the GP-SANS and TOPAZ data reduction workflow. NDIP now allows raw data to be automatically ingested into the platform as data come off a number of beamlines, and the platform also automatically executes workflows when this happens.

The Neutrons Advanced Visualization and Analysis (NOVA) framework couples with the NDIP platform to provide a next-generation visualization and analysis framework for neutron scattering science. This project developed a visualization architecture and framework that enables cutting-edge technologies to be used for delivering visualization tools to instrument scientists and end users. NOVA integrates seamlessly with the tools and services already provided by NDIP. This project provided the necessary infrastructure as code to enable the framework to be transitioned into a production environment with minimal effort and disruption.

Mission Relevance

Data analysis and visualization are important for extracting scientific knowledge from neutron scattering experiments. Neutron scattering enables work in a wide range of scientific fields such as materials science, advanced manufacturing, chemistry, battery science, and others that advance DOE's mission to address the nation's energy, environmental, and nuclear challenges. This project helped address visualization and analysis needs in neutron scattering by creating a modern, modular framework for developing and deploying web-based applications that provide advanced user interfaces and visualization capabilities.

Results and Accomplishments

This project completed the development of NOVA, which is a framework for building integrated web-based visualization and analysis tools. The NOVA framework provides the following key benefits:

- A single common framework for neutron data processing tools and visualization applications
- A uniform mechanism for all users to access reduction, analysis, and visualization tools
- The ability to use specialized hardware resources for demanding visualization tasks
- A scalable, modular, reusable, and maintainable design methodology for visualization applications
- The ability to leverage advanced visualization techniques provided by PyVista, VTK, and other open-source libraries
- The ability to share and reuse analysis code that is developed for user interfaces (e.g., validation of input parameters) and backend tools (e.g., computational algorithms)

The NOVA framework supports legacy and third-party visualization applications as well as new neutron analysis applications being developed for specific techniques or instruments by ORNL's in-house team. New visualization applications are predominantly (but not exclusively) based on a boilerplate template that provides a basic model–view–viewmodel design pattern, along with all the components necessary for advanced visualizations, and the ability to interact with NDIP, Mantid, and other libraries commonly required for neutrons data analysis. Existing web-based applications will continue to coexist with NOVA without requiring any modification; however, these platforms can also use NOVA functionality via Python libraries should that be of benefit to the application. Legacy and third-party applications that were not designed for web-based deployment (e.g., Amira, Jana2020, SasView) are supported using containers

that incorporate a Virtual Network Computing client or server. Legacy and third-party applications that are web-based (e.g., Refl1D) have been containerized and deployed as web applications on the NOVA framework.

NONEQUILIBRIUM AND EMERGENT TRANSIENTS IN ADVANCED AND SOFT MATERIALS (NEAT)

11720: Floquet Quantum Physics Studied by Neutrons

C. Hua

Project Description

Most current state-of-the-art technologies are dominated by equilibrium macroscopic phenomena, in which thermodynamic quantities such as temperature and pressure are well-defined. In the realm of Floquet physics, in which systems are driven to out-of-equilibrium states by a time-periodic external source, macroscopic properties such as temperature, conductivities, and diffusivities are no longer relevant descriptors of the systems. Floquet engineering of quantum materials will enable programmable quantum materials for quantum computing, genuine nonequilibrium quantum phases without an equilibrium counterpart such as discrete time crystals, and coherent transport in quantum many-body systems for low-loss and low-power electronics. Switching electronic states with light excitation is reversible and controllable and can be used in light-controllable electronic devices. Despite the experimental advances in systems such as cold atoms and trapped ions, nonequilibrium phenomena in condensed matters in general are not well-studied. Therefore, the rich physics of quantum Floquet systems offer plenty of opportunities to explore and advance quantum information technologies.

Quantum spin systems offer a method to uncover the laws of nonequilibrium physics in condensed matter. The ubiquitous presence of integrability in 1D spin chains—a property of dynamical systems with an infinite set of conserved commuting quantities—leads to intriguing nonequilibrium physics such as quantum wake dynamics, finite Drude weight, and relaxation to nonequilibrium steady states. Highly frustrated magnetism in quantum spin ice can cause bottlenecks in the nonequilibrium dynamics of magnetic monopole excitations. These nonequilibrium excitations form Floquet states under a periodic external excitation.

This project investigated the decay behaviors of Floquet quantum states, which exist in nonequilibrium quantum spin systems subject to a periodic external source. This work used a unique laser-pump neutron-probe capability. Because of symmetry restriction of the spin operator, spin–spin interactions alone in quantum magnets are insufficient to bring nonequilibrium spin states back to equilibrium. Interaction with the environment (i.e., thermal baths) leads to the ultimate decay but occurs on a timescale much slower than that of the spin–spin interaction at low temperatures. With the control of such quantum systems using periodic external driving at a rate faster than the spin–environment relaxation, Floquet spin states are formed and maintained at a timescale that should be accessible by neutrons. To characterize the decay behavior of the Floquet states, this project directly observed the actual quasiparticle flow and collision processes using inelastic neutron scattering and neutron diffraction measurements under a periodic laser excitation. This study pioneered a new research direction in neutron science to study transient phenomena in condensed matter and expanded the pump–probe capabilities in ORNL's neutron sources beyond the existing excitation schemes and sample environment. The insight gained from this project narrowed the knowledge gap of nonequilibrium physics in emerging spin systems that impeded the progress in spintronics based on magnetism by enabling Floquet engineering of quantum states.

Mission Relevance

Enabling Floquet engineering of quantum materials could lead to new tuning mechanisms to create quantum coherence for technologies such as quantum computing, data storage, information transport, and wireless communication. Because the control of quantum states is important to these applications, fundamental insights into nonequilibrium quantum physics leading to exotic quantum properties and functionalities support America's energy technologies leadership and promote DOE missions for energy and national security.

Results and Accomplishments

A next-generation millikelvin optical cryostat optimized for neutron scattering environments was procured and entered formal construction in FY 2024. The system, built by ICE Oxford Ltd., was designed to minimize parasitic neutron scattering contributions from structural and cryogenic components, thereby enabling high signal-to-noise measurements on small-volume samples.

The cryostat was configured for standard operation as a 1.5 K variable temperature insert supported by a closed-cycle liquid helium recondensing system, which ensured stable thermal performance with minimal cryogen consumption. Integrated within the variable temperature insert environment was an ultralow-temperature probe that employed advanced dilution refrigeration technology to achieve and maintain continuous base temperatures near 300 mK. The system was engineered to deliver ≥100 μW of cooling power at the sample position, which provided a critical thermal budget for experiments requiring optical, electrical, or radio frequency excitation. These external driving fields introduced dissipative heating that would otherwise limit thermal stability or attainable base temperature in conventional cryogenic platforms.

Upon deployment, this instrument will support elastic and inelastic neutron scattering studies on quantum materials, frustrated magnets, and driven quantum systems in nonequilibrium regimes previously inaccessible in conventional neutron scattering experiments.

NONPROLIFERATION SCIENCE INITIATIVE (NSI)

11103: Characterizing Indoor Pattern-of-Life and Anomalies for Secure Facilities

D. De, J. Baldwin, C. Gunaratne, S. Guggilam, C. Young, T. Waddle, T. Stinnett, H. Niu, D. Mann, G. M. Thakur, C. Christopher

Project Description

This project designed and deployed a data-driven, AI/machine learning–based, intelligent pattern-of-life (PoL) characterization system that can analyze regularities and normalcies of indoor environments, as well as automatically and dynamically detect anomalous activities in secure facilities. The project specifically focused on facilities with physical security needs, such as nuclear reactor facilities or facilities with radiological materials. This work enabled a unique PoL characterization solution for indoor (and campus-level) environments of secure facilities (i.e., facilities of national interest and mission-critical operations), especially at an unprecedented, spatially fine-grained manner that also preserves the privacy of facility workers. The three main objectives of the project were to (1) design a cost-effective, easily deployable and scalable, multisensor data collection system capable of passively monitoring indoor spatiotemporal mobility and activity patterns of facility workers at such a fine-grained spatial resolution that is not available in state-of-the-art, real-world solutions and that simultaneously preserves privacy; (2) model regularities and anomalies in secure facilities' indoor PoL using the designed AI/machine learning–based algorithms to considerably improve current state-of the-art research; and (3) showcase the solution with a real-world test bed at an ORNL indoor facility and integrate it with a designed metaverse (virtual reality). Together, these goals have several downstream applications for safety, security, and safeguards in comprehensive nuclear governance, PoL, and national security to proactively prevent the threat of nuclear proliferation.

Mission Relevance

This project was relevant to ORNL's mission for science advancements and applications for national security. It focused on advancing current state-of-the-art research in PoL domains and its application in facility safeguard technologies in nuclear nonproliferation, addressing DOE's mission to ensure America's security and prosperity by addressing its nuclear challenges.

Results and Accomplishments

The project had three main aims: (1) design and deploy a multisensor data collection system for sensing PoL in secure facilities and collect multiple months of PoL data from a nuclear reactor facility, (2) design novel AI/machine learning–based PoL modeling and anomaly detection algorithm(s), and (3) design a virtual reality–based immersive environment for emulating normal PoL as well as anomalous behavior in nuclear reactor workers (and generating corresponding data).

First, an extensive compliance and approval process occurred for deploying a PoL sensing system in a nuclear reactor facility and collecting long-term data from the facility. Simultaneously, the multisensor PoL (and anomalies) sensing and data collection system was designed, which is capable of detecting spatially (indoor) very fine-grained movement patterns of facility users (but in a privacy-preserved way) who wear a custom-designed PoL smart badge. Also, passive motion sensors were used to capture movement of nonbadged users (i.e., not wearing the PoL smart badge) and material movements. Additionally, fine-grained indoor microclimate and structural vibrations were measured using the multimodal PoL sensing system. This system was deployed in ORNL's High Flux Isotope Reactor nuclear reactor facility, and continuous PoL sensing data was collected for an extended 10-month period.

Second, two different AI/machine learning–based PoL modeling and anomaly detection algorithms were designed, which, in an unsupervised way (i.e., no labeled training data required), can detect various

normal operating patterns in nuclear facilities, as well as various anomalies from spatial-temporal movements and other activities data of facility indoor environments. These multimodal spatiotemporal PoL and anomalies data was generated from efforts in aim 1 from deploying the system in a real-world nuclear reactor facility (as well as simulated and emulated data).

Third, an immersive virtual reality environment was designed with AI model-driven nonplayer characters posing as normal facility workers and an interactive user of the virtual reality emulating anomalous activities and behavior (i.e., enacting insider threat scenarios).

11131: Next-Generation Particle Mapping for Isotopic, Chemical, and Elemental Analysis

B. T. Manard, T. Spano, T. Beiswenger, H. Andrews, C. Hexel, D. Dunlap, V. Bradley, A. Zirakparvar

Project Description

Material analysis for nuclear nonproliferation is a challenging problem. Extracting maximum analytical value from finite quantities of material with robust and knowable uncertainties is a critical concern of stakeholders. Emphasis has been placed on detecting extremely small quantities of critically important collected materials among a plethora of background matrices (e.g., dust, soil). Currently, lack of throughput and sensitivity limitations restrict the speed of analysis and hinder comprehensive particle characterization. The primary aim of this work was to leverage ORNL's innovative, state-of-the-art equipment to solve these limitations. Micro-Raman spectroscopy, scanning electron microscopy with energy dispersive x-ray spectroscopy, and simultaneous laser-induced breakdown spectroscopy (LIBS)/laser ablation (LA)–inductively coupled plasma mass spectrometry (ICP-MS) were used in tandem to provide comprehensive Mapping of Isotopic, Chemical, and Elemental (MICE) information for samples containing an abundance of particles. The overall MICE workflow involved high-throughput chemical mapping using micro-Raman spectroscopy and elemental mapping using automated scanning electron microscopy with energy dispersive x-ray spectroscopy to locate and characterize particles of interest. LIBS/LA–ICP-MS was then employed to provide simultaneous elemental (e.g., carbon, oxygen, fluorine) and isotopic (e.g., uranium-234, -235, -236, and -238) analysis for whole samples, including background matrices and preidentified radiological particles of interest. The result of this effort was a state-of-the-art methodology for particulate analysis, which will be injected into a high–technology readiness level workflow that will provide the ability to obtain high-throughput particle characterization in the presence of complex matrices.

Mission Relevance

Although isotopic, chemical, and elemental mapping capabilities exist through various (separate) analytical techniques, correlated, spatially resolved deployment of disparate analytical techniques on samples with significant background noise are difficult to obtain. Thus, combining spatially resolved (i.e., mapped) isotopic, chemical, and elemental information on samples of interest among high-background matrices, coupled with advanced data analytics, represents a new capability for the nuclear nonproliferation community. Specifically, this work applied to DOE and National Nuclear Security Administration missions of national security and nuclear nonproliferation.

Results and Accomplishments

This project focused on developing and integrating new analytical approaches for particle MICE information. This project first used LIBS to rapidly analyze surrogate fuel particles. Tristructural-isotropic (TRISO) particles are a proposed nuclear fuel alternative for high-temperature reactors. These particles consist of a zirconium dioxide kernel (as a surrogate to uranium) surrounded by an inner pyrolytic carbon layer and an outer carbide layer (zirconium carbide was used in this project) to act as a barrier to fission products generated during nuclear reactions. These particles are embedded within a graphite compact and housed within the reactor core. Because of their robust nature, performing elemental analysis of these

particles poses a challenge. In this project, LIBS was explored as a method for characterizing elemental constituents of these particles, focusing on rapid elemental mapping and depth profiling. Different from traditional elemental analysis techniques (e.g., inductively coupled plasma–based methods), LIBS is advantageous because it can directly analyze the sample surface and detect light elements such as carbon and oxygen, making it a viable technique for analyzing small, multilayered particles to obtain spatial elemental information for producing these particles. In this work, LIBS was successfully used for discerning small layers (30–50 μm), detecting the location of carbon and oxygen layers, and providing fast 2D mapping and rapid depth profiling.

In parallel, work was performed with LA–ICP-MS to rapidly map uranium inclusions in a complex matrix. This project developed a unique method to initially screen large samples via LIBS to find the inclusions and subsequently employed LA–ICP-MS to scrutinize the inclusion for its isotopic composition. This work leveraged the high-speed applicability of LIBS to screen the samples for their elemental constituents. After determining the location of the uranium inclusions (via LIBS), high-resolution LA–ICP-MS was employed to further characterize the inclusions. The high-resolution (submicrometer) capabilities of LA–ICP-MS were able to extract important information from the uranium minerals, including discerning its chemical form (e.g., finchite from carnotite, strontium- and potassium-bearing uranyl vanadates) as well as their uranium-235/-238 isotopic composition. This approach significantly reduced the analysis time (approximately 95%) compared with employing an approach using only LA–ICP-MS. Furthermore, this work presented a novel approach to analyzing inclusions via a particle/inclusion analysis tool, which is commercially available within the iolite 4 software and allows for a more accurate characterization of the isotopic distribution of the inclusions and rapid sizing of the inclusions. This analytical approach can be applied to other sample types in which the target species (e.g., micrometer-sized inclusions) are embedded in complex matrices (e.g., centimeter-sized samples). Other work in this project focused on optimizing multicollector ICP-MS for LA of particles. The motive was to determine the performance of a multicollector ICP-MS when using all–Faraday cup detectors.

To expand on the multicollector ICP-MS work, this project also combined the approach with LIBS. This method allowed the team to simultaneously measure fluorine and uranium isotopics in single laser pulses on single uranium particles. This work highlighted how the LIBS/LA–multicollector ICP-MS could detect fluorine loss through calcination within particles and provide high-precision uranium isotopic measurements. Next, the project aimed to improve throughput for isotopic determination within particles. This work combined ultrafast LA with time-of-flight ICP-MS to detect and isotopically characterize ruthenium particles collected from air samples deployed in the field. The results demonstrated high-throughput particle analysis. Additionally, two unique data analysis software platforms were generated during this project to handle LIBS-based data and time-of-flight LA–ICP-MS.

11397: Development of Cosmic Radiation Noise Cancellation Method

J. Bae, R. Bean, A. Enders, S. Tognini

Project Description

Radiation portal monitors (RPMs) serve as a critical first line of defense for detecting illicit trafficking of nuclear and radioactive materials at borders and secure facilities. Most deployed RPMs rely on polyvinyltoluene plastic scintillators because of their low cost, mechanical robustness, and high light yield. However, their limited energy resolution and susceptibility to background fluctuations constrain detection sensitivity. Among all background sources, cosmic ray muons represent the most challenging noise contributor because they are unshieldable, highly energetic, and produce signal signatures in scintillators that are difficult to distinguish from legitimate gamma events.

This project developed and demonstrated a first-of-its-kind cosmic radiation noise cancellation framework for RPM applications. The approach combined analytical modeling, high-fidelity Monte Carlo

simulations, and data measured from ORNL's cosmic ray muon detector and RPM systems to characterize muon flux, energy deposition, and temporal behavior. Based on these insights, advanced signal processing algorithms were developed to suppress muon-induced noise and preserve valid radiation signatures. In parallel, a complementary hardware solution—a metal-latticed collimator—was designed to reduce background gamma contributions from the surrounding environment. Together, these software and hardware innovations significantly improved RPM signal stability and detection sensitivity without requiring changes to existing detector materials or operating concepts.

Mission Relevance

This project advanced DOE and ORNL missions by strengthening national security and nuclear nonproliferation capabilities related to radiation detection. RPMs are widely deployed to prevent unauthorized movement of nuclear and radioactive materials, yet their effectiveness is fundamentally limited by high–background noise environments. By addressing cosmic radiation—one of the most persistent and least tractable background sources—this work enhanced the reliability and sensitivity of existing RPM infrastructure.

Results and Accomplishments

Comprehensive characterization of cosmic ray muon interactions in RPM scintillators was completed using analytical models and GEANT4-based Monte Carlo simulations informed by CORSIKA air-shower calculations. These efforts quantified muon energy deposition as a function of energy, angle, and detector geometry to establish a physics-based understanding of how muons contribute to RPM background fluctuations. Time-dependent muon simulations enabled realistic reproduction of stochastic muon-induced signal behavior observed in operational RPM data. Measured RPM background signals were systematically analyzed in the time and frequency domains. Fast Fourier transform techniques were applied to identify dominant sources of periodic and long-term fluctuations, including temperature-driven baseline drift and cyclic background components. Frequency-domain filtering successfully reduced baseline instability and interdetector variability. This work provided a first stage of noise suppression.

Building on this foundation, an advanced noise cancellation algorithm based on Kalman filtering concepts was developed to explicitly account for muon-induced contributions to gamma signals. By fusing measured RPM data with modeled muon behavior, the algorithm reduced uncertainty in gamma background, scanning, and alarm channels without sacrificing valid gamma counts. Application to simulated and live RPM datasets demonstrated a marked reduction in signal variance and a corresponding increase in the effective alarm threshold margin, which indicated enhanced detection sensitivity. The associated computational latency was shown to be negligible relative to RPM scanning times, which confirmed operational feasibility. To complement software-based noise reduction, a metal-latticed lead collimator was designed and evaluated to suppress external gamma background from naturally occurring radioactive materials during object approach and exit. Simulation and system-level assessments showed that the collimator stabilized background baselines and improved sensitivity to radiation sources located within the RPM passage.

Collectively, these advances resulted in a validated, deployable framework for cosmic radiation noise mitigation in RPMs. The project supported invention disclosures for the noise cancellation algorithm and the metal-latticed collimator and established a strong technical basis for enhancing RPM performance across existing installations.

11442: Vehicular Pattern of Life Analysis from Uncontrolled Multiple Views

S. Ghanem, J. Holliman II, R. A. Kerekes, A. Duncan, K. J. Dayman

Project Description

Vehicle re-identification (re-ID) across disparate sensing modalities remains a fundamental challenge for transportation research. In this work, the project team introduced a deep multiview vehicle re-ID framework that leverages Siamese networks to compare pairs of vehicle images and produce matching scores to enable robust association across drastically different viewpoints such as those from UAVs, surveillance cameras, and ground sensors. The model exploits convolutional neural networks to learn features that remain discriminative under changes in angle, distance, and illumination, which supports more generalizable re-ID performance. As part of this effort, the team also developed an automated pipeline to synchronize roadside and UAV video streams to provide a multiperspective dataset that complements preexisting and synthetic collections. Together, these contributions advanced the capability to re-identify vehicles across wide viewing baselines; established a foundation for scalable, reproducible research in vehicle re-ID; and opened pathways for future applications such as conducting pattern-of-life (PoL) analyses.

Mission Relevance

Unconstrained multiview vehicle re-ID is an important capability for DOE's mission in nuclear nonproliferation because it supports persistent situational awareness, cross-sensor data fusion, and security operations across complex environments. Analysts, facility operators, and others that use PoL analysis often rely on heterogeneous sensing platforms, including fixed cameras, mobile systems, and UAVs, where vehicles are observed under varying viewpoints, distances, and environmental conditions. The ability to reliably associate these observations to the same vehicle, even in the absence of ideal imagery or overlapping camera coverage, enhances infrastructure protection, movement analysis, and anomaly detection. Moreover, the approach offers broad applicability across facilities because it does not involve controlled measurements, specific views, or specialized data collection. This adaptability makes it particularly suitable for retrofitting existing nuclear facilities, which often exhibit significant variations from one another in their design and operation.

Results and Accomplishments

In this work, the described matching concept was expanded to re-identify vehicles across a wide variety of viewing angles and distances, such as those that could be acquired from satellite imagery, UAVs, surveillance cameras, and ground sensors. To further enhance the robustness and efficiency of the proposed re-ID approach, synthetic data were integrated into the training set. Appropriate inclusion of synthetic data constructed under a wide variety of simulated viewing angles and environmental conditions can enhance the overall re-ID accuracy beyond what could be achieved with only collected data. The team developed software to render synthetic high-resolution and photorealistic vehicular images. The generated vehicular synthetic images spanned various conditions of lighting, elevation angles, and camera positions. Synthetic images combined with existing datasets of real imagery formed a robust training dataset. Additionally, an algorithm that matches multiview vehicle imagery was developed. The team assessed the performance of this vehicle-matching algorithm against simulated and real data. Several parameters were evaluated, including disparate view angles, illumination, and vehicle type.

UAV platforms were leveraged to collect multiperspective, real-imagery vehicle data for training and validation. This data acquisition simulated surveillance cameras and overhead assets over a wide variety of viewing configurations. On-demand UAV-based data acquisition is more flexible than fixed-sensor installations in terms of image resolution, direction, distance, and angle, and it can provide a continuum of different viewpoints for evaluation. Furthermore, UAVs can be used to simulate surveillance cameras and overhead assets. When used for wide-area surveillance, UAV imagery also provides the opportunity to

observe a given vehicle from multiple angles as it traverses the field of view, which enables techniques that can increase identification accuracy by leveraging additional visual information. This approach significantly stretches the limits of vehicle re-ID across disparate sensor views to provide a more robust capability for tracking.

The training dataset used in this study comprised three components: preexisting datasets, a synthetic dataset generated using Blender, and a collected dataset obtained during this project. This task employed publicly available surveillance benchmark datasets of real imagery. These datasets included VeRi-776, VehicleID (Peking University), and VeRi-Wild, as well as two UAV-collected vehicular datasets—namely, VRAI (Vehicle Re-identification for Aerial Image) and UAV-VeID. The team also used the PRIMAVERA (Profile Images and Annotations for Vehicle Re-identification Algorithms) dataset. This dataset contains a comprehensive collection of side-view vehicular images collected at ORNL and employed in previous research on vehicle re-ID. This dataset contains 636,246 images collected for 13,963 distinct vehicles, including daytime and nighttime captures, gathered over several years.

In addition to real imagery, the team generated a synthetic dataset of high-resolution, photorealistic vehicle images. The project team selected random pairs from all datasets for each training epoch, and the selection included equal numbers of positive/negative matches for vehicular pairs. The model in this study was trained across 90,000 epochs, and its performance was assessed on a fixed validation set of 10,000 varied vehicle pairs every 200 epochs. Training accuracy was measured by averaging results across 100 batches. Matching accuracy was evaluated with a threshold of 0.5 on the similarity scores, and scores below this threshold were considered true matches. The model accuracy was computed as the fraction of model predictions that were correct. To maintain consistency, a fixed validation set was used to evaluate the model's performance. The fixed validation set was selected randomly from data that were not a part of the training set. After 79,000 epochs, the training pairs of vehicle images that were mistakenly re-identified were saved in a separate folder. These image pairs were subsequently integrated into the retraining dataset used to fine-tune the network parameters. The performance of the refined model was evaluated on the same validation set employed in the previous phase to ensure consistency in comparison. After retraining across all combinations of viewpoints that shared common visible regions of the vehicle, the model achieved a matching accuracy of 97.7%.

11446: Materials Signatures of Metallic Phases

Z. Brubaker, J. Neu, C. Parish, D. Moseley, A. Wood, R. Petrova, J. Roach

Project Description

This project advanced several nanoscale characterization and synthesis capabilities important for understanding and interpreting metallic particulates. Key achievements included the development of high-throughput nanocrystallography using transmission Kikuchi diffraction, the implementation of nanotomography for 3D particle reconstruction, and the demonstration of atomic force microscopy modalities for physical property mapping—each contributing to a more comprehensive understanding of particle formation and behavior. To address the throughput limitations of advanced characterization techniques, the team developed and demonstrated particle relocation and multiparticle sectioning workflows. These methods enable simultaneous preparation and analysis of multiple particles, significantly reducing time and cost per sample. For 316L stainless steel, the team achieved five times better sectioning throughput and two to three times better electron diffraction throughput.

In parallel, analog synthesis capabilities were matured. Controlled synthesis methods now enable the production of high-fidelity uranium alloy samples, and combustion synthesis provides a flexible platform for studying oxidation behavior under extreme conditions.

Collectively, these efforts represent a step forward in the ability to characterize, interpret, and replicate the formation conditions of metallic particulates. The tools and workflows developed under this project could be used to support mission needs in nonproliferation, forensics, and materials attribution.

Mission Relevance

This project was initiated to advance high-resolution characterization of metallic particulates and analog synthesis, with the goal of enabling more robust attribution and supporting needs in nuclear security. Particulates are pervasive across fuel cycle operations and can provide valuable insights into conversion and production activities. For isotopically unperturbed particles, analysis generally must rely on chemical and physical characterization—often requiring micrometer- or submicrometer-scale spatial resolution. This project demonstrated advanced characterization techniques and novel synthesis approaches. These efforts enhance the ability to analyze and interpret metallic particulates, which supports DOE's mission in nuclear nonproliferation.

Results and Accomplishments

Advancing Nanoscale Particulate Analysis

Nanocrystallography. Crystallography is a powerful tool for understanding particle formation history, but traditional approaches using transmission electron microscopy (TEM) are slow, expensive, and limited in spatial coverage. To address these limitations, the team explored the use of transmission Kikuchi diffraction for nanoscale crystallography. In this diffraction method, particles were thinned in a manner similar to that of TEM preparation but were analyzed in a scanning electron microscope by positioning the sample above the electron backscatter diffraction detector. This approach significantly increased the mapping throughput compared with that of TEM and improved spatial resolution by approximately an order of magnitude over conventional electron backscatter diffraction.

Nanotomography. Conventional electron microscopy provides only 2D views of particles, focusing either on the surface for as-collected particles or the interior for focused ion beam–sectioned samples. Although effective for homogeneous materials, these methods can oversimplify the structure of heterogeneous particles. Additionally, they do not support volumetric quantification of crystal phases or void content. To overcome these limitations, the team developed nanotomography capabilities that enabled 3D reconstruction of micrometer-scale particles with nanoscale resolution. This was achieved by serial sectioning with nanometer-scale steps and reconstructing the slices using Avizo software. The process was optimized for use with a plasma focused ion beam system and included fiducial marking for accurate alignment. Nanotomography was successfully demonstrated on particles from applied projects. Using backscattered electron image contrast correlated with atomic number, the team reconstructed particles to reveal compositional heterogeneity. This technique enabled quantitative assessment of crystal phase fractions and void content, which are important for understanding particle transport behavior.

Throughput optimization. Nanocrystallography and nanotomography offer powerful characterization capabilities, but their throughput is limited by the need to individually locate, extract, and prepare each particle for TEM analysis. Preparing and analyzing a single particle typically requires 1 day of full-time effort and costs $4,000–$5,000 per particle. To address this bottleneck, the team developed methods for particle relocation and simultaneous sectioning, initially focusing on 316L stainless steel. Using either a probe tip or microgripper, particles were repositioned into engineered clusters based on substrate and particle characteristics. These clusters were then lifted out and sectioned together. This approach significantly reduced sample handling and registration time, resulting in five times better sectioning throughput and two to three times better TEM diffraction throughput for 316L stainless steel.

Enabling Analog Synthesis

A key objective of this project was to mature laboratory-scale synthesis capabilities for uranium-based alloys, enabling more routine and tunable production of analog materials. Two complementary synthesis approaches were pursued: controlled and combustion synthesis. These efforts focused on maturing early-stage synthesis capabilities into reliable, repeatable methods for producing and characterizing uranium-based alloys. Controlled synthesis was demonstrated on several uranium alloy systems; UNi_2 and UFe_2 + Ni are representative examples.

High-quality UNi_2 samples were produced by arc melting and annealing and were characterized at the Spallation Neutron Source and by using a physical property measurement system capable of handling radiological materials. This combination of advanced synthesis and characterization enabled direct comparison with computational models and demonstrated excellent agreement with prior computational predictions. In the case of UFe_2 + Ni, gram-scale ingots were produced via arc melting and subjected to systematic annealing to investigate the effects of thermal treatment on structural and physical properties. Multiple annealing protocols were explored to assess their influence on crystal structure, lattice constants, magnetic behavior, and phase fractions. This study helped establish processing–structure–property relationships critical for interpreting phase stability and material history. In both cases, characterization leveraged cross-directorate collaborations and represented a significant step toward routine, tunable synthesis of uranium alloys.

To complement controlled synthesis, the team extended combustion synthesis techniques from uranium oxide systems to uranium–steel compounds to enhance their relevancy for applied programs. Precursor materials from the annealing studies were used as inputs, which enabled the team to target specific compositions and observe their oxidation behavior under extreme conditions. This capability allowed the team to identify oxide phases that may form under rapid or high-temperature formation environments. This work ultimately informed material history and degradation pathways.

For example, combustion synthesis targeting $U_2NiCrFe_4$ resulted in rapid surface oxidation, which was confirmed by electron microscopy and Raman spectroscopy. Notably, the resulting oxides were predominantly iron-based, with no detectable UO_2—an unexpected outcome that highlighted the complexity of multicomponent oxidation behavior.

11517: Frictionless Knowledge Injection for Few-Shot Models

K. Dayman, J. Stomps, B. Welch, M. Morrell, C. Braden., C. Randolph

Project Description

The nuclear nonproliferation community commonly focuses on rare events, whose scarcity limits the applicability of modern data analytic models. Cutting-edge approaches to data analysis could enhance the national security enterprise if properly adapted to accommodate these ubiquitous sparse datasets. Although few-shot models have been shown to be reasonably effective in developing high-capacity models in limited-data scenarios, few-shot model performance is not comparable with models built on large datasets. Incorporating physics knowledge could improve the performance of few-shot models, but existing approaches to inject subject matter expert knowledge require data science expertise on the part of users, are limited in the types of features that may be injected, and produce one-off custom models that are not generalizable to different applications.

To remedy these shortcomings, this project adapted a recently proposed model explainability technique called the concept activation vector to allow a non-data-science expert to easily inject any type of domain knowledge into models and identify higher-level features of the data (as opposed to low-level features such as pixels in an image) learned by the model that may be more informative to a human analyst. To demonstrate the viability and potential efficacy of this new physics-informed, sparse data analysis workflow, these techniques were applied to two nonproliferation-relevant problems: discerning uranium minerals from Raman spectra and differentiating uranyl by-products via morphology information provided by scanning electron microscopy imagery. In both cases, few-shot learning improved analysis capabilities versus canonical supervised learning, and incorporating this project's additional knowledge injection methods further improved analyses. This study quantitatively showed the potential of these data modality agnostic analysis methods.

Mission Relevance

The capabilities developed under this project expanded the reach of modern data analytics and the effective data exploitation built on such models to previously untreated challenging national security problems, which is of relevance to DOE missions in nuclear security and nonproliferation. The project allowed models to be built on limited amounts of known data and easily incorporated complex subject matter expertise to increase the efficacy of these models and their trustworthiness. Related to DOE and ORNL missions in computing, this technology will enable rapid development of effective machine learning models by reducing the volume of data required to train models. This work will also allow nonexperts to easily inject proliferation subject matter expertise to improve performance and verify the models to leverage salient features identified by users.

Results and Accomplishments

The primary technical accomplishment of this project was the successful development of an easy-to-use knowledge injection capability using the concept activation vector, which used an example-driven approach to encoding arbitrary and potentially abstract knowledge into the embedding space of a neural network. Two methods for using the concept activation vectors for knowledge injection were developed and tested. The first method implicitly modified the embedding function by projecting encoded onto the space spanned by the vectors associated with the total collection of vectors and encoded concepts. The second approach modified the loss function used to train the model to encourage the model to respect the vectors associated with injected knowledge.

The efficacy and potential of these methods for improving model performance were demonstrated on two mission-relevant datasets. The first analyzed scanning electron microscopy data to differentiate uranyl phases via morphology, and the second differentiated between uranium minerals using Raman spectra. In both cases, data analytic models were developed using three to five samples per data class, and in both cases, each of this project's knowledge injection methods improved model performance over conventional few-shot learning methods.

Additionally, a user interface was developed to help general analysts navigate the workflow developed under this project. This user interface will increase the pool of potential users and increase the overall impact of the work.

SELF-DRIVEN EXPERIMENTS FOR SCIENCE/INTERCONNECTED SCIENCE ECOSYSTEM (INTERSECT)

10751: Enabling Adaptively Controlled Additive Manufacturing Through Automation

S. DeWitt, K. An, B. Turcksin, J. Haley, Y. Lee, T. Feldhausen

Project Description

The project team developed and demonstrated an autonomous, closed-loop system for adaptively controlling metal additive manufacturing (AM) using simulations and in situ monitoring. This system plans and adaptively controls an AM build process for targeted material properties that was codesigned with the INTERSECT crosscutting projects. To form a closed-loop control, this system leverages the INTERSECT ecosystem for automated coordination among the AM build system, simulations, and in situ characterization. The system consists of the AM build system, simulations, in situ characterization, a data assimilation tool, and optimal process planning software. This system was tested on two AM platforms: one under development at the Spallation Neutron Source (SNS) and another at the Manufacturing Demonstration Facility.

Mission Relevance

This project supported DOE and ORNL missions through developing and demonstrating a new approach for controlling metal AM processes using advanced simulation and in situ monitoring techniques. This technical breakthrough brought the team one step closer to meeting the grand challenge of born-qualified AM, which is important for enabling robust supply chains for US energy infrastructure as well as for US manufacturing dominance. By codeveloping the broader INTERSECT ecosystem and serving as an early exemplar of a successful interconnected science ecosystem, this project laid the groundwork for future innovation in AI-driven autonomous science.

Results and Accomplishments

This project resulted in the development of a new approach for using faster-than-real-time simulations to control metal AM and the demonstration of this new approach. Alongside this primary accomplishment, the project had a variety of secondary accomplishments related to the development of improved simulation and characterization tools and workflows.

Through this project, the team developed a new scheme for adaptive control of metal AM: sample-based model predictive control with a model reference. This new approach leveraged faster-than-real-time thermal simulations to plan, monitor, and control the printing process. The control scheme had two phases: an offline planning phase and an online control phase. In the planning phase, simulations were run in an optimization loop using the INTERSECT Dial microservice[1] for Bayesian optimization to find layer-wise printing parameters (e.g., wait time between layers and laser power) that yielded desired printing conditions (e.g., microstructure distribution). This optimization loop output the planned parameters for printing each layer and the reference solution—the target thermal evolution throughout the entire part. In the online control phase, the print started using the planned parameters. Because the part was printed, infrared camera data were collected to obtain surface temperature information. Pixels from the infrared images were cast onto the simulation mesh, and a Bayesian data assimilation method was used to incorporate information from those observations into a faster-than-real-time finite element simulation to form a real-time digital twin. The real-time digital twin provided a state estimate for the part as it was printed. Using the state estimate as initial conditions, a series of forward-looking simulations were initiated to investigate potential printing parameters for upcoming layers. Those forward-looking simulations were scored against the reference solution from the offline planning stage, and the parameters

[1] L. Drane, et al., *Dial* (Zenodo, 2025), https://doi.org/10.5281/zenodo.14872254.

from the best-scoring simulation were sent to the printer. This new control scheme was demonstrated during a printing campaign at the SNS VULCAN (Engineering Materials Diffractometer) beamline.

In the course of developing the new adaptive control capability, numerous enhancements were made to ORNL's Adamantine open-source simulation code.[2] Through a series of performance enhancements, Adamantine's performance on CPUs improved significantly and thus permitted faster-than-real-time simulations using only a few CPU cores (in some cases, a single CPU core). The team also improved GPU support in Adamantine, such as incorporating the Kokkos[3] performance portability library to support Advanced Micro Devices Inc. GPUs on the Oak Ridge Leadership Computing Facility's Frontier supercomputer. Adamantine was also modified to be more flexible, with added support for a de facto standard toolpath format and improved pipelines for meshing and toolpath generation. Initial support for thermomechanics was added for Adamantine, as well for thermoelasticity and thermoplasticity. Moreover, Adamantine's data assimilation capabilities were extended to include improved support for distributed memory parallel calculations.

Improvements to the Stereo Correlated Optical and Pyrometric System for in situ monitoring were also made over the course of the project. These improvements included interfaces to two printers to ferry control messages between the simulation-based controller and the printer, automated calibration routines, and enhancements to the user interface.

As part of the new control scheme demonstration, the team redesigned the OpeN-AM system for operando experiments at the SNS VULCAN beamline. In an initial experiment at VULCAN, SNS safety staff determined that an enclosure was needed to contain the activated oxide debris that flakes off the printed specimens. The redesigned version of OpeN-AM added an enclosure that was approved for use by the SNS safety team. Furthermore, the redesign incorporated a rotating sample stage, which enabled multiple printing experiments without opening the VULCAN beamline cave and dramatically accelerated experiment throughput.

The experiments at VULCAN included ex situ lattice strain characterization of the printed specimens to better understand the formation of residual strains that result from heating and cooling cycles during printing. This project involved collaboration with other INTERSECT teams to deploy a new INTERSECT-based adaptive sampling method using the Dial active learning microservice and the SNS External Instrument Control application programming interface. This new approach replaced an earlier method that involved installing machine learning libraries on the VULCAN control computer, which required privileged access and navigating the complex, immutable software environment on that computer. This new setup allows users to drive the beamline from their laptops using secure tokens for authentication and authorization. Users see projected strain maps update in real time as the experiment proceeds instead of waiting for postprocessed results after the experiment. The team also developed a new method for leveraging simulation results to accelerate lattice strain mapping.[4]

[2] B. Turcksin and S. DeWitt, "Adamantine 1.0: A Thermomechanical Simulator for Additive Manufacturing," *Journal of Open Source Software* 9, no. 102 (2024), https://doi.org/10.21105/joss.07017.
[3] H. Edwards, C. Trott, and D. Sunderland, "Kokkos: Enabling Manycore Performance Portability Through Polymorphic Memory Access Patterns," *Journal of Parallel and Distributed Computing* 74, no. 12: (2014), https://doi-org.ornl.idm.oclc.org/10.1016/j.jpdc.2014.07.003.
[4] S. Venkatakrishnan, et al., "Simulation Driven Adaptive Sampling for Neutron-Diffraction Based Strain Mapping of Additively Manufactured Parts," *Machine Learning: Science and Technology* 6, no. 4 (2025): 04037. https://doi.org/10.1088/2632-2153/ae0ff2.

11421: Self-Driven Experiments for Science/Interconnected Science Ecosystem Initiative Radiochemical Innovating Separations

A. AlNajjar, C. Greulich, P. Thakur, M. Du, L. Sadergaski

Project Description

Since 1946, ORNL's Isotope Program had maintained a distinguished legacy of supplying isotopes for medical, industrial, and research applications. However, the design, production, and purification of radiochemical products remain labor-intensive and involve multiple stages, including target design, chemical dissolution, and multistep separation processes, to achieve the desired product purity. Meeting stringent product specifications requires precise control over numerous process parameters. Factors such as pellet quality, separation efficiency, and product cut decisions significantly influence yield. Although traditional statistical methods have been employed to improve yields across various subprocesses, they often fail to fully exploit the high-dimensional data generated during production. As a result, valuable insights remain underused, even in advanced manufacturing environments. This project developed an instrument computing elemental separation ecosystem to transform isotope production by integrating several advanced technologies and methodologies: (1) data analysis to optimize radiochemical separation processes, (2) real-time optical and gamma spectroscopy for in situ monitoring within shielded hot cells, (3) digitization and structuring of historical and real-time data to enable predictive modeling and autonomous decision-making, and (4) software–hardware integration to enable automated control of processing conditions.

Mission Relevance

This project was relevant to advancing isotope production at ORNL. It addressed national needs related to DOE missions in medical, industrial, and research applications of radioisotopes; national security and space exploration, in which isotopes such as ^{238}Pu are used; reduction of reliance on foreign suppliers by enhancing domestic production capabilities; modernization of legacy processes by integrating AI/machine learning, real-time monitoring, and automation; and improvement in safety, efficiency, and scalability of radiochemical separations in extreme environments (e.g., shielded hot cells).

Results and Accomplishments

During this project, several tasks were accomplished toward supporting autonomous isotope separations.

Isotope Separation Ecosystem Design

This project designed an ecosystem to support autonomous and remote control of chemical separations across multiple facilities. In this ecosystem, the instruments for optical spectroscopy and the column separation workstation are in a user facility, and the remote control and analytics of spectra are conducted remotely from a computing facility located in a different facility separated by multiple domain firewalls.

The ecosystem relies on ZeroMQ in a request-reply paradigm, in which a client (i.e., the user on their remote computer) issues commands via specially crafted messages to a server in the lab that executes those commands on the connected scientific equipment. A relay server was used so that the instrument control computer can be disconnected from the external network to maintain instrument control safety and reliability.

Remote Control Plane for Chemical Equipment

A software stack was developed to remotely acquire and process optical spectra from Ocean Optics–based spectrometers. Based on the signature of spectrum of the metals, different actions can be performed autonomously.

Another software stack was developed that controls a dual eight-port syringe pump (J-Kem). The stack works by using Python to pass serial messages to a program that J-Kem developed and allows for selecting the port to use, the seed of the pump, and the amount of fluid to move.

A final software stack was developed that controls an 80-vial fraction collector (J-Kem) that, in conjunction with the pump, stores and processes 20 mL samples of various reactants and elutes. Importantly, the combination of the two J-Kem equipment enables arbitrary dilutions to control the acidity of the solution before it is pumped into a column.

A series of workflow scripts was developed that performs an experiment and displays the results to the user via an interactable Jupyter lab notebook.

This project also developed a first-of-its-kind, open-access digital platform designed to function as a centralized repository for radiochemical separations data. This platform integrated a range of data sources, including experimentally derived results, curated datasets from peer-reviewed literature, and preexisting open-access fundamental databases. The dataset contains the results of nearly 10,000 individual experiments.

11521: Multi-Workflow Orchestration and Lightweight Integrated Data Analysis Across Facilities

R. Santos Souza, L. Casey, M. Coletti, R. Ferreira da Silva, O. Kuchar, T. Skluzacek, F. Suter, T. Bullock, R. Moore II, J. Woods, T. Naughton III, M. Brim, T. Jones, M. McDonnell, R. Prout, A. Ruckman, G. Wiggins, D. Rosendo, K. Evans

Project Description

Researchers in many fields increasingly rely on end-to-end campaigns that connect experimental instruments, edge resources, and leadership computing facilities, often across organizational and geographic boundaries. A representative pattern is a feedback loop in which an instrument generates high-rate data streams that are preprocessed at the edge and transferred to high-performance computing (HPC) systems for large-scale AI analysis; results are used to guide and retune the experiment in near–real time. Today, these campaigns are typically assembled as disconnected workflows, and scientists must manually coordinate execution, data movement, and context across steps and facilities.

This project advanced a system-of-systems approach aligned with INTERSECT by combining (1) multi-workflow orchestration across facilities and (2) lightweight integrated data analysis at runtime. The technical approach leveraged the Zambeze framework as a control-plane controller for cross-facility workflow automation and FlowCept as a complementary capability that interfaces with data and control planes to capture runtime provenance, metadata, and references to large data objects. Together, these capabilities enable a multi-workflow to be treated as a single, automated science campaign that provides an integrated view of what has run, where it ran, what data were produced, and how outputs related while the campaign was executing.

Mission Relevance

This project supported DOE and ORNL missions by strengthening the technical foundation for trusted, scalable, and automated cross-facility science. DOE's mission emphasizes addressing national challenges through transformative science and technology solutions, and ORNL's mission emphasizes scientific discoveries and technical breakthroughs that accelerate the development and deployment of energy and global security solutions. Cross-facility campaigns that integrate instruments, edge resources, HPC, and AI are important to achieving these outcomes, but they require robust orchestration, reproducibility, and responsible data management across heterogeneous environments.

By enabling a multi-workflow campaign to run as a systematic, automated process and providing lightweight runtime data and provenance capture to improve transparency and reproducibility, this project

helped reduce friction in AI-enabled experimental science. This capability can directly improve the reliability and throughput of computationally accelerated discovery workflows that depend on ORNL user facilities and leadership computing systems.

Results and Accomplishments

This project delivered and validated key elements of an INTERSECT-aligned cross-facility workflow capability that spanned orchestration, runtime data capture, and applied demonstrations on leadership computing systems.

A major outcome was integration into the INTERSECT ecosystem, including development of a secure scientific service mesh application programming interface adapter to streamline task orchestration between domain science labs and leadership computing facilities. This work enabled early end-to-end communication patterns and helped establish practical pathways for composing cross-facility campaigns using INTERSECT message layers and application programming interfaces. The project achieved successful integration of the Oak Ridge Leadership Computing Facility with the INTERSECT ecosystem using the appropriate interfaces and messaging layers, which provides a concrete foundation for future multi-workflow campaigns that span facilities.

The project also demonstrated progress in cross-facility orchestration in a manufacturing context by establishing an initial INTERSECT-based communication path between a manufacturing use case and Defiant (Integrated Research Infrastructure test bed).

In parallel, the project advanced large-scale validation activities on Frontier that demonstrated the feasibility of combining distributed automation with scalable AI workflows. The project completed an evolutionary algorithm–based hyperparameter tuning workflow for AtomAI, which was executed on a large-scale experiment on the Frontier HPC system. This capability demonstrated the project's ability to coordinate distributed AI workflows at scale. The project also explored approaches to assess trade-offs between energy efficiency and model quality during large-scale deep learning development on leadership computing systems and conducted early experiments in runtime provenance and performance data management for large model training runs. Together, these activities strengthened the practical basis for responsible and efficient AI workflows in HPC contexts while aligning with the project's integrated control and data campaign view.

Overall, the project demonstrated that combining multi-workflow orchestration with lightweight, integrated runtime data capture is a viable and effective strategy for reducing manual coordination across facilities.

11540: Scientific Data Layer for Self-Driven Experiments for Science/Interconnected Science Ecosystem Initiative

S. Boehm, P. Widener, D. Stansberry, S. Hitefield, A. May, T. Jones, C. Engelmann, O. Kuchar

Project Description

This project defined an open scientific data layer (SDL) for the laboratory of the future. The SDL prototype leveraged the INTERSECT microservice specification using uniform application programming interfaces to support cross-facility and edge computing scientific use cases as well as reduce the time, effort, and complexity associated with producing and consuming scientific data regardless of data location or format. When possible, the implementation supported automatic capture and association of metadata and provenance with data sources to aid in discoverability and interoperability. A minimum viable implementation provided data services to find and retrieve data. The overall objectives for the SDL were to enable simple; uniform; efficient; and findable, accessible, interoperable, and reusable access to scientific data sources across facility boundaries and at all stages of the data life cycle.

Mission Relevance

As AI becomes increasingly used for scientific research, the quality of the data underlying AI models is a focus. Scientific data management helps ensure data quality and suitability for AI use. Although the quality of data is highly domain-dependent, many of the steps of managing the data are not. Providing a framework to quickly implement and deploy an SDL for any type of research advances DOE's mission across its energy, environmental, and national security mission spaces.

Results and Accomplishments

Aim 1 had the goal of producing a catalog of scientific use pattern, which was planned to be based on surveys of the data use patterns of the INTERSECT domain science projects. Aim 2 focused on the architectural framework of the SDL. Aim 3 focused on a prototype implementation for the SDL.

Aim 1: Catalog of Scientific Usage Patterns

The goal for aim 1 was to catalog the different ways in which scientific data are used by the INTERSECT domain science projects. Understanding how data are created and used across the different projects helps to identify commonalities and differences across projects. It was anticipated that this would reveal insight into the more common usage patterns and help with designing a modular system that can support science use cases across all core areas of expertise. To develop a catalog of usage patterns, this project created two questionnaires. The first was aimed at understanding the system generating the data, and the second aimed at understanding data usage. The results of the platform questionnaire were used to create the ontological description, based on the Semantic Sensor Network Ontology, of the platform generating data, which becomes an integral part of the provenance data for any dataset generated by a given platform.

Aim 2: Design of SDL Services Within the INTERSECT Architectural Framework

The SDL services were designed as part of the INTERSECT architectural framework and implemented using a microservice architecture. The SDL relied on established ontologies to support data cataloging with the Data Catalog Ontology and tracked provenance data using the Provenance Ontology. To capture the origin of the data, the provenance records pointed to the platform that generated the data using the Semantic Sensor Network Ontology.

The SDL was based on a layered architecture consisting of three tiers. The first tier abstracted the infrastructure layer and consisted of a storage management service, data adapter services, and data movement services. The storage management service was responsible for abstracting storage infrastructure, such as file systems and object storage, and managing the corresponding metadata. Data adapter services were used to abstract structured data in relational or document databases, and the data movement services were responsible for moving data between different systems. The second tier consisted of data registry and data deployment services. Data registry services stored additional metadata such as provenance and platform descriptions. The third tier provided data catalog and data repository services. Although data repositories manage data assets and the provenance records, the data catalog service was responsible for providing metadata for the data assets. Extensive documentation is available.[5]

Aim 3: Proof-of-Concept Prototype

The prototype was implemented within the Autonomous Chemistry Laboratory. The Autonomous Chemistry Laboratory is developing capabilities to automate solid-state chemistry. It can synthesize compounds and automatically analyze samples of the compound using various instruments, such as an x-ray diffractometer or mass spectrometer. The prototype provided a data registration service/data catalog service for the different files produced during a campaign. It also provided the means to track samples and

[5] "INTERSECT Scientific Data Layer (SDL)," Oak Ridge National Laboratory, GitHub, https://intersect-sdl.github.io/.

associate the data produced with a particular sample. Finally, the prototype provided a simple web frontend that allowed a user to browse the available raw data. The code is available.[6]

11739: Large Language Models for Scalable Design of Gradient Copolymers

I. Ivanov, R. Moore II, F. Wang, M. Goswami, R. Vasudevan, J. Yin, R. Advincula, P. Balaprakash

Project Description

This project integrated large language models (LLMs) with large-scale molecular dynamics simulations to rapidly iterate on chemical synthesis recipes for a continuous-flow chemistry approach, called AuotFlowS. This work aimed to enable autonomous and scalable generation of new gradient copolymers with desired properties. The hypothesis was that LLMs trained on appropriate polymeric scientific literature could be used in conjunction with molecular dynamics simulations to better predict the outcome of real experiments on AutoFlowS and therefore significantly accelerate searching through the synthesis trajectory space. To demonstrate the impact, this project targeted gradient copolymers because of their ability to show unique properties such as interfacial behavior and thermal and mechanical properties in bulk and solution. Traditional synthesis of gradient copolymers is complicated and generates very limited amounts of material, which is not optimized for a particular property. Furthermore, because gradient copolymers possess a gradual change in composition along the chain backbone, they have varied monomer sequence distributions. This variety means that a vast monomer sequence design space exists, which is impractical to probe using traditional synthesis. Using this project's approach with LLMs, molecular dynamics, and AutoFlowS, synthesis was elucidated to enable precision gradient copolymers at scale with optimized mechanical and thermal properties.

Mission Relevance

This project solved a challenging problem in synthesis science. It also provided a means for facilitating practical applications of gradient copolymers to large classes of monomers and properties. The developed approach can be applied to multiple DOE mission-relevant areas, including catalysis, energy storage, chemical separations, sensing, and transformative manufacturing.

Results and Accomplishments

The general task of integrating an LLM into an organic synthesis laboratory with targeted synthesis of gradient copolymers and predictive modeling of their thermal properties changed with the fast evolution of new LLMs and retrieval-augmented generation LLMs that connect LLMs with external knowledge bases (e.g., documents, databases) for referenced, context-specific information retrieval while overcoming hallucination problems. In the 2 years of this project, the AI/machine learning approaches to R&D evolved toward the concept of agentic AI, in which instead of integrated, complex structures comprising interdependent units, each subtask is accomplished by an independent agent. This approach is equivalent to a subroutine practice in programming in which independent subroutines could be reused or reintegrated in different workflows. This project adapted program efforts to keep the work in line with the development of the field. Each autonomous agent retrieval-augmented generation LLM or tools agent integrated Python software and relevant knowledge data (publications) to enable contextual reasoning.

This project identified publications and Reaxys Elsevier protocols of 5,820 relevant copolymerization reactions (extracted from the publications and patents) from the database, including 247 copolymerization reactions of styrene and butyl acrylate. This collection of literature and extracted copolymerization was used to generate a training dataset based on question/answer to train the FORGE LLM and incorporate into an OpenAI ChatGPT Polycookbook chatbot. The performance of the chatbot was tested against

[6] "intersect-sdl," GitHub, https://github.com/intersect-sdl.

100 questions on copolymerization, which were generated separately using the most recent publications not included in the training dataset. This effort also tested prompt engineering.

This project then designed a digital twin of a copolymerization reaction by integrating the thermally initiated Mayo–Lewis reaction mechanism as a physics-based model that integrated kinetics of propagation, termination, and chain transfer. The digital twin was used to generate an extended dataset for machine learning–based modeling of reactions and extract dependency of the monomer concentration and temperature to maximize the yield and efficiency of the reaction. For incorporating molecular weight modeling into the digital twin, this project used the kinetic Monte Carlo stochastic modeling approach. This work also incorporated special cases of copolymerization; for gradient, random, and ladder polymers, the selection was based on the reactivity ratio consideration. The model considered a monomer acting as a solvent. The project team observed a net effect of the chain transfer agent reducing or controlling the average molecular weight of the polymer because it increased the number of chain termination events relative to propagation events. The results of this digital twin correlated with the results previously published in the literature. This model allowed visualization of the interdependency of reactor parameters in the form of a correlation matrix for each copolymer structure. The copolymerization digital twin was further extended to incorporate the difference in the reaction design using continuously steering tank reactor (batch copolymerization) and plug flow reactor models along with the Mayo–Lewis reaction mechanism; this work had a final goal of maximizing yield and efficiency using Bayesian optimization.

This project designed a continuous flow system called Train of Reactors (TOR), which enabled carrying out one or multiple reactions on a series of milliliter-scale samples. Two liquid plugs were separated by inert gas as a spacer. The flow was enabled by inert gas as well as opened or closed valves. TOR incorporated external and internal measurement units that measure properties of a liquid in a traveling drop. The internal unit used a conductivity flow cell sensor, and the external unit could be any diagnostic instrument. The function of the measurement units was to identify if a specific condition (e.g., concentration of product) was reached. After achieving this condition, the reaction train was guided to the autosampler for collection of each drop into a separate vial. If the condition was not achieved, the reaction train was guided back into the reactor. The external measurement unit, whose operation was enabled by the Python script, processed the signal and identified if the condition was achieved. If positively identified, the software sent a trigger signal to general-purpose input/output pins of the data acquisition device, which generated a digital electrical signal (Transistor-Transistor Logic) pulse of 5 V and triggered a sequence of closing or opening valves for redirecting liquid. Gas and liquid flows were enabled by two syringe pumps, with an option to incorporate up to three solvents.

11819: Edge-to-Ecosystem Connect Integration

A. AlNajjar, N. Rao, T. Naughton, G. Wiggins

Project Description

Instrument-computing ecosystems (ICEs) are increasingly being deployed and have enabled scientific discovery, particularly by employing AI-driven workflows to steer experiments and provide real-time analytics. Complex ICEs for DOE applications typically span experimental facilities with instruments and high-performance computing (HPC) systems, which are dispersed across various geographical locations and separated by network domains. The most effective use of ICE resources requires end-to-end integration to support real-time instrument control and measurement transfers to and from HPC facilities for simulations and analytics. Several ICE integration issues must be addressed to support end-to-end self-driven experiments and real-time analytics, including addressing the multidomain network firewall to support transferring data and control messages, integrating local instruments with programmable interfaces to support workflow automation, addressing job scheduling for on-demand analytics on HPC systems, and developing workflow orchestration to support resilient end-to-end workflow execution.

This project used the Edge-to-ecosystem ConNect Integration (Eco-NetI) framework that addresses end-to-end ecosystem integration of science instruments and HPC resources to enable autonomous scientific workflows across ICE facilities. Eco-NetI realizes advanced ICE infrastructures with novel solutions to enable cross-facility workflows that span science instruments and leadership computing facilities across ORNL and beyond across DOE facilities. This project team targeted the computing platforms at the Oak Ridge Leadership Computing Facility (OLCF) at ORNL and Argonne Leadership Computing Facility (ALCF) at Argonne National Laboratory as use cases for building end-to-end workflows with resilient computation.

Mission Relevance

Using Eco-NetI, this project addressed ICE infrastructure integration aspects of software and hardware requirements to support end-to-end autonomous workflows—in particular, across multisite DOE facilities. The integration solutions offered by Eco-NetI are essential to accelerate the development and deployment of various science discoveries and energy missions across different science fields relating to energy, environmental, and security challenges. Eco-NetI contributed to DOE's mission in (1) supporting building multisite ICEs of heterogenous instrument and computing platforms by providing expertise and software and hardware solutions to design and develop these ICEs and (2) demonstrating the capability to integrate scientific laboratories as part of generic and large infrastructure to expedite the speed of discovery.

Results and Accomplishments

In this Eco-NetI project, end-to-end integration hardware and software solutions were designed and developed to integrate edge systems (e.g., instruments, control platforms, computing workstations, and clusters). The developed solutions are necessary to enable self-driving experiments and computations for science discovery.

Edge-to-ICE Infrastructure Integration and Support

The Eco-NetI team designed and developed integration solutions to support instrument edge integration with the domain network. This task involved developing instrument control planes of programmable Python wrappers to support remote instrument control as part of closed-loop workflow orchestration. It also included designing and developing two-tier edge-to-ICE integration, in which the instruments are indirectly controlled from a remote site through a gateway system deployed at the instrument facility. The Eco-NetI team also designed and developed data transfer solutions as part of instrument integration, in which the instrument measurements are transferred to computing platforms in real time.

Interdomain Integration with Leadership Computing Facilities

The Eco-NetI team designed and developed a scientific workflow framework that integrates instruments with leadership computing platforms (such as the OLCF Advanced Computing Ecosystem test bed at ORNL and the ALCF Polaris at Argonne National Laboratory) for on-demand analysis to support autonomous workflow tasks of instrument steering and workflow orchestration. In Eco-NetI, the ORNL–OLCF–ALCF domain network firewalls and network traffic access were addressed to enable real-time instrument control and data messages passing, as well as workflow tasks execution.

Support for a Variety of Domain Projects

Other than the main tasks of Eco-NetI, the team members worked closely with domain projects to support their workflow integration with the ICE. Particularly, the team developed customized integration solutions based on edge systems and instruments anticipated in the domain workflows. The solutions included developing instrument control wrappers for remote control, developing machine learning for domain analysis, and building ecosystems infrastructure to support instrument integration in outdoor fields.

TRANSFORM INITIATIVE

11429: Energy-Efficient Removal of Oxygen from Carbon Dioxide Streams

H. Chen, T. Toops, S. Sinha Majumdar, Y. Park, S. Islam, F. Boukouvala, C. Sievers

Project Description

This project aimed to develop energy-efficient and cost-effective technologies to purify carbon dioxide streams to meet the recommended impurity specifications for safely transporting carbon dioxide through pipelines. Specifically, the project focused on removing oxygen and water impurities (typically found in captured carbon dioxide streams) to below the limits of 10 and 500 ppm, respectively. A variety of catalytic oxygen reduction (COR) approaches using different reducing agents, including hydrogen, carbon monoxide, methanol, and methane, were evaluated for the deep removal of oxygen for comparison with commercial automotive exhaust emission control catalysts. The results demonstrated that with precise control of the amount of the reductant added into the feed, hydrogen, carbon monoxide, or methanol could effectively lower oxygen from 1.5% to less than 10 ppm with a single-reactor design. This would achieve above 99.93% oxygen removal efficiency and simultaneously meet the impurity limits for the added reagent or their reaction intermediates. With methane as the reductant, the single-reactor design could not meet all the targets. Thus, a dual-reactor design was invented. Detailed techno-economic analysis (TEA) and life cycle assessment (LCA) studies showed that the methane route offers the lowest purification cost and highest carbon dioxide recovery, followed by the methanol, carbon monoxide, and hydrogen routes. This order is mainly affected by the material cost of the reductants. The TEA and LCA results also indicated that water removal has a significant effect on process economics and environmental footprint. Therefore, it is most cost-effective to perform dehydration after COR, where compression and condensation steps can be leveraged for simultaneous water removal and compression.

Mission Relevance

Carbon dioxide is a valuable resource for chemical synthesis, the food industry, and enhanced oil recovery. High-purity carbon dioxide is a requirement for these applications as well as for safe transport of carbon dioxide through pipelines. The invention of a dual-reactor system in this project is a technical advancement in this research field. It enables the use of methane, a readily available and low-cost material, to purify carbon dioxide streams to meet the impurity specifications required for pipeline transport and offers the lowest environmental impact. This technology may accelerate the deployment of carbon dioxide purification processes, creating economic opportunities for the nation.

Results and Accomplishments

Evaluation of COR Approaches with Dry Feeds

This project initially assessed a variety of technologies for deep oxygen removal from carbon dioxide streams. COR was identified as the most promising lead that could potentially be applicable to a wide range of applications. In this study, a set of commercial automotive exhaust emission control catalysts were selected for their technical maturity, high activity, and excellent stability. Four different reagents—hydrogen, carbon monoxide, methanol, and methane—were tested as the reductants.

With hydrogen, the COR reaction could be initiated at an inlet temperature of approximately 60°C, making it highly compatible with low-temperature carbon dioxide capture processes. By controlling the amount of hydrogen added to the feed at the stoichiometric point or with a slight excess, the oxygen impurity level can be lowered to less than 10 ppm. A major concern about using hydrogen as the reductant is the associated material cost. This was determined and evaluated in the TEA/LCA analysis, which is discussed below.

With carbon monoxide, the COR reaction occurred around 150°C. Controlling the amount of carbon monoxide added to the feed at the stoichiometric point was important to simultaneously meet the oxygen and carbon monoxide specifications. Deviations from the stoichiometric point can lead to oxygen or carbon monoxide breakthrough that exceeds the limits.

With methanol, the COR reaction occurred at approximately 250°C. The amount of methanol introduced to the system must be carefully controlled so as not to exceed the stoichiometric point to ensure that its concentration in the treated carbon dioxide stream remains below 1 ppm. Catalysts with oxygen storage components can mitigate small fluctuations and meet the specification targets.

With methane, a much higher reaction temperature of approximately 650°C was required to activate the reagent and maintain conversion stability. Under these reaction conditions, methane can react with carbon dioxide and generate a considerable amount of carbon monoxide. Consequently, with a single-reactor design, the system cannot simultaneously meet the oxygen and carbon monoxide targets. A dual-reactor system was thus invented in this study. In the first reactor, methane removes the majority of oxygen in the feed while a portion also reacts with carbon dioxide, producing carbon monoxide. By adjusting the amount of methane added to the system, the concentrations of the produced carbon monoxide and the residual oxygen in the outlet gas from the first reactor can be tuned to achieve a 2:1 molar ratio. The outlet gas is subsequently cooled and fed to a second reactor, in which carbon monoxide and oxygen react with each other to reach the levels below the pipeline specifications.

TEA and LCA Studies of the COR Approaches and Water Removal Processes

TEA and LCA analyses of the four COR pathways were performed by collaborators at the Georgia Institute of Technology. Process flow sheets, including water removal, were developed and simulated in Aspen Plus for carbon dioxide streams of varying water amounts and feed gas temperatures that were representative of low- and high-temperature capture processes. The results showed that operating expenses dominated the total cost, accounting for an average of 89% of the levelized purification cost across all routes. The oxygen removal step accounted for 15% to 39% of the overall levelized costs, with the remaining 61% to 85% attributed to the dehydration section.

Among the four COR routes, the methane route achieved the lowest purification cost and highest carbon dioxide recovery, followed by the methanol, carbon monoxide, and hydrogen routes. For the hydrogen route to be competitive, the prices for hydrogen must be below approximately $0.56 or $0.84 per kilogram, depending on the conditions of the carbon dioxide feed. As a reference point, the current price of hydrogen is approximately $4 per kilogram, making the hydrogen route the least economically favorable.

Capital costs were comparable across all routes, ranging from $2.2 to $2.5 per ton of carbon dioxide. Of the total capital investment, the compressor costs associated with the drying step accounted for 87% of the bare module cost, highlighting that the compression required to meet water content specifications constitutes the largest capital investment rather than the catalytic reactor for oxygen removal. Additionally, the capital and operating costs for the dehydration step were relatively insensitive to the variation in the water content in the feed gas (within the 5% to 15% range). Therefore, from a process design standpoint, performing drying after catalytic oxygen removal is most cost-effective because compression and condensation steps can be leveraged for simultaneous water removal and compression.

Evaluation of COR Technologies with Wet Feeds

Although the choice of reducing agent affected oxygen removal costs, the downstream water removal section remained the primary cost driver for carbon dioxide purification in capital and operating costs. Because most COR approaches generate additional water (except the carbon monoxide route), removing oxygen before the dehydration processes is more energy-efficient and cost-effective. This means that COR reactions should be carried out in the presence of water. Thus, additional COR evaluations with wet feeds that contained 15% water were conducted. The results showed that the COR processes were very

robust and insensitive to the water content in the feed. All four COR approaches can meet pipeline specifications.

11462: Electrochemical Seawater Decarbonization Advancement

J. Szybist, F. Dal Forno Chuahy, T. Toops, A. Serov, X. Lyu, V. Lobodin, S. Whitted, J. Willocks, M. Hatzel, H. Choi, H. Song, J. Scott, Y. Kim

Project Description

This effort evaluated the ability to scale up electrochemical processes to remove inorganic carbon from seawater. State-of-the-art electrochemical technologies for removing inorganic carbon from seawater show promising results in terms of energy input required[1] but are at a low technology readiness level. Many questions remain about the ability to scale these technologies, their durability, and improvement opportunities using different materials, designs, or operating techniques. This project established a facility at ORNL that allowed a tenfold scale-up in electrode area and experimental duration to evaluate durability. Complementary low-dimensional (1D and 0D) modeling focused on developing a better understanding of the chemistry, including the related pH of each process stream, and the balance-of-plant and economics. The project focused on a three-chamber, dual-membrane electrochemical cell with a hydrogen looping strategy in which the form of inorganic carbon was shifted from bicarbonate to gaseous by shifting the pH acidic. Progress was made in terms of understanding the challenges associated with precipitation from alkaline seawater, the carbon removal process from acidic seawater, and the balance-of-plant costs.

Mission Relevance

This project advanced the DOE mission of ensuring America's security and prosperity by addressing its energy and environmental challenges. Sequestering carbon from seawater will improve the health of the ocean by increasing its pH, which has dropped 0.1 due to the absorption of carbon dioxide from the atmosphere.

Results and Accomplishments

A laboratory at ORNL was established to investigate the electrochemical removal of inorganic carbon from seawater at scales that are an order of magnitude larger than the where the initial proof-of-principle work is conducted. Conducting these experiments is useful for understanding scaling challenges and energy inputs associated with the balance-of-plant. The laboratory facility consisted of an 1,100 L supply tank with synthetic seawater, formulated in accordance with the ASTM D1141-98 synthetic seawater specification. The technology that was investigated used a three-chamber electrochemical cell: the anode compartment consumed hydrogen, which was fed by a mass flow controller. The center and cathode compartments contained seawater and were fed by separate parallel metering pumps so that their respective flow rates could be controlled separately. A potentiostat was used to supply electricity to the cell in a constant current configuration with a high voltage shutoff of 5 V direct current.

An electrochemical cell for this application was designed and fabricated from titanium to avoid any possible corrosion issues that could arise with ferrous metals. The cell was designed to have an electrode area of 400 cm^2, which represented an order of magnitude scale-up from the demonstrated technology in the literature. A provisional patent was filed for the design of the center compartment of the cell, which included a novel seawater static mixing technology. The cell was constructed using treated Nafion membranes to separate the anode compartment from the center compartment and the center compartment from the cathode compartment. For the anode, the platinum/carbon catalyst was applied to carbon paper, and for the cathode, the platinum/carbon catalyst was applied to nickel foam. Titanium felt was used to provide structural support for the Nafion membrane in the interface with the central compartment. The

[1] L. Yan, et al., "An Electrochemical Hydrogen-Looping System for Low-Cost CO_2 Capture from Seawater," *ACS Energy Letters* 7 (2022): 1947–1952, https://doi.org/10.1021/acsenergylett.2c00396.

cell was assembled so that no conductivity occurred through the metal casing by using Viton gaskets to seal and electrically isolate the compartments and using nylon inserts for the bolts that fasten the cell together.

The operating principle of the electrochemical cell is that hydrogen ions are produced from the hydrogen in the anode and transported into the center compartment, which makes the water acidic and changes the form of the inorganic carbon, making it removable. Correspondingly, hydroxide ions are produced and transported to the cathode compartment, making it alkaline. During the initial experiments of the electrochemical cell, the operating principle of the electrochemical cell was successfully demonstrated. The pH of the center compartment decreased to a value of <4, and the pH of the cathode compartment increased to a value of >9. However, the voltage potential across the cell did not stabilize but instead progressively increased for a given amount of current. After removing the voltage potential, allowing the system to rest for several minutes, and then reapplying the voltage, the voltage returned to a lower value but then increased again. After a number of attempts, the electrochemical cell was disassembled, and the project team found that although the anode and center compartments were nearly free of deposits, excessive precipitation occurred from the seawater in the cathode compartment. Chemical analysis revealed that brucite (magnesium hydroxide) was the primary component of the precipitate. Further investigation revealed that brucite precipitation began at a pH of roughly 9.0.

The first option to control the alkalinity of the cathode compartment to mitigating brucite precipitation was to implement an alternative flow path for the electrochemical cell. In this configuration, the acidic stream was removed from the center compartment, the inorganic carbon was removed under vacuum, and the acidic seawater was then recycled back to the cathode compartment. The hypothesis was that starting with a lower initial pH at the inlet to the cathode would prevent high alkalinity. This process was tested experimentally, and this work found that the hypothesis was incorrect: starting with a lower pH did not successfully attenuate the alkalinity or the brucite precipitation. Modeling was conducted to understand why, and the project team found that the inorganic carbon content, which was removed from the recycled stream, was critical to limiting the pH increase. A series of buffering reactions involve the inorganic carbon–containing species that consume the free hydroxide ions in the cathode, and without these species present, the increase in the pH occur much more rapidly. Thus, this alternative flow path was determined not to be a feasible method of attenuating alkalinity.

The second option to control the alkalinity of the cathode compartment was to increase the flow rate of seawater in the cathode compartment to a higher value than the center compartment. Although this technique is effective at modulating the alkalinity, this work found that the buffering reactions make it such that the pH rise with the amount of seawater is nonlinear. Thus, the seawater flow rate through the cathode compartment would have to increase by more than a factor of five over the seawater flow rate through the center compartment to prevent precipitation, which plays into the economic feasibility.

In the process of determining how to attenuate the alkalinity increase of the seawater in the cathode compartment, this project also addressed challenges associated with removing the inorganic carbon from the acidified seawater from the center compartment. Specifically, even though the form of the carbon had been shifted in the electrochemical cell to something that was more easily removed from the seawater, gas/liquid equilibria dictated that this was not a species that was readily separated into the gas phase. Instead, a vacuum of 0.04 atm had to be applied to the acidified seawater to recover 80% of the inorganic carbon. Even under those circumstances, the purity of the inorganic carbon was approximately 20%, with the balance being water. For the purposes of this project, a single-stage vacuum separation system was used, but to inform the economic analysis, a multistage separation process was designed for the economic analysis that included purification of the inorganic carbon.

Finally, the techno-economic feasibility of this electrochemical technology was assessed on an industrial scale. The techno-economic model considered seawater pumping, operation of the electrochemical cell, and separation processes to remove the inorganic carbon from water that had been acidified in the

electrochemical cell. The model also included capital costs. It was found that the operation of the electrochemical cell amounted to 20%–25% of the overall cost. In contrast, in a plant design in which the cathode and center compartment had equal flows, the pumping cost would amount to 40% of the overall cost, and another 15% of the cost would come from operating the inorganic separation processes. The balance of the overall costs can be attributed to capital costs. Thus, the balance-of-plant costs that were identified in this project comprise the majority of the operating expenditures for operating such a plant on an industrial scale.

11498: Coupled Ecosystem-and-Engineering Decision-Making Framework for Enhanced Weathering

J. Mao, Y. Wang, M. Uddin, F. Yuan, X. Gu, D. Ricciuto, X. Shi, I. Busch, D. Lu, N. Singh

Project Description

Enhanced rock weathering (ERW) is a promising carbon dioxide removal (CDR) strategy that accelerates natural carbon dioxide capture by applying finely ground silicate rocks over land surfaces. Previous assessments of ERW potential and economics were limited by incomplete treatment of geochemical–biological feedbacks and weak coupling between ecosystem processes and supply-chain logistics. This project developed a scalable modeling framework integrating Earth system and engineering approaches to optimize long-term ERW deployment across the contiguous United States (CONUS). The research team implemented a new ERW module within the Land Model of the DOE Energy Exascale Earth System Model (ELM) to represent key geochemical and biological feedbacks and built a machine learning–based surrogate model to accelerate ensemble simulations of the coupled ELM–ERW system. In parallel, the research team developed a source–sink optimization model to account for application rate, areal extent, grain size, and transport matching. Two-way coupling between the surrogate and optimization models established a unified decision-making framework. Using this integrated system, the team quantified the CDR potential and economic costs of ERW under multiple environmental and socioeconomic scenarios. The resulting models and decision-support tools enable improved assessment of ERW feasibility and provide a foundation for broader CDR strategy evaluation at regional to global scales.

Mission Relevance

This project supported DOE's mission of advancing transformative science and technology solutions for environmental challenges. By integrating geochemical, biological, and socioeconomic feedbacks within the DOE Earth system modeling framework, this work enhanced national capacity to evaluate large-scale CDR–environment feedbacks across interconnected land, water, atmosphere, and human systems. The resulting tools provide data-driven insights to guide ERW implementation.

Results and Accomplishments

The project developed an ERW module for ELM, with two-way coupling between the reactive-transport processes relevant to silicate rock weathering in the ERW module and the ecosystem processes in the rest of ELM at subdaily radiative time steps. The coupled ELM–ERW model was validated at point-scale with publicly available data on three field sites located in the Northeast, Midwest, and West US. The validated ELM–ERW was applied to simulate CDR rates under ensemble scenarios of grain size, application rate, frequency, and start years over the CONUS at 0.5° resolution under future climate forcing. A surrogate model was developed for the ELM–ERW model based on the ensemble simulations, enabling extrapolation of the gridded CONUS outcomes to unsimulated scenarios.

The ERW module considers all key processes in ERW: (1) silicate rock powder dissolution as affected by mineral type, grain size, pore water ion concentrations, soil moisture, soil temperature, and growth of passivation layer; (2) soil pH dynamics, as a result of newly released cations from dissolution, carbon dioxide equilibrium between soil air and pore water, and cation exchange equilibrium between the pore

water and adsorbed phase on clay minerals; (3) the influence of pH on soil total cation exchange capacity; (4) secondary mineral precipitation; (5) vertical leaching and horizontal runoff of solutes; (6) release of phosphorus from the dissolved rock powder; and (7) impact of soil pH dynamics on soil nitrous oxide emission. The consideration of passivation layer and two-way coupling represents an advancement over the prevalent lack of passivation process and offline coupling between reactive-transport and ecosystem processes. In addition to the physical processes, several technical developments were implemented in ELM–ERW to fulfill the data and computation speed demands of the CONUS ensemble simulations. New CONUS maps of soil cation exchange properties and secondary mineral contents were interpolated from thousands of soil pedons for model inputs. A self-calibration method was developed to estimate background weathering fluxes into soil in the absence of CONUS-level maps of soil mineral compositions. A grid-search algorithm was developed to quickly solve the carbon dioxide and cation exchange equilibrium. A semianalytical algorithm was developed to enable accurate simulation of vertical solute transport despite the long, inflexible time step (1 hour or longer) and uneven vertical discretization in ELM.

The ELM–ERW simulations showed that the passivation mechanism slowed the dissolution rate, but contrary to what was expected in the proposal, did not affect the total CDR at the decadal timescale. Instead, two previously underrecognized mechanisms were found to be important. The first factor was soil liquid water content. At the snow-affected Northeast site, the simulated soil cation concentrations reproduced the observed persistence only if dissolution rates were scaled down by soil liquid water content. Additionally, the CONUS simulations showed that CDR values differed considerably between cropland and needleleaf forest land use types, even when the two were subjected to identical climate and soil forcings, suggesting that differences in evapotranspiration affected CDR through soil moisture levels. The second factor was soil cation exchange. In general, a large fraction of dissolved cations from the applied silicate rock are adsorbed into the cation exchange phase and are prevented from contributing to CDR through the riverine and oceanic processes. Feedback effects of ERW on ecosystem processes were found to contribute to CDR levels. Aside from the promotion of vegetation growth by phosphorus release, soil pH changes due to the silicate rock application generally increased soil nitrous oxide emissions, which reduced the net CDR. Those findings demonstrated the significant value of coupled ecosystem–ERW simulations.

Related to supply chain modeling, the research team produced a comprehensive national-scale analysis of basalt-based ERW for the CONUS. A biobjective optimization model was developed to jointly minimize system-wide costs and maximize net CDR while incorporating the full logistics of mining, grinding, and transport operations. The modeling framework integrates detailed geological, land use, and transportation datasets, enabling simulation of rock movements from more than 200 potential basalt sources to thousands of application sites via both road and rail networks. Spatial optimization results indicated that grasslands and croplands in the southern and southwestern United States offer the most favorable combination of accessibility, weathering reactivity, and land availability for large-scale basalt application.

The supply-chain optimization model was extended to incorporate coupling with a surrogate representation of ELM–ERW. This integration linked the decision-making optimization framework with high-resolution, process-based estimates of weathering rates, soil carbon dynamics, and carbon dioxide uptake efficiency derived from Earth system simulations. The surrogate ELM–ERW model provided spatially explicit environmental parameters, including soil pH, temperature, and moisture, as inputs to the optimization algorithm. This coupling enabled the selection of deployment strategies that reflect mechanistic biogeochemical processes and site-specific environmental conditions.

11502: Soilcosm Phenotyping to Counteract Priming for Sequestration

L. York, M. Cregger, B. Sulman, F. Santos, E. Webb

Project Description

Soil health is an important component of sustainable agriculture and ecosystem resilience, but research is hampered by the lack of high-throughput automation for soils. The intricate balance of soil microorganisms plays a vital role in maintaining soil fertility, structure, and nutrient cycling. Although many researchers propose enhancing soil health through practices such as crop rotation, cover cropping, and organic amendments, fewer researchers address the complex interactions within the soil microbiome that significantly influence these outcomes. Understanding the variety, functions, and dynamics of soil microbes is essential for developing effective soil health strategies. The rhizosphere, which is the narrow region of soil influenced by root secretions and associated microbial activity, is a hot spot for microbial interactions that can enhance or impede plant growth and soil health. This project aimed to develop advanced soil phenotyping technologies to study these interactions in real time using automated systems to monitor microbial activity, nutrient flows, and soil properties. This understanding is crucial for understanding bioenergy feedstock systems as well as managed systems such as grazing lands and forests. The SoilCosm system leverages real-time isotopic carbon dioxide analysis and soil moisture sensors to create an AI-ready, dense time-series dataset of soil dynamics as well as forecasting and simulation modeling. Methods were created to analyze raw data directly from the instruments with batch processing, which allowed the system to be further tuned for real-time analysis and insights.

Mission Relevance

Soils are important for food security and for national defense. For example, bioenergy feedstock agricultural systems can leverage marginal lands not used for food production to increase American capacity to maintain strategic reserves of fuels. Soil health can increase water holding capacity to buffer droughts and increase the utility of associated microbiomes in the soil to benefit plants. AI-enabled autonomous labs help maintain US competitiveness in science and technology and accelerate knowledge gains across domains.

Results and Accomplishments

The SoilCosms project focused on codesigning hardware and software to optimize a novel experimental platform for soil science, as well as applying this platform to understand how soil microbes are influenced by nutrients and water. Multiple instruments were integrated, including a LI-COR Environmental multiplexer, four extension manifolds, a novel carbon dioxide isotopic analyzer, a methane analyzer, and a nitrous oxide analyzer. These were set up in a flask configuration with continuous pumping to purge gases. The new 7825 analyzer uses principles of optical feedback cavity enhanced absorption spectroscopy to resolve isotopes of carbon dioxide in real time. After substantial research and trial, industrial weatherproof polycarbonate electrical enclosures were identified as machinable and airtight vessels for the SoilCosm chambers and were modified to accept the two air tubes and gland for the soil sensor cable. Meter volumetric and tensiometer based sensors were used with online data loggers so that data streams could be combined via programmatic control. Additionally, microdialysis ports were tested to allow either sampling or synthetic application of substrates, similar to root exudation.

Although the hardware integration allowed continuous monitoring, the vendor software was too limited for real-time analysis and decision-making. Therefore, this project developed syncing methods to pull raw data files from the LI-COR multiplexer; each file represented a single measurement time of a single SoilCosm chamber. Python code was developed to batch analyze these files by extracting the time-series accumulation data and fitting models to determine flux. The flux data were combined and stored along with the experimental metadata using "barcodes" and "description" files, which were developed originally for the Advanced Plant Phenotyping Laboratory, that associate unique IDs for each sample with

the treatments applied as needed for downstream statistics and analysis. Additionally, the METER Group ZENTRA application programming interface was leveraged to programmatically download the soil moisture data to store in a similar format. In this way, AI-ready data were continuously generated.

Several pilot studies were conducted to test the system's functioning and its dependence on design choices, such as the amount of soil used and the volume of the chambers. Soil mass and jar volumes (in early work, the project used glass jars) were factorially combined, which demonstrated an optimum mass-to-volume ratio that was more important than the soil mass used. Decreasing the headspace made the system more sensitive. This project also tested the dependency of the system on soil nutrients and water. The system was sensitive to these amounts but typically produced quick flushes of flux that differed in peak and duration. This work typically used dried, sieved soils but discovered that this method possibly should be studied further to compare with soils handled more naturally without drying.

The project also tested the ability of the system to use isotopic determination of carbon dioxide by using labeled glucose in solution at various ratios. This study found that the system was able to distinguish the ratios correctly and could therefore be used to measure how microbes use the added carbon source versus native carbon.

Connecting to simulation modeling allowed the model-experiment (ModEx) approach to be used to fully understand biological systems. Simulations can embody known assumptions and processes to test experimental results and help guide the necessary experiments by filling simulation knowledge gaps. This project used the simple pool-based CENTURY model for the initial work and found that this study could parametrize the model to fit the data reasonably well regarding the amount of substrate added. However, the model struggled with water and nutrient loads, which revealed promising avenues to increase model accuracy. The CENTURY model is used in global weather models, so increasing its accuracy has direct implications for global predictions.

By codeveloping hardware and software infrastructure of automated soil experiments, this project laid the groundwork to make general-purpose frameworks for autonomous labs in biological and environmental sciences. This work showed how previous batch processing development can be used to develop real-time monitoring capabilities.

11508: Point Source Carbon Dioxide Capture Using Biphasic Solvents

C. Tsouris, A. Wiechert, G. G. Jang, M. Cordon, R. Custelcean, D. Stamberga, A. Jackson, G. S. Jung, A. Meyer

Project Description

Established solvent methods for point source carbon dioxide capture are severely limited by their high regeneration energies. For example, in the state-of-the-art monoethanolamine (MEA) capture process, costs associated with solvent regeneration account for more than 50% of the process's operating costs. Biphasic solvents are considered one of the most promising alternatives to MEA for point source capture because of their excellent carbon dioxide recovery rates and relatively low regeneration energies. These benefits are largely tied to the phase separation process that splits the solvent into two distinct phases when it is loaded with carbon dioxide: carbon dioxide–rich and carbon dioxide–lean. Because the lean phase does not require regeneration, only the rich phase is thermally treated, which greatly reduces the total solvent volume that needs to be regenerated and meets or exceeds the recovery rates of MEA. Moreover, using a biphasic solvent in point source capture allows for process intensification. Thus, the work performed in this project focused on first identifying and then analyzing promising biphasic solvent formulae. This project investigated solvent thermal regeneration performance, carbon dioxide flux characteristics, phase separation mechanisms, and process capture costs. Additionally, continuous-flow phase separation via centrifugation was examined.

Mission Relevance

This project supported DOE's mission of advancing transformative science and technology solutions for environmental challenges.

Results and Accomplishments

Preliminary experiments with pure carbon dioxide were performed to assess the phase change behavior, absorption capacity, and regeneration performance of various biphasic solvent solutions. These tests examined a number of different amine blends and organic phase-change agents in aqueous and nonaqueous biphasic formulations. Solvent blends that exhibited undesirable characteristics (e.g., solid precipitation after phase separation, inconsistent or incomplete rich–lean phase recombination following regeneration, low carbon dioxide capacities) were removed from consideration. From this test, two aqueous formulations were identified as promising candidates because of their superior performance characteristics. The first solution used a blend of triethylenetetramine (TETA) and diethylaminoethanol (DEAE), and the second used a blend of diethylenetriamine (DETA) and dimethylaminoethanol (DMAE). Both formulations used triethylene glycol dimethyl ether (TEG-DME) as their phase change agent. An initial cyclic test performed with DETA:DMAE over five loading/regeneration cycles in pure carbon dioxide suggested that the regeneration energy demand of the DETA blend was 67% lower than the MEA benchmark per mole of carbon dioxide recovered. Similarly, the TETA blend's estimated regeneration energy demand per mole of carbon dioxide recovered was 50% lower than that of MEA in preliminary testing with pure carbon dioxide.

The selected blends (2 M DETA, 3 M DMAE, and 20 vol % TEG-DME; 1.5 M TETA, 2.5 M DEAE, and 15 vol % TEG-DME) were examined using a natural gas flue gas (NGFG) simulant stream (i.e., 5 vol % carbon dioxide and 95 vol % nitrogen gas). In the NGFG simulant, rich-phase-only regeneration of the biphasic solvent was only slightly less energy-intensive than the MEA benchmark because the biphasic solvents needed to be regenerated at a higher temperature (i.e., 125°C vs. 110°C) to ensure consistent phase recombination. To reduce the regeneration temperature, a portion of the solvent's lean phase was used as a homogenous regenerant. The lean phase was expected to assist in regeneration because it was assumed to be relatively rich in tertiary amines (i.e., DMAE or DEAE), which is beneficial because tertiary amines facilitate the otherwise very energy-intensive deprotonation step of the carbon dioxide desorption process. Lean phase–assisted regeneration was first tested using the DETA solvent blend with a rich–lean volume ratio of 9:1. At this volume ratio, the DETA solvent was successfully regenerated at 110°C across a total of four loading/regeneration cycles. Overall, DETA's energy demand per mole of carbon dioxide recovered was 23.6% ± 17.8% lower than MEA using this method.

The rich–lean volume ratio was then increased to 9:2 for another four-cycle experiment. This increased rich–lean volume ratio yielded greater energy savings (i.e., energy demand per mole was 31.7% ± 5.0% lower than MEA). Because the energy efficacy of lean phase–assisted regeneration increased when the rich–lean volume ratio increased, additional experiments were performed to determine the optimal volume ratio for the DETA solvent by using rich–lean volume ratios of 3:1, 9:4, 9:5, and 3:2. The regeneration energy demand was greatly reduced when the rich–lean volume ratio was increased from 9:2 to 3:1, but subsequent volume ratio increases did not result in statistically significant changes to the regeneration energy. For these high rich–lean volume ratio experiments, lean phase–assisted regeneration of the DETA solvent averaged an energy demand per mole of carbon dioxide recovered that was 45.5% ± 3.8% lower than that of MEA. Based on these results, cyclic testing with the TETA blend (eight total cycles with a rich–lean volume ratio of 3:1) showed that TETA's average per mole of energy demand was 35.3% ± 5.2% lower than that of the MEA benchmark. Because of the superior regeneration performance of the DETA solvent, the TETA formula was not investigated further. Following these tests, the DETA blend's performance was also examined using coal flue gas (CFG) simulant loaded solvent samples (i.e., 13 vol % carbon dioxide and 87 vol % nitrogen gas). Rich-phase-only regeneration of the CFG-loaded DETA was found to be energy-efficient (i.e., energy demand per mole of carbon dioxide

recovered was 35.7% ± 1.1% lower than that of MEA). Further energy savings could also be obtained for the CFG DETA using lean phase–assisted regeneration (i.e., energy demand per mole was 43.3% ± 1.9% lower than that of MEA). Additionally, these energy savings from lean regeneration could be obtained at a much lower rich–lean volume ratio (10:1) in the CFG-loaded DETA.

Although lean regeneration is efficient, it increases the volume of solvent requiring regeneration. As such, heterogeneous regeneration of CFG- and NGFG-loaded DETA using various solid materials (i.e., Amberlyst-15, titania, quartz powder, and silica gel) was examined. Amberlyst-15 and titania were selected based on their Brønsted and Lewis acidity, respectively. Quartz powder and silica gel were used to assess the practical impact of surface area on regeneration because of the low surface acidity of pure silica and their significant difference in surface area (i.e., 1.3 m^2/g for quartz and 290 m^2/g for silica gel). Standard heterogeneous regeneration used a solvent–solid mass ratio of 200:1.

For NGFG-loaded DETA, the regeneration energy demands per mole of carbon dioxide recovered were 31.9% ± 1.9%, 36.3% ± 2.3%, 37.8% ± 2.2%, and 43.6% ± 1.4% lower than that of the MEA benchmark for Amberlyst-15, quartz, titania, and silica gel, respectively. Conversely, the regeneration energy demand per mole for CFG-loaded DETA was 47.7% ± 1.7%, 48.1% ± 2.3%, and 46.8% ± 1.5% lower than MEA when using Amberlyst-15, titania, and silica gel, respectively. In all these experiments, heterogeneous regeneration greatly increased the rate of DETA regeneration (e.g., maximum carbon dioxide desorption rate increased from 8.4 to 10 mol/[L·h] in CFG-loaded DETA when silica gel was added). Nevertheless, heterogeneous regeneration did not significantly affect the regeneration efficiency. Previous studies argued that surface acidity is central to the efficacy of heterogeneous solvent regeneration. In particular, earlier work suggested that Brønsted and Lewis acid sites on the regenerant's surface catalyze deprotonation and carbamate decomposition. However, based on the performance of silica gel, the results obtained in this work indicated that the energy efficiency of solvent regeneration is more closely tied to the regenerant's surface area. Based on these results, additional experiments with titanium foil were performed. The foil was anodized using an established method to produce a superficial layer of high–surface area titania nanotubes. A control experiment with foil that had not been anodized yielded a regeneration energy demand per mole that was 18.7% lower than that of the MEA benchmark. However, solvent regeneration with the anodized foil consumed significantly less energy than the control case (i.e., energy demand per mole was 50.3% lower than that of the MEA benchmark).

Flux experiments were also performed using DETA solvent to determine the carbon dioxide uptake rate under various circumstances. For NGFG, DETA solvent samples were loaded at 25°C, 40°C, or 60°C for periods of up to 30 h. In this case, the average fluxes of the DETA solvent before phase separation in NGFG simulant were 3.07 ± 0.24 mmol/(m^2·s) at 25°C, 4.16 ± 0.10 mmol/(m^2·s) at 40°C, and 4.01 ± 0.25 mmol/(m^2·s) at 60°C. The flux remains constant from 40°C to 60°C because of competing kinetic and thermodynamic effects. On one hand, higher temperatures increased the rate of mass transfer because of the reduced viscosity of the solvent. On the other hand, the carbon dioxide absorption reaction was exothermic, which negatively impacted the driving force at higher temperatures. In all experiments, the DETA solvent remained in a single phase until reaching a loading threshold of 0.85 mol of carbon dioxide per liter of solvent. After phase separation, the fluxes in the NGFG simulant dropped to 2.66 ± 0.22 mmol/(m^2·s) at 25°C, 2.73 ± 0.23 mmol/(m^2·s) at 40°C, and 2.89 ± 0.26 mmol/(m^2·s) at 60°C. This drop was expected because of the mass transfer limitations imposed by phase separation; absorption proceeded mostly in the rich phase, and carbon dioxide had to diffuse through a layer of the lean phase to reach the rich phase. DETA flux experiments in CFG simulant at 40°C were used to evaluate the effect of suspended solids (i.e., with and without silica gel) on carbon dioxide uptake. These experiments showed that suspended solids had no effect on the carbon dioxide flux.

Molecular dynamics simulations were used to investigate the phase separation behavior of DETA at various loadings (0–3 mol of carbon dioxide per mole of DETA). Initially, simulated DETA formed a single, well-mixed phase. Two phases were formed when carbon dioxide was added to the simulated solution. One phase contained water and absorption products (i.e., the rich phase), and the other contained

TEG-DME and unreacted DMAE (i.e., the lean phase). Additionally, each phase became more distinct as carbon dioxide loading increased. These results aligned with expectations and conformed to this project's experimental findings. Based on experimental findings, a techno-economic analysis of the DETA solvent indicated that a DETA-based capture process in NGFG would cost $55 per metric ton of carbon dioxide captured ($85 per metric ton for MEA). A life cycle assessment revealed that the DETA process had a net carbon dioxide capture rate (i.e., mass captured less process emissions) of up to 0.83 in NGFG (0.62 for MEA in CFG). Centrifugation required multiple steps to achieve continuous-flow phase separation. Instead, gravity-driven phase separation in a settling tank was found to be more energy and cost effective.

11519: Carbon Dioxide Capture from Seawater by a Novel Contactor

A. Rahman, S. Islam, A. Ivanov, T. Saito

Project Description

Seawater is a critical natural sink for carbon dioxide. Current state-of-the-art seawater–carbon dioxide separation technologies are mainly based on an electrochemical approach involving an energy-intensive process, high capital, and operating costs, which significantly hinder their large-scale deployment. The overarching goal of this project was to design novel, scalable, energy-efficient, and cost-effective polymeric sorbent–grafted membrane contactors for highly efficient carbon dioxide separation from seawater or aqueous sources. This project designed responsive polymer sorbents that selectively bind bicarbonate ions and release under pH swing. Additionally, the project integrated these ligand-grafted polymer sorbents onto the surface of polyvinylidene fluoride (PVDF) hollow fiber membranes to enable scalable process development. Achieving this goal involved understanding selective capture and controlled release of bicarbonate ions, which led to creating novel, energy-efficient, and cost-effective carbon capture technologies.

Mission Relevance

This project supported DOE's mission by advancing transformative science and technology solutions for environmental challenges.

Results and Accomplishments

This project designed and demonstrated polymer sorbent–grafted hollow fiber contactors as a promising approach for bicarbonate ion removal from aqueous sources. The initial step of this project performed a qualitative computational screening to guide ligand selection. Density functional theory calculations were employed first to identify ligands with high affinity for bicarbonate ions under aqueous conditions. With the information from computational results, a guanidinium-functionalized polyguanidinium (PGA) polymer was synthesized for bicarbonate ion capture and was grafted to commercially available PVDF hollow fibers. The ligand-derived polymer was synthesized through a reversible addition–fragmentation chain transfer polymerization reaction with *N*-(3-aminopropyl)methacrylamide monomer and 4-cyano-4-((phenylcarbonothioyl)thio) pentanoic acid as the chain transfer agent. The resulting polyamine was converted to PGA through postpolymerization modification with 2-ethyl-2-thiopseudourea hydrobromide to yield a PGA bromide salt. Conversion of the amine functionality to guanidinium was monitored by nuclear magnetic resonance spectroscopy. To evaluate the intrinsic bicarbonate ion binding capability of the guanidinium-functionalized PGA polymer before integrating it into the PVDF hollow fiber contactor, this project first examined bicarbonate ion uptake by neat PGA bromide polymer in aqueous sodium bicarbonate solution. Using a total inorganic carbon analyzer, this project quantified a removal efficiency of 63%, which translates to an adsorption capacity of 0.36–0.38 mg of bicarbonate ions per 1 mg of PGA polymer.

The next step in contactor development involved covalent attachment of PGA bromide to solid supports. PVDF hollow fiber was selected as the physical support because of its prolific use in water treatment, largely because of its cost-effectiveness and chemical or material robustness. A bench-scale prototype

contactor module was constructed from the functionalized PGA polymer sorbent–grafted PVDF fiber to demonstrate carbon capture from water. Grafting PGA polymer onto PVDF (PGA–PVDF) was confirmed by Fourier transform infrared spectroscopy, x-ray photoelectron spectroscopy, and thermogravimetric analysis. After characterization of the PGA–PVDF, this project sought to identify the overall binding capacity of the resultant fiber. The system operated through a two-step process: selective ion exchange–enabled adsorption of bicarbonate ions, followed by desorption triggered by pH swing. In the first phase, adsorption occurs where bicarbonate ion–rich water was loaded into the PGA-functionalized PVDF hollow fibers containing a sealed module. During this phase, ion exchange occurs, which swaps the counterions (e.g., halogens) bound to the guanidinium for bicarbonate anions dissolved in the aqueous feedstock. To investigate the binding capacity of this PGA–PVDF fiber contactor module, different compositions of aqueous solutions were assessed with and without the presence of sodium chloride. To assess bicarbonate binding in the absence of competing ions, the fibers were first exposed to 100 mL of a 200 ppm sodium bicarbonate solution prepared in Milli-Q deionized water. Although conductivity measurements provided preliminary evidence of ion exchange, quantitative verification of bicarbonate ion removal was obtained through total inorganic carbon analysis of the same solutions. The bench-scale prototype achieved an initial bicarbonate ion removal of 34%, which increased to 98% after four adsorption–desorption cycles.

To evaluate the regeneration capability of the PGA–PVDF sorbent via a pH swing process, bicarbonate-saturated fibers were introduced to a slightly acidic solution of hydrochloric acid (pH 4.5) in a flow process. In this process, the guanidinium ligands are regenerated through a mild pH swing, effectively releasing bound bicarbonate anions as gaseous carbon dioxide and restoring the guanidinium salt. The pH was intermittently adjusted back to 4.5 when the solution approached a pH of 7 to maintain desorption conditions. The regenerated anion in guanidinium salt is determined by the conjugate base of the inorganic acid used in the desorption step. PGA sorbent stability was evaluated in multiple cycles using fresh 200 ppm sodium bicarbonate solutions. Across these cycles, bicarbonate ion removal efficiency was maintained around approximately 39%. Notably, the fibers consistently demonstrated reversible bicarbonate ion capture and release over eight cycles, which confirmed the feasibility of the PGA–PVDF contactor design for repeated operation.

After establishing bicarbonate ion capture performance in pure sodium bicarbonate solutions, adsorption cycles 9–13 evaluated the selectivity of PGA–PVDF fibers in increasingly saline conditions. Sodium chloride was introduced at mass ratios of 1:1 up to 10:1 relative to sodium bicarbonate. At a 1:1 sodium chloride:sodium bicarbonate ratio, the inorganic carbon decreased by 22.0 ± 2 ppm (16% removal), which is lower than the 34.0 ± 2 ppm (25% removal) observed in the salt-free cycle 8. As sodium chloride concentration increased, bicarbonate ion removal efficiency declined because of competitive occupation of guanidinium binding sites by chloride ions. However, even at the highest salinity tested (10:1 sodium chloride:sodium bicarbonate), the PGA–PVDF fiber contactor achieved 11% bicarbonate ion removal, which indicated that passive inorganic carbon capture in brackish water remains feasible. To place these results in an environmental context, this project compared the contactor performance with salinity levels representative of natural water bodies. Extrapolation of multicycle operation suggested that more than 25% of dissolved inorganic carbon could be removed after four adsorption–desorption cycles even under salinity similar to seawater. These findings highlighted the robustness of guanidinium-based ion exchange under competitive ionic conditions and underscored the potential of PGA–PVDF contactors for carbon removal across a wide range of aqueous environments, from freshwater to brackish and marine systems.

This study demonstrated the potential of guanidinium-functionalized PGA–PVDF hollow fiber contactors as a viable and regenerable platform for selective capture of bicarbonate ions from aqueous environments. This prototype system achieved an initial bicarbonate ion removal efficiency of 34%, which increased to 98% after four adsorption–desorption cycles. In the presence of seawater, the system successfully removed 10% of bicarbonate ions, which highlighted the ion selectivity of the guanidinium-functionalized polymer even in the presence of sodium chloride—the dominant ion in seawater. The functionalized

fibers showed excellent aqueous stability by maintaining performance over at least 13 cycles in model sodium bicarbonate solutions. Furthermore, the system demonstrated robust regeneration via a mild pH swing, which allowed for complete sorbent recovery with minimal acid input. This process converted bound bicarbonate ions to carbon dioxide and restored the active guanidinium sites, which reinforced the durability of the functionalized hollow fiber sorbent contactors. Notably, the system maintained its selective performance even in the presence of competing chloride ions. Carbon removal remained above 10% at a 10:1 sodium chloride:sodium bicarbonate ratio, which underscored the relevance of this technology for real-world applications in natural water bodies. Overall, this work established the feasibility of using surface-modified contactors for carbon capture from aqueous sources.

VALIDATED DESIGN AND EVALUATION OF FUSION WALL COMPONENTS

11163: Helium-Cooled, Additively Manufactured, Refractory Metal Plasma Facing Component Response to Thermal Transients

T. Gray, T. Seibert, C. Wiggins, M. Kirka

Project Description

Plasma-facing components (PFCs) are required in fusion reactors to protect radio frequency antennas (limiters), first wall components, and divertors, where much of the plasma heat and particle exhaust are focused. However, some initial design studies of compact fusion reactors that use high magnetic fields and high plasma densities predicted steady-state heat fluxes to divertor PFCs greater than 100 MW/m^2 assuming a significant amount of injected power is radiated away via impurity seeding (f_{rad} approximately equal to 0.9).[1,2] Plasma transients such as vertical displacement events and edge-localized modes further increase PFC heating. Vertical displacement events are disruptions that terminate the plasma and result in most of the plasma's stored energy being deposited onto the PFCs, which can potentially melt and destroy them in the process. Edge-localized modes are periodic relaxation events that result in cyclic heating of the first wall and divertor PFCs. Designing, testing, and validating the design of actively cooled PFCs that can survive the steady-state and transient heat fluxes are necessary. This project combined expertise in additive manufacturing (AM) of refractory-metal PFCs, high–heat flux (HHF) testing, and active cooling to prototype PFC designs against transient heat fluxes and validate thermomechanical models of the design.

Mission Relevance

Fusion power has the potential to unlock significant energy resources for the United States. This project aimed to address an important obstacle to developing a fusion power plant. The goal of the work was to use AM to simplify designing and manufacturing actively cooled PFCs.

Results and Accomplishments

The goal of this work was to use AM to simplify designing and manufacturing actively cooled PFCs. The helium-cooled, multiple-jet concept was chosen as a well-established, complex PFC design for this trial.[3] The first objective was to create a finite element model of an AM tungsten PFC. This goal included modeling water and helium as coolants. AM tungsten has been shown to have an approximately 30% reduction in ultimate tensile strength.[4] The results of finite element analysis using water as a coolant were recently published;[5] that study found that the simplified helium-cooled, multiple-jet concept could reach up to 10 MW/m^2 at the steady state, which is similar to the original, more complex design with water cooling.[6]

Additional analysis was performed with helium as the coolant in anticipation of using the PREFACE HHF facility at Tokamak Energy. PREFACE has a limited helium mass flow rate of ≤9 g/s but operates at

[1] A. Q. Kuang, et al., "Divertor Heat Flux Challenge and Mitigation in SPARC," *Journal of Plasma Physics* 86 (2020): 86586050.
[2] J. E. Menard, et al., "Fusion Pilot Plant Performance and the Role of a Sustained High Power Density Tokamak," 29th IEEE Symposium on Fusion Engineering, 2021.
[3] P. Norajitra, et al., "Status of Development of the EU He-Cooled Divertor for DEMO," 22nd International Atomic Energy Agency Fusion Energy Conference, Geneva, Switzerland, 2008.
[4] C. Ledford, et al., "Microstructure and High Temperature Properties of Tungsten Processed via Electron Beam Melting Additive Manufacturing," *International Journal of Refractory Metals and Hard Materials* 113 (2023): 106148, https://dx.doi.org/10.1016/j.ijrmhm.2023.106148.
[5] A. Bhat, "Design and Analysis of an Integrated Additively Manufactured Test Article for Plasma-Facing Components," *Fusion Science and Technology* 81, no. 7 (2025): 661–670, https://doi.org/10.1080/15361055.2025.2476850.
[6] M. Yoda and S. I. Abdel-Khalik, "Overview of Thermal Hydraulics of Helium-Cooled Solid Divertors."

a pressure of 9 MPa. A parametric scan of incident heat fluxes ranging from 2 to 6 MW/m^2 was carried out at helium mass flow rates of 2, 3, 4 and 9 g/s. The surface temperature of the test article, T_{surf}, was predicted to reach the steady state in 1–2 s. However, the test article's performance was limited during the parametric scan because the test article's stainless steel inlet and outlets exceeded the melt temperature of the steel at the helium flow rates at which PREFACE is designed to operate.

To create better digital models for analysis, a series of AM tungsten tensile test samples were fabricated. These were similar to those previously produced but contained 3% by weight of rhenium to improve the ductility of the AM tungsten.[7] The tungsten–rhenium alloyed samples were tensile-tested at temperatures up to 900°C to measure the stress–strain relationship of the alloy. Samples of multiple thicknesses, including 3, 6, and 10 mm, were printed, and so-called "dog bones" were removed from the samples using a wire electrical discharge machining process. However, most of the 10 mm samples did not survive the electrical discharge machining process. This failure was believed to be caused by the 10 mm samples having significantly different crystalline structure compared with the thinner samples, which led to increased brittleness. Tensile testing of the 3 and 6 mm samples showed that adding 3% rhenium to the tungsten increased the ductility of AM tungsten by as much as 50% compared with pure tungsten prints.

Once tensile testing data were available, the third objective of this project was to build an actively cooled, AM tungsten PFC test article that could be used for benchmarking and validation of the models under reactor-relevant HHF conditions. The integrated test article was designed and digitally iterated upon during the finite element analysis and parametric design studies. The AM tungsten–rhenium alloy was printed on top of a solid tungsten disk as a solid hexagonal piece. This disk would also form the pressure boundary during HHF testing with either water or helium as the coolant. A strong metallurgical bond was verified at the additive–solid interface. Subtractive manufacturing was then used (sinker electrical discharge machining) to remove the inner portion of tungsten–rhenium material and install the coolant inlet and outlet, which was machined out of stainless steel.

11412: Machine Learning Informed Multiscale Simulation of Plasma-Facing Materials

R. Juneja, M. Alnaggar, C. Parish, S. Reeve, P. Seleson, J. Weinmeister, Y. L. Pape, J. Rapp

Project Description

Plasma-facing materials (PFMs) are the primary armor in fusion devices; they protect components from intense heat fluxes, energetic particles, and high neutron doses. Under these extreme conditions, however, they undergo severe thermomechanical stresses, irradiation-induced defects, and cumulative degradation that limit component lifetime and reactor availability. To address the need for robust, physics-based tools to assess damage accumulation and fracture in PFMs and to guide materials selection for future fusion power plants, this project developed and demonstrated a machine learning–informed multiscale modeling framework to quantify crack resilience in tungsten PFMs exposed to reactor-relevant high heat fluxes. The framework was rooted in experimentally observed microstructures, starting from a curated database of 2D scanning electron microscopy images of pure tungsten from the JUDITH electron beam facility under varying heat fluxes and base temperature conditions. Generative adversarial networks (GANs), including conditional GAN and StyleGAN2-ADA architectures, were trained on these images to generate realistic synthetic 2D microstructures conditioned on reactor-relevant parameters, thereby augmenting the experimental dataset and enabling improved statistical sampling of microstructural variability. Real and synthetic microstructures were reconstructed into 3D grain-resolved representations using DREAM3D and passed to the Multi-Physics Lattice Discrete Grain Model (MP-LDGM), where mesoscale thermoelastic and elastic perfectly plastic simulations were used to determine effective constitutive responses as functions of grain orientation, temperature, and loading. Mesoscale stress–strain curves were

[7] C. Ledford, et al., "Microstructure and High Temperature Properties of Tungsten Processed via Electron Beam Melting Additive Manufacturing."

generated and used to inform the parameters of macroscale constitutive relations, including elastic modulus, yield stress, and ultimate strain for longitudinal and transverse grain orientations. The macroscale constitutive relations were then implemented in CabanaPD, which is a peridynamics-based fracture mechanics code, to simulate crack initiation and propagation in tungsten monoblocks under spatially and temporally varying heat fluxes. The coupled, microstructure-informed multiscale framework quantified thermomechanical damage thresholds, and it provided a basis for model-informed performance indices and classification of damaged versus undamaged material states, thereby directly addressing the challenge of predicting PFM reliability and lifetime in fusion-relevant extreme environments.

Mission Relevance

DOE's mission is to ensure America's security and prosperity by addressing its energy, environmental, and nuclear security challenges. This project developed quantitative tools to evaluate fracture in tungsten PFMs under reactor-relevant conditions. Predictive modeling of PFMs underpins safe, reliable, and economically viable deployment of fusion as a future energy source, which supports this mission.

Results and Accomplishments

This project established an end-to-end, data-driven, multiscale fracture modeling framework for tungsten PFMs under high heat fluxes. The major technical accomplishments are summarized here.

Machine Learning–Based Synthetic Microstructure Generation

This project assembled and curated a dataset of approximately 3,200 scanning electron microscopy images of pure tungsten subjected to eight different combinations of surface heat flux and base temperature. These images captured a rich variety of microstructures, including variations in grain size, grain shape, and crack morphology, across loading conditions. Using this dataset, this project developed and trained conditional GAN and StyleGAN2-ADA models to generate synthetic tungsten microstructures conditioned on heat flux and temperature. Model quality was evaluated using image quality and physics-aware metrics. The project used the Fréchet Inception Distance, which is a standard image quality metric quantifying the distance between feature distributions of real and synthetic images. Lower Fréchet Inception Distance values indicate better visual similarity. The physics-aware metric employed grain area histograms extracted from real and synthetic images to assess whether generated microstructures preserved grain size distributions. The StyleGAN2-ADA model achieved the lowest Fréchet Inception Distance value of 46.4, which substantially outperformed the conditional GAN model with a best Fréchet Inception Distance value of approximately 280. The StyleGAN2-ADA model also showed strong physics-aware metrics, with grain area histograms of real and synthetic microstructures exhibiting significant overlap. These results demonstrated that the StyleGAN2-ADA architecture can generate high-fidelity, physics-consistent synthetic tungsten microstructures across a range of heat flux and temperature conditions. This dataset was used as input for subsequent mesoscale simulations to enable systematic exploration of how microstructural variability influences damage.

Mesoscale MP-LDGM Simulations and Macroscale Constitutive Parameter Extraction

To incorporate 2D microstructures into 3D mesoscale simulations, this project employed DREAM3D to reconstruct 3D grain geometries from the 2D scanning electron microscopy images. The resulting 3D microstructure meshes resolved grain morphology and connectivity, which provided suitable input for MP-LDGM simulations. Within MP-LDGM, an elastic perfectly plastic constitutive law was formulated at the grain level. Mesoscale simulations were performed over a range of temperatures and loading conditions to characterize the effective thermomechanical response of pure tungsten for longitudinal and transverse grain orientations relative to the surface. From these simulations, stress–strain curves were generated and used to calibrate key macroscale constitutive parameters, including effective elastic modulus, yield stress, and ultimate strain. The temperature dependence of these parameters was quantified by fitting appropriate functional forms to the mesoscale simulation outputs. The resulting parameter sets provided a compact representation of microstructure-dependent mechanical behavior under cyclic thermal

loading for use in macroscale fracture simulations. Crucially, the parameter differences between longitudinal and transverse grain orientations encoded the effects of microstructural anisotropy observed in the mesoscale simulations. The generated mesoscale outputs further demonstrated how physics-based simulations can augment limited experimental data.

Macroscale Peridynamics Damage Simulations

The mesoscale-derived macroscale constitutive parameters were implemented in CabanaPD, which is an ORNL-developed peridynamics fracture mechanics code that naturally represents complex crack phenomena, including initiation, growth, and branching, without requiring predefined crack paths. To this end, the available CabanaPD brittle elastic-only constitutive models were extended to model thermoelasticity and plasticity, and a heat transfer solver coupled with the mechanics solver was also implemented in CabanaPD. Using CabanaPD, this project performed macroscale thermomechanical simulations of tungsten PFMs subjected to reactor-relevant heat fluxes. The simulations employed multiphysics coupling of heat transfer and thermomechanics to capture the thermal stresses resulting from the evolving temperature gradients. The results qualitatively reproduced key experimentally observed behaviors for both grain orientations—namely, no damage at low heat fluxes, the transverse grain material failing before the longitudinal grain material, and both materials ultimately failing at very high heat fluxes. These outcomes highlighted the importance of explicitly accounting for microstructural anisotropy in the macroscale damage response of tungsten PFMs. The consistency between the simulated trends and experimental observations provided confidence that the microstructure-informed multiscale framework can capture essential features of tungsten damage under high–heat flux conditions.

Performance Index and Damage Classification

To complement the physics-based simulations, this project developed a machine learning–driven performance index to classify tungsten microstructures as damaged or undamaged under specified heat flux and base temperature conditions. This work trained a classification model that ingested macroscale simulation results and output a damage label, and the results demonstrated this method's effectiveness for tungsten with longitudinal and transverse grain orientations. The performance index thus provided a rapid, data-driven indicator of damage and undamaged states that can be useful for screening microstructures and identifying regimes of high damage risk and integrating microstructure-based metrics into higher-level design or control frameworks for PFMs.

This project delivered an integrated, machine learning–enabled multiscale modeling workflow for tungsten PFMs. Despite the simplifying modeling assumption of an elastic perfectly plastic constitutive law and the lack of explicit treatment of the brittle-to-ductile transition, the framework successfully captured microstructure-dependent damage evolution and laid the groundwork for a quantitative assessment of crack resilience in tungsten PFMs under extreme conditions. These results enhanced modeling capabilities for plasma-facing components and established the foundation for incorporating additional plasma–material interaction phenomena.

11420: Design of a Lead–Lithium Corrosion Testing Loop and Experimental Planning

S. Smolentsev, M. Romedenne, J. Fan, Y. Jiang, S. Aduloju, E. Burkhardt

Project Description

Lead–lithium is the worldwide front-runner for liquid tritium breeder materials that will be used in fusion reactors. Among the most serious limitations in designing a robust blanket system is corrosion of structural blanket materials such as reduced activation ferritic-martensitic steel in the flowing hot lead–lithium in the presence of a strong plasma-confining magnetic field. The available phenomenological models have not been validated, and the experimental data on the corrosion of candidate steels in lead–lithium are scarce and often contradictive. The precipitation of corrosion products

in the cold (i.e., nonradioactive) section of the lead–lithium loop is another feasibility issue. The United States currently does not have an experimental facility to address these issues. This project developed a simulation workflow and designed a flowing lead–lithium corrosion loop with a magnetic field and surface heating in the test section to reach prototypical blanket conditions. The proposed lead–lithium loop will be able to show the corrosion performance of candidate fusion materials in an experimental fusion prototypical environment and establish their corrosion mitigation. The developed computational tools were used in preexperimental simulations to carefully plan future experiments on corrosion and precipitation to deliver data to validate phenomenological models and computer codes and help design and analyze real blanket systems.

Mission Relevance

This project supported DOE's mission to ensure America's security and prosperity by advancing transformative science and technology in the fields of energy and nuclear systems. The development of a flowing lead–lithium corrosion loop with an integrated magnetic field and surface heating addresses one of the challenges in fusion energy research: the corrosion and compatibility of structural materials under reactor-relevant conditions. By enabling experimental validation of corrosion and precipitation models, the project provided foundational understanding necessary for designing reliable, safe, and efficient fusion blanket systems, which is a key component for achieving sustainable fusion power generation.

Results and Accomplishments

A new flowing lead–lithium corrosion loop, called CoUPLE (Corrosion Upgrade Lead–Lithium Experiment), with a magnetic field and surface heating of the test-section was designed at ORNL. The proposed ORNL lead–lithium loop reaches prototypical blanket conditions (lead–lithium temperature up to 650°C, velocity up to 20 cm/s, magnetic field up to 4 T) and is able to show the corrosion performance of candidate structural steels and functional electroinsulating materials as well as establish their corrosion mitigation. The developed computational model in COMSOL Multiphysics was used in this study as a design tool as well as for the preexperimental analysis to plan future experiments on corrosion and precipitation. COMSOL also provided data to help design and analyze real blanket systems. Components of an experimental corrosion program enabled by this work are the following:

- Materials to be tested include bare reduced activation ferritic-martensitic, coated reduced activation ferritic-martensitic, new perspective steels (oxide dispersion strengthened), silicon carbide (composite and foam), alumina, and preoxidized steels to form alumina protective layers.
- The flow, heat transfer, and magnetohydrodynamic effects, as well as the effect of lead–lithium alloy chemistry and microstructure, can be accessed in the experiments for various lead–lithium velocities (up to 20 cm/s), temperatures (up to 650°C), and magnetic fields (up to 4 T).
- Two test blanket modules, horizontal and vertical, can be sequentially used to reproduce flow regimes predicted in the blanket studies—forced, mixed, and natural convection.
- Corrosion of extended surfaces (walls), submerged miniature specimens, and welded parts can be studied in monomaterial and dissimilar material environments.
- Redeposition (precipitation) studies in the cold (i.e., nonradioactive) leg of the loop can be performed along with the corrosion studies in the hot (i.e., radioactive) leg.

11441: Engineering Optimization of First Wall Protection Limiters

K. Borowiec, F. Oz, J. Nichols, J.-W. Bae, A. Sircar, J. Rapp

Project Description

Even for the relatively well-studied tokamak concept, additional challenges remain to develop a fusion pilot plant. The operating conditions of a fusion pilot plant are challenging, and a robust first wall system is essential to withstand high heat and particle fluxes under irradiation conditions. Balance must be maintained between irradiation resistance, temperature limits, and influence on tritium breeding ratio.

One concept being considered is introducing protection limiters on the first wall, which concentrate plasma fluxes on armored components to reduce plasma fluxes to fragile breeding surfaces and still maintain small gaps between the plasma and first wall (this is an especially important consideration for compact fusion pilot plant designs). The limiter concept mitigates the heat flux from ionized particles during steady-state operations and disruptions. In the event of a disruption, the limiter can serve as a sacrificial component to protect the first wall from severe damage because the limiter is easily replaced. However, the presence of a limiter has important implications for steady-state operations that must be considered during the design process. This project emphasized limiter performance during steady-state operations. This work focused on engineering modeling of the limiters considering the multiphysics aspects of the design under realistic boundary and source conditions from plasma physics codes. The developed tools can be used to investigate the heat mitigation capability and engineering feasibility of different protection limiter designs. The optimization workflow created as part of this project allows rapid iteration of the limiter design for any tokamak. The coupled workflow was executed driven by an optimizer to improve selected performance metrics.

Mission Relevance

This project supported DOE and ORNL missions by addressing energy challenges using cutting-edge science and engineering. This project strengthened modeling and simulation for fusion energy research by incorporating tight coupling between plasma physics and engineering calculations. The workflow generated in this project was designed to easily interface with larger whole-device modeling efforts, which can provide insight and independent design review for conceptual designs considered for a future fusion pilot plant.

Results and Accomplishments

A fully integrated multiphysics modeling and simulation framework was developed. This new framework allowed rapid iteration of limiter design. The workflow included implementation of the limiter geometry into parametric geometry generation software TRACER, which enabled investigation of limiters for many tokamak designs and allowed changes to the limiter dimensions, profile, and cooling solutions. The analysis-specific workflows were developed to enable neutronics, thermal, and plasma physics analysis. The neutronics analysis was developed to assess the influence of limiters on the tritium breeding ratio. Because of the complexity of the design, geometry homogenization was necessary. The cooling channels for the limiter and first wall were treated as one with the surrounding structural material. To increase the accuracy of the solution, this project developed a correction based on microscopic cross sections and showed significant reduction in homogenization error. The thermal analysis was initially planned to include conjugate heat transfer to explicitly model the cooling channels using computational fluid dynamics. However, the computational resources required for this approach proved to be prohibitive for design optimization. To address this issue, a more robust and efficient method was developed in which 3D heat conduction was coupled to 1D Navier–Stokes equations. Specifically, a 1D fluid flow solution was implemented to replace the computational fluid dynamics simulations of the cooling channels. This solution was then coupled to the 3D conduction equation in OpenFOAM to reduce the simulation time to a fraction of that required for a full conjugate heat transfer analysis. The approach relied on established correlations for the heat transfer coefficient, which are well-characterized for helium cooling flows. In cases in which these correlations were unavailable, reduced-order models could be developed using higher-fidelity conjugate heat transfer simulations and could be integrated into the workflow. Additionally, a mesh-based geometry specification was developed to support design optimization.

The coupled workflow was integrated into an optimizer to enable automated optimization of the limiter design. The workflow's integration with the optimizer allowed seamless input generation and code execution. Once the specific metrics of interest were identified, the optimizer drove the simulation to minimize prescribed loss function without user input.

The fully coupled workflow connected to the optimizer was demonstrated on one of the ORNL clusters. The figures of interest included the (1) maximum temperature of the first wall, (2) maximum temperature of the limiter, and (3) tritium breeding ratio. This project successfully developed and demonstrated automated multiphysics optimization of the limiter design.

11555: Modeling for Lithium Fluoride and Beryllium Fluoride (FLiBe) Fusion Tritium Breeding Compatability

M. Gehrig, P. Humrickhouse, R. Pillai, M. Romedenne, S. Cambier, A. Willoughby, T. Lowe

Project Description

The molten salt FLiBe ($2LiF–BeF_2$) has gained interest as a tritium breeder for a fusion pilot plant, mainly because of its low electrical conductivity. FLiBe is less influenced by high magnetic fields and has a smaller effect on tokamak plasma control systems. A critical challenge for this blanket concept is the lack of compatible low-activation material. FLiBe melts at 459°C, which is too high for operation with reduced activation ferritic-martensitic steel (550°C maximum). Corrosion is also a concern; nickel alloys have superior corrosion resistance in salt but are not feasible in a fusion reactor because of their neutron activation characteristics, which lead to long-lived radioactive waste and decay heat production exceeding fission reactors. This work pursued a two-fold design strategy to address these shortcomings: (1) an alternate salt mixture (52% lithium fluoride/48% beryllium fluoride) that melts at a lower temperature of 392°C and (2) thin nickel or tungsten coatings on reduced activation ferritic-martensitic steel to mitigate corrosion. To assess this strategy, this project (1) conducted a staged thermal design of a FLiBe/reduced activation ferritic-martensitic steel blanket to identify an operating temperature window using computational fluid dynamics (CFD) and magnetohydrodynamic (MHD) models; (2) used neutronic and activation analyses to determine the effect of coating materials and salt composition on decay heat, waste production, and tritium breeding ratio; and (3) conducted a series of capsule experiments in which coated reduced activation ferritic-martensitic steel specimens were exposed to FLiBe to demonstrate their performance.

Mission Relevance

This project addressed DOE's energy mission by supporting the advancement of tritium breeder blanket science and technology. This project examined a pathway to enable a more rapid timeline for deployment of a viable molten salt tritium breeding blanket technology by exploring a relatively simple change in FLiBe composition and studying the compatibility effects. It leveraged experience in material science, nuclear engineering, and thermal hydraulic analysis to address a technological challenge by using FLiBe as a tritium breeding blanket material. The goal of this project was to find a solution that improved the compatibility of FLiBe with other materials to allow for more near-term and available fusion energy deployment.

Results and Accomplishments

This project identified a FLiBe composition that may be viable to serve as a coolant in a tritium breeding blanket first wall component. Although a range of FLiBe compositions with different mole fractions of beryllium fluoride were known to have different melt temperatures, it was unknown if the changes in other properties such as viscosity might have invalidated the use of FLiBe as a first wall coolant with the project's candidate structural materials. This result was achieved by sweeping the range of beryllium fluoride properties and melt temperatures with a series of constraints for heat transfer, temperature, and pressure drop to find the most suitable composition. This outcome enabled the continued exploration of a 48% mole fraction beryllium fluoride FLiBe for an outboard breeding blanket configuration with poloidally oriented first wall channels, neutronic compatibility analysis, and corrosion mitigation. This analysis also was used to determine the channel dimensions in the first wall, with selected channel dimensions being 0.03×0.05 m rectangular poloidal channels.

To address thermal compatibility, this project employed CFD and thermal modeling of the outboard blanket with the first wall, bulk blanket, and added volumetric heating caused by the nuclear reactions in the blanket. This activity used nuclear heating data retrieved from the nuclear analysis specific aim and incorporated the data into CFD models in STAR-CCM+ to assess the 3D geometry of a representative slice of the outboard blanket. This analysis found that with the conditions of a 0.25 MW/m^2 heat flux and a maximum of approximately 9 MW/m^3 volumetric heating, the 48% mole fraction beryllium fluoride FLiBe was able to maintain the maximum temperature of the first wall at the limitation of approximately 550°C in steady-state conditions with a flow rate of 500 kg/s across eight channels. The overall simulation showed that across the entire outboard blanket, heat from 313 MW of thermal power can be removed with a pressure drop less than 2 MPa, which provided a case for the viability of the 48% mole fraction beryllium fluoride FLiBe as a candidate for a coolant in a liquid breeder blanket.

In addition to the CFD modeling, MHD modeling was performed to assess the effects of magnetic field on the overall thermal and fluid performance of the 48% mole fraction beryllium fluoride FLiBe. For analysis in the first wall, calculations examined the 48% mole fraction beryllium fluoride FLiBe with magnetic fields of 0, 10, and 20 T to determine changes in friction factor and heat transfer. For the cases studied, this project found that even at the maximum of 20 T magnetic field, the maximum temperature was less than 5°C higher than that for 0 T. For the fluid analysis, the friction factor in the 20 T magnetic field case was reduced relative to the 0 T case (0.00629 vs. 0.0068, respectively). This reduction in the friction factor was caused by turbulence suppression in the MHD flow. The MHD calculations found that for a straight first wall channel, the maximum velocity could be reduced from the originally calculated velocity of 16 m/s to 12 m/s to maintain the maximum temperature of the structural material at 550°C.

The project then assessed the nuclear compatibility of the 48% mole fraction beryllium fluoride. This activity was performed using Monte Carlo N-Particle (MCNP) and FISPACT II to calculate the waste disposal rating and decay heat expected from a tritium breeding blanket with a first wall coated with nickel and one coated with tungsten, with thicknesses up to 400 μm. These neutronic and activation analyses were used to determine that neither coatings would exceed waste disposal rating limitations. However, to avoid excessive decay heat, the coating thicknesses should be limited to only that which would be required for corrosion mitigation.

Another goal was to assess tritium breeding of the 48% mole fraction beryllium fluoride FLiBe. The MCNP and FISTPACT II calculations were also used to ensure that the tritium breeding ratio with 48% mole fraction beryllium fluoride FLiBe would not be ruled as insufficient for useful tritium production. With each using the same configuration, the tritium breeding ratios for 48% mole fraction beryllium fluoride FLiBe and the standard 2:1 FLiBe were calculated. The ratio of the 48% mole fraction beryllium fluoride FLiBe was only marginally lower than that of the standard 2:1 FLiBe tritium breeding blanket (1.3 vs. 1.34, respectively). Notably, this calculation was performed to serve as a comparison with standard 2:1 FLiBe and not as a full assessment of the design of the blanket; the geometry lacked some structural components whose inclusion may result in a lower calculated tritium breeding ratio. This calculation also provided information for nuclear heating in the blanket, which was used for the aforementioned thermal analyses.

To address material compatibility and corrosion, results were obtained from the coating process and corrosion experiment. The first of these results identified a coating thickness and the beginning of the experimental campaign for specific aim three, which involved coating samples of F82H, a reduced activation ferritic-martensitic steel, with nickel, tungsten, and a nickel–tungsten alloy. Coating thicknesses were determined to be sufficient at 50 μm to allow for a 20 μm thick interdiffusion zone between the coating and substrate materials. Two different coating mechanisms were used: thermal spray and electroplate. Both coating mechanisms produced intact coatings, although the nickel–tungsten coating contained cracks and multiphase precipitates.

For baseline comparison, static capsule tests were performed with a standard 2:1 FLiBe ratio for uncoated F82H, nickel-coated F82H, and nickel–tungsten-coated F82H at two temperatures (450°C and 550°C) and three times (100, 55, and 1,000 h). For the bare F82H, significant metal loss was observed after 1,000 h. For the coated specimen, both coatings mitigated corrosion compared with the bare sample, and no change was observed in chromium concentration in either coating. The nickel–tungsten-coated F82H showed the presence of chromium-rich precipitates in the coating. The corrosion experiment for the 48% mole fraction beryllium fluoride FLiBe was performed with the same temperatures and times at those for uncoated F82H, nickel-coated F82H, and tungsten-coated F82H. The results were available for the mass change in the 100 h samples, with mass loss observed in the uncoated F82H. Mass gain was observed in some tungsten-coated specimens, which indicated an update from salt impurities or salt entrapment in coatings. One of the tungsten-coated specimens showed significant mass loss, which suggested coating spallation. The nickel-coated specimen showed little mass change. The samples for all the 48% mole fraction beryllium fluoride FLiBe corrosion tests were cleaned and prepared for characterization. Additionally, the third specific aim included demonstrating the recent implementation of new beryllium handling practices.

In general, the results of the thermal, CFD, MHD, and neutronic modeling suggested that the 48% mole fraction beryllium fluoride FLiBe may be theoretically feasible as a liquid breeder material in tritium breeding blankets, but this material warrants further investigation to better assess parameters such as feasible pressure drop in molten salts. The corrosion experiments showed that the coatings mitigated corrosion against the standard 2:1 FLiBe composition, with the nickel-coated F82H showing the most intact coating following the static capsule tests and minimal formation of precipitates in the interdiffusion zone.

VERTEX: ADVANCED MULTIPHYSICS SIMULATIONS FOR CORE APPLICATIONS

11058: Plasma-Eroded Material Migration for Vertex: Advanced Multiphysics Simulations for Core Applications Initiative

T. Younkin, C. Hauck

Project Description

This project focused on advancing capabilities for simulating bulk plasma, plasma sheath, and plasma–material interactions relevant to isotope enrichment and fusion energy systems. Specifically, it addressed gaps in modeling plasma-eroded material migration and the integration of eroded material into the plasma using enhanced self-consistent algorithms and hybrid particle–fluid methods. The plasma separation process (PSP), which is a promising isotope enrichment mechanism and applies to linear plasma devices such as the Material Plasma Exposure eXperiment (MPEX), provided the foundational motivation for this effort. The project sought to build upon ORNL's global impurity tracking code (Global Impurity Transport Code, or GITR) by porting relevant operations and expanding physics capabilities in the Vertex framework. This project was driven by the following three primary objectives:

- **Incorporation of self-consistent effects**—The project aimed to relax the trace approximation used in the GITR model by tracking plasma and eroded species together. A combination of these self-consistent effects and time-scale separation for tracer particles was developed. Therefore, this project used a combination of time-marching and iterative solutions for source terms to account for source terms, perturbed electric fields, and evolving plasma profiles.
- **Application to the PSP challenge problem**—This work leveraged project improvements to simulate key parameters of PSP devices and provide design insights for future facilities. Starting with the basic trace impurity migration model and then moving to higher-fidelity physics gave insight into PSP parameters and trade-offs in model fidelity.
- **Integration with the Vertex framework**—Finally, the project focused on infrastructure development and porting of physics to Vertex to ensure full compatibility of particle tracking with shared data structures. This work enabled scalable simulations for particle–plasma dynamics.

Mission Relevance

This project advanced plasma–material interaction modeling through algorithmic improvements to GITR, applications to the PSP, and integration with Vertex. These developments contribute to deeper scientific insights and transformative enrichment technologies that address America's energy and security needs.

Results and Accomplishments

Advancing the GITR Model in Vertex

Through algorithmic advancements, the GITR model was progressively enhanced to incorporate self-consistent plasma effects into the Vertex framework. The iterations included impurity–impurity collisions, self-consistent plasma effects by solving electron fluid equations, and tracer particles for electrons and ions that can supplement solution procedures. This work resulted in a multispecies (neutrals, multiple ion species, and electron) plasma solver that captures plasma sheath formation at surfaces, tracks ions to the surface in those regions, and addresses the resulting plasma–material interaction. Those interactions include deposition, sputtering, reflection, and integration of those interactions back to the plasma. Additionally, this work modeled ion–ion collisions through a binary-pair collision algorithm that replicates characteristics of Coulomb collisions from the Landau Fokker–Plank equation. This model is critical to assess the temperature and density of ion cyclotron heated particles in enrichment devices. Finally, this effort created structured and unstructured mesh high-performance computational capabilities

with various backend solvers as well as the ability to couple with additional finite element fluid and multiphysics models.

PSP Challenge Problem

Initial simulations of PSP environments using the updated GITR version explored erosion dependencies on baseline parameters such as plasma density, device geometry, and magnetic field strength. Subsequent iterations analyzed scaling laws for impurity deposition and assessed the impact of Coulomb collisions under nonlinear plasma conditions. Results included refined guidelines for balancing feed erosion and downstream transport to maximize isotope throughput, predictions of optimal cyclotron resonance conditions for minimal cross contamination at the collector, and insights into preserving separative work efficiency despite potential re-erosion effects.

Additional fusion-relevant demonstration problems were developed to drive model development and library implementation. These domain-focused demonstrations included the following:

- **Source problem**—The main problem involved a plate where plasma contact shows plasma sheath, sputtering, and ionization of surface material. This problem is relevant to plasma–material problems in tokamak divertors, MPEX, and isotope enrichment.
- **Resonance problem**—In a magnetized section of plasma inflow and outflow, an injected wave excites a resonant species, which interacts with itself and other ion species to show the competition of heating versus collisional energy loss.
- **Collector problem**—This problem concerns complex geometry in which particles deposit and interact with leading edges of the geometry. Angles are varied to demonstrate how heat flux and erosion can be effectively modified.

Integration with the Vertex Ecosystem

Developing hybrid particle–fluid models and interfacing GITR with the Vertex framework were integral to this project's success. Realizing spatially and temporally scalable particle-based simulations was important for advancing predictive modeling capabilities for PSP and MPEX. Achievements included full compatibility with Vertex's shared grid structures and adaptive meshing tools through the Trilinos library, successful deployment of the GITR model as a component in the Vertex framework for large-scale high-performance computing simulations, and enhanced runtime performance for particle-bin physics operations enabled through advanced node architectures.

11228: Vertex: Advanced Multiphysics Simulations for Core Applications Initiative Core

F. Curtis, A. Almeldein, D. Arndt, V. Badalassi, J. Bludau, K. Borowiec, F. Brandao, L. Bullerwell, K. Chong, M. Delchini, W. Ge, R. Glasby, K. Gottiparthi, S. Hamilton, I. Jarrah, K. Kincaid, F. Oz, C. Plechaty, S. Reeve, R. Savery, A. Sircar, S. Slattery, S. Smolentsev, D. Stefanski

Project Description

The Vertex Core project created a new multiphysics simulation framework supporting ORNL's mission-critical challenges in modeling and simulation of gas dynamics, rarefied flow, plasma–surface interaction, electromagnetics, magnetohydrodynamics (MHD), and thermal hydraulics for conducting fluids and collisionless and collisional plasma. The team established a coupled physics framework enabling predictive simulations for (1) next-generation enrichment science for applications including the plasma separation process (PSP), gas centrifuge, and electromagnetic isotope separation; (2) fusion science, including blanket design and evaluation of novel fusion concepts; and (3) advanced scientific computing workflows integrating data-based models. This project created a computational and data science–ready framework that simulates coupled multiscale/multiphysics phenomena in core ORNL mission spaces. Advanced predictive simulation capabilities were developed, meeting ORNL

science-driver needs. Strategic program challenge problems were demonstrated, emphasizing the PSP and fusion blankets.

Mission Relevance

Vertex Core contributed to DOE's mission in energy and nuclear innovation by delivering high-fidelity modeling capabilities for fusion and fission reactor design, including isotope enrichment and plasma–surface interactions, which are central to advancing nuclear energy technologies. This project also advanced DOE's mission of transformative science and technology by integrating cutting-edge computational tools such as Trilinos, TensorFlow Lite, and Hypre and supporting GPU acceleration via Kokkos, which enables scalable, high-performance simulations that push the boundaries of current modeling capabilities.

Results and Accomplishments

The project successfully addressed the needs in modeling and simulation of gas dynamics, rarefied flow, plasma–surface interaction, electromagnetics, MHD, and thermal hydraulics for conducting fluids and collisionless and collisional plasma. The team established a coupled physics framework that enables predictive simulations for these core science areas. The project focused on next-generation enrichment science for applications including the PSP, gas centrifuge, and electromagnetic isotope separation; fusion science, including blanket design and evaluation of novel fusion concepts; and advanced scientific computing workflows integrating data-based models. The team created a computational and data science–ready framework that simulates coupled multiscale/multiphysics phenomena in core ORNL mission spaces. Advanced predictive simulation capabilities were developed, meeting ORNL science-driver needs. Strategic program challenge problems were demonstrated that emphasized the PSP and fusion blankets.

The VERTEX-CFD framework is a multiphysics solver designed for fusion applications. It is built on a finite element foundation and supports CPU and GPU architectures through Kokkos.[1] The solver integrates key libraries such as Trilinos,[2] TensorFlow Lite,[3] and Hypre,[4] which enables advanced modeling capabilities, including external field coupling and automated differentiation. Some noteworthy features are its open-source accessibility and continuous integration, verification and validation against benchmark cases, and modular design for extensibility across fusion-relevant physics.

Two primary turbulence modeling strategies were explored: (1) large eddy simulation, which compares algebraic wall-adapting local eddy-viscosity models and implicit large eddy simulation approaches and focuses on transient, incompressible Navier–Stokes flows and highlighting differences in numerical dissipation and time integration schemes, and (2) Reynolds-Averaged Navier–Stokes, which implements multiple Reynolds-Averaged Navier–Stokes models such as Spalart–Allmaras, k-ε, and Shear Stress Transport. These models are tailored for finite element methods and validated using canonical cases such as the NACA 0012 airfoil.

The solver includes robust MHD capabilities that support full induction and inductionless formulations. Validation was performed using analytical solutions for Hartmann and Shercliff flows, with results demonstrating high accuracy in mass flow rate and pressure drop predictions. Applications include multiregion simulations for fusion blanket designs and integration of MHD effects into broader thermal-fluid simulations.

[1] Trott, C., et al., "The Kokkos EcoSystem: Comprehensive Performance Portability for High Performance Computing," *Computing in Science and Engineering* 23, no. 5 (2021): 10–18. https://doi.org/10.1109/mcse.2021.3098509.
[2] *Trilinos* (Github.io), https://trilinos.github.io.
[3] Abadi, M., et al., "TensorFlow: Large-scale machine learning on heterogeneous systems," arXiv (2016), https://doi.org/10.48550/arXiv.1603.04467.
[4] "Hypre: Scalable Linear Solvers and Multigrid Methods," Lawrence Livermore National Laboratory, https://llnl.gov/casc/hypre.

Adaptive mesh refinement was introduced to enhance simulation accuracy and computational efficiency. The framework leverages NASA's Refine[5] software to dynamically adjust mesh resolution based on scalar field gradients. This approach is particularly effective for fusion blanket problems with localized phenomena. The mesh refinement technology ensures that computational resources are used efficiently by focusing detail where it is most needed and allowing coarser meshes in regions of lesser importance, thereby optimizing overall computational effort.

The framework is optimized for high-performance computing environments, with a focus on GPU acceleration. Challenges related to memory allocation and automated differentiation were addressed, and scaling studies demonstrated strong performance for large-scale 3D turbulent flow simulations. This includes strategies for efficient GPU memory management and performance benchmarks across various hardware configurations. The integration with modern GPU technologies ensures that the framework can handle increasingly complex simulations, maintaining high resolution and accuracy even as problem sizes grow.

A data-driven approach was employed to develop turbulence closures for MHD flows in molten salt blankets. High-fidelity large eddy simulation and direct numerical simulations were used to generate training datasets for machine learning models. These models were then embedded into the solver using TensorFlow Lite for real-time Reynolds-Averaged Navier–Stokes simulations. Innovations include physics-informed machine learning techniques, integration of uncertainty quantification, and enhanced predictive capabilities for fusion-relevant flows. The use of machine learning allows the solver to adapt more flexibly to a wide range of flow conditions, which improves accuracy and efficiency.

A hybrid particle mesh solver was developed to address complex phenomena such as plasma–surface interactions, turbulent liquid metal flows, and isotope tracking. The framework supports curved mesh geometries and was designed for seamless integration with ORNL's data workflows. Its capabilities include multiphysics coupling for fusion and materials science, adaptive meshing and particle tracking, and a modular design for future expansion. This hybrid approach combines the strengths of particle-based and mesh-based methods to provide a versatile tool for simulating a broad spectrum of physical phenomena.

[5] NASA, *Refine* (GitHub), https://github.com/nasa/refine?tab=readme-ov-file.

SUMMARIES OF PROJECTS SUPPORTED BY THE SEED PROGRAM

BIOSCIENCES DIVISION

11776: Protein Engineering to Multiplex Gas Biosensors

I. Del Valle Kessra, E. Teixeira Prates, T. Walker, P. Kruse, M. Shah, D. Jacobson

Project Description

Gas biosensors are genetically engineered microorganisms that respond to an environmental input by producing an indicator gas. They have emerged as a new tool to understand microbial processes belowground, where using fluorescent proteins or other visual reporters can be challenging.[1,2] However, studies using gas biosensors are still limited because only one gas-producing enzyme has been identified as an output reporter that works under aerobic and anaerobic conditions and is suitable for plant studies.[3] This enzyme, called methyl halide transferase (MHT), uses as substrates S-adenosyl methionine and halide ions (Cl^-, Br^-, and I^-) to produce methyl halide gas.[4] Engineering MHT variants with specificity toward one halide ion would increase the number of gas biosensors that could be used simultaneously in soil studies and would open the possibility of detecting multiple environmental inputs or tracking the growth of multiple microbes within hard-to-image samples. This project used a protein rational design framework to reprogram MHT halide preference and generate a Br^--specialist variant. This approach integrated targeted library construction, structural modeling, computational screening, and experimental validation using gas chromatography–mass spectrometry to quantify methyl halide gas production across halide types and concentrations. Through this Design–Build–Test–Learn cycle, this project engineered three highly active and halide-specific MHTs.

Mission Relevance

By establishing a versatile enzyme design platform, this project advances the DOE mission to address national challenges in the nation's energy, security, and critical materials and minerals stewardship. Engineered gas biosensors enable in situ monitoring of microbial activity and physicochemical conditions in soils and subsurface environments. By providing real-time detection of microbially mediated processes, they can be coupled to report on nutrient availability, redox transformations, and pathogen activity. These biosensors help better predict, monitor, and reduce biological risks that could affect bioenergy crops.

This integrated workflow combines protein engineering, structural modeling, and hybrid machine learning/physics-based approaches to build infrastructure for engineering specialized biosensors and ion-responsive systems. This capability supports DOE priorities in the critical materials and minerals space by enabling future adaptation of the workflow to rare earth element biomining, including engineering proteins with enhanced rare earth element specificity and sensitivity, and using MHT-based biosensors to assess microbiome conditions favorable for rare earth element mobilization.

Results and Accomplishments

This effort engineered three highly active and halide-specific MHTs. The project used a Design–Build–Test–Learn method that integrated a high-throughput sequence- and structure-based clustering algorithm to guide the selection of natural variants for experimental testing, which was complemented by comparative structural analysis and molecular dynamics simulations to inform rational design of MHTs with enhanced halide specificity.

---

[1] H.-Y. Cheng, et al., "Ratiometric Gas Reporting: A Nondisruptive Approach To Monitor Gene Expression in Soils," *ACS Synthetic Biology* 7, no. 3 (2018): 903–911, https://doi.org/10.1021/acssynbio.7b00405.
[2] I. Del Valle, et al., "Artificial Soils Reveal Individual Factor Controls on Microbial Processes," *mSystems* 7, no. 4 (2022): e00301-22, https://doi.org/10.1128/msystems.00301-22.
[3] I. Del Valle, et al., "Translating New Synthetic Biology Advances for Biosensing Into the Earth and Environmental Sciences," *Frontiers in Microbiology* 11 (2021), https://doi.org/10.3389/fmicb.2020.618373.
[4] H.-Y. Cheng, et al., "Volatile Gas Production by Methyl Halide Transferase: An In Situ Reporter Of Microbial Gene Expression In Soil," *Environmental Science and Technology* 50, no. 16 (2016): 8750–8759, https://doi.org/10.1021/acs.est.6b01415.

In the Design phase, the project used ORNL's clustering workflow to identify chemically distinct binding pockets to employ structural modeling and molecular dynamics simulations to propose targeted mutations to natural and engineered MHTs. In the Build phase, candidate enzymes and mutants were transformed and expressed under a tightly inducible system in *Escherichia coli* for gas production. In the Test phase, gas chromatography–mass spectroscopy was used to evaluate indicator gas production of designed variants compared to the wild-type. Each halide specialist's activity was characterized across a range of temperature and salt conditions. In the Learn phase, structural predictions, molecular dynamics simulations, and quantum–chemical methods were used to interpret activity differences emerging from the proposed mutations and select new point mutants for subsequent design rounds. In two cycles, this project tested 36 mutations in proteins from 23 organisms, including plants, fungi, and bacteria. The following subsections briefly outline these infrastructural developments and accomplishments.

Cross-Kingdom Sequence Clustering Infrastructure Guides Experimental Screening

To prioritize candidate sequences for experimental assay, this project developed a clustering workflow for cross-kingdom screening candidates from 8,466 putative MHT sequences using feature embedding of the amino acids forming their active binding site. This set of 17 features included physicochemical properties such as volume, quantum–chemical features such as polarizability, and sequence-based structural predictions of disordered regions in the protein. Using *k*-means clustering, this project identified 10 clusters. Comparing the halide activity of representative sequences from each cluster, the project team found that this approach improved upon phylogeny alone, identifying enzymes with chemically similar binding pockets or residues even when overall sequence similarity was low. It also did the reverse: cases of low binding-site similarity were identified even when sequences were phylogenetically closely related. These results show the relevance of including structural and quantum–chemical features in screening workflows to capture binding site similarity at single-residue resolution.

Predictive Structural Analysis and Deep Learning Models Assist Design of Halide Specialist Enzymes in Three Species

To complement the sequence clustering approach, template-based deep learning approaches in Boltz v. 1.0.0.[5] were used to generate structural predictions from enzyme sequences. Predictions used multiple sequence alignments from HHsuite v. 3.3.0[6] against the UniRef100 database.[7,8] Comparative structural analysis with these models provided insights for rational design of MHTs and allowed the researchers to identify point mutations that conferred halide specificity in proteins from *Arabidopsis thaliana*, *Batis maritima*, and *Burkholderia xenovorans*.

Experimental Assays Characterize Halide Specificity, Confirm Halide Specialist Activity, Across a Range of Temperatures and Salt Conditions

This project successfully implemented a highly sensitive protocol to quantify microbial gas production. Using this approach, this study found point mutations of key residues that conferred 100–300-fold halide specificity, with enzyme specialists tailored for various salt conditions (20 mM, 100 mM, and 10:10 mM and 50:50 mM salt admixture) and temperature ranges (30°C–37°C). The particular mutations and each halide specialist's activities are the subjects of invention disclosures and will be included in a future publication.

---

[5] J. Wohlwend, et al., "Democratizing Biomolecular Interaction Modeling," bioRxiv (2024): 2024.11.19.624167, https://doi.org/10.1101/2024.11.19.624167.
[6] M. Steinegger, et al., "HH-Suite3 for Fast Remote Homology Detection and Deep Protein Annotation," *BMC Bioinformatics* 20, no. 1 (2019): 473, https://doi.org/10.1186/s12859-019-3019-7.
[7] B. E. Suzek, et al., "UniRef: Comprehensive and Non-redundant UniProt Reference Clusters," *Bioinformatics* 23, no. 10 (2007): 1282–1288, https://doi.org/10.1093/bioinformatics/btm098.
[8] B. E. Suzek, et al., "UniRef Clusters: A Comprehensive and Scalable Alternative for Improving Sequence Similarity Searches," *Bioinformatics* 31, no. 6 (2014): 926–932, https://doi.org/10.1093/bioinformatics/btu739.

Molecular Dynamics and Quantum Chemical Studies Contribute Mechanistic Insights for Interpreting Halide Specificity

Simulations with classical molecular dynamics, coupled with metadynamics, identified the most energetically favorable conformations to enable a mechanistic study of the specialists. This project tested the hypothesis that interactions between the halide and tightly coordinated water molecules are important for conferring halide specificity using quantum–chemical calculations to refine the findings. The results highlighted the roles of ion hydration and pocket hydrophobicity in halide specificity.

BUILDINGS AND TRANSPORTATION SCIENCE DIVISION

11603: Enhanced Tunnel/Cavity Detection via Through-Transmission Signature Analysis

M. Kass, C. Finney, O. Marcillo, D. Splitter, M. Maceira

Project Description

The objective of this effort was to improve the detection of underground tunnels. The distinguishing feature of the proposed concept was placing an acoustic source at a depth below the tunnel so that the generated signal created a through-transmission signal. This approach differs from the standard surface-based methodologies currently in use (e.g., ground penetrating radar, electrical resistance profiles, surface seismic/acoustic generators) The hypothesis was that the propagation and attenuation of a through-transmission acoustic signal would provide critical information to facilitate tunnel detection. This project's approach was to bore a horizontal tunnel (using an approved directional drilling contractor) at a suitable site on the ORNL reservation. Before tunnel construction, baseline measurements were performed by placing an array of acoustic surface sensors along the surface at a preapproved site. A vertical tunnel was bored to a depth of 20 m (approximately 60 ft). At selected depths, an acoustic transmitter was lowered into the tunnel and coupled to the soil for source signal generation at multiple frequencies. Upon completion of the baseline tests, a horizontal tunnel (0.5 m diameter) was bored at a depth of 3 m, and the process was repeated. The project deliverable was a report demonstrating the effectiveness of integrating through-transmission with pulse-echo-based detection methods to improve detectability.

Mission Relevance

The technology and approach of this project are applicable toward preventing nonproliferation of nuclear material transport, as well as other illicit items such as narcotics. As such, the project directly related to the National Nuclear Security Administration's mission to prevent nuclear weapon proliferation.

Results and Accomplishments

The experiments were undertaken in late August 2024 in the 7000 area at ORNL. The testing grounds had been previously mowed to facilitate the experiments, and the ground was dry; no rain had recently fallen. In total, 24 geophones were placed on the ground surface at various locations. Data were collected continuously by the geophone array at a frequency of 1,000 Hz. A commercial vertical boring machine was used to drill the vertical bore at depths of 1, 5, 10, 15, 20, and 30 ft. The drilling rate was around 200 rpm, and the bore diameter was 6 in. At each depth, an electronic trumpet/snail horn was lowered into the bore. The horn bell was oriented in the direction of the tunnel. The first day of testing involved taking baseline measurements before the horizontal tunnel was put in place. The horn produced a 510 Hz signal at 105–118 dB. The horn was operated by a battery, and the signal was pulsed at 1 s intervals. The second day involved tunnel placement. This was accomplished using a jack-and-bore operation. The tunnel consisted of two 20 ft sections of steel pipe that had a wall thickness of 1 in. The tunnel diameter was 3 ft. After the first 20-foot section was installed, the second pipe was welded onto the first, and then drilling continued until both sections (totaling 40 ft) were placed at a depth of 10 ft below the surface. The final day consisted of repeating the vertical bore and signal generation but with the tunnel in place.

The results confirmed the hypothesis that by placing an acoustic source below the tunnel (at 20 or 30 ft), a characteristic subharmonic was observed that was not present for transmissions that originated above the tunnel. This characteristic subharmonic was not observed for sound transmissions placed at depths approaching the tunnel depth. Interestingly, these subharmonics were detected by the surface geophones 60 ft from the end of the tunnel bore. Even though the soil was a settled tailing pile, the sensors were able to pick up the acoustic frequencies.

The research team was able to (1) demonstrate that acoustic frequences around 500 Hz were detectable using a standard geophone at distances of 60 ft; (2) observe that a characteristic subharmonic occurred, indicating the presence of a tunnel; and (3) determine that this subharmonic only occurred when the source signal was placed at depths well below the tunnel.

12036: An Intelligent Metasurface-Based Wireless Power Transfer System

E. Aydin, M. Mostak, E. Ali Riza

Project Description

This project focused on designing, optimizing, and experimentally validating an intelligent, metasurface-assisted wireless power transfer (WPT) system to mitigate performance degradation caused by coil misalignment and increased transfer distance. This work was completed by implementing an optimized 3 × 3 controllable metasurface coil array for a 1 kW WPT system and systematically investigating coil activation strategies under spatial misalignment. Measured datasets corresponding to four anchor misalignment conditions were preprocessed and consolidated, followed by visualization of average per-coil efficiency to identify spatial activation patterns. Based on these data, baseline and supervised learning–based models were developed to predict optimal metasurface coil activation for a given receiver position. The results demonstrated that data-driven metasurface control can effectively redirect magnetic field density, improve transfer efficiency, and increase receiver positioning tolerance. This project established a practical framework for adaptive and scalable metasurface-assisted wireless charging systems.

Mission Relevance

This project advanced DOE and ORNL missions by enabling a resilient, energy-efficient, and intelligent WPT technology to address critical challenges in energy delivery and advanced manufacturing. By developing and experimentally validating a metasurface-assisted, machine learning–enabled WPT system, the project delivered a scalable solution that improved power transfer efficiency and robustness under misalignment conditions, which are considered challenging for real-world applications of WPT systems, including high-power applications. These advances supported DOE's mission by contributing to secure, efficient, and adaptable energy technologies.

Results and Accomplishments

This project established a comprehensive performance baseline for conventional two-coil, three-coil, and metasurface-based WPT systems operating at 85 kHz and up to 1 kW. This work explicitly quantified the effects of ground clearance, load resistance, and spatial misalignment. One method of improving this misalignment tolerance is using intermediary components—either a single intermediate coil or a more complex multielement metamaterial—placed between the transmitter and receiver. With the proper configuration of these intermediary coils and metamaterials, usually in the form of metasurfaces, the transmitter magnetic field can be focused and directed.

Extensive parametric sweeps demonstrated that misalignment and increased transfer distance significantly degrade efficiency and power output in traditional two-coil configurations, but three-coil and metasurface-based architectures showed improved robustness at the cost of increased system complexity and higher source voltage requirements. These results clarified the fundamental trade-offs between efficiency, misalignment tolerance, and circuit stress and provided a rigorous reference against which advanced metasurface-assisted designs could be evaluated.

Building on this foundation, the project delivered a fully parameterized and optimized 3 × 3 controllable metasurface coil array, including systematic optimization of coil turns, pitch, side length, interelement gap, resonant frequency, and spatial placement. Optimization studies identified a dominant metasurface

configuration—activating centrally located elements while suppressing corner elements—that consistently maximized coupling and efficiency. The optimized metasurface geometry achieved efficiencies exceeding 95% at kilowatt-level power, with peak performance obtained when the metasurface was tuned slightly off the primary or secondary resonance.

A major outcome of the project was demonstrating that selective metasurface coil activation enabled active control of the magnetic field distribution, which allowed the WPT system to compensate for receiver misalignment in a controlled and repeatable manner. By exhaustively evaluating all 512 activation states of the 3 × 3 metasurface array, this study established clear relationships between activation patterns and induced secondary voltage, output power, and efficiency. The results showed that activating metasurface elements nearest the receiver while suppressing weakly coupled elements concentrated magnetic flux in the desired region, which led to substantial performance improvements under current- and voltage-source excitation.

Across multiple anchor misalignment conditions, the metasurface-assisted system consistently maintained higher efficiency and reduced sensitivity to lateral displacement compared with conventional fixed-coil designs. Rather than experiencing sharp performance degradation with misalignment, the system adapted by reconfiguring active metasurface elements to redirect magnetic fields toward the displaced receiver. These findings validated the core hypothesis that controllable metasurfaces boost the magnetic field, which established a practical and scalable approach for misalignment-tolerant, intelligent WPT systems.

Finally, comparative analyses confirmed that the optimized metasurface-assisted WPT system offers superior robustness at larger ground clearances and under geometric scaling of the primary coil. The system outperformed conventional two-coil designs as separation distance increased. The metasurface design preserved efficiency advantages at elevated ground clearances and enabled stable 1 kW operation across multiple misalignment cases with moderate source voltage requirements. Collectively, these outcomes demonstrated that metasurface-based control is not merely an incremental improvement but also a qualitative advancement in WPT system capability that enables intelligent, scalable, and misalignment-tolerant wireless charging architectures suitable for demanding applications such as autonomous UAV platforms.

In summary, this project successfully completed an intelligent metasurface-based WPT system by designing, optimizing, and developing the metasurface; implementing an AI-based activation strategy based on receiver position, and performing experimental tests and validations. In the first phase of the project, the metasurface for a 1 kW wireless charging system was designed, analyzed, and optimized. Following the design and metacoil optimization stage, this project developed a practical pipeline to determine the on/off states of each coil switch in the array based on the receiver coil's alignment. This pipeline consolidated data to learned coil selection under misalignment and incorporated interpretable baselines and a multilayer perceptron model trained on soft, augmented labels. To validate the approach, the system was fabricated and experimentally tested in the laboratory. The results demonstrated that the metasurface-based WPT system outperforms a conventional two-coil system under misalignment conditions and at larger air gaps, which enables efficient power transfer over longer distances.

12053: Hydrogen Analysis by Gas Chromatography–Mass Spectrometry

V. Lobodin, Y. Park, C. Finney

Project Description

Detecting hydrogen in complex gas mixtures, including volatile and semivolatile organic and inorganic compounds, is important for developing and optimizing hydrogen production but is difficult due to its low molecular weight, high diffusivity, and the presence of interfering species such as oxygen, nitrogen, carbon monoxide, carbon dioxide, and methane. Conventional methods, including residual gas analyzers,

thermal conductivity detectors, and helium ionization detectors, are limited by poor selectivity, nonlinear response, or inadequate separation of hydrogen from other gases. However, in the absence of chromatographic separation, hydrogen quantification and comprehensive characterization of other species in a mixture are not reliable. Furthermore, quadrupole mass spectrometers, commonly used as residual gas analyzers and in gas chromatography (GC)–mass spectrometry (MS) configurations, are subject to the zero-blast effect. Traditionally, chromatographic analysis of hydrogen would require a two-column configuration for achieving its separation from other gases (e.g., oxygen, nitrogen, carbon monoxide, carbon dioxide, methane, and other volatile hydrocarbons). The columns that are used are packed or molecular sieve porous layer open tubular (PLOT) columns connected to a particle trap on the detector side to prevent sorbent particles from entering the ion source of the mass spectrometer. This project accomplished separation of hydrogen and other gases in a refinery gas test mixture using helium as a carrier gas with a single capillary column. The MS detection method involved simultaneous acquisition in selected ion monitoring mode at a mass-to-charge ratio (m/z) of 2 for hydrogen detection and scan mode within the range of m/z = 10 to 160 for other gases. Detection of hydrogen in selected ion monitoring mode at m/z = 2 significantly improved the signal-to-noise ratio and detection limit. The baseline separation of all gases in the refinery gas test mixture was achieved. One of the most critical regions for separation and detection of hydrogen was 2–7 min, and hydrogen was completely separated from oxygen, nitrogen, carbon monoxide, and methane, thus minimizing interference from the zero-blast effect for accurate quantitation of hydrogen. Additionally, all three isobar gases (nitrogen, carbon monoxide, and ethene), which have the same molecular weight of 28 Da but different compositions, were reliably separated chromatographically.

Mission Relevance

This project supported DOE missions by advancing innovation in energy abundance, availability, and security. This effort aligned with DOE's mission to ensure America's security and prosperity by addressing its energy challenges through transformative science and technology solutions. Developing a quantitative, highly sensitive, and repeatable method for hydrogen analysis strengthens the foundation for hydrogen production, storage, and distribution, which are key components of the hydrogen energy utilization sector. Reliable hydrogen quantification is important for scaling hydrogen technologies in power generation, fuel cells, and energy storage, as well as supporting emissions reduction, leak detection, and environmental safety. This project's methodology had strong commercial and cross-sector applicability, including hydrogen production and distribution quality control, fuel cell vehicle infrastructure, industrial process optimization, grid-scale energy storage monitoring, environmental leak detection, geological hydrogen exploration, and aerospace safety.

Results and Accomplishments

Development of a GC–MS Hydrogen Detection Method

To address the limitations of detecting hydrogen, this project first developed and validated a novel analytical method for direct hydrogen analysis that employs a single capillary GC column and an electron ionization mass spectrometer. Several innovations were implemented.

Modified sampling procedure. This project designed an improved gas introduction system that allows efficient and controlled injection of gas mixtures into the GC–MS instrument. Using precision sample loops (ranging from 250 µL to 2 mL) enabled consistent and reproducible introduction of hydrogen-containing mixtures and minimized sample loss and contamination.

Single-column GC configuration. Traditional GC systems for hydrogen analysis often require two-column configurations with molecular sieve or packed PLOT columns, accompanied by particle traps to protect the ion source. This study achieved complete separation of hydrogen and other gases using a single 60 m GS-GASPRO capillary column operated under subambient conditions. Helium was used as a

carrier gas to provide optimal peak shape and resolution. A split ratio in the range of 5:1 to 100:1 was tested to deliver the best signal-to-noise ratio and chromatographic performance across the tested range.

Modified ion source for enhanced sensitivity. The electron ionization ion source of the mass spectrometer was modified with a low-gauss magnet (220 G) to improve focusing of low-mass ions. This adjustment, combined with optimized tuning macros for hydrogen, significantly enhanced sensitivity and signal stability. The system achieved reliable detection of hydrogen down to 5 ppm.

Chromatographic and Mass Spectrometric Performance

Hydrogen and other gases in a refinery gas test mixture were successfully separated using helium as the carrier gas. The MS detection method involved simultaneous acquisition in selected ion monitoring mode at $m/z = 2$ for hydrogen detection and in scan mode at $m/z = 10$ to 160 for other gases. The selected ion monitoring mode dramatically improved the signal-to-noise ratio and detection limit for hydrogen. Baseline separation of all gases in the mixture was achieved, with hydrogen eluting between 2 and 7 min and completely resolved from oxygen, nitrogen, carbon monoxide, and methane, thereby eliminating interference from the zero-blast effect and ensuring accurate quantitation.

The method also achieved chromatographic separation of isobaric species such as nitrogen, carbon monoxide, and ethene (each with a nominal molecular weight of 28 Da), which are often challenging to distinguish. This capability allows for more comprehensive compositional analysis in complex gas mixtures and provides qualitative and quantitative insight in a single analytical run.

Repeatability and Quantitative Accuracy

The method demonstrated strong repeatability and precision. Triplicate analyses of standard gas mixtures yielded consistent hydrogen peak areas with standard deviations below 5%. Calibration curves were linear across the examined concentration range (5 ppm: 5%), which confirmed the method's suitability for trace and bulk hydrogen quantification. This robustness established the technique as a dependable analytical platform for hydrogen process monitoring, quality assurance, and research applications.

CHEMICAL SCIENCES DIVISION

11684: Novel Concept of Polymer-Based Solid-State Batteries

C. Gainaru, Md. A. Rahman, G. Polyzos, A. Sokolov

Project Description

Solid-state batteries are of interest because of their potential for safety and high energy density. However, their deployment is currently limited by the absence of solid electrolytes that satisfy all requirements for conductivity and mechanical and electrochemical stability. This project aimed to develop fully functioning solid-state batteries based on single-ion (Li^+) conducting polymers rather than traditional batteries that use ceramics, doped poly(ethylene oxide), or doped poly(ethylene oxide)–ceramic composites, which so far have been proven successful only for niche applications. This project developed several Li^+-conducting electrolytes to enable in situ (i.e., directly in the assembled cell) polymerization and crosslinking. Without developing new chemistry or addressing the fundamentals of charge transport in conducting polymers, this project aimed to optimize the single-ion conducting polymer concept devised at ORNL for its implementation in fully operating batteries.

Mission Relevance

The work aligned with the DOE's mission by delivering breakthrough science for an efficient and secure energy future by developing a novel concept for solid-state batteries.

Results and Accomplishments

This research introduced a new concept of doped quasi-single-ion conducting copolymers that can be polymerized and crosslinked in situ and enabled (1) good adhesion with the electrodes, which mitigated delamination and interfacial resistances; (2) homogeneous distribution of charge carriers near the electrode surfaces to prevent dendrite formation; (3) interchangeable transport of lithium and sodium (and possibly other cations) using simple chemistry procedures at variance with superionic ceramics in which switching cations is obstructed by structural stability; (4) relatively large lithium and sodium conductivities (approximately 0.1 mS/cm at ambient conditions) and a decent transport number (approximately 0.5), which hampers polarization effects; and (5) mechanical robustness with safety benefits in preventing electrolyte spilling upon battery damage. The proposed lithium electrolyte demonstrated good performance in Li^+/Li symmetric cells, as well as in lithium metal cell with a nickel manganese cobalt or lithium iron phosphate cathode. Additionally, the proposed electrolyte concept was easy to implement with current battery technologies and provided a viable route for designing high–energy density, safe, polymer-based batteries capable of a long duty cycle.

11691: Novel Alkaline Exchange Membranes for Economic Water Electrolyzer

X.-G. Sun, I. Popovs, X. Lyu

Project Description

Electrolyzers using alkaline exchange membranes (AEMs) can significantly reduce overall production cost because of the low cost of the membranes and enhanced reaction kinetics, as well as the use of nonprecious metals as catalysts. In comparison, electrolyzers using proton exchange membranes must use precious noble metals (e.g., platinum) as a catalyst, and perfluorinated proton exchange membranes have a high cost. To date, the application of AEMs in electrolyzers has been limited by its low ionic conductivity and low chemical stability. The overarching goal of this project was to improve the chemical stability and ionic conductivity of AEMs with low swelling for application in AEM electrolyzers. In this project, the chemical stability of AEMs was improved by eliminating β-hydrogen atoms, and low swelling was achieved by using controlled substitution on the hydrophobic polymer main chains. The

project also attempted to enhance ionic conductivity of the AEMs by increasing the flexibility of the cationic group and optimizing the charge density at the side chains to achieve ionic conductivities higher than 100 mS/cm at 80°C.

Mission Relevance

The stable AEMs obtained in this research can be used for AEM electrolyzers and fuel cells. AEM technology enables the transition from titanium to stainless steel for accessory parts, which could eliminate 75% of the stack part cost. This accessory cost reduction could significantly reduce the cost of hydrogen as an energy source, thus supporting DOE's mission to ensure America's security and prosperity by addressing its energy challenges.

Results and Accomplishments

The goal of the project was to synthesize an all-aromatic polymer backbone on which the styrene monomer had a specific functional group (i.e., without active β-hydrogen atoms to improve its stability in a strong basic environment). The initial obtained copolymer failed to obtain a stable membrane because of insufficient copolymerization. Next, the project assessed whether the targeted functional monomer would lead to a polymer that is more stable than the analog polymer with active β-hydrogen atoms. The obtained functional styrene polymer demonstrated better stability than its analog polymer in the strong basic environment but suffered much lower ionic conductivity caused by the bulky hydrophobic environment around the ammonium cation.

COMPUTATIONAL SCIENCES AND ENGINEERING DIVISION

11757: Computational Error Model for a Fault-Tolerant Quantum Processor

T. LeBlond, R. Bennink, C. Seck, P. Groszkowski

Project Description

This project modeled logical errors for surface code circuits implemented on a grid-based trapped-ion quantum charge-coupled device (QCCD) architecture. The project computed logical quantum channels and examined error metrics under a realistic hardware error model involving both stochastic and coherent errors. To simulate surface code circuits at practical code distances in a manner that is asymptotically exact, a Monte Carlo technique was used to sample from an underlying quasiprobability distribution of Clifford circuits. The project verified the fault tolerance of the surface code at coherent dephasing rates relevant to current trapped-ion quantum computers. This work laid the foundation for the first realistic computational error model for a fault-tolerant quantum processor, which will enable the most accurate quantification to date of resource requirements for practical circuits running on these processors.

Mission Relevance

Quantum computing promises to rapidly accelerate scientific progress in fields such as chemistry, cryptanalysis, and optimization, which are relevant to DOE's mission to ensure America's security and prosperity by addressing its energy, environmental, and nuclear challenges. However, efficient fault tolerance is essential for quantum computing to deliver new scientific or cryptanalytic capabilities. This project's focus on the efficacy of fault tolerance for realistic hardware models and when it may be achievable are, therefore, key issues.

Results and Accomplishments

This project simulated realistic trapped-ion QCCD hardware mapping for surface code under an accurate noise model that was developed in consultation with Quantinuum. The trapped-ion QCCD noise model was implemented within the Oak Ridge Quasi-Clifford Simulator (ORQCS), a state-of-the-art simulator developed at ORNL that was improved within the context of this project. In particular, ORQCS is now able to process time-resolved trapped-ion surface code circuits generated by the Trapped-Ion Surface Code Compiler (TISCC) by inserting appropriate instructions corresponding to noise processes relevant to the QCCD architecture. During simulation, ORQCS tracked the time accumulated during idling and movement for each ion to rebuild noise channels on the fly according to this information. Additionally, the infrastructure for defining errors in the simulator was generalized to allow multiqubit error processes, and the ability to track sampled errors at circuit locations was implemented. Furthermore, extensive testing of ORQCS within this project verified the correctness of its implementation as well as its scalability.

The ORQCS simulation method involved an up-front decomposition of all non-Clifford hardware operations as superpositions of Clifford operators. The Clifford operators were then probabilistically sampled and efficiently simulated using stabilizer propagation to yield expectation values for Pauli observables at the circuit output. Optimal decompositions were implemented for single-qubit Z-axis rotations and rotations about an axis in the X–Y plane, which are primitive instructions in trapped-ion quantum computers. By minimizing the L1 norms of the Clifford decompositions, the statistical variance of the simulation outputs was minimized.

This project implemented all the infrastructure necessary to perform trapped-ion surface code simulations using (1) TISCC to generate circuits, (2) Stim to generate a description of how errors (if treated stochastically) affect otherwise-deterministic combinations of measurement outcomes in what is called a detector error model, (c) ORQCS or Stim to simulate these circuits, and (d) PyMatching to decode

simulated measurement outcomes using the detector error model. The resulting workflow allows researchers to efficiently obtain decoded expectation values for the Pauli-basis logical operators of a surface code patch following noisy rounds of error correction.

Using this workflow, process tomography was simulated to compute effective logical channels for error-corrected surface code patches of various sizes and various amounts of hardware error. Logical channels were computed in the form of Pauli transfer matrices and Chi matrices. These characterizations provide much more information about the behavior of the surface code than typical approaches that estimate logical error rates from stochastic Pauli noise. Thus, the team was able to quantify coherent logical error buildup in the surface code under coherent physical errors. The team was also able to calculate the diamond error, which is the most rigorous known definition for the logical error rate.

This project accurately simulated the impact of coherent dephasing errors—a form of qubit error that is difficult to simulate but suspected to be more detrimental than other forms of noise that are commonly simulated. Coherent dephasing, which accumulates during physical qubit idling and movement, has been suspected in previous studies to contribute to the emergence of coherent logical error. This project fully characterized such error using process tomography and demonstrated the surface code sizes and coherent dephasing rates at which coherent logical error emerges.

These simulations provided the first detailed logical error model for the surface code that included coherent physical errors within a circuit-level noise model. A key conclusion was that dephasing error has a negligible impact on surface code logical error rate for a plausible model of near-future trapped-ion QCCD architectures. Secondly, the project confirmed that the coherent component of logical error is negligible at large code distances (the coherent component was negligible for code distance above 5). Finally, the project also found that accurate simulation of dephasing generally leads to larger logical error values than simulations that approximate the dephasing error using Pauli twirling.

In summary, this project established a capability for performing simulations of the surface code under realistic hardware mappings and error models and for extracting logical error models and rigorous logical error rates from these simulations. As a proof of principle, this project implemented trapped-ion QCCD mapping for the surface code and simulated it under coherent dephasing noise, which has been conjectured to be more detrimental than other error sources in that hardware platform. This work verified the feasibility of storing quantum information using the surface code on this architecture by showing that logical error rates decrease with increasing code distance up to a code distance of 11 at experimentally relevant values of hardware error.

11794: Complementary Metal–Oxide–Semiconductor Hyperentangled Sources for Quantum Networking

H.-H. Lu, A. Miloshevsky

Project Description

Across the ORNL campus, quantum photonic integrated circuits deployed in multiple buildings use existing fiber and wavelength-division multiplexing to distribute entangled states. This project aimed to design and fabricate (using commercial foundries) complementary metal–oxide–semiconductor (CMOS)-compatible photon sources that produce high-flux polarization–frequency hyperentangled photons. This work aimed to enable seamless integration with a variety of quantum devices and pave the way for robust and scalable quantum networks and increasingly capable quantum photonic integrated circuits for advanced quantum information processing. This project developed simulation and photonic integrated circuit testing to design good-quality photon sources.

Mission Relevance

This research focused on understanding how CMOS-compatible photon sources can be improved to deliver robust and scalable quantum networks. This work aligned with DOE and ORNL missions in energy security and energy economy by combining basic energy science with circuit design.

Results and Accomplishments

This project employed a simulation workflow using Lumerical MODE/finite-difference time-domain to design and validate several core components: edge couplers, polarization splitter-rotators, Y-branches (Y-junction splitters), and micro-ring resonators on a silicon nitride platform. The project also developed analytical microring resonator models to guide optimal design. This study systematically varied waveguide widths, gaps, and lengths to minimize loss, maximize polarization-splitting extinction, or meet target quality factors. The project included additional test structures to measure loss, coupling, and quality factors to guide the next fabrication run; training the larger dataset required approximately 1 week. The project also found that tuning the learning rate can help strike a balance between accuracy and training time. In addition to simulation, this project characterized each photonic chip on an automated silicon photonics probe station, which rapidly scanned many optical structures to perform linear measurements, including coupling and propagation.

COMPUTER SCIENCE AND MATHEMATICS DIVISION

11724: Federated Learning in Heterogeneous Environments for Science Applications

O. Kotevska, R. Prout

Project Description

Federated learning (FL) offers a decentralized approach to model training that enables data owners to collaboratively train machine learning models without sharing raw data, which is an essential feature for privacy-preserving applications in domains such as energy systems. Although prior studies have largely focused on FL in simulated environments, its deployment on supercomputing platforms and advancement in privacy-preserving methodologies remain underexplored. This project addressed these gaps by developing two open-source software packages: PETINA (Privacy prEservaTIoN Algorithms) and PRESTO (Privacy REcommendation and SecuriTy Optimization). A novel privacy-preserving algorithm was introduced to advance the state of the art in secure FL. Additionally, NVFlare was customized for the Frontier supercomputer by integrating PETINA, and a user manual was produced to facilitate broader adoption. These contributions establish a foundation for scalable, privacy-aware FL in heterogeneous high-performance computing (HPC) environments and demonstrate its potential in scientific applications.

Mission Relevance

This project supported DOE's mission by advancing transformative science and technology solutions that address critical challenges in data privacy, distributed intelligence, and scalable computing, which are key enablers of national energy and security priorities. By enabling FL on HPC platforms such as the Frontier supercomputer, the project outcomes facilitate secure, collaborative model training across sensitive data sources without compromising privacy.

Results and Accomplishments

The project produced a comprehensive software and methodological stack for privacy-preserving FL at scale, spanning algorithmic innovation, system-level integration, and practical tooling. Two software packages were developed: PETINA, which enables privacy-enhancing training across distributed scientific data sources, and PRESTO, a privacy evaluation and recommendation tool. Together, these tools support adaptive privacy configuration, systematic privacy–utility analysis, and trust-aware FL model development.

PETINA is a modular Python library that implements a broad suite of differential privacy (DP) mechanisms for machine learning and federated workflows. It provides standard noise mechanisms (e.g., Gaussian, Laplace, Rényi–Gaussian); projection- and sketch-based perturbation methods; and adaptive clipping, pruning, and utility-preserving transformations that can be directly integrated into PyTorch- and NumPy-based pipelines. PETINA was designed to support flexible privacy accounting and mechanism composition, enabling privacy-preserving training under heterogeneous data distributions and system constraints. Example workflows demonstrated integration with deep learning models and federated training pipelines, making PETINA a reusable foundation for deploying DP in scientific FL settings.

PRESTO complements PETINA by providing an automated framework for privacy mechanism evaluation and recommendation. PRESTO leverages Bayesian optimization to analyze dataset characteristics, rank privacy mechanisms and hyperparameters, and quantify privacy–utility trade-offs with confidence and reliability metrics. Rather than requiring users to manually tune DP parameters, PRESTO systematically recommends mechanisms that best satisfy user-defined utility and privacy objectives and provides visualization tools for comparing original and privatized data distributions. Its

extensible architecture enables new privacy algorithms to be incorporated as plugins, supporting deployment in regulated scientific and enterprise environments.

At the system level, the project customized NVIDIA's NVFlare framework for deployment on the Frontier exascale supercomputer to enable large-scale FL orchestration under realistic system and data heterogeneity. NVFlare's client–server architecture was adapted to Frontier's HPC environment, where the FL Server orchestrates training workflows, manages job life cycle and scheduling, and securely distributes global models across federated rounds. FL Clients were mapped to isolated GPU or multi-GPU node resources to ensure that local training occurs on private datasets and only model updates are communicated back to the server. The Aggregator module was configured to efficiently combine updates using federated optimization algorithms (e.g., FedAvg variants) while accounting for heterogeneous compute performance and partial client participation. An integrated Admin Server (FLARE Console) enabled secure job submission, monitoring, and experiment control. This integration was operationalized through a user manual and configuration toolkit, including job templates and deployment scripts. Ongoing collaboration with the NVFlare team continues to extend this work toward larger client counts and more complex FL workloads.

In parallel, the project introduced a novel projection-based privacy-preserving algorithm for FL to extend the ProjUnit framework with mechanisms tailored for cross-node and large-scale FL deployments. The method applied structured random projections to client gradient updates before DP noise injection, which reduced gradient dimensionality and sensitivity while preserving essential learning signal. This project also introduced a dynamic projection dimension (K) scaling strategy, which enabled the algorithm to adaptively balance privacy, accuracy, and computational efficiency based on training dynamics. Extensive experimental evaluation against state-of-the-art baselines, including SRHT-based projections, CountSketch-based methods, and Gaussian DP, demonstrated competitive accuracy, stable performance across increasing client counts, and improved communication efficiency. Validation on publicly available benchmark datasets confirmed the method's suitability for privacy-preserving FL in heterogeneous, real-world HPC environments.

11727: Verifying Quantumness: Developing Algorithms for Nonlocal Games

E. Wong, S. Chehade, A. Delgado

Project Description

Nonlocal games (NLGs) have rich mathematical theory with untapped practical applications because of the complexity of optimal strategy constructions. This project incorporated variational quantum algorithms into strategy computation. Variational quantum algorithms encoded game rules and strategies into a Hamiltonian optimized through a classical–hybrid loop. This project aimed to rigorously develop the mathematical theory to enable algorithm design for verifying the quantum behavior of systems without relying on external benchmarks—a concept referred to as self-testing. In particular, this project explored nonlocality through the magic squares game (MSG). Within the project's framework, the five key objectives were (1) analyzing how to generalize the MSG explicitly (if possible), (2) designing variational quantum algorithms to compute quantum strategies, (3) demonstrating end-to-end NLG strategy construction, (4) automating this project's implementation to facilitate broad adoption of this methodology, and (5) establishing theoretical guarantees for multigame scenarios by combining algebraic and experimental methods.

Mission Relevance

This project focused on finding theoretical setups for entanglement certification in quantum systems, as well as enabling multisystem verifications under this framework. The project aimed to provide mathematical guarantees for the existence of perfect strategies and algorithmic primitives for multigame

scenarios. This work could potentially lead to experimental designs of quantum-accelerated high-performance computing systems, which aligns with DOE's mission in quantum computing and high-performance computing integration. These results apply to quantum device verification in a potentially scalable way.

Results and Accomplishments

A Variational Quantum Algorithm for Playing NLGs

Variational quantum algorithms offer a promising near-term approach to finding optimal quantum strategies for playing NLGs. These games test quantum correlations beyond classical limits and enable entanglement verification. This work presented a variational framework for the MSG: a two-player NLG with perfect quantum advantage. The project constructed a value Hamiltonian that encoded the game's parity and consistency constraints, then optimized the parameterized quantum circuits to minimize this cost. This approach built on the stabilizer formalism, leveraged commutation structure for circuit design, and was hardware-efficient. Compared with existing work, this project's contribution emphasized algebraic structure and interpretability. The project validated this method through numerical experiments and outline generalizations to larger games and automated the process via scripts and notebooks.

Hilbert Space Compression for Parallel NLGs

In finite dimensions, the standard method for playing several games in parallel requires a tensor product of the local Hilbert spaces, which scales additively in the number of qubits. This project showed that this additive cost can be reduced by exploiting algebraic embeddings. This work introduced two forms of compressions. First, when a referee selected one game from a finite collection of games at random, the game quantum strategy could be implemented using a maximally entangled state of dimension equal to the largest individual game, thereby eliminating the need for repeated state preparations. Second, the project established conditions under which several games could be played simultaneously in parallel on fewer qubits than the tensor product baseline. These conditions were expressed in terms of commuting embeddings of the game algebras. Moreover, this work provided a constructive framework for building such embeddings. Using tools from Lie theory, this project showed that aligning the various game algebras into a common Cartan decomposition enables such a qubit reduction. Beyond the theoretical contribution, this framework cast NLGs as algebraic and algorithmic primitives for distributed and resource-constrained quantum computations and suggested NLGs as a comparable device-independent dimension witness.

Identification of Infinite Dimensional Barriers for Larger MSGs

Previous work from Paddock et al. characterized finite graph incidence groups, which is a class of generalized MSGs, via two disjoint connected components of cycles (Theorem 1.4).[1] A 4 × 4 MSG is isomorphic to $K_{4,4}$, which is easy to construct two disjoint 4-cycles, implying that the graph incidence group is not finite and thus it is infinite. The key technical result from Paddock et al.[2] was that the presence of two disjoint cycles in a graph (as in $K_{4,4}$) caused the associated solution group $\Gamma(G, b)$ to be infinite. Physically, this implied that any perfect quantum strategy cannot be faithfully represented using only finite-dimensional unitary representations. Consequently, the measurement operators involved may not close under any finite Lie algebra, and they may not respect any finite-dimensional Cartan subalgebra structure.

What this means for entanglement verification is that unlike the 3 × 3 MSG, which needs two Bell states (a finite amount of entanglement) or three for the variational version, the 4 × 4 case has an infinite solution group. Therefore, no known finite-dimensional state implements a perfect quantum strategy.

[1] C. Paddock, et al., "Arkhipov's Theorem, Graph Minors, and Linear System Nonlocal Games," *Algebraic Combinatorics* 6, no. 4 (2023):1119–1162.
[2] Ibid.

Thus, an experiment would need an arbitrarily high amount of entanglement—if such a state can be approximated. In principle, the amount of entanglement is unbounded, which is not practical in reality.

11822: Discontinuous Basis Heirarchy for Kinetic Plasma Equations

M. Stoyanov, E. Endeve, S. Schnake, C. Hauck

Project Description

This project developed a new sparse grid (SG) method for high-dimensional, high-fidelity kinetic simulations of fusion plasma equations. These methods are available to the broader scientific community through the ORNL-developed scalable open-source framework Adaptive Sparse Grid Discretization (ASGarD).[3] The SG methods create multidimensional approximation schemes induced by single-dimensional hierarchies of basis functions. The main advantage of the hierarchical representation in SG comes from the much better scaling in terms of degrees of freedom per dimension compared with the standard full tensor approach. However, major challenges to the SG approach came from the computational cost of constructing the hierarchical basis and manipulating the resulting matrices with Kronecker block structure. This project developed a fast method for basis construction that exploited the sparsity structure and managed to overcome the previous bottleneck in the process. The project also implemented GPU-accelerated kernels for fast Kronecker multiplication that use the throughput of modern accelerated devices. Combined, these two results allowed this project to reduce the overall computational cost, which enabled the deployment of an interpolation strategy for dealing with nonseparable and nonlinear operators. This project was successful in performing 5D simulations of drift-kinetic equations and describing plasma confined in a magnetic mirror. With those capabilities, ASGarD can now simulate challenging high-dimensional problems of interest in plasma physics.

Mission Relevance

ASGarD is an open-source software package designed to simplify the development of solvers for high-dimensional partial differential equations, with special relevance to plasma physics applications. Its improved computational efficiency enables researchers to tackle complex problems related to the design, optimization, control, and uncertainty quantification of magnetic fusion devices such as tokamaks and stellarators. By facilitating high-fidelity simulations, such as 5D drift-kinetic models, ASGarD allows detailed exploration of plasma behaviors.

Results and Accomplishments

This project developed several numerical techniques that enabled simulations of high-dimensional kinetic equations that represent models of interest to the field of plasma physics. Kinetic equations are the standard method for modeling the physical phenomena associated with nonequilibrium plasma, such as those found in magnetic fusion devices like tokamaks and stellarators. The solution to kinetic equations (i.e., the kinetic phase-space density) is defined over spatial and velocity domains, which collectively have twice the dimensionality of real space (e.g., 6D instead of 3D). Drift and gyrokinetic models simplify these equations by averaging out the gyro motion of particles along magnetic field lines, resulting in a reduced 5D representation. Despite this dimensional reduction, the challenge of high dimensionality persists.

The traditional approach for solving multidimensional equations involves tensor product grids, which are straightforward in terms of implementation and analysis. However, the drawback of this approach is the exponential growth of degrees of freedom with increasing dimensionality. For dimensions of 4D and higher, even storing the solution on leadership-scale computing facilities (e.g., Oak Ridge Leadership Computing Facility's Frontier) becomes infeasible. Consequently, simulating and evolving the field in time are computationally prohibitive. Moreover, mathematical error estimates indicate that the full tensor

[3] *Asgard*, GitHub, https://github.com/project-asgard/asgard

basis space is highly suboptimal, with significant computational work yielding minimal improvements in numerical approximation accuracy.

SGs offer a modern alternative by leveraging the simplicity of tensor product constructions and mitigating the limitations of traditional approximation spaces. The fundamental idea behind SG is to use hierarchical 1D basis functions with varying resolutions and combine tensors of differing resolutions to form an optimal approximation space. However, deploying SG strategies in practice presents several challenges: (1) transforming between hierarchical and regular basis representations is computationally intensive, (2) operator approximation matrices exhibit high fill-in and often deviate from standard conventions supported by existing linear algebra libraries, and (3) inherent difficulties exist in applying the underlying tensor structure to nonseparable and nonlinear operators, which are key areas of interest in plasma physics.

To address the first challenge, this project developed an algorithm for hierarchical basis transformation that used the intrinsic sparsity of the problem. The transformation process was structured as a product of three sparse matrices with a predefined fill pattern. Although performing general sparse matrix–matrix multiplication remains an open problem, this project was able to analyze the specific graph structure of this case and determine the minimal pattern required to efficiently compute the product. This approach bears similarities to fast Fourier transform methods because it employs a hierarchical decomposition of the data that reduces computational complexity to $O(N\log(N))$ compared with brute-force alternatives, which can scale quadratically or even cubically with the problem size N. This algorithm has significantly improved the ability to handle partial differential equations involving complex operator terms.

The challenge of increased fill in matrices associated with SG methods can be addressed by harnessing the computational capabilities provided by modern GPU accelerators. GPUs are optimally suited for high-performance computations because of their massive parallelism and high memory bandwidth. However, the block-sparse structure of these matrices poses unique challenges that are not well-supported by standard linear algebra libraries. These libraries typically rely on explicit indexing of sparse matrix elements, do not account for Kronecker product structures inherent in the matrices, and are often optimized for matrices with significantly fewer nonzero entries per row or column, which lead to inefficiencies when applied to this project's problem.

To overcome these limitations, this project developed specialized GPU kernels tailored to the unique characteristics of SG computations and their associated data structures. Rather than explicitly storing the sparsity pattern of the matrices, this approach avoided storage overhead by implicitly handling indexing within the kernels. This design directly interfaced with SG representation, which streamlined calculations and reduced the memory footprint. Additionally, this strategy fully exploited the Kronecker block structure of the matrices. By leveraging this structure, this project achieved dual benefits: a reduction in overall computational cost and an increase in flop intensity and device throughput. ORNL developed open source library for ASGarD that provides implementation of the accelerated kernels tuned for the current generation accelerators provided by AMD and Nvidia and deployed in DOE's leadership computing facilities.

The models most relevant to plasma physics often involve nonlinear and nonseparable operators, which introduce significant complexity into computational simulations. A common approach to handling nonlinear operations is to decompose them into a sequence of linear steps to simplify their treatment. By leveraging the increased computational capabilities achieved through hierarchical basis transformations and accelerated GPU kernels, this project implemented an interpolation strategy that uses three computationally expensive but separable operations. The combination of these advancements allowed this study to extend the SG methodology beyond the limitations of linear-separable cases. Using the ASGarD library, this project successfully performed fully 5D simulations of a drift-kinetic model.

As a benchmark, this project simulated a plasma field confined within a magnetic mirror system by employing strategies similar to those used in tokamaks and stellarators. This benchmark demonstrated the

applicability of drift-kinetic models to a broad range of plasma regimes and highlighted their importance in advancing high-fidelity simulation capabilities. This work marks a step forward in the development of advanced plasma physics modeling tools to enable more accurate, multidimensional representations of plasma behavior in fusion devices.

CYBER RESILIENCE AND INTELLIGENCE DIVISION

11748: Human–Artificial Intelligence Collaborative Decision-Making Facilitated by Interactive Visual Interfaces

S. Monadjemi, A. Ghosh

Project Description

Experts across domains increasingly leverage large amounts of noisy data to guide researchers toward precise decisions with critical real-world implications. Because the volume and complexity of datasets surpass cognitive abilities to extract knowledge, researchers have designed AI techniques that learn from historical data and make decisions. However, these AI techniques often suffer from undesirable characteristics, including (1) a lack of domain-specific knowledge, (2) a lack of explainability because of their black box design, and (3) inaccessibility to domain experts because of the computational knowledge required to deploy them. These concerns lead to domain experts hesitating to adopt AI techniques in their workflows, particularly in safety-critical settings involving national security and critical infrastructure. To address these problems, this project created a human–AI collaborative paradigm for model steering and decision-making. This work exploited the highly interpretable nature of causal AI techniques and designed interactive visual analytic systems that seamlessly augment human expertise with underlying causal models. Through a set of novel interaction techniques and automatically discovered causal relationships, this project empowers users to ask what-if questions before taking a costly action with potentially adverse outcomes.

Mission Relevance

This work resulted in novel data visualization and analytic capabilities. The results enhance decision-making in domains that address DOE's mission, such as material science (e.g., designing less costly scientific experiments) and critical infrastructure (e.g., maintaining electric grid reliability by analyzing sensor data).

Results and Accomplishments

This project conducted a scoping review of how humans and artificial agents share agency while conducting analytic tasks. These settings, referred to as mixed-initiative visual analytics, involve humans and artificial agents making decisions and initiating actions. The project characterized this human–AI collaboration in light of existing taxonomies and typologies, including Parasuraman's 10 levels of automation[1] and Horvitz's principles of mixed-initiative user interfaces.[2]

Next, the project narrowed the focus to causal modeling techniques for data analysis and decision-making. These techniques are more robust and reliable compared with traditional modeling techniques in that they unravel the underlying causal relationships (as opposed to correlations) by combining observational data and domain expertise. This causal characteristic allows researchers to design interventions that lead to desired outcomes in a complex system. Examples include designing scientific experiments and configuring the electric grid. The causal modeling pipeline is generally divided into two steps: causal discovery and causal inference. This project built capabilities to enhance the causal discovery phase of this workflow. In doing so, the project considered algorithms such as the Peter–Clark algorithm and the Linear Non-Gaussian Acyclic Model (LiNGAM) for automated extraction of causal relationships from observational data. Then, this work built a visual interface to enable seamless user interaction with these algorithms to inject domain expertise and iteratively refine the causal graph.

[1] R. Parasuraman, T. B. Sheridan, and C. D. Wickens, "A Model for Types and Levels of Human Interaction with Automation," *IEEE Transactions on Systems, Man, and Cybernetics A* 30, no. 3 (2000): 286–297, https://doi.org/10.1109/3468.844354.
[2] E. Horvitz, "Principles of Mixed-Initiative User Interfaces," *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems the CHI is the Limit - CHI '99* (ACM Press, 1999), Pittsburgh, Pennsylvania, 1999, 159–166, https://doi.org/10.1145/302979.303030.

Finally, this project validated this approach on recovering known laws of physics from noisy, synthetic data points. The following subsections expand on each of these steps.

Modified LiNGAM Algorithm for Causal Discovery

This project considered the set of openly available causal discovery algorithms, including LiNGAM and the Peter–Clark algorithm. Their outputs were compared, especially in cases in which the ground truth for the data generating process was known. Ultimately, the project team decided to focus on the LiNGAM algorithm because it also quantifies the strength of causal relationships. However, the commercial off-the-shelf version of this algorithm did not consider injected domain knowledge in its output. Therefore, this project modified the algorithm to accurately incorporate domain expertise provided by the users while discovering causal relationships from data.

Visual Analytic Tool to Steer Causal Graphical Models

This project involved designing and implementing a web-based system that provides a user interface for mixed-initiative curation of causal graphical models. This system seamlessly integrates observational data, domain expertise from the user, and algorithmic approaches for discovering and quantifying causal relationships. Because some analytic tasks such as deriving new attributes and data manipulation are more efficiently performed through code, the project built this tool to leverage the Jupyter ecosystem as its backend. This design choice enabled the interplay between code and no-code interaction while curating causal models.

Validation

The project validated this system for causal discovery with two distinct approaches: (1) open-source datasets without ground truth on the causal structure and (2) synthetic datasets generated from known laws of physics with ground truth causal structures. For each case, the datasets were loaded into the system, and the project team examined the causal graphical model produced by an algorithm with no background knowledge. Then, this work was iterated over the graph by making incremental corrections (thus injecting domain expertise) and updating the algorithm's causal model in light of the injected knowledge. This process continued until the researchers arrived at the ground truth (when it was known) or were satisfied with the resulting graph (when ground truth was not known).

11834: Practical Three-Dimensional Curvilinear Synthesis Aperture Radar

M. Story, T. Allen, T. Deller

Project Description

Synthetic aperture radar (SAR) typically forms a 2D image of the ground using radar data collected along a straight-line path. This technology is not impeded by clouds, smoke, or darkness, which makes it ideal for chaotic environments such as disaster recovery from damage to critical energy infrastructure.[3] However, the appearance of SAR imagery is quite different from optical imagery, and it is difficult to interpret. Adding 3D information can aid in interpretation, but the most mature SAR 3D techniques require two or more carefully matched imaging paths.[4] The logistics of this matching process are daunting under normal circumstances and infeasible in chaotic scenarios. Using curvilinear paths, some amount of 3D information can be extracted from a single SAR collection using methods less thoroughly studied. SAR collections using such curvilinear paths are becoming more feasible as SAR hardware is miniaturized to fit on a small, agile uncrewed autonomous systems. Although most curvilinear SAR

[3] S. Plank, "Rapid Damage Assessment by Means of Multi-temporal SAR—A Comprehensive Review and Outlook to Sentinel-1." *Remote Sensing*, no.6 (2014): 4870–4906, https://doi.org/10.3390/rs6064870.
[4] C. V. Jakowatz Jr., et al., *Spotlight-Mode Synthetic Aperture Radar: A Signal Processing Approach* (Springer Science and Business Media, 1996), https://link.springer.com/book/10.1007/978-1-4613-1333-5.

techniques are computationally intensive, a recent paper showed how to compute elevation of a single point in a computationally feasible manner using techniques related to monopulse direction finding.[5] This technique is referred to as Curvilinear Aperture Monopulse (CLAM). However, the initial development opened several questions of immense importance to the practical application of the technique. This project continued recent work to address these deficiencies. The project extended the CLAM technique to realistic platform paths and showed how to mitigate interference between different bright points in the image by identifying which pixels have valid measurements. Finally, the project showed that CLAM works when autofocus is included properly. These developments advanced the technology to the point at which it is ready to demonstrate the techniques using real hardware.

Mission Relevance

This project addressed DOE and ORNL missions in energy security by advancing America's capability to respond to disasters impacting critical energy infrastructure by assessing damage quickly after disaster occurs, even in low-visibility conditions.

Results and Accomplishments

The key accomplishment of this project was a demonstration via simulation that CLAM provides sufficient accuracy on sufficiently many points to aid in operator interpretation of SAR imagery. These points were identified using a screening criterion developed for this project, and their height was measured using computations developed for this project. Previous research showed that interference can corrupt height measurement for nearby image points with similar brightness, and the key accomplishment of the current project can be seen as vindication that some points are sufficiently bright to overcome this effect and can be used as a means of mitigating the interference. The simulation emulated pixel brightness and scattering center height distribution of a real SAR image by using a 3D point cloud derived from a high-resolution interferometric sample pair of real SAR images made publicly available by the National Geospatial-Intelligence Agency.[6] This project established a strong case via simulation that CLAM will work in real data collection.

This work also developed the capability to use CLAM with practical platform paths. Specifically, using new CLAM computations developed for this project, platform paths may have curvature in the downrange dimension, which enables a wider variety of realistic paths. The project further developed the capability to identify and avoid paths that result in ambiguous answers and choose paths with good performance.

Finally, this project showed that autofocus is compatible with CLAM. In typical collections with real hardware, sufficient error occurs in phase measurement to necessitate automatic correction of phase errors. This automatic correction is known as autofocus. This project showed that when the phase is corrected uniformly across the image, the error introduced in CLAM height measurement is manageable.

Together, these accomplishments show through simulation that CLAM is a practical technique for supplementing SAR imagery with sufficient height information to aid in operator interpretation.

[5] M. Story, "Curvilinear Aperture Monopulse," *Proc. 2023 IEEE Radar Conference (RadarConf23)* (IEEE, 2023), San Antonio, Texas, May 1–5, 2023, https://doi.org/10.1109/RadarConf2351548.2023.10149596.
[6] "Sample SICDs," (National Geospatial-Intelligence Agency, 2022), https://github.com/ngageoint/six-library/wiki/Sample-SICDs.

12059: Exploring Electromagnetic Signals Through Sonification

A. Rice, L. Thompson, J. Maddox, T. Potok, A. Albright, T. Karnowski

Project Description

Electronic devices emit electromagnetic signals during operation that reveal a variety of operational characteristics and provide insight into the operational state of electronic systems. Traditional signal analysis methods are based on statistics, specific to the use case, largely rely on visual interpretation, and are not transferable across domains. However, it is often difficult to visually identify weak or sporadic signals that do not greatly differ from the neighboring data. Additionally, visual analysis through tools such as spectrograms is sensitive to parameter values and requires some a priori knowledge of which frequency bands may be of interest so that the parameters can be tuned appropriately. Otherwise, users risk producing spectrograms that are visually tuned in a way that does not display relevant or insightful information. Furthermore, spectrogram analysis requires a more sedentary environment to allow users to work from a single position to interpret the visualization, which limits the operator's ability to move and engage in other tasks. Spectrograms are useful for analyzing the emitted signals' frequency content, but additional information surrounding the underlying processes driving signal emission may be gained through studying the raw signals. In this scenario, unsupervised learning algorithms can be employed to cluster or reduce the dimensionality of signal observations to gain insight into the underlying processes. However, thorough analysis of the resultant clusters to infer the driving forces is still a tedious and visual process with limited capability to be conducted peripherally or in near–real time.

To address the shortcomings and complications associated with visual-only analysis, this project augmented traditional analysis methods with an auditory component through sonification, which is a method of converting data into sound that leverages the human auditory system to gain insight into data that would otherwise be challenging to derive with visual cues alone. Traditional approaches, or parameter-mapped sonification, extract raw signal characteristics and use them to modulate auditory parameters such as pitch, tempo, or volume to convey information about the signal. The output sound is intended to be intuitive and interpretable by a user so that the user is aware of what various fluctuations in the sound correlate to and can make inferences about the signal based on the sound. Sonified representations of complex data have advantages in temporal, spatial, dynamic range, and frequency resolution that may be exploited as strategic enhancements to visualization methods.

However, parameter-mapped sonification has notable issues, such as a lot of trial and error and developer- or user-dependent decisions on which signal parameters to map to which auditory characteristics. This may lead to differing interpretations of the underlying signal. This project aimed to implement and test the effectiveness of using a generative adversarial network (GAN) for sonification applications using two sonification approaches—parameter-mapped and GAN—to transform unintended radiated emissions from common electronic devices to audible sounds. A user study assessed the quality of the sonified representations of the electromagnetic signals, wherein human users were tasked with distinguishing between electronic devices based on listening to pairs of sounds.

Mission Relevance

This project aimed to ensure America's national security by exploring novel ways of analyzing and interpreting electromagnetic signals. System security is multifaceted, and awareness of the electronic environment is a necessity. Sonification allows for mobile assessment, which frees operators from a desk space. Additionally, sonification uses auditory senses, which reducing fatigue from constant visual stimuli and thus reduces the likelihood of missed detections or human oversight of important information. Because of the increased discriminatory capability with auditory interpretation, device characterization and anomaly detection within the electromagnetic spectrum become easier tasks that are able to be performed rapidly. Fluctuations in the spectrum can occur quickly, and it is important to have the capability to interpret these changes as quickly as possible.

Results and Accomplishments

Unintended radiated emissions from the Flaming Moe[7] dataset were used as signals to demonstrate sonification. This dataset contained emissions from 14 electronic devices commonly found in homes and office spaces such as internet routers, computer monitors, and corded telephones. Emissions were captured in a shielded environment while the devices were powered on and off. In the parameter-mapped approach, dominant frequencies extracted with a Fourier transform were mapped to pitch, and amplitudes of the dominant frequencies were mapped to volume. Postprocessing methods were performed to create music-like sounds from these mappings. In the GAN approach, emissions from each device were segmented and embedded into a fixed-length vector. Each device was associated with a minimalist music sample. The embedding vector was used as input to the GAN, along with a clip of the minimalist music sample to condition the output. At the end of the sonification effort, each device had a parameter-mapped and a GAN-produced sonification.

In total, 70 participants across ORNL were recruited to participate in the comparative user study to evaluate the sonifications. The user study comprised two phases, the first of which involved the participants taking a baseline hearing test at ORNL Health Services. In the second phase, participants reported to the Human and Biometrics Laboratory, where they completed the sonification evaluations through a computerized user interface that the team developed. During the evaluations, users were presented with 74 pairs of sounds of four varieties. In some cases, the pairs were different device emissions sonified with either parameter-mapped or GAN sonification approaches. In other cases, the same device emissions were sonified with either parameter-mapped or GAN methods. Users were asked to answer whether the two sounds were generated from the same or different devices and then rank on a scale of 1 (easy) to 5 (hard) the difficulty they experienced in making that determination. Overall, this project noted markedly better performance at determining when the devices were the same or different with the GAN method, along with higher confidence in the user responses with the GAN method.

This project explored two methods of sonification of electromagnetic signals to augment visual analysis with auditory representations. This work found better differentiability between devices and overall better performance with GAN sonifications than those produced with a parameter-mapped strategy.

[7] T. Karnowski, et al., *Flaming Moe* (Oak Ridge Leadership Computing Facility, 2021), https://doi.org/10.13139/ORNLNCCS/1784765.

ENVIRONMENTAL SCIENCES DIVISION

11635: Tracking Disturbance Signals Along River Networks

J. D. Gomez Velez, P. Le

Project Description

Hydrologic disturbance driven by anthropogenic and natural processes has far-reaching and long-lasting impacts on the watershed's physical and biogeochemical functions. Disturbance, either abrupt (e.g., wildfire) or progressive (e.g., land cover/land use change), changes hydrologic systems and can potentially drive them to new equilibrium states. Significant observational and modeling efforts have focused on characterizing the magnitude of disturbance and its implications for the system's evolution (e.g., quantifying recovery rates or transitions). Typically, these efforts implicitly assume that observational networks capture signals from the disturbed parts of the landscape. However, the complex interplay between the river network topology and geometry, the spatiotemporal dynamics of precipitation, and the spatial heterogeneity of disturbance can result in situations in which such signals are only partially or episodically captured or wholly missed, potentially resulting in incorrect conclusions. This project proposed a novel and parsimonious framework to track the propagation of hydrologic disturbance signals along river networks. This framework was used to develop a rigorous approach to interpret observations that capture the effects of hydrological and biogeochemical disturbances. This project validated and tested the framework in synthetic watersheds and in the East Fork Poplar Creek watershed, which is the focus of long-term DOE and ORNL studies.

Mission Relevance

High-resolution integrated watershed modeling and hydrologic disturbance analysis support DOE priorities by strengthening the resilience, reliability, and security of energy infrastructure systems under changing environments and extreme event conditions. Thus, this research advanced DOE's energy and environmental missions.

Results and Accomplishments

This project implemented a parsimonious modeling framework to track the propagation of hydrologic disturbance signals along river networks. This framework's core was the ability to capture the interplay between the spatiotemporal dynamics of precipitation, river network topology and geometry, and spatial heterogeneity of disturbance. This model was tested in synthetic watersheds generated with Optimal Channel Network algorithms as well as in a real-world watershed from the East Fork Poplar Creek. To validate the framework, Advanced Terrestrial Simulator simulations were performed for all watersheds. This work provided a high-fidelity benchmark to assess the accuracy of the system's responses. Key findings from model comparison demonstrated the proposed framework's capability to capture hydrologic variability by accurately simulating changes in streamflow and water depth across multiple scales. This project also successfully evaluated environmental disturbance effects and showed how the model captures the impacts of varying storm intensity and direction on watershed responses. Finally, the project characterized cascading effects, tracking how hydrologic signals move and dissipate as they propagate downstream through the network.

This project finished an initial version of the new transfer-function modeling framework, which included developing generic Python classes equipped with routines for geomorphological analysis and flow estimation. The team also completed consolidating observations for the East Fork Poplar Creek watershed model. This work included datasets for the ATS-PFLOTRAN model and new routines to process atmospheric drivers from the Analysis Of Record for Calibration gridded dataset. The Analysis Of Record for Calibration routines were essential for this analysis because of their higher resolution—for instance, they provided a 1 km spatial resolution with hourly data. In contrast, the previous datasets had a 4 km resolution and were updated daily. The ATS-PFLOTRAN model for flow and transport was implemented

and tested. This study started with an initial East Fork Poplar Creek model that only included flow and conservative solute transport. Based on this model, the project developed a strategy for signal tracking that employed selective injections of multiple conservative tracers.

11680: Impact of Carbonaceous Particles on Mixed-Phase and Ice Cloud Processes

A. Alsante, M. D. Cheng

Project Description

Carbonaceous aerosol, which is produced by combustion from engines, wildfires, and explosives, is ubiquitous in the atmosphere. The ice nucleating properties of carbonaceous aerosol sources are poorly constrained, which leads to large uncertainties in aerosol–cloud interaction models.[1,2] The goal of this project was to develop cutting-edge experimental aerosol technology that could be used to understand the ice nucleating properties of carbonaceous aerosol, which is necessary for regional and global modeling. The data may also be used to better understand the atmospheric fate and transport processes of anthropogenically emitted aerosol. The ability to identify physical and compositional features of single particles and subsequently determine their ice nucleation efficiency under a wide range of atmospheric conditions enables investigation of aerosol composition and ice formation under mixed-phase and cirrus cloud formation conditions in a controlled laboratory setting. This project investigated the ice nucleation efficiency of three distinct carbonaceous aerosol sources: simulated aviation-like soot obtained from ETH Zürich and diesel combustion soot collected at the ORNL National Transportation Research Center (NTRC). The primary dataset of this project was obtained during the field campaign in collaboration with NTRC. Soot aerosol was collected from a diesel engine using exhaust gas recirculation (EGR) and non-EGR combustion strategies. EGR is a technique used to reduce nitrogen oxide emissions that may also alter soot emissions and composition.[3] Particle size distributions were measured with an electrical low-pressure impactor, and data were collected about size-classified aerosol for ice nucleation and chemical characterization. Scanning electron microscopy coupled with energy dispersive x-ray spectroscopy was used for morphological and elemental characterization. Optical photothermal infrared spectroscopy coupled with Raman spectroscopy was used for single-particle chemical characterization to understand the composition that controls ice nucleation of soot emissions.

Mission Relevance

This project advances the DOE mission of ensuring America's energy security and prosperity by addressing its environmental challenges using transformative science and technology solutions to understand anthropogenic cloud formation and properties that affect aerosol–cloud interactions as well as atmospheric fate and transport of aerosol.

Results and Accomplishments

This project developed the ice nucleation cold stage and successfully implemented it to capture the cloud formation conditions of aerosol in a temperature- and humidity-controlled environmental cell. The cold stage was shown to have an uncertainty of ±0.1°C determined by calibration over the experimental temperature range. The development of this system provides the capability to investigate aerosol–cloud interactions.

---

[1] A. N. Alsante and M. D. Cheng, "Aerosol-Cloud Interactions from Aviation Soot Emissions," *Journal of Geophysical Research: Atmospheres* 129, no. 12 (2024): e2023JD040277, https://doi.org/10.1029/2023JD040277.
[2] Z. A. Kanji et al., "Overview of Ice Nucleating Particles," *Meteorological Monographs* 58, no. 1 (2017): 1.1–1.33, https://doi.org/10.1175/AMSMONOGRAPHS-D-16-0006.1.
[3] M. Zheng, G. T. Reader, and J. G. Hawley, "Diesel Engine Exhaust Gas Recirculation—A Review on Advanced and Novel Concepts," *Energy Conversion and Management* 45, no. 6 (2004): 883–900, https://doi-org.ornl.idm.oclc.org/10.1016/S0196-8904(03)00194-8.

Simulated aircraft-like soot emitted from spray combustion had a wide nucleation temperature range spanning more than 20°C from −35.0°C to −10.9°C. Scanning electron microscopy images demonstrated a wide range in particle size, and large agglomerated soot (>1 μm) may have contributed to the variability. Most previous studies used lab-generated soot with different physical and chemical characteristics from realistic soot emissions,[4] which resulted in a wide range of reported ice nucleating properties. Therefore, in this study, field campaigns to collect realistic soot aerosol were shown to be crucial to account for their contribution to atmospheric ice formation.

This project successfully collected representative anthropogenically emitted soot aerosol at the ORNL NTRC user facility. NTRC provided the aerosol source that was needed to deliver on the promised dataset of this project to investigate the ice nucleation conditions of characterized soot. This study determined that the EGR combustion strategy resulted in higher soot emissions (4.4×10^5 cm^{-3}) when compared with the non-EGR strategy (3.0×10^5 cm^{-3}). Compositional differences in emissions were found between EGR and non-EGR techniques that depended on particle size across aerodynamic cut sizes ranging from 15.5 to 94.9 nm. Most notably, the infrared spectrum indicated that soot emitted during EGR combustion conditions with an aerodynamic cut size of 94.9 nm had a distinct composition with the presence of heterogeneous organic compounds in the carbon–hydrogen region (2,800 to 3,000 cm^{-1}). Aerosol collected from these conditions had a significantly warmer nucleation temperature (−26.7°C ± 3.0°C) compared with the remaining soot samples collected using the lowest recorded median freezing temperature at −29.0°C ± 2.4°C. Additionally, the number of ice nucleation active sites was highest for EGR conditions: 1×10^{-5} cm^{-3}, compared with the lowest observed number of active sites for non-EGR conditions at 1×10^{-6} cm^{-3} at a nucleation temperature of −25°C. In comparison, the non-EGR soot at 94.9 nm was more homogeneous with the presence of graphitic carbon. This work resulted in the first known study to provide single-particle physicochemical measurements used to constrain the ice nucleating properties of diesel soot.

[4] U. Trivanovic, G. A. Kelesidis, and S. E. Pratsinis, "High-Throughput Generation of Aircraft-Like Soot," *Aerosol Science and Technology* 56, no. 8 (2022): 732–743, https://doi.org/10.1080/02786826.2022.2070055.

FUSION ENERGY DIVISION

11704: Development of Cryogenic Shell Pellets for Disruption Mitigation

J. Herfindal, D. Shiraki, T. Gebhart, A. Dvorak, T. Ha

Project Description

The goal of this project was to rapidly cool plasmas from the inside out within magnetic fusion devices using cryogenic shell pellets (CSPs). The project also aimed to validate the physical processes for disruption mitigation. An efficient rapid shutdown method is essential for next-generation burning plasma devices, in which unmitigated loss of plasma confinement poses a significant risk of major damage to plasma-facing components. The inside out mitigation technique cools the core of the plasma first, where most of the energy resides, by encapsulating impurities within a minimally radiating shell that burns through as it progresses into the plasma. Simulations of this technique showed a high assimilation fraction and stochastization of the flux surfaces, leading to complete loss of confined electrons (no runaway electron seed), which was not observed for mitigations from the outside in. This project encompassed (1) developing the CSP technology, (2) modifying ONRL's pellet system at the DIII-D National Fusion Facility, and (3) performing plasma experiments using this new mitigation scheme.

Mission Relevance

This project addressed the DOE mission of developing fusion energy resources for energy security and abundance. Significant hurdles still must be overcome before fusion energy can be deployed on the electrical grid. Magnetically confined fusion concepts rely on a clear separation between the plasma and surrounding structure (e.g., vacuum vessel, first wall, cooling lines). However, crossing certain operations limits can trigger a disruption—a total loss of plasma confinement—that can severely damage plasma-facing components and cause lengthy repair downtime. Minimization of this damage using a promising reactor-relevant disruption mitigation technique was the primary goal of this project.

Results and Accomplishments

Leveraging ORNL's expertise in pellet injection technology and existing laboratory infrastructure, this project developed and captured images of a CSP moving at various speeds between 300 and 600 m/s and performed initial experiments of firing the CSP into a plasma at the DIII-D National Fusion Facility. The CSP consisted of an encapsulated core of a high–atomic number radiating impurity (neon) surrounded by a nonradiative deuterium shell. CSP shell–core formation was verified through images that showed a dark striation within the pellet's core in an hourglass shape, which clearly distinguished it from a pure deuterium pellet that showed no dark striation. Core and shell parameter spaces were explored by varying freezing temperatures along the cold zone and adjusting the injected gas quantity, with the goal of minimizing neon exposure at the pellet's edge.

The CSP technique was transferred to the DIII-D pellet system, with hardware modifications and device repositioning in the machine hall successfully completed. An experiment applied CSP injection to tokamak plasmas at DIII-D. Initial tests achieved reliable CSP formation. The subsequent experiment campaign, after additional testing, showed that approximately 90% of pellets fractured before entering the plasma. Some intact CSPs were injected into the plasma, but they used a recipe with higher neon content, which produced a conventional outside-in mitigation because more neon was exposed at the edge of the pellet. The issue is believed to be related to too much propellent delivery during the initial launch of the pellet.

The experimental run day yielded several notable results. The CSPs that were broken or had high neon content were treated as proxies for single pellets because they were not ideal layered pellets used in this work's simulations. The plasma density rise was the same order of magnitude as in prior shattered pellet

injection mitigated plasmas (this study analyzed approximately 80 shots). Moreover, comparable current quench durations were observed when CSP injections with neon inventories of 0.1 $Pa \cdot m^3$ were compared with shattered pellet injection discharges with 10 times the neon inventory. These similarities indicated that comparable mitigation can be achieved with substantially less injected material when delivered as a solid pellet rather than many fragments. This method had a variety of benefits, such as using additional mitigation techniques after the thermal energy was dissipated that require neutral pressures to be at low levels and that can ease the burden on the cryogenic pumps, which may shorten downtime for the device.

MANUFACTURING SCIENCE DIVISION

11407: Highly Wear-Resistant Nanostructured Bainitic Steels via Powder Metallurgy

R. Kannan, Y. Wang, P. Nandwana, T. Grejtak, J. Rakhmonov, B. Lim

Project Description

This project combined computational thermodynamic alloy design with advanced machine learning approaches to develop highly wear-resistant, low-cost nanostructured bainitic steels with properties superior to existing commercial steels. Historically, nanostructured bainitic steels have seen limited industrial use because of poor weldability and machinability stemming from their high carbon content and hardness. Additionally, these steels typically require aging for days to weeks to obtain the desired nanostructured bainite, which further constrains their commercial adoption. Despite these challenges, their attractive wear properties and significantly lower cost (approximately 90 times lower) make them strong candidates to replace current steels used in gearboxes, bearings, and cutting tools. Their superior wear resistance can substantially extend the service life of critical components in energy technologies as well as in automotive applications. This project aimed to develop a new faster, less computationally intensive approach to designing these advanced alloys using CALPHAD, machine learning, and data mining. The approach was validated through powder metallurgy processing used to fabricate nanostructured bainitic steels. This work resulted in a successful new approach and demonstrated an alloy with some superior characteristics to state-of-the-art materials such as AISI 52100 steel.

Mission Relevance

This project pertains to DOE's mission by empowering domestic manufacturing to strengthen US supply chain resilience and energy security.

Results and Accomplishments

This project developed and validated an accelerated alloy design approach that coupled CALPHAD with machine learning and data mining. This alloy design approach can be run on consumer-grade laptops, which enables sampling over a larger composition space. The alloy design approach developed in this project reduced the time for new alloy discovery from nearly a month to just a few hours without the need for high-performance computing.

Using the newly designed methodology, the ORNL-designed alloy exhibited seven times faster onset kinetics for bainite formation and three times faster kinetics for complete bainite transformation. Furthermore, this project demonstrated that the ORNL-developed alloy possesses superior wear resistance compared with state-of-the-art gearbox materials such as AISI 52100 steel.

This project also showed that the designed alloy can be successfully atomized into powders—specifically, prealloyed nanostructured bainite powders can feasibly be atomized. The ability to atomize the alloy confirms that powder metallurgy is a viable route for fabricating these steels and indicates potential for large-scale, near net shape production. This project also achieved sintered components with higher density and hardness than those produced by the conventional route.

This project demonstrated accelerated alloy discovery enabled by coupling computational thermodynamics with machine learning on consumer-grade computers.

11568: Thermochemistry and Kinetics for Electroslag Additive Manufacturing

A. Stevens, V. Singh, R. Kannan, Y. Wang

Project Description

To be cost-effective at the infrastructure scale, an additive manufacturing process must use high-volume, low-cost feedstocks. One such method is electroslag processing combined with feedstocks consisting of direct reduced iron (DRI) and ferroalloys. Electroslag processing methods commonly use a consumable feedstock electrode (e.g., metal strip or ingot) but can be modified to operate with nonconsumable electrodes (e.g., graphite) paired with discontinuous feedstocks. To establish the feasibility of an electroslag additive manufacturing process combining DRI and ferroalloys in a refining slag bath, this project studied the thermochemical and kinetic aspects of the slag–metal–gas interactions using thermomechanical computational and physical simulation under varying chemistry and process conditions.

Mission Relevance

Based on historical trends, large iron and steel castings, which are key needs to support domestic energy generation, are increasingly difficult to source domestically. These components rely on overseas manufacturing and long supply chains (>18 months). Development of a metal additive manufacturing method capable of using low-cost and readily available feedstocks stands to accelerate domestic energy component production and enable reshoring of associated manufacturing infrastructure.

Results and Accomplishments

Electroslag steelmaking was simulated at the laboratory scale by combining 5 g DRI with commercial CaF_2–SiO_2–Al_2O_3 fluxes in an argon atmosphere at approximately 1,600°C. The solidified, phase-separated product was characterized via x-ray computed tomography to quantify the relative fraction of metal and slag phases produced, with approximately 4.5 g of iron refined from the DRI. Further characterization via electron backscatter diffraction of the refined iron demonstrated purity above 99.9%. Miniature cylindrical specimens extracted from the metallic product were tested in compression and were found to perform similarly to pure iron values reported in the literature. Computational thermodynamics simulations were performed for the same slag–metal system, with the predicted iron purity agreeing with the simulation to within 0.1%. These results demonstrated a proof of concept for additive manufacturing of metallic iron from low-cost inputs (<$0.50/kg DRI).

11645: Synthesis of Monophosphides for Extreme Environments

T. Aguirre, B. Lamm, T. Muth

Project Description

Transition metal and rare earth element monophosphides have the potential to expand the operational window of what is considered an extreme environment. These materials do not occur naturally on Earth and are shown to form under difficult conditions—traits that indicate high chemical stability in terrestrial applications. However, these materials are not well-understood outside the semiconductor industry, which conducted most investigations into these materials in the 1960s and 1970s. Extreme-environment materials are relevant to the energy, defense, and nuclear industries. Monophosphide development can unlock pathways for bulk material or coating development for improved material performance in these sectors. Specifically, this project sought to help mature molten salt reactor technology by developing salt-compatible materials for use as molten salt heat exchangers—specifically, printed-circuit heat exchangers. The project focused on the processing science, such as chemical vapor deposition, to develop titanium phosphide and zirconium phosphide coatings as well as characterization of the mechanical, thermophysical, and thermomechanical properties of these materials using x-ray fluorescence and x-ray

diffraction. These materials have a unique combination of properties that result in compounds that are oxidation-resistant, corrosion-resistant, and stable in liquid metals and, possibly, halide salts.

Mission Relevance

The significance of this research is in advancing the field of chemical vapor deposition–coated rare earth element monophosphides. These materials can possibly advance energy technologies by helping enable molten salt reactor technology and high-temperature heat exchange in corrosive environments. These applications support DOE's mission to address America's energy challenges through transformative science and technology solutions.

Results and Accomplishments

Project results included a Taguchi design of experiments to explore the process parameter space for chemical vapor deposition of zirconium phosphide coatings. The process parameters were the experimental temperature and the flow rates of the phosphorous- and zirconium-containing precursors. This project found that a temperature of 1,250°C, a zirconium pentachloride flow rate of 10 standard cubic centimeters per minute, and a phosphorous trichloride flow rate of approximately 12.78 standard cubic centimeters per minute produced samples with a zirconium-to-phosphorus ratio closest to 1, which is ideal for a stoichiometric combination of zirconium and phosphorous to make zirconium phosphide. The zirconium-to-phosphorus ratio was assessed by x-ray fluorescence. Furthermore, x-ray diffraction showed that for all nine experiments within the Taguchi design of experiments, zirconium phosphide was the dominant zirconium- and phosphorus-containing phase that was deposited during chemical vapor deposition.

12045: Printing Trees: Applying Wood Science to Advanced Manufacturing

K. Copenhaver, B. Sanders, Y. Pu, M. Lamm, A. Hubbard, S. Wasti

Project Description

This study combined research into the chemistry, microstructure, and adhesion of lignocellulosic materials with extrusion-based manufacturing technologies to create a process for 3D printing 100% bio-based wood products. The project developed an adhesive-free formulation of wood flour (WF), lignin particles, cellulose nanomaterials, and enzymes that is extrudable and crosslinkable in situ. Relying on the natural molecular structuring of plant cell walls within wood, the process used enzyme treatments to initiate crosslinking of lignin among wood particles, essentially reconstructing wood from its discrete constituents. WF and lignin have a low cost and are typically waste or by-products from lumber processing and papermaking, and WF is currently used with formaldehyde-based adhesives in fiberboard and other construction materials. The technology developed in this project enabled continuous production of petroleum-free wood products via planar extrusion and additive manufacturing (AM), with extensive potential applications in the building and construction industry as continuously extruded fiberboard and directly printed objects or structures with complex geometries. At a large scale, the materials and processes developed herein could function in the same applications as 3D printed concrete with a substantially lower carbon footprint. If scaled, the developed technology could aid in addressing the US housing crisis by enabling production of wood-based construction materials from low-cost, environmentally friendly sources.

Mission Relevance

The development of a wood-based formulation suitable for extrusion-based 3D printing has primary applications in the building and construction industry. The formulation and processes developed in this work could be scaled up to a system resembling concrete AM. AM is generally less wasteful, requires less labor, and is more sustainable than conventional construction techniques. Developing a 3D printable

wood could significantly cut down on the material, time, and cost necessary for home construction. These efficiency improvements could lead to significant reductions in housing costs.

Results and Accomplishments

Formulations for an adhesive-free, extrudable, wood-based feedstock material were first developed comprising WF, lignin, enzyme suspensions, cellulose nanofiber (CNF) suspensions, rheological modifiers, and water. The resulting materials were dough-like and could be extruded to form freestanding structures that were solid upon drying at ambient conditions. One type and size of WF were used, and two types of lignin were investigated: kraft lignin powder and sodium lignosulfonate suspension. This project found that the WF/lignin ratio is crucial in the flow and final performance of the material. Ratios ranging from 70/30 to 90/10 (WF/lignin) were investigated.

Laccase was the primary enzyme considered in this work. This enzyme represents a class of copper-containing oxidases, found in fungi and certain bacteria, that catalyze the oxidative coupling of phenolic substrates, including lignin monomers and fragments, to higher–molecular weight polymers under mild, aqueous conditions. The amount and source of enzyme used to polymerize lignin were varied, as well as the presence of pH modifiers and buffers. This project showed that in isolated studies of lignin/enzyme solutions, enzyme treatments can cause precipitation of solid particles from previously homogeneous solutions, which indicates polymerization of lignin fractions to form insoluble particles. Small-molecule mediators or catalysts are also often used in similar systems with bulky molecules such as lignin and laccase to facilitate in charge transfer. A range of mediators were identified for trial in the proposed systems, but so far, only one radical accelerator has been tried.

CNFs were also added to the mixtures in small quantities. CNFs were obtained from mechanical and/or chemical treatments to biomass, which isolate and purify the nanofibrils. CNFs have extremely high surface area, are covered in hydroxyl groups, and have been used as binders in other lignocellulosic systems. The high hydroxyl content makes CNFs extremely prone to hydrogen bonding with surrounding moisture or other lignocellulosic media. Therefore, CNFs are excellent rheological modifiers in aqueous systems, with a low weight percentage being sufficient to form freestanding gels in water and provide a strong shear thinning effect. Although CNFs have been shown to achieve structural performance when used as a binder in fiberboard products, they are prohibitively expensive for use in commodity products. This project thus employed CNFs in small quantities (<5 wt % in final materials) to modify the flow properties of the wood doughs and augment the adhesive effects of the enzymatic oxidation. To the researchers' knowledge, until this project, the production of wood products using lignin crosslinking after enzymatic oxidation had not yet been used with the addition of CNFs.

Additional modified celluloses were also used as rheological modifiers and water retention aids. Methylcellulose (MC) and hydroxypropyl methyl cellulose (HPMC) were tested, and this project found that adding the modified cellulose was critical to enable any flow of the wood doughs. Without MC or HPMC, the mixtures immediately formed a solid plug in the extrusion nozzles, and further addition of pressure would effectively squeeze out excess water around the plug. Adding very small amounts of MC or HPMC (<1 wt % in final materials) enabled homogeneous flow of the doughs through nozzles of varying sizes, making these modifiers crucial for extrusion. Finally, the water content of all formulations was closely monitored. WF and lignin are extremely hydrophilic and need high concentrations of water to form a dough. The water content was found to be one of the most important factors in processing and crosslinking, with both positive and negative effects. From a structural or mechanical performance and process speed standpoint, low water content is desirable to increase the density of the dried product and reduce drying time. However, a very dilute suspension (i.e., high water content) is desirable for the enzymatic treatment to enable diffusion of the large laccase and lignin molecules and generally enable hydration of the WF and lignin powders to form moldable and extrudable doughs. The most easily processible water contents in the doughs were found to be 65%–75%, and increasing the concentration of MC or HPMC enabled reduction of the overall water content.

In the first round of preparing formulations and trial extrusion with handheld syringes, the dough material was very stable after extrusion, which indicated a high degree of modulus and viscosity recovery upon cessation of shear. However, sagging was noticed on lower layers when test walls or boxes were additively manufactured, which indicated some degree of creep in the material. The extrudate also exhibited little to no die swell, which facilitated near net shape AM. Interestingly, this project found that the rheological behavior of the dough changed from day to day; the doughs felt stiffer or more viscous if they were allowed to dwell longer in a sealed container (to mitigate water evaporation), even without adding enzyme to the dough. This could have been caused by further hydration or water uptake of the WF and lignin powders, which reduced the free water in the system over time. Ultimately, this study found that adding CNF led to a system with high modulus and viscosity with a lower degree of creep over time. The addition of CNF to the system also reduced its yield stress and increased its degree of shear thinning, further facilitating extrusion.

Print trials were conducted initially on enzyme-free systems. Material was prepared with WF, lignin, CNF, HPMC, and water and was printed directly after mixing. The extruded material ultimately flowed out of the nozzle but exhibited a high degree of cracking on its exterior akin to solid breakup. An approximately 3 × 3 in. box was printed that destabilized after approximately 1 in. of print height, with the walls beginning to sway and cause uneven deposition. The same material was prepared again and allowed to dwell for 3 days in a sealed container at room temperature. Upon printing, the material was noticeably smoother and more stable, indicating that the longer dwell time is essential for a more homogeneous extrudate. A final batch of material was prepared without adding CNF and was allowed to dwell for 3 days. Extrusion of the CNF-free material proved unsuccessful; the extrudate exhibited an even higher degree of cracking and wall instability than the CNF-containing material printed immediately after preparation. Overall, the print trials suggested that adding CNF and a dwell/hydration period after mixing are crucial for a high-quality extrudate.

Finally, samples for mechanical testing were prepared using the WF/lignin/CNF/HPMC/water formulations with the addition of enzymes and other synthetic binders. Prepared doughs were molded to plaques and allowed to dry in ambient conditions, after which samples were cut for flexural testing. Adding laccase and a radical accelerator was shown to nearly triple the flexural modulus and increase the flexural strength over 85% compared with an enzyme-free system. Furthermore, the flexural properties of the doughs crosslinked with enzymatic treatments were found to be superior to those prepared with the addition of acrylate or polyurethane binders.

MATERIALS SCIENCE AND TECHNOLOGY DIVISION

11370: Time- and Space-Resolved Deformation of Advanced Nuclear Fuels

E. Cakmak, D. Arregui-Mena, G. Helmreich, A. Kercher

Project Description

This work focused on a proof-of-principal study to demonstrate that the deformation behaviors of advanced nuclear fuel pellets can be resolved as a function of time and space using x-ray computed tomography while loading in situ. Aluminum(III) oxide, or alumina, was used as a fuel surrogate to facilitate this study to reduce environment, safety, and health concerns. The project's goal was to create a framework that can help accelerate qualification efforts for new fuel concepts and lead to faster progression up the technology readiness level ladder.

Although the thermomechanical deformation that composite fuel forms experience under neutron irradiation is complicated, the problem can be compartmentalized into isolated focus areas (e.g., mechanical deformation, neutron irradiation, thermal expansion). In the case of composite structures (e.g., tristructural-isotropic [TRISO] fuel compact), the constituents react differently to the applied loads when interacting with each other. Thus, the strain/stress concentrations must be mapped throughout the structure as a function of applied load to determine the effect of inherent defects, as well as crack initiation and progression to document the failure mechanism. Rapid testing with maximum data output while also using the vast database of physics-based properties that are already well-known for typical nuclear materials such as graphite, silicon carbide, zirconium, and more are key to achieving faster qualification.

Normally, localized fractures are impossible to detect from the outside, and postmortem characterizations offer limited information. In this project, high-resolution x-ray computed tomography with an in situ mechanical loading capability was used to track crack and defect formation, progression and distribution, and eventual failure of the fuel surrogates. The application of deformation representative of a reactor accident scenario in situ with high-resolution x-ray computed tomography can produce correlated datasets that are information-dense and time-conserving. This approach enables determination of the location of preexisting defects in the as-received state, shows their evolution with deformation, and detects the formation of new defects (e.g., cracks) and their interaction within the microstructure. The two sets of qualitative and quantitative data can be matched to accurately describe the effect of load on individual components of the system with the observed deformation and failure mechanisms. When established, such a framework can aid the efforts of accelerated nuclear qualification for faster deployment of new technologies.

In this study, miniature representative fuel surrogate pellets containing 20 TRISO particles were manufactured and tested under compression in situ until complete failure. Deep learning–aided segmentation was used to differentiate the different layers of the TRISO particles and discern defects and cracks within the pellets from zero load to failure. The strain distributions were calculated and visualized in 3D using a novel digital volume correlation approach. A finite element modeling mesh was fitted to the 3D structure to calculate the stresses present. This step was important because traditionally, safety models are based on perfect geometries or structures that are unrealistic. Overall, a framework that relies on information-dense, correlated datasets with added temporal information was established that can be extended to other advanced fuel forms.

Mission Relevance

This project addressed DOE missions related to energy security and abundance. Although nuclear energy remains a viable energy source, newer nuclear technologies are slow to reach the market because of the lengthy certification processes. Nuclear qualification efforts have historically relied on the iterative

generation of large datasets that can be time- and cost-intensive to collect and analyze. This project's goal was to create a framework that can help accelerate qualification efforts for new fuel concepts and lead to faster progression from concept to prototype to production via rapid prototyping and testing on scaled-down versions that are representative of the whole.

Results and Accomplishments

The first objective of this exploratory work was to produce a representative miniature fuel surrogate because the in situ testing required using a specialized miniature test setup. This study successfully manufactured mini-compacts containing 20 TRISO particles compacted in a carbon matrix of a cylindrical shape using alumina surrogates for the fuel. These were coated in series with a porous carbon buffer layer, an inner pyrolytic carbon layer, a silicon carbide layer (acting as the main pressure vessel), and an outer pyrolytic carbon layer before being overcoated in a graphite/binder mixture for final pressing and sintering.

The specimens were tested in situ under compressive deformation (which could occur under swelling conditions). With the help of deep learning–aided segmentation, this study digitally separated all the layers constituting the fuel compacts and detected as-built defects and the evolution of cracks during deformation. In all compacts, the deformation followed the overcoating boundaries of the TRISO particles, and damage was mainly confined to the graphitic matrix. This was a result of the TRISO particles being loosely bound to the matrix (i.e., the cracks did not penetrate into the TRISO particles), keeping them intact until the matrix completely failed.

Following matrix failure, the load was transferred to the TRISO particles. Crack propagation within the layers of the TRISO particles was clearly observed, including deflections and delaminations, due to the deep learning–aided segmentation and high-resolution x-ray computed tomography data.

Digital volume correlation enabled high-resolution visualization of displacement fields within the bulk of the specimens and revealed the strain fields present as a function of applied load. Furthermore, depending on how the TRISO particles were distributed within the graphitic matrix, the strain fields could greatly differ. The combination of such time-resolved qualitative and quantitative deformation datasets offered a unique perspective for the development efforts of advanced nuclear fuels.

Finally, the meshes and strain fields obtained from the tested samples were fed into a finite element modeling framework to map the stresses present in the sample. This information is important particularly for safety models because typical modeling work relies on perfect geometries and mechanical properties. However, this exploratory yet systematic study showed that ideal geometries are not the norm but are instead the exception.

11599: Origin of Potential Strong Diamagnetism in Copper-Doped Lead Oxyapatite

J. Gardner, D. Parker, M. McGuire, G. Eres

Project Description

This work investigated claims of superconductivity at room temperature and ambient pressure in $Pb_9Cu(PO_4)_6O$. Since the claims were made in early 2024, several other groups around the world started producing this compound, as well as variants of it. This project included making high-quality samples in various forms and measuring the bulk transport and magnetic properties with high accuracy and fidelity. The crystal structure was analyzed using x-ray scattering, and the diamagnetic nature of these samples was probed with neutrons and muons. Using a model driven approach, this project made several samples with the published composition and several that were doped in various ways that would increase the metallic nature of the sample. The project also used different routes to produce $Pb_9Cu(PO_4)_6O$. In all

cases, this work found no proof of a superconductor—that is, zero resistance and magnetic expulsion at any temperature above 4 K.

Mission Relevance

This research focused on the synthesis of quantum materials for advancing technological exploitation of quantum phenomena. Specifically, an ambient-pressure, room-temperature superconductor is highly sought after in the world of quantum computing and quantum sensors. Coils of such a material could revolutionize applied magnetism. Even if this is a low-current-carrying compound, it will provide fundamental materials physicists a new pathway to superconductivity. If materials could be found that can carry current over long distances with no losses, the electrical grid would dramatically change. Thus, by pursuing production of such materials, this work supported DOE's energy mission.

Results and Accomplishments

This project reported on the structural, electrical, and magnetic measurements in as-grown polycrystalline samples of $Pb_{10-x}Cu_x(PO_4)_6O$. This project's as-grown specimen had excellent x-ray diffraction matching with the original submission.[1,2] This sample had up to 1.5% of Cu_2S as an impurity phase, but the project team was able to reduce this to 0% with better control of the synthesis. A resistive transition around 380 K, possibly corresponding to structural transitions of Cu_2S, was observed in those samples where the impurity phase was visible by x-ray diffraction. Magnetization measurements showed linear diamagnetic behavior that cannot be associated to the superconducting state. In summary, this work found no evidence for room-temperature, ambient-pressure superconductivity in copper-doped lead apatite $Pb_9Cu(PO_4)_6O$.

First-principles calculations of the electronic structure of the parent lead apatite $Pb_{10}(PO_4)_6O$ compound and related materials were conducted using the all-electron density functional theory code WIEN2K.[3] In general, the calculations confirmed the insulating nature of the parent compound with a band gap on the order of 3 eV.

Because the primary challenge in this series of oxide materials was inducing metallization (a generally necessary precondition for the emergence of the desired superconductivity), additional calculations were performed with one of the lead atoms by either copper (as was claimed to produce superconductivity[4]) or bismuth, which did not appear to have been previously reported. Although the copper-substituted calculations did not show evidence of metallization, which was consistent with experimental reports,[5] the bismuth substitutions showed evidence of metallization based on direct substitution within the unit cell and the use of the virtual crystal approximation.

Experimentally, this project developed an alternative, rapid (<2 h), scalable reaction route toward generating large quantities of the polycrystalline parent phase using a mechanochemical method. This method was significantly easier than that originally reported in Lee et al. (2023).[6] The resulting powder was black-gray and 95% phase-pure lead apatite. A Pawley refinement was created of the powder produced in the ORNL lab. Impurities observed in the powders included other ternary Pb–P–O phases

[1] S. Lee, et al., "Consideration for the Development of Room-Temperature Ambient-Pressure Superconductor (LK-99)," *Journal of the Korean Crystal Growth and Crystal Technology* 33, no. 2 (2023): 61–70, http://doi.org/10.6111/JKCGCT.2023.33.2.061.
[2] S. Lee, J.-H. Kim, and Y.-W. Kwon, "The First Room-Temperature Ambient-Pressure Superconductor," arXiv:2307.12008, https://arxiv.org/abs/2307.12008.
[3] P. Blaha, et al., "WIEN2k: An Augmented Plane Wave + Local Orbitals Program for Calculating Crystal Properties," presented at Technische Universität Wien, Austria, 2001.
[4] S. Lee, J.-H. Kim, and Y.-W. Kwon, "The First Room-Temperature Ambient-Pressure Superconductor," arXiv:2307.12008, https://arxiv.org/abs/2307.12008.
[5] P. Puphal, et al., "Single Crystal Synthesis, Structure, and Magnetism of $Pb_{10-x}Cu_x(PO_4)_6O$," *APL Materials* 11, no. 10 (2023): 101128, https://doi.org/10.1063/5.0172755.
[6] S. Lee, et al., "Consideration for the Development of Room-Temperature Ambient-Pressure Superconductor (LK-99)," *Journal of the Korean Crystal Growth and Crystal Technology* 33, no. 2 (2023): 61–70, http://doi.org/10.6111/JKCGCT.2023.33.2.061.

with slightly different stoichiometries, which can be further suppressed by the addition of small amounts (approximately 5%) of excess P_2O_5 during the reaction.

11719: Computational Framework for Modeling Titanium Diboride Chemical Erosion by Plasma

Y. Osetskiy, M. Eisenbach, G. Samolyuk, L. Nuckols, E. Zarkadoula

Project Description

This project developed a multiscale, machine learning–based computational framework that integrated density functional theory, large-scale atomistic modeling, and rate theory to predict surface erosion, modification, and plasma–material interaction in ultrahigh-temperature ceramics, as well as the behavior of these materials under radiation. Plasma-facing materials are crucial components of any fusion device. Ultrahigh-temperature ceramics have excellent high-temperature properties: strength, thermal and electrical conductivity, tunable neutronics, may be composed of low- to mid-atomic number transition metals, and are refractory. Titanium diboride is good for plasma (low quenching), complies with plasma boronization for coating plasma-facing materials, has fewer transmutation problems (unlike nitrides), is expected to have good chemical erosion, retains tritium, and facilitates structurally stable thermal transport (unlike carbides). However, experimentally generated knowledge on the behavior of ultrahigh-temperature ceramics under the relevant plasma irradiation conditions is very limited. Direct modeling of all plasma–material interaction (PMI) mechanisms using accurate first-principles methods is currently computationally impractical, but classical atomistic modeling requires accurate interatomic force fields (IFFs), which are not yet available for most plasma-contacted materials. This project focused on developing, testing, and validating high-fidelity machine learning IFFs for large-scale atomistic modeling of plasma-facing titanium diboride. The results of the atomistic modeling were used to parameterize the chemical erosion rate model. The developed platform enabled predictive modeling of plasma–material interaction and radiation effects in the application of ultrahigh-temperature ceramics.

Mission Relevance

Fusion power has the potential to provide unprecedented energy abundance and security for the United States. This project addressed multiple areas in the validated design and evaluation of fusion wall components to develop further understanding of material erosion in the context of plasma–material interaction.

Results and Accomplishments

Ab Initio Modeling

A database of titanium diboride properties, including bulk crystalline and liquid phases, point defects, and defect clusters, was created for training interatomic potential models using machine learning. Modeling of hydrogen–titanium and hydrogen–boron reactions on the crystalline titanium diboride surfaces showed that these reactions are unlikely to contribute to chemical erosion because of the high binding energy (>>1 eV) between products and surfaces.

A wide range of interaction energies with deuterium and the component atoms titanium and boron was found on the amorphous titanium diboride surfaces when assuming a possibility of chemical erosion for low–binding energy products. Deuterium transport in bulk titanium diboride was highly anisotropic, predominantly along the boron-occupied (0001) planes. At comparable temperatures, deuterium transport in amorphous titanium diboride was approximately 20 times slower than in the crystalline state. This result suggested a high concentration of deuterium within the amorphous layer.

Machine Learning IFFs

Three different machine learning IFFs were trained using a dataset of approximately 22,000 structures along the trajectories of the ab initio molecular dynamics simulations:

- ChIMES (Chebyshev Interaction Model for Efficient Simulation)—Not GPU accelerated for LAMMPS
- SNAP (Spectral Neighbor Analysis Potential)—GPU accelerated but not flexible enough
- HIP-NN (Hierarchical Interatomic Potential Neural Network)—Most accurate and GPU accelerated

Multiple training loops with a few hundred of new, uncorrelated configurations were used at each cycle to improve accuracy.

Experiment Summary

The goal of this experiment was to characterize titanium diboride's temperature-dependent erosion behavior over a wide range of temperatures to generate curves of erosion yield versus temperature to validate chemical erosion formation energy predictions from the model. Fully dense titanium diboride samples were exposed to low-energy deuterium plasma from the PISCES-RF linear plasma device. Plasma conditions were selected to evaluate temperature dependencies of titanium diboride's erosion behavior. The target temperature varied between 70°C and 850°C. Exposures were performed so that the average energy of bombarding deuterium ions to the target materials were below physical sputtering thresholds for titanium diboride (approximately 19.9 eV). Other plasma parameters, such as electron density and temperature, ion density and temperature, and plasma particle flux and fluence, were held constant throughout all sample exposures. Postmortem characterization following plasma exposures included (1) mass loss measurements to characterize erosion yields of the materials and (2) x-ray photoelectron spectroscopy to characterize depth-resolved chemistry of the materials to capture any preferential erosion processes.

The results confirmed temperature-dependent erosion occured at elevated (>800°C) temperatures. However, direct chemical erosion of deuterated boron hydride and boron dihydrate is an unlikely erosion mechanism at these high temperatures. It significantly exceeds the temperature regime for deuterated boron hydride and boron dihydrate chemical erosion that is reported for boron-doped carbide materials or boron-coated materials.[7,8]

Conclusions

Density functional theory–based analysis of titanium–boron–hydrogen surface reaction rates showed that chemical erosion on the crystalline titanium diboride surfaces is unlikely because of the high binding energies of boron and titanium hydrides to the titanium diboride surfaces. Lattice structural instability and amorphization can occur because of deuterium atoms accumulating beneath the free surface. The low mobility of deuterium in the amorphous phase increased the deuterium concentration. Reactions on the amorphous surface had a wide range of energies (and, therefore, rates). Low-energy reactions were hypothesized to be the mechanism of chemical erosion.

The best accuracy was demonstrated by the HIP-NN force fields. Sputtering yield curves and depth-resolved surface chemistry measurements of titanium diboride targets after exposure to 20 eV deuterium plasma showed enhanced boron depletion at the highest tested temperatures (825°C) over the lower temperatures. This result possibly occurred because of chemical erosion based on the boron–deuterium–oxygen.

[7] A. Annen and W. Jacob, "Chemical Erosion of Amorphous Hydrogenated Boron Films," *Applied Physics Letters* 71 (1997): 1326–1328, https://doi.org/10.1063/1.119885.
[8] D. G. Whyte and R. P. Doerner, "Observations on Chemical Erosion in DIII-D and PISCES," *Physica Scripta* 2001, no. T91 (2001): 7, https://doi.org/10.1238/Physica.Topical.091a00007.

The findings in this project suggested the following possible mechanism of chemical erosion under low-energy deuterium plasma irradiation:

- No direct damage of titanium diboride occured because of the subthreshold ion energy of 20 eV.
- No chemical erosion was possible at the titanium diboride crytalline surface because of the high binding energy of hydrides.
- At high dose, accumulation of deuterium atoms within the narrow surface layer, <2–5 nm, led to structural instability and amorphization of the titanium diboride surface layer. The amorphization rate depended on the exposure time, plasma density, and temperature.
- A wide range of deuterates' binding energies with the amorphous surface assumed the possibility of chemical erosion because of an inate probability of low-energy (high-rate) reactions under the plasma experiment conditions.

12067: Thermomagnetic Processing of CeCo-Based Gap Magnets

M. Kesler, M. Thompson, M. McGuire

Project Description

Relying on critical materials for key technologies poses a national security threat. Of particular concern is the need for rare earths for fabricating magnets. Currently, NdFeB-based magnets are used in multiple applications that do not require the high performance of sintered NdFeB magnets (up to approximately 54 MGOe), so the material is diluted to form a bonded magnet with polymers such as nylon. These bonded magnets generally fill an intermediate performance regime; the substitute magnet material candidates are referred to as gap magnets. As of yet, no practical substitute gap magnet materials can significantly affect criticality. However, CeCuCo-based alloys are a promising rare earth–free magnet material system. This material forms a hard magnet upon casting, which is an important factor for scalability and market impact. However, the lack of crystallographic and magnetic alignment in these cast structures limits performance. These magnets appear to display an intrinsic coercivity that is not linked to the conventional ideas of the emergence of coercivity—magnetic and insulating phases. Thus, improving the coercivity is elusive. This project successfully built off promising preliminary results using thermomagnetic processing during casting and heat treatment of these alloys. Improvements in remanence magnetization and coercivity were realized, and the cast microstructure was fully aligned in the *c*-axis crystallographic direction parallel with the applied field direction. This project demonstrated a processing route for a viable, noncritical gap magnet by improving crystallographic and magnetic alignment as well as coercivity.

Mission Relevance

This research outlined a cost-effective and scalable process to fabricate gap magnets without using rare earth elements. In doing so, this work advanced and addressed DOE missions related to national and energy security as well as producing economic opportunities for the nation. Global security and supply chain flexibility will be advanced by limiting the reliance on critical minerals imported from other countries. This technological breakthrough can give US magnet manufacturers a competitive advantage by giving them a way to produce magnets with materials sourced from US mines to address broad technological and energy needs of the country. This process provides a low-cost path for rare earth–free permanent magnet production.

Results and Accomplishments

Two magnet systems were explored during this project: $Ce(Co_{1-x}Fe_x)_{5.7}Cu_{0.7}Zr_{0.1}$ (x = 0.20, 0.25, or 0.30), referred to as 1:7, and $Ce(Co_{0.65}Cu_{0.22}Fe_{0.12})_{5.2}$, referred to as 1:5. The project initially focused on the 1:7 system because this system has great potential for high remanence magnetization and, thus, a higher energy product. The 1:5 composition resulted in significant crystallographic texture and enhanced magnetic properties.

The 40 g ingots cast in the 9 T applied magnetic field resulted in apparent full *c*-axis crystallographic alignment in the direction of the applied magnetic field, which was elucidated by x-ray diffraction, and yielded significant improvements in magnetic saturation, M_s (6.3 kG), and remanence, M_r (5.3 kG), compared with the 0 T cast specimen (M_s = 5.6 kG and M_r = 3.8 kG). Notably, the magnetization measurements were conducted on monolithic 3 × 3 × 3 mm cubes, whereas conventionally, these tests would be conducted on crushed and magnetically aligned powders to orient the crystallites to maximize the magnetization. The high M_r/M_s ratio of a monolithic specimen for the in-field cast ingot was a significant finding. Generally, the only way to achieve this in a material, if possible, is by complex directional solidification techniques or single crystal growth furnaces.

After heat-treating the 0 T cast specimens at 0 and 9 T, the expected improvements in coercivity were observed, and the expected reduction in M_r was also observed in the 0 T heat treatment (2.6 kG, 68% retained). However, M_r was nearly fully retained in the 9 T heat-treated specimen (3.7 kG, 97% retained). Furthermore, the coercivity improved 5% in the 9 T heat-treated specimen (8.4 kOe) compared with the 0 T heat-treated specimen (8.0 kOe). The 0 and 9 T heat-treated specimens were then crushed and aligned. The magnetization results showed a moderate improvement in M_s (5.2 kG for 0 T and 6.1 kG for 9 T) and a significant improvement in M_r (4.9 kG for 0 T and 5.7 kG for 9 T), as well as a reduction in coercivity, which was expected after crushing. However, the 9 T heat-treated specimen retained more coercivity than the 0 T heat-treated specimen (7.3 vs. 5.5 kOe). Notably, none of the 0 T cast specimens were able to attain the M_s of the 9 T cast specimen even after crushing and aligning. This result implies that the 9 T casting achieved a greater degree of magnetic alignment than what is achievable through conventional casting.

Heat treatment of the 9 T cast ingot at 0 T resulted in a surprising increase in M_r from 5.3 to 5.7 kG, with an accompanied expected increase in coercivity from 3.6 to 5.7 kOe. Generally, M_r is expected to decrease during heat treatment with the trade-off of gains in coercivity. In this case, gains occurred in both. This result might be due to the columnar morphology of the original cast grains in the 9 T specimen, in which during grain growth upon heat treating, grain rotation is constrained by this morphology. Alternatively, in the near-isotropic polycrystalline 0 T casting, grains are free to rotate, thereby increasing isotropy. Notably, similar energy products were attained by the 0 T cast/9 T heat-treated specimen (7.5 MGOe) and as those in the 9 T cast/0 T heat-treated specimen (7.4 MGOe), suggesting that both routes could potentially be viable depending on if an application required higher coercivity (0 T casting).

Finally, the specimen cast and heat-treated at 9 T yielded significant gains in M_s (6.6 kG), M_r (6.3 kG), coercivity (7.0 kOe), and energy product (9.6 MGOe). These results show that when a magnetic field is applied during casting and heat treatment, the previous trade-off of M_r and coercivity gives way to significant simultaneous improvements in both. The remanence magnetization was retained, but it was also improved by 29% (6.3 vs. 4.9 kG) over the crushed and aligned specimen that was conventionally processed at 0 T. Saturation magnetization was also maximized by 27% (6.6 kG after 9 T for cast and heat treatment vs. 5.5 kG after 0 T for conventional process).

NEUTRON TECHNOLOGIES DIVISION

11632: Validated Predictive Capability for Neutron Instrument Background

A. Ibrahim, R. Archibald, K. Grammer, F. Gallmeier, B. Winn

Project Description

Fully exploiting the range of advanced capabilities at ORNL's Spallation Neutron Source (SNS), High Flux Isotope Reactor, and the future SNS Second Target Station requires overcoming the inability to predict and control instrument background. Background arises from spurious scattering within the beamline, leakage through shielding and neighboring beamlines, photon-induced neutron production and emission after radioactive decay, capture and decay photons that trigger detector readouts, and cosmic rays. This project addressed this challenge by integrating two complementary capabilities: (1) very high-fidelity, cradle-to-grave radiation transport (RT) simulations using full-scale models enabled by the hybrid Monte Carlo/deterministic acceleration techniques and (2) advanced machine learning (ML) analysis of the extensive scattering data and facility control metadata recorded at SNS. By restricting the scope to an isolated beamline, the project demonstrated that accurate background prediction is achievable. It established the groundwork for full-scale background simulations of SNS, the High Flux Isotope Reactor, and the Second Target Station, which would enable a true digital twin framework that supports automated instrument operation, real-time decision-making, and transformative advances in new and upgraded instrument design and performance.

Mission Relevance

By integrating very high-fidelity, full-scale, cradle-to-grave RT simulations with cutting-edge ML analysis of operational and scattering data, this project enabled accurate prediction and control of instrument background. Neutron instrument background is difficult to predict because it cannot be simulated using simple ray-tracing techniques. Thus, this project provided the missing component required to realize true digital twins of neutron instruments. These digital twins will enable automated instrument control, faster calibration, improved experiment planning, optimized user time, and a higher scientific return from DOE neutron user facilities. For future facilities, this capability enables predictive evaluation of shielding, bunker design, instrument placement, and beamline configuration. The integration of high-fidelity modeling with ML therefore advances multiple elements of the DOE mission: improving the performance of national scientific infrastructure; enabling next-generation, AI-driven experimental science; and strengthening America's leadership in neutron-based research and technology.

Results and Accomplishments

For the first time, this project demonstrated that neutron instrument background can be accurately predicted by combining (1) cradle-to-grave RT simulations using very high-fidelity, full-scale models and (2) ML analysis of instrument signals and facility operation data. Although the methods developed in this project can apply to any neutron instrument background, the SNS Hybrid Spectrometer HYSPEC's prompt pulse background feature was selected for the proof-of-principle demonstration in this project. This background is quantified as the integral of the detectors' count rate over the first 3 ms following the proton-on-target time stamp.

Machine Learning

The ML analysis used a comprehensive, time-resolved dataset representing the HYSPEC scattering data and the SNS facility operational metadata. The scattering data captured the prompt pulse recorded by the HYSPEC detector tubes, and the operational metadata recorded the instantaneous facility power, proton beam energy, SNS shutter, location of the HYSPEC detector vessel within its room, and other data. More than 450,000 HYSPEC runs from the years 2013 to 2025 were examined. Preprocessing of these data

filtered all runs in which the prompt pulse was obscured by elastic neutron scattering peaks, inelastic neutron scattering peaks, or gamma-ray peaks, along with runs affected by erroneous detector location metadata, insufficient accumulated proton charge, abnormally short acquisition duration, incomplete detector frames, or effectively zero facility beam power. Together, these filtration steps rejected approximately 73% of runs, yielding a final dataset of roughly 117,000 high-quality prompt pulses suitable for ML analysis.

To isolate shutter contributions across different operational periods, nonnegative linear factorization was applied. During 2013–2016, when the average prompt pulse was approximately 43.0 counts per microampere-hour, nonnegative linear factorization showed that the opening of shutter 13 contributed to approximately 25% of the shutter-attributable background. Beginning in late 2021, the HYSPEC team installed boron-rich Mirrorbor panels and cadmium shielding around the detector vessel, especially on the side facing shutter 13. Following these installations, the average prompt pulse between 2022 and 2025 decreased to approximately 28.8 counts per microampere-hour, and shutter 13's contribution dropped to approximately 3.0 counts per microampere-hour. In this most recent period, shutters 11, 13, and 10 were the dominant contributors, accounting for approximately 27.5%, 10.4%, and 6.7%, respectively, of the nonnegative linear factorization–decomposed prompt pulse.

To highlight the effects of HYSPEC's physical location within its instrument room, the detector vessel location was discretized into eight configuration indices. After 2022, configurations farther from shutters 11, 13, and 10 exhibited substantially lower prompt pulse levels—approximately 38.9% lower than other configurations. Supervised ML was used to predict the prompt pulse background using an L_1 sparse autoencoder, which was chosen for its ability to capture nonlinear dependencies and enforce sparsity in the latent representation. The model inputs included the HYSPEC detector vessel's location within the room, a vector representing the states of the 18 SNS primary shutters, the integral of the first 100 μs of the HYSPEC prompt pulse, the SNS beam power, the proton energy, and a step function equal to 0 before March 1, 2022, and 1 thereafter to distinguish operational regimes before and after shielding installation. The dataset was divided into 70% for training, 15% for validation, and 15% for testing. The autoencoder employed an overcomplete latent space eight times the size of the input dimension with L_1 regularization to suppress nonessential features and noise. Training was performed in PyTorch using the Rectified Adam optimizer with a learning rate of 10^{-4}, batch size of 32, maximum of 3,000 epochs, and early stopping criterion with a patience parameter of 50 epochs to prevent overfitting and reduce computational time. Over all variations of the inputs, the model was used to predict the prompt pulse integrals over 0–1, 1–2, and 2–3 ms after the proton pulse, yielding root mean square error values of approximately 3.8%, 9.3%, and 18.9%, respectively. This result illustrated the ability of the supervised ML model to reduce the time-of-flight region obscured by the prompt pulse. Although the absolute error decreased for later time intervals, the relative root mean square error increased because those later portions contained a larger fraction of time-dependent background components not considered in this analysis.

Unsupervised ML using clustering algorithms based on density-based spatial clustering of applications with noise with cosine and Euclidean similarity measures revealed eight distinct clusters of prompt pulse shapes, each characterized primarily by its decay behavior. These clusters appeared consistently in every operational period but differed in prevalence, with clear transitions between 2013 and 2016, 2017 and 2021, and 2022 and 2025, which reflected shielding installations and broader facility changes. This project devised a reconstruction algorithm of the full 3 ms prompt pulse signal using the least squares fit of a linear combination of the time-dependent functions representing the eight clusters during the first 100 μs. The reconstruction prompt pulses achieved errors 20 to 40 times smaller than those obtained through supervised prediction alone. This reconstruction capability reduced the time-of-flight region obscured by the prompt pulse by approximately 96.7% because it enabled subtraction of the cluster-based approximation from the measured scattering data, which revealed any inelastic scattering features obscured by the 2.9 ms prompt pulse tail.

Monte Carlo Simulation

A full-scale Monte Carlo N-Particle (MCNP) model was developed to simulate HYSPEC's prompt pulse, beginning from protons on the target and going all the way to its detectors. The geometry for this simulation integrated two detailed MCNP models: the HYSPEC instrument model, containing the curved supermirror guides, chopper housing, drum-shield structure, detector vessel, and internal shielding, and the SNS core-vessel and monolith model, which includes moderators, shutter assemblies, bulk shielding, target cell walls, and the surrounding concrete and steel structures. These geometries were integrated using the GeomWriter tool developed at ORNL. Because the prompt pulse at HYSPEC is produced by an extremely small fraction of the high-energy particles generated at the target and transported through tens of meters of shielding, the calculation required advanced variance reduction techniques. The simulation employed the Consistent Adjoint Driven Importance Sampling hybrid put into the ORNL ADVANTG (Automated Variance Reduction Generator) code. The resulting prompt pulse prediction of 49.14 ± 2.7 counts per microampere-hour fell within the distribution derived from historical prompt pulse data and agreed with the ML predictions for pulses recorded under comparable operational states.

Conclusion

This project fully achieved and exceeded its goals by demonstrating the ability to accurately predict neutron instrument background. The project combined cutting-edge ML analysis with high-fidelity, full-scale RT simulations. ML validated the RT computational methods and data and guided the development of the simulation models. Conversely, full-scale RT simulations supplied synthetic data that were needed to complete and optimize ML models, which extends predictive capability to configurations and materials not present in existing facilities as well as to future facilities. Scaling this approach to full-facility simulation on high-performance computing systems will lay the foundation for a practical digital twin framework that can optimize and automate instrument operation, improve measurement efficiency, and guide the design of future neutron facilities.

11758: Workflow from Atomistic Simulation to Monte Carlo Ray Tracing: An Innovative Approach to Enhance Diffuse Neutron Scattering Interpretations

F. Islam, M. McDonnell, S. Seal, T. Huegle

Project Description

This project outlined a novel workflow designed to accurately model diffuse elastic scattering data to ensure comparability with experimental results. Central to this approach was the calculation of the diffuse scattering kernels using atomistic modeling. This kernel was then integrated into a Nanoscale Ordered Materials Diffractometer (NOMAD) instrument ray tracing simulation, aiming to produce data that closely mirror experimental findings. As a proof of concept, this project focused on a liquid argon system. Simulated results were compared with existing experimental data to validate this model's accuracy and reliability. By achieving a high degree of comparability with experimental observations, this research provided a new method to process and analyze neutron diffuse scattering. This project enhanced the accuracy for interpreting neutron diffuse scattering data by creating an end-to-end workflow of a fully virtual diffuse scattering instrument simulation using atomistic models as input.

Mission Relevance

This project provided a new method to process and analyze neutron diffuse scattering as well as a comprehensive simulation analysis workflow tool. This tool enables more precise material characterization to support DOE's mission in addressing America's energy and nuclear challenges.

Results and Accomplishments

This project developed and deployed an end-to-end virtual neutron experiment workflow that integrates instrument simulation, sample modeling, and data analysis. This workflow enables users to design, execute, and analyze virtual neutron scattering experiments in a manner that closely mirrors real beamline operations to support informed experimental decision-making before beam time.

NUCLEAR ENERGY AND FUEL CYCLE DIVISION

11773: Advancing Noninvasive Ultrasonic Flow Measurement Capabilities

D. Orea, T. D. Nguyen, H. Sun

Project Description

For acquiring experimental measurements characterizing fluid flow in detail, invasive and noninvasive measurement techniques can be applied. Because invasive methods alter the flow, noninvasive methods are usually preferred. This project sought to demonstrate the application of a noninvasive diagnostics technique for single and multiphase flow featuring ultrasound imaging and velocimetry. The goal was to establish the capability to detect and accurately measure particulates and gas bubbles in liquids. The 2D ultrasound imaging system acquired flow measurements in a small water loop test facility designed for developing multiphysics codes.

Mission Relevance

This project developed the capability to detect and track solid particles and gas bubbles in a working liquid using ultrasonic imaging velocimetry. This project addressed DOE missions by providing a method to perform advanced measurements to further inform complex fluid flow dynamics pertaining to nuclear power and advanced manufacturing.

Results and Accomplishments

The first task of this project was to establish a small water loop to be used to test the ultrasonic sensor against conventional optical techniques. The loop consisted of a circulation diaphragm pump and high-accuracy flow meter capable of flowing up to 10 L/min. The main horizontal and vertical test sections of the loop were constructed from an optically transparent square acrylic tube, which allowed researchers to visualize the internal fluid flow.

A Vantage 32 LE system from Verasonics was used to control the ultrasonic transducers for data acquisition and image reconstruction. The system had 64 independently controlled transmit channels and 32 receiving channels. Two linear phased array transducers (64 and 128 elements for increased field of view) with a center frequency of 5 MHz were connected to the Vantage system using a Universal Transducer Adapter. The system's MATLAB program interface provided user control of the transducer and image reconstruction. The transducer could operate in a variety of modes, such as a single element activation, element scan/sweep, or full matrix acquisition.

The ultrasound images were compared with the optical technique commonly referred to as shadowgraph. The transducer and digital camera were placed on the vertical leg of the flow loop. Water filled the loop past the area of interest before helium gas bubbles were injected. Ultrasonic and optical images of the bubbles were captured at a constant gas bubble injection rate but varying water flow rate (i.e., 0, 2, and 7 L/min). Both imaging techniques demonstrated the ability to detect and track bubbles and particulates in the flow.

Open-source tools such as Fiji for image processing and Python for data processing were employed for comparing results from both techniques. The frame acquisition for optical and ultrasonic imaging was 30 and 67 Hz, respectively. Conversions of pixels to physical units were determined for both imaging techniques using wire or tubing of known dimensions. The bubble size distribution and bubble velocities resulting from optical and ultrasonic data were compared. The distribution functions showed that bubbles of larger diameters (>2.5 mm) increased the noise of the ultrasound signal, which resulted in greater uncertainty. Conversely, bubbles of smaller diameters (<2 mm) and particulates as small as 10 μm were clearly observed and tracked using ultrasonic imaging. The helium bubble velocity distribution showed good agreement with that of the shadowgraph results for all flow conditions.

Because this study observed that micrometer-sized particles could be detected and tracked with a large signal-to-noise ratio, particle tracking velocimetry and particle image velocimetry were applied to determine the liquid phase flow velocity. This was performed by seeding the liquid with hollow glass spheres approximately 50 μm in diameter, which acted as tracers in the fluid that followed the fluid streamlines. Using the ultrasound with an image acquisition rate of 100 Hz, the particle displacements were determined via image cross-correlations, resulting in a 2D fluid flow velocity field.

To determine the limitations of the ultrasound system on a high–acoustic impedance wall material and fluid, a benchtop test stand made from quartz glass was designed to hold approximately 500 g (77 mL) of liquid metal eutectic composed of gallium (68.5 wt %), indium (21.5 wt %), and tin (10 wt %). A series of initial tests using a single element were conducted to determine the penetration depth for multiple frequencies (0.5 to 5 MHz) of the ultrasound waves before moving on the full matrix mode. This study concluded that the allowable frequencies from the linear phased array did not allow the gas bubble to be resolved because of the large differences in acoustic impedance of the helium and surrounding liquid metal. However, by applying the Doppler acquisition mode, the frequency shift caused by the moving liquid allowed qualitative measurements of the bubble activities.

NUCLEAR NONPROLIFERATION DIVISION

11633: Three-Dimensional Flow Measurements with Magnetic Resonance Imaging

M. Benson, T. Nguyen

Project Description

Dispersion of contaminant materials advecting through air, water, or other media—both passive and reacting—involves detailed interactions between complex wind (water) fields, topography, and thermal effects within a variety of atmospheric boundary layer regions. Complexity increases when this is coupled with uncertainties in weather, including any realistic and spatially varying topographical effects. Improving the understanding of the interaction between transport and dispersion united with weather and topography is a current research topic.

Experimental diagnostics in scenarios involving contaminant plumes typically require facilities that are expensive, large, and have multiple personnel to safely operate and conduct high-quality experiments. However, cost and time requirements can be reduced by using magnetic resonance imaging (MRI) systems. When used successfully, an MRI diagnostic for a 3D velocity field, a technique referred to as magnetic resonance velocimetry, usually yields a dataset of 6–10 million measurements at submillimeter isotropic resolution in a few hours of scan time. Similarly, the concentration field measurements, known as magnetic resonance concentration, can provide the species concentration or mixture fraction at similar resolution as for magnetic resonance velocimetry.

This project centered on measurements of the complex turbulent fluid flow fields within an urban-inspired scenario, including contaminant plume tracking, and used MRI systems at Vanderbilt University in Nashville, Tennessee, and the Mayo Clinic in Rochester, Minnesota. The project consisted of equipping and conducting a matched series of experiments and computational studies measuring the 3D, three-component, time-averaged velocity and concentration field data in a test section. The experimental data were compared with matched computational fluid dynamics simulations as well as studies using traditional 2D laser diagnostics for velocity and concentration field comparisons.

Mission Relevance

Whether for vulnerability or hazard assessments within or outside the United States, decision-makers require improved, fast-running modeling capabilities that are accurate and validated. Reliable, robust 3D experimental data are important to that improvement. The demonstration of a well-qualified 3D experimental capability within a challenging geometry without optical accessibility supports multiple aspects of DOE's security mission.

Results and Accomplishments

Four different experiments were conducted on a model geometry and compared for velocity and concentration field data. Additionally, two simulations—Reynolds-averaged Navier–Stokes and Large Eddy Simulations using commercial software—were conducted for comparisons. The experiments comprised two MRI-based measurements: one each for velocity (magnetic resonance velocimetry) and concentration (magnetic resonance concentration). Additionally, two laser diagnostic experiments were conducted, including particle image velocimetry and planar laser-induced fluorescence to capture 2D velocity and concentration fields, respectively.

PHYSICS DIVISION

11728: Benchmarking Computational Primitives for HPC-QC Systems

A. Delgado, P. Date

Project Description

This project introduced innovative research aimed at advancing the integration of high-performance computing (HPC) and quantum computing (QC) systems, particularly in the realms of quantum machine learning and Hamiltonian simulations (HSs). Although several projects have focused on integrating HPC and QC systems, limited efforts have identified and tailored suitable applications for such integrations. This project addressed this lack of research through three primary objectives. The first objective was to identify computational primitives that are optimally suited for HPC-QC settings, with a special emphasis on those essential for quantum machine learning and HSs. The outcome was a comprehensive catalog of these primitives that assesses their potential to enhance efficiency in HPC-QC systems. In the second objective, the project undertook benchmarking these identified primitives on HPC and QC platforms. This analysis revealed each platform's strengths and limitations to guide the effective optimization of these primitives for HPC-QC environments. The third objective was the application of these primitives to practical problems in quantum machine learning and HSs to demonstrate their practical utility in real-world scenarios. This included implementing the primitives in complex scientific and engineering problems to showcase their ability to improve problem-solving efficiency and accuracy.

Mission Relevance

This project advances the field of QC, which has the potential to enable breakthroughs in various scientific domains in which quantum-enhanced computing could be used to solve problems that are currently intractable for classical computing systems alone. These problems include energy, environmental, and nuclear challenges that are part of DOE's mission.

Results and Accomplishments

As computational demands in scientific applications continue to rise, hybrid HPC systems integrating classical and QC are emerging as a promising approach to tackling complex computational challenges. One area of application is HS, which is a fundamental task in quantum physics and other large-scale scientific domains. This project investigated strategies for quantum–classical integration to enhance HS within hybrid supercomputing environments. By analyzing computational primitives in HPC allocations dedicated to these tasks, this project identified components in HS workflows that stand to benefit from quantum acceleration. To this end, the project systematically broke down the HS process into discrete computational phases, which highlighted specific primitives that could be effectively offloaded to quantum processors for improved efficiency. These empirical findings provided insights into system integration, potential offloading techniques, and the challenges of achieving seamless quantum–classical interoperability. This project assessed the feasibility of quantum-ready primitives within HPC workflows and discussed key barriers such as synchronization, data transfer latency, and algorithmic adaptability. These results contributed to the ongoing development of optimized hybrid solutions to advance the role of quantum-enhanced computing in scientific research.

11786: Particulate Effluent Characterization

A. Sandifer, E. Kabela, W. McCarter, M. Willis

Project Description

This project demonstrated the viability of ORNL as a test bed for monitoring studies of particulate emissions from nuclear facilities. A robust understanding of these emissions is desirable to help differentiate between normal nuclear operations and potential nuclear accidents or nuclear weapons tests.

ORNL is unique because it has several types of nuclear facilities on-site: an operating production reactor, radiochemical separation facilities, and a spallation neutron source. Many of these types of facilities have not been studied by the monitoring community and have only been looked at from a broad regulatory point of view, so their detailed emissions are not well-documented. Distances between samplers and facilities are also comparable with some samplers within the International Monitoring System network and civilian facilities. The stacks at ORNL's High Flux Isotope Reactor, Radiochemical Engineering Development Center, Spallation Neutron Source (SNS), and several of the smaller radioisotope facilities on-site have monitors installed for regulatory reasons. To provide ground truth for detections, this project deployed a high-volume particulate air filter sampler to collect airborne particulate effluent from the nuclear facilities on-site at ORNL. Collections occurred regularly from November 1, 2024, through June 30, 2025, and were analyzed via gamma spectroscopy in the laboratory. The radioisotope iodine-123 was detected in several samples throughout the collection period. Detailed atmospheric transport modeling was performed on all detections for multiple sources throughout the Knoxville, Tennessee, area for source attribution, and the most likely source of the iodine-123 was determined to be SNS.

Mission Relevance

Effluent from nuclear facilities can affect measurement results for nonproliferation and nuclear test treaty monitoring missions. Although campaigns are underway to understand radioxenon backgrounds, no current understanding of anomalous particulate detections at radionuclide monitoring stations has been elucidated. Understanding particulate emissions from nuclear facilities could help differentiate between normal nuclear operations and potential nuclear accidents or nuclear weapons tests, thus advancing DOE's mission and ensuring America's security by addressing nuclear challenges.

Results and Accomplishments

Most samples contained only naturally occurring radionuclides such as beryllium-7, lead-212, and radon daughters. Initially, longer-lived isotopes with small concentrations were expected, so samples were allowed to cool for 24 hours before counting to let short-lived radon daughters decay and, thus, improve the background and minimum detectable concentration in the sample. However, iodine-123 was detected in March 2025 and then regularly on a roughly monthly cycle thereafter. Half-life analysis was performed on several samples to confirm the isotope's identity. Once this shorter-lived isotope was identified, the counting strategy was altered to count samples immediately after collection for 24 hours and then recount for another 24 hours. This allowed confirmation and detection of the short-lived isotope.

Detailed atmospheric transport modeling was performed on all detections for source attribution. To support source attribution of airborne radionuclide detections, two complementary modeling approaches were employed: back trajectory analysis and release time analysis. These methods leveraged meteorological model data and atmospheric transport modeling to infer the most probable origin of detected isotopes observed during this project. SNS is a known emitter of xenon-123, which is the parent isotope for iodine-123, so SNS was considered to be the primary source. However, consideration needed to be given to other potential emitters, such as the Watts Bar Nuclear Plant. Three on-site locations at ORNL were considered as potential source locations because of nuclear and radiochemistry processes that occur at these locations. Finally, because iodine-123 is a medical isotope, oncology offices and facilities in the broader Knoxville, Tennessee, area were also included as potential sources.

Release time analysis simulates hourly unit releases from candidate source locations and tracks their transport to the receptor site. This method enables identification of sources that directly contribute to a collected sample, contribute via plume wraparound, or do not contribute at all. Release time analysis was applied to all elevated iodine-123 samples collected during the study period. The analysis revealed that for the sample with a collection stop date of March 10, 2025, SNS was the most likely source of the detected iodine-123.

Further release time analysis and nonparametric wind regression analyses reinforced the conclusion that SNS was the dominant source of iodine-123 for most detections. SNS is known to regularly report xenon-123 stack emissions, which supports this attribution. A comparison of SNS xenon-123 emissions with iodine-123 detections at the sampler showed strong temporal alignment, further strengthening the case for SNS as the primary source. However, at least one sample collected in May 2025 showed transport patterns inconsistent with SNS emissions, suggesting a possible alternate source—potentially from the central ORNL campus.

The combined results from observed meteorology and transport modeling provided a consistent and compelling narrative regarding the likely sources of the iodine-123 detections observed during the project. Analyses using wind direction and speed data from ORNL's meteorological towers indicated that elevated activity levels were most frequently associated with winds originating from the west-northwest and northwest, which are directions that align with the location of SNS. Nonparametric wind regression and weighted polar frequency plots further supported this directional bias, suggesting that SNS was a persistent and probable contributor to the observed radionuclide activity.

An comprehensive dataset of nuclear facility stack emissions, long-range gamma collections, and atmospheric transport and meteorological conditions was generated for the 8-month collection period.

RADIOISOTOPE SCIENCE AND TECHNOLOGY DIVISION

11616: Theoretical Calculations for Emerging Targeted Alpha Therapy Actinide Radionuclide Couples

C. Celis-Barros, M. Martelles

Project Description

Targeted alpha therapy (TAT) and radiotheranostics offer powerful alternatives to treat late-stage cancers for which other therapies are limited. Delivering alpha particle radiation directly to cancer cells generally requires a radiometal, a targeting vector, and a metal chelator. The actinium-225/bismuth-213 couple has been shown to be a promising candidate for TAT; studies are at an advanced stage. However, their progress has been hindered by limited availability for widespread application. These limitations can be overcome with the new and emerging uranium-230/thorium-226 radionuclide couple, which is being considered as a novel alternative for TAT. Chelators are key to safe implementation of TAT. They must show fast metal complexation kinetics, high thermodynamic stability, and high in vivo stability. The lack of knowledge regarding the coordination chemistry of actinium and, to a lesser extent, thorium makes it difficult to accurately foresee which ligands will form stable complexes in vitro and in vivo without running experiments in the laboratory. Another challenge is that because of the larger ionic radius of these ions compared with other radionuclides, kinetics and thermodynamics are decreased when using current chelators for TAT. To overcome these challenges, this project took a theoretical and computational approach to understand the nature of metal–ligand interactions with chelating agents that have been explored for use in actinium-225/bismuth-213 TAT. The information collected in this work provided a starting point for researchers to gain an understanding to develop more suitable chelators for new and emerging uranium-230/thorium-226 TAT.

Mission Relevance

This project advanced DOE's mission by addressing nuclear challenges through transformative computational science. TAT relies on radionuclides such as actinium-225 and bismuth-213, whose production, handling, and chemical behavior are central to the nation's nuclear science and isotope programs. By developing a predictive, theory-driven understanding of radiometal–chelator interactions, this work aimed to reduce reliance on extensive experimental trial and error with scarce and highly radioactive isotopes, thereby improving safety, efficiency, and stewardship of nuclear materials.

The project also supported DOE objectives in nuclear security by strengthening the fundamental chemical knowledge required for the safe use, containment, and control of heavy radionuclides. The results demonstrated that thermodynamic stability and bonding behavior cannot be inferred from surrogate metals or simplified covalency metrics, which provided essential guidance for rational design of robust coordination environments for actinides. This insight is broadly relevant to isotope production, nuclear medicine, and radiochemical processing workflows supported by DOE.

Results and Accomplishments

This project established a comprehensive theoretical framework for evaluating the thermodynamic stability and bonding characteristics of emerging radiometal–chelator systems relevant to TAT. Using relativistic density functional theory, 20 metal–ligand complexes involving actinium(III), bismuth(III), lanthanum(III), and thorium(IV) coordinated by five clinically relevant chelators were systematically investigated in aqueous solution:

- 1,4,7,10-tetraazacyclododecane-1,4,7,10-tetraacetic acid (DOTA)
- 1,4,7-tris(carboxymethyl)-10-[2-(bis(carboxymethyl)amino)ethyl]-1,4,7,10-tetraazacyclododecane (DEPA)
- Diethylenetriaminepentaacetic acid (DTPA)

- Pyridine-based tetrapicolinic acid (Py4PA)
- Macropa

Complexation free energies, energy decomposition analysis, and topological analysis of the electron density (quantum theory of atoms in molecules) were jointly employed to assess trends in stability and bonding across metals and ligands. This unified computational approach enabled direct comparison of macrocyclic and acyclic chelators for radionuclides that are experimentally challenging to study due to radioactivity and limited availability.

A central outcome of this work was demonstrating that thermodynamic stability cannot be predicted from covalent bonding metrics alone. Although energy decomposition analysis revealed that orbital (covalent) interactions were primarily metal-dependent and that macropa exhibited the strongest orbital stabilization across all metals, these interactions did not correlate with complexation free energies. Instead, electrostatic contributions, which are strongly influenced by ligand charge, were found to dominate overall thermodynamic stability. As a result, macropa consistently emerged as the least thermodynamically stable chelator for all four metals studied in this work despite its pronounced covalent character. This finding provided a mechanistic explanation for why covalency-based descriptors alone are insufficient for ranking chelators in TAT applications.

The project further showed that favorable chelator performance is highly metal-specific and cannot be inferred from surrogate elements. Although lanthanum(III) is often used experimentally as a nonradioactive analogue for actinium(III), the calculated stability trends revealed substantial differences between these metals. For example, DEPA and Py4PA were identified as the most thermodynamically stabilizing ligands for actinium(III), lanthanum(III), and thorium(IV), whereas DOTA was the most favorable chelator for bismuth(III). Notably, metals with nearly identical ionic radii, such as bismuth(III) and lanthanum(III), displayed markedly different ligand preferences, which underscored that ionic size alone does not dictate chelation behavior. These results demonstrated that extrapolating chelator suitability from chemically similar metals can lead to incorrect conclusions for radionuclide design.

Finally, quantum theory of atoms in molecules analysis confirmed that increased electron density at metal–ligand bond critical points does not translate into enhanced thermodynamic stability. Although some ligands exhibited higher electron density and more covalent bonding character, these did not correspond to more favorable complexation free energies. Together with the energy decomposition analysis results, this result established a key conclusion of the project: thermodynamically stable radiometal–chelator systems for TAT must be evaluated using a balance of electrostatic and covalent interactions rather than covalency alone.

SUMMARIES OF PROJECTS SUPPORTED BY THE STRATEGIC HIRE PROGRAM

FUSION ENERGY DIVISION

11713: Theory of Tokamak Pedestal Stability

H. Wilson, M. Anastopoulos, M. Cianciosa, O. Lopez Ortiz, P. Snyder, G. Staebler, M. Yang

Project Description

A tokamak confines deuterium–tritium fusion fuel in the plasma state using a toroidal configuration of a magnetic field. In H-mode, high confinement is achieved because of a narrow insulating region that spontaneously forms near the plasma edge as the heating power is increased through a threshold. This region is called the pedestal. As insulating properties of the pedestal in a fusion pilot plant improve, the fusion power increases. Two properties influence this: the pedestal width and the pressure gradient it supports. These are determined by two pieces of physics: the transport of heat and particles through the pedestal and its vulnerability to eruptions called edge-localized modes (ELMs). Both are a consequence of plasma instabilities. The ELMs are driven by large-scale magnetohydrodynamic (MHD) instabilities called peeling–ballooning modes, which affect the full pedestal. Transport is influenced by plasma turbulence, which is driven by fine-scale kinetic instabilities. This project advanced the theoretical understanding of classes of instability and used that understanding to provide a predictive model for a pedestal that is consistent with fusion pilot plant requirements. The aims included understanding how (1) the MHD instabilities can be influenced to avoid or mitigate ELMs that otherwise cause excessive damage; (2) the pedestal stability of the spherical tokamak can be efficiently analyzed, and (3) kinetic instabilities transition between the pedestal and the exhaust (scrape-off layer) plasma that surrounds it.

Mission Relevance

Understanding how to optimize the stability of the pedestal region of the tokamak plasma will lead to high plasma confinement in tokamaks, which will advance their viability as a pathway to commercial fusion power. The science undertaken in support of this goal therefore aligned with the DOE and ORNL missions to accelerate the development and deployment of energy solutions.

Results and Accomplishments

The central hypothesis of this project was that externally applied 3D magnetic perturbations modify the peeling–ballooning instabilities that underlie ELM dynamics. The primary objective of this work was, therefore, to develop a reliable and computationally efficient theoretical framework for pedestal stability in the presence of 3D magnetic geometry, with particular emphasis on calculating the energy functional, low-n modes and spherical tokamak configurations.

This project's approach built on the ELITE code, which is a standard tool for axisymmetric (2D), ideal MHD pedestal stability analysis. Progress has been made toward developing predictive theoretical and computational tools to understand and quantify the impact of externally applied 3D magnetic perturbations on tokamak pedestal stability and ELM physics. Improvements to the low-n version of the ELITE pedestal stability code have improved robustness for experimentally relevant equilibria. A new Grad–Shafranov equilibrium solver was developed and coupled to ELITE to refine experimental EFIT equilibria. This work enabled highly resolved inputs suitable for axisymmetric and 3D stability calculations.

The reconstruction of the axisymmetric energy functional was progressively improved, overcoming substantial numerical challenges, and was successfully demonstrated for realistic DIII-D equilibria. In parallel, ELITE was optimized for spherical tokamak applications; it achieved nearly a threefold speedup for a representative Mega Ampere Spherical Tokamak Upgrade equilibrium. Analytical derivations of the leading-order 3D corrections to the ideal MHD energy functional were completed and are ready for numerical implementation and testing. These developments established a foundational toolset for future studies of low-n stability and 3D effects in conventional and spherical tokamaks.

GEOSPATIAL SCIENCE AND HUMAN SECURITY DIVISION

11291: Environmental Anomaly Detection for Biopreparedness

A. Anyamba, S. Guggilam, H. Tubbs

Project Description

The emergence and spread of vector-borne diseases have the potential to result in outbreaks and epidemics across the world, which threaten US national and global health security. Vector-borne and zoonotic pathogens comprise at least two-thirds of the 57 top infectious disease threats to US Department of War personnel and account for more than 17% of all infections, with more than 700,000 deaths globally per year. The incidence of vector-borne diseases is increasing within the US Armed Forces and the global public, especially since chikungunya and Zika viruses emerged in the western hemisphere in 2013 and 2015, respectively. The epidemiology of many vector-borne pathogens is driven by climate and environmental conditions that critically influence vector survival, reproduction, biting rates, feeding patterns, pathogen incubation and replication periods, and the efficiency of pathogen transmission among multiple hosts. Although specific shifts in patterns of climate and weather anomalies are known to precede certain vector-borne disease outbreaks (e.g., dengue, chikungunya, Rift Valley fever, Zika, and hantavirus, among others), no study has comprehensively estimated driving environmental and climatic thresholds, temporal persistence of the anomalies, and generalization of these metrics at a global scale. This project undertook such a systematic evaluation of various satellite remote sensing and climate data in combination with disease outbreak datasets to determine various metrics preceding and during selected vector disease outbreaks and epidemics.

Mission Relevance

Vector-borne and zoonotic pathogens result in outbreaks and epidemics, which threaten US national and global health security. This project addressed DOE's mission to ensure America's security and prosperity by developing methods, understanding data sources and uncertainties, and proposing appropriate solutions based on environmental and climate anomaly metrics for vector-borne disease threats.

Results and Accomplishments

This project used several assets: harmonized multidecadal climate and environmental time series, georeferenced outbreak curation, and a semisupervised multivariate anomaly scoring framework. The project applied these assets in retrospective case studies that demonstrated context-consistent associations (e.g., persistent precipitation or vegetation anomalies preceding Rift Valley fever activity, climate-linked clustering of outbreaks during El Niño, and regional climate correlates for plague and hantavirus). Within the project period, however, heterogeneous reporting, spatial and temporal resolution mismatches, and ecosystem-specific dynamics limited the ability to declare stable, generalizable anomaly magnitude or persistence thresholds across pathogens and regions.

NUCLEAR ENERGY AND FUEL CYCLE DIVISION

11697: Accelerated Nuclear Fuel Development and Qualification

D. Adorno Lopes, R. Juneja, M. Kurley III, W. Cureton

Project Description

Accelerated fuel qualification (AFQ) frameworks for nuclear fuel have been proposed that leverage advanced experiments, materials modeling, and computational tools to minimize the need for integral fuel irradiation experiments. However, AFQ still lacks a clearly defined pathway for how to combine such a broad range of tools to ultimately receive regulatory approval. This project proposed to demonstrate the ability to reliably predict nuclear fuel properties and performance under various environmental conditions. Various fuel performance phenomena (e.g., solid swelling and thermal conductivity degradation) were predicted using machine learning and density functional theory and were validated via separate effects testing. Various methods of data analysis were employed, such as convolution of crystal lattices in descriptors and transfer learning. Model validation consisted of two parallel tracks: (1) fabrication of uranium dioxide, uranium mononitride, and uranium monocarbide materials, each with an array of solid fission products, followed by out-of-pile characterization, and (2) comparison with historic and ongoing in-pile (long-term) fuel performance data, leveraging High Flux Isotope Reactor MiniFuel irradiation campaigns and data from ongoing projects.

Mission Relevance

This project advanced DOE's mission by developing scientific foundations, analytical methods, and experimental capabilities required to accelerate the qualification of nuclear fuels for advanced reactors. By integrating machine learning, density functional theory, and targeted separate-effects experiments, this work supported DOE's objective to ensure America's energy security through transformative science and technology solutions.

Results and Accomplishments

The project delivered scientific and technical advances toward an integrated AFQ framework. This project developed physics-informed machine learning models capable of predicting thermophysical degradation in uranium dioxide, uranium mononitride, and uranium monocarbide as a function of fission product chemistry and defect types. Integrated density functional theory workflows were used to quantify defect formation energies and solid-solution behavior across uranium dioxide, uranium mononitride, and uranium monocarbide (each with an array of solid fission products) systems, which directly informed machine learning feature sets and experimental targets.

Progress was achieved in producing and characterizing single-effect surrogate samples to isolate the roles of individual fission products. This project fabricated controlled-chemistry uranium dioxide and cerium dioxide samples doped with neodymium(III) oxide and yttrium(III) oxide, which represented soluble fission product analogs. These samples allowed the team to disentangle single-solute effects that could not be resolved in SIMFUEL or irradiated fuels in which multiple fission products coexisted. High-resolution x-ray diffraction and Rietveld refinement demonstrated measurable lattice parameter evolution with dopant concentration to confirm solid-solution incorporation behavior and enable calibration of machine learning–derived descriptors. Nanoindentation measurements revealed systematic changes in elastic modulus and hardness with increasing dopant concentration, which validated model predictions of mechanical degradation mechanisms driven by local defect chemistry. Total scattering and pair distribution function analysis at the Spallation Neutron Source (using the Nanoscale-Ordered Materials Diffractometer, NOMAD) provided the first direct quantification of oxygen sublattice distortions, vacancy clustering, and local structural disorder in doped uranium dioxide and cerium dioxide systems. These neutron data were essential for resolving oxygen stoichiometry variations, redox-driven defect configurations, and nanoscale strain fields that cannot be captured with laboratory x-rays.

For integration with an AFQ framework, this project established a quantitative workflow linking experimental measurements, density function theory energetics, and machine learning models to evaluate how individual fission products modify fuel performance–limiting properties. The project also built a comparison methodology incorporating High Flux Isotope Reactor MiniFuel data, historic long-term irradiations, and advanced modeling outputs to evaluate accelerated burnup sensitivities.

SUMMARIES OF PROJECTS SUPPORTED BY DISTINGUISHED STAFF FELLOWSHIPS

RUSSELL FELLOWSHIP

No projects in this fellowship were completed in FY 2025.

WEINBERG FELLOWSHIP

11284: Precision Polymer Synthesis via Dynamic Catalysis

J. Foster

Project Description

Synthetic polymers can be found in virtually every consumer product and comprise structural materials, food packaging, textile fibers, and medical devices. The function of a polymer is related to its molecular structure. Polymer structure emerges from its molecular weight, chemistry, tacticity, and architecture. Polymer scientists now understand how polymer molecular weight and architecture influence the behavior of polymeric materials in solution or in bulk. Thus, strategies have been developed to precisely control these parameters. However, in contrast to the ability to control polymer molecular weight, the methods to precisely define polymer sequence are meager. The chemical sequence of a polymer is as important as its length. Proteins are perhaps the best exemplar of this concept. Polymer scientists have attempted to mimic protein biosynthesis; however, developed approaches are limited in terms of scale and efficiency. The inability to precisely control polymer sequence has severely restricted the design space of synthetic polymer materials and precluded understanding of how sequence can be leveraged to target desired material properties. Thus, the overarching goal of this research was to unravel the fundamental principles for sequence-controlled polymerization and understand how sequence influences the bulk properties of synthetic polymer materials. This project investigated common classes of cyclic olefin monomers for ring-opening metathesis polymerization (ROMP) to identify monomers for single-unit addition during sequential monomer addition synthesis. The project discovered that oxanorbornene imide monomers are suitable for these types of reactions. Findings on monomer kinetic trends in copolymerizations were then applied for plastic waste upcycling. This work revealed that a variety of butadiene sources can produce alternating copolymers under mild conditions.

Mission Relevance

This project targeted control over the process by which reactants (i.e., monomers) assemble into products (i.e., polymers) and leveraging the high-information content of sequence-defined macromolecules to accelerate materials discovery. This research is relevant to catalysis science and transformative manufacturing, with novel polymerization catalysis potentially enabling the creation of new manufacturing technologies. Enhanced control over polymer structure can be directly leveraged to enable polymer upcycling by creating materials with site-specific degradation behavior. This project contributed to the objective of being able to design and make materials with precisely defined functions and properties, with implications for America's security and prosperity.

Results and Accomplishments

This project found that easily synthesized oxanorbornene imide monomers are suitable for single-addition reactions. With the identified monomers, this project demonstrated the synthesis of multiblock copolymers containing up to three precise functionalization sites and singly cross-linked four-armed star copolymers.[1]

A distinct strategy to upcycle waste polyalkenamers such as polybutadiene into new, performance-advantaged materials was developed by using them as drop-in additives for ROMP. The polyalkenamers served as competent chain-transfer agents in ROMPs of common classes of cyclic olefin monomers to facilitate good molecular weight control, allow low ruthenium catalyst loadings, and enable efficient incorporation of the polyalkenamer into the synthesized polymeric material. This project successfully demonstrated ROMP using model polyalkenamers and translated these results to leverage

[1] J. C. Foster, J. T. Damron, and H. Zhang, "Simple Monomers for Precise Polymer Functionalization During Ring-Opening Metathesis Polymerization," *Macromolecules* 56, no. 19 (2023): 7931–7938, https://doi.org/10.1021/acs.macromol.3c01068.

commercial polybutadiene and acrylonitrile butadiene styrene as chain transfer agents for ROMP copolymerizations. The developed strategy was shown to be highly efficient and operationally simple by quantitatively incorporating the polyalkenamer and inheriting aspects of its thermomechanical performance.[2]

Dynamic catalysis was further exploited to fine-tune polymer microstructure. This project studied the capability of ring-opening cross-metathesis polymerization to produce alternating copolymers. By treating commercial polybutadiene with bulky oxanorbornene monomers and ruthenium-based olefin metathesis catalysts, alternating copolymers were produced under mild conditions with high sequence fidelities. This project found that alternating copolymers could be produced starting from a variety of butadiene sources, including polybutadiene, cyclooctadiene, or *t*,*t*,*t*-1,5,9-cyclododecatriene. This result highlighted for the first time the kinetic pathway independence of this process. Kinetic copolymerization analysis of an oxanorbornene monomer with *t*,*t*,*t*-1,5,9-cyclododecatriene revealed that much higher monomer conversions were obtained compared with the analogous homopolymerizations and showed evidence of alternating monomer incorporation. Copolymerization of these monomers also enabled good control when targeting different molecular weights. Copolymer thermal analysis revealed a strong correlation between thermal behavior and alternating sequence fidelity, which provides a second lever beyond composition to tune thermal behavior. These data demonstrated that a broad variety of polymer microstructures can be accessed via ring-opening cross-metathesis polymerization and highlighted the potential of *t*,*t*,*t*-1,5,9-cyclododecatriene in alternating copolymer synthesis.[3]

[2] J. C. Foster, et al., "Polyalkenamers as Drop-In Additives for Ring-Opening Metathesis Polymerization: A Promising Upcycling Paradigm," *Journal of the American Chemical Society* 146, no. 48 (2024): 33084–33092, https://doi.org/10.1021/jacs.4c10588.
[3] J. C. Foster and I. T. Dishner, "Tuning Copolymer Microstructure Using Ring-Opening Cross-Metathesis Polymerization," *Macromolecules* 58, no. 10 (2025): 5170–5176, https://doi.org/10.1021/acs.macromol.5c00659.

WIGNER FELLOWSHIP

No projects in this fellowship were completed in FY 2025.

SUMMARIES OF PROJECTS SUPPORTED BY THE EARLY CAREER COMPETITION

BIOSCIENCES DIVISION

11349: Harnessing Fungi to Enhance Soil Carbon Sequestration

J. Tannous, C. Salvador, R. Ployet, K. Carter, K. Davis, T. Rush, L. York

Project Description

Ectomycorrhizal fungi (EMF) play a vital role in soil carbon sequestration. These fungi enhance plant growth, which contributes to faster carbon dioxide capture through biomass production, and help stabilize carbon in soil by integrating it into their hyphal networks, thereby slowing its release back into the atmosphere. The underlying mechanisms of fungal soil carbon sequestration, particularly in EMF, are not completely understood. This project aimed to better understand how specific EMF communities influence carbon cycling and ecosystem function. The project established a compatible synthetic fungal community, developed a lab-scale mesocosm system with real-time carbon dioxide monitoring, and used it to study the effect of EMF on plant performance, soil carbon dynamics, and volatile organic compound emissions. The results revealed that these fungi significantly influence plant traits linked to carbon sequestration and modulate soil volatile organic compound profiles.

Mission Relevance

This project aligned with DOE's mission to enhance environmental resilience through transformative science. By investigating the role of EMF in soil carbon dynamics and plant–soil interactions, this research contributed to a deeper understanding of terrestrial ecosystem processes. Insights gained from this work support the development of sustainable land management practices and strategies to maintain ecosystem stability.

Results and Accomplishments

This project made advancements in understanding the role of EMF in soil carbon dynamics and plant–soil interactions. The project conducted in vitro coculture experiments with six EMF species across four synthetic media to identify a compatible community without antagonistic interactions. Four species—*Laccaria bicolor*, *Rhizopogon salebrosus*, *Hebeloma brunneifolium*, and *Cenococcum geophilum*—were selected for further study. The project also established and optimized a lab-scale mesocosm system integrated with a LI-COR Environmental multiplexer for real-time carbon dioxide measurement. Using *Populus* genotype 717 and sterilized topsoil, this system provides a controlled platform for studying plant–soil–microbe interactions.

This project leveraged the mesocosm system to evaluate the effect of the EMF community on carbon dioxide emissions, volatile organic compounds, and plant development. Findings indicated that EMF inoculation enhances plant traits associated with carbon sequestration and alters soil volatile organic compound profiles, which suggests influence on ecosystem functions.

Finally, the project designed and validated species-specific quantitative polymerase chain reaction primers for the selected EMF strains, which enabled precise quantification of fungal colonization in soil and root samples.

BUILDINGS AND TRANSPORTATION SCIENCE DIVISION

11352: Cross-Sector Spatiotemporal Energy Analysis Framework

R. Wang, M. Pan, W. Li

Project Description

Previous research often modeled transportation, building, and grid energy in isolation, which limited insight into the interactions among these sectors and their combined influence on energy use. To address this limitation, this project developed a cross-sector energy analysis framework that integrated mobility, building, and grid datasets to evaluate how different transportation modes and travel behaviors affect urban energy systems. The developed framework incorporated the Model America building stock database, the National Household Travel Survey, origin-destination trip data, and Alternative Fuel Data Center charging infrastructure. These inputs supported trip chain simulations and EnergyPlus models to jointly estimate building and charging demand, which were mapped to the IEEE 37-bus distribution feeder. OpenDSS simulations were then applied to assess grid performance under varying levels of electric vehicle adoption and charging infrastructure expansion. The resulting system illustrated how mobility-driven occupancy patterns shape spatiotemporal energy demand and grid reliability, which provides planners with a scalable tool to evaluate technology adoption pathways.

Mission Relevance

This project advanced DOE's mission by providing a transformative, data-driven framework for analyzing the connections among transportation, buildings, and the power grid—sectors that are traditionally studied in isolation. By integrating multiple national datasets and applying advanced simulation tools, the framework improved the understanding of how mobility and occupancy patterns drive energy demand and grid performance in urban environments. This knowledge supports America's energy security and prosperity by equipping planners and decision-makers with scalable tools to guide infrastructure investments and system modernization.

Results and Accomplishments

This project delivered a scalable simulation framework that integrated mobility and energy domains to better capture the interplay among travel behavior, building occupancy, and electricity demand. The framework generated realistic daily travel patterns, including charging trajectories, through trip chain simulations based on the National Household Travel Survey and origin-destination data. Building energy use was modeled using the Model America dataset and EnergyPlus to capture occupancy-driven loads, which were then spatially allocated to census block groups and mapped to the IEEE 37-bus feeder. OpenDSS power flow simulations were used to evaluate voltage stability and unbalance under different levels of market penetration.

The framework was demonstrated in a case study of 51 census block groups in downtown Atlanta, Georgia. A baseline scenario simulated more than 68,000 vehicle trajectories, generating over 240,000 daily trips, or approximately 3.5 trips per person. Charging demand was linked to 55 Alternative Fuel Data Center–identified stations, primarily Level 2 ports, with additional infrastructure expansion scenarios tested. Results showed that higher-power chargers (150–400 kW) shifted demand peaks earlier in the day as vehicles charged and departed quickly. Although this compressed the load into shorter time windows and raised peak demand, scenarios in which only one new port was added per station revealed that 150 kW chargers produced higher simultaneous energy use than 400 kW chargers because of longer dwell times and overlapping demand. Strategically expanding ports at already-busy stations amplified local demand more than evenly distributing them, which highlighted the importance of siting strategies. In fact, only five high-power ports met most charging demand, with limited benefits gained from adding more.

The simulations also revealed important building–charging interactions. Residential buildings saw strong evening peaks as people returned home and recharged vehicles, making these buildings the most sensitive to future technology adoption. Commercial and industrial buildings displayed steadier load profiles, with industrial sites contributing relatively less to charging demand. Comparisons with workplace charging datasets from the National Laboratory of the Rockies confirmed that the model reasonably reproduced daily charging trends.

A grid impact analysis showed that at 10% market penetration, voltage stayed within ±5% of the acceptable range, with the lowest levels occurring in the late afternoon. At 40% market penetration and above, sharper voltage drops emerged, particularly in the evening when residential charging was concentrated. At 80% market penetration, voltage unbalance rose significantly in dense residential areas and propagated instability into surrounding zones, which underscored the risks of localized charging clustering.

Overall, this work demonstrated how incorporating mobility-informed occupancy improves accuracy when estimating building and charging loads. It highlighted how charger power levels shift the timing of demand and how residential density drives spatial clustering effects. These insights provided guidance for utilities, planners, and policymakers and emphasized the need for time-of-use pricing, targeted siting of infrastructure, and consideration of future vehicle-to-grid solutions to mitigate evening peak stress.

CENTER FOR NANOPHASE MATERIALS SCIENCES

11354: Stacks to Synapses—Layering Materials for Neuromorphic Systems

S. Neumayer

Project Description

The increasing importance of AI and big data applications demands the development of high-performance, energy-efficient computers. A promising approach is to combine neuromorphic computing with neuromorphic hardware. This brain-like method is inspired by the fast and energy-efficient way the human brain processes and stores information using synapses and neurons. This project aimed to develop functional material stacks that act as artificial synapses for neuromorphic hardware elements and provide a more energy-efficient method than conventional computing architectures. To achieve neuromorphic functionality, hardware elements must have multiple states, output signals dependent on past and present input, and tunability of physical parameters. Stacking materials creates interfaces that can modulate the local functional properties through (electro)chemical interactions, electronic correlations, and charge migration. These interactions affect functional material properties, which can be used to store and process information. To achieve the desired neuromorphic signature, it is necessary to understand the interactions arising at the interfaces between top and bottom building blocks and their effect on the stack functionality. This project found that $PdSe_2$ nanosheets showed enhanced memresistive and memcapacitive properties when stacked due to trapped charge carriers at the interfaces. This behavior was revealed through a newly developed nanoscale electrical characterization protocol based on atomic force microscopy. Moreover, this project discovered that the type of metal contact interfacing with van der Waals layered $CuInP_2S_6$ strongly affects the measured polarization as well as the switching dynamics between these states. This phenomenon is governed by the stabilization of different structural phases through the metal–$CuInP_2S_6$ interfaces.

Mission Relevance

The project advanced DOE's mission by addressing the energy and national security challenges posed by the high power demands of AI and big data applications that currently rely on conventional computing technologies. In particular, the project addressed outstanding needs in the fields of advanced manufacturing, microelectronics, and neuromorphic computing.

Results and Accomplishments

Tunable impedance states in electronic materials are key to neuromorphic computing. To detect such tunability, this project developed a nanoscale characterization method based on scanning microwave impedance microscopy. By applying triangular, bipolar voltage sweeps in a grid while simultaneously recording the scanning microwave impedance microscopy capacitance and scanning microwave impedance microscopy resistivity channels, capacitance–voltage and resistivity–voltage curves could be locally obtained. The characteristics of these curves revealed charge accumulation, depletion, inversion, memcapacitance, memresistivity, and signal ranges, which could be directly linked to semiconducting behavior and memimpedance characteristics. This microscopic method provided crucial insights into local variations in electrical and electronic properties that underpin device performance.

Using this atomic force microscopy–based technique, this project detected electrical memory effects in van der Waals layered $PdSe_2$ nanosheets. Stacking nanosheets enhanced memresistivity and memcapacitance—an effect attributed to trapped charge carriers at the interfaces. This project also observed the existence of p-type $PdSe_2$ nanosheets despite prior reports showing only n-type dominated ambipolar characteristics in macroscopic measurements. Further interfaces between $CuInP_2S_6$ and $PdSe_2$, graphene, and $AgBiP_2Se_6$ were studied using electrical atomic force microscopy modes.

Beyond interfaces between 2D layered materials, this project discovered that metal electrodes can significantly influence the functional properties of van der Waals materials. Ferroelectric $CuInP_2S_6$ exhibited different behavior depending on the specific metal deposited. In the absence of an applied direct current voltage, the silver electrode increased the piezoresponse, whereas the copper interfaces suppressed it. Polarization switching probed with direct current voltage spectroscopy likewise showed strong differences between the two metal–$CuInP_2S_6$ interfaces. Under silver electrodes, $CuInP_2S_6$ displayed polarization switching behavior similar to uncovered $CuInP_2S_6$ regions. In contrast, $CuInP_2S_6$ under copper electrodes exhibited hysteresis loops with orientations opposite to the surrounding uncovered $CuInP_2S_6$. These distinct polarization dynamics can be attributed to the stabilization of different, energetically close polar and nonpolar structural phases that can exist in $CuInP_2S_6$. Because integrating van der Waals layered materials into microelectronics and neuromorphic devices requires metal contacts for reading, writing, and transmitting signals, these findings highlight that the electrode material should be carefully selected. Moreover, the electrode material can serve as a tuning parameter for optimizing device performance.

COMPUTATIONAL SCIENCES AND ENGINEERING DIVISION

11355: Widespread and Persistent Climate Extremes: Identification and Impact Assessment

D. Rastogi

Project Description

Weather extremes, including heat waves, storms, and hurricanes, result in noticeable and even devastating impacts on infrastructure and human health. This project identified widespread and persistent weather extremes associated with heat waves and cold waves using an ensemble of high-resolution Earth system datasets with a detection and tracking algorithm deployed on high-performance computing. Using high-performance computing resources at the Oak Ridge Leadership Computing Facility, this project developed AI-based methods to generate high-resolution datasets from coarse Earth system model outputs. It also used an ensemble-based approach to evaluate differences across dynamical, statistical, and AI-based approaches to improve the understanding of the impact of methodological choices on data quality and demand assessments. This project also demonstrated the impact of weather extremes on a selected human system. As a proof of concept, this project evaluated the impact on electricity consumption. This work set up an analytical workflow that can be expanded to other types of weather extremes (e.g., storms, hurricanes) and assess impacts on other human systems.

Mission Relevance

This project supported DOE's mission to advance national energy security by developing a high-resolution, data-driven framework to identify and analyze extreme weather events and their implications for grid reliability. The resulting workflow provided a robust framework to further evaluate a range of extreme events, including temperature and precipitation extremes, and assess long-term water availability and its impacts on critical infrastructure.

Results and Accomplishments

A methodology was established to detect extreme weather events across the conterminous United States using the TempestExtremes algorithm implemented on the Andes Oak Ridge Leadership Computing Facility cluster. Grid-based magnitude and minimum area thresholds were applied to identify historical weather extremes, including heat waves and cold waves, in the Daymet (1980–2019) and Livneh (1980–2018) observational datasets. This approach enabled consistent identification of large-scale, multiday weather extremes and provided a foundation for evaluating regional sensitivity and historical impacts.

The weather extremes identification workflow was extended to a high-resolution, 42-member ensemble consisting of simulations from six Coupled Model Intercomparison Project Phase 6 Earth system models, which were downscaled using dynamical and statistical techniques. Two AI-based downscaling models—Super-Resolution Convolutional Neural Network and Super-Resolution Generative Adversarial Network—were trained using Daymet data. This multimethod ensemble supported a systematic assessment of projected changes in weather extreme event intensity, spatial footprint, and persistence. To better understand the sources of projection spread, variance decomposition was performed for key weather extreme metrics. Variance decomposition showed that Earth system model choice was the primary source of uncertainty in projected changes, especially for intensity metrics. Downscaling methods contributed modest additional spread mainly for spatial footprint, whereas the reference datasets played a minimal role. Additionally, to quantify method-specific disagreement, the Method Variability Index was developed, which confirmed generally low method disagreement for heat wave temperature and duration metrics and had higher sensitivity for affected-area changes and cold wave characteristics. These diagnostics identify where methodological choices matter most for downstream impact assessments.

To evaluate energy system consequences, historical extreme weather events were linked with observed electricity demand to develop statistical demand models that were applied to future projections. Two indices—Compound Extremes Stress Index and Cumulative Load Stress Index—were introduced to quantify combined extremes and demand stress. Compound Extremes Stress Index showed moderate and regionally consistent increases, but Cumulative Load Stress Index rose much more sharply and unevenly across the regions. These results showed that even modest, near-term warming can drive disproportionately large increases in electricity demand during extreme weather events, which underscores emerging risks to grid reliability.

NUCLEAR ENERGY AND FUEL CYCLE DIVISION

11350: f-Element Structural Engineering Through Orbital Perturbation

K. Lawson

Project Description

The charge density of the actinide elements, which can vary given their unique 5f orbital interactions, is a key driver of actinide material properties. With complex 5f orbitals, actinides possess rich chemistry, including accessing many oxidation states. Developing tailored actinide materials is hindered by a lack of sufficient understanding of the effects of perturbations in the f-orbitals. Systematic manipulation of material properties can provide the foundation for high-fidelity structure–property relationships and, subsequently, for the design of new actinide materials with computationally predicted properties.

The overarching goal of this project was to control actinide dioxide thermochemical and thermophysical properties through changes in the oxidation states of the starting materials and their f-electron charge density. The work used actinide oxalate and oxide systems as training grounds for f-element structural engineering because of their importance and universality throughout the nuclear fuel cycle. Synthesis of oxalate precursors and oxide solids was performed for the tailored development of cerium-, thorium-, uranium-, and neptunium-containing materials. Within the series of elements proposed, accessible oxidation states included 3+, 4+, 5+, and 6+. The logical progression of this work was to establish methods to stabilize molecular precursors with uncommon f-element oxidation states as well as synthesize varying actinide oxalate compounds, develop a mechanistic understanding of uncommon oxalate decomposition reactions toward actinide dioxide formation, and characterize and compare thermochemical and thermophysical properties of synthesized actinide solids using advanced solid-state characterization techniques for correlating structure–property relationships to tailor actinide dioxide properties. Characterization of solids included scanning electron microscopy coupled with energy dispersive spectroscopy, powder x-ray diffraction, thermogravimetric analysis, and differential scanning calorimetry. The application of state-of-the-art solid-phase characterization, which is still limited in the field of actinide chemistry, was the foundation for characterizing key thermochemical and thermophysical properties of these systems. This work provided a basis to assert systematic control over f-element chemical phenomena in many actinide material systems.

Mission Relevance

This project advances DOE's mission through basic science research related to nuclear fuel cycle materials. Fundamental research on actinide elements and how to tailor key actinide material phases, such as oxides, applies to nuclear fuel development, nuclear fuel cycle processes, and nuclear forensics and nonproliferation efforts.

Results and Accomplishments

Results of this work included expanding and refining oxidation state control of f-elements, developing a novel synthesis technique for oxalate solids at higher oxidation states, developing thermal analysis techniques for transuranic materials, and growing the body of available imaging for f-element oxalates and oxides. Oxidation state control for the f-elements studied in this project was a key requirement to perform syntheses. Initially, the project hypothesized that oxidation state control for the elements of interest—cerium, uranium, and neptunium—would be possible through electrochemistry alone. Although electrochemistry was successful in stabilizing some oxidations, the electrochemical technique had shortcomings for solutions with high concentrations of these f-elements. Specifically, high-concentration solutions, which were needed for many of the synthesis techniques, could not be stabilized electrochemically because of the surface area limitations of available electrodes. Honing the electrochemical cell, including designing custom electrodes, proved fruitful to manage some of these effects. Additional techniques for oxidation state control were also pursued, including photochemistry. A

small photochemical reactor allowed for rapid control of the oxidation states of uranium and neptunium—even at higher concentrations of these elements in solution—and opened new pathways to achieve oxidation state control of f-elements in the absence of chemical additives.

Using the initial chemical manipulations performed through electrochemistry and photochemistry, synthesis of an array of f-element oxalates was carried out over the course of this project, including cerium(III), cerium(IV), uranium(IV), uranium(VI), neptunium(IV), and neptunium(V) oxalates. An area of advancement in this project was the application of a solvent displacement crystallization technique to grow crystals of uranium(VI) oxalate. The uranium(VI) oxalate complex in solution is reported to have remarkable stability and, correspondingly, high solubility. Consequently, it is difficult to achieve the supersaturation state that is required to precipitate uranyl oxalate phases from solution. In this project, solvent displacement crystallization was tested and successfully synthesized uranium(VI) oxalate from methanol, ethanol, 1-propanol, and isopropanol as additive solvents. Characterization of the crystallized phases showed an absence of solvent influence on the bulk structure but indicated a variability in observed morphology and particle size based on solvent employed in synthesis. Clear trends emerged regarding solvent polarity used in solvent displacement crystallization for uranium(VI) oxalate precipitation. Specifically, the total yield of solid was a function of solvent polarity, with lower solvent polarity increasing yield. This approach can be useful for other high-solubility actinide complexes, primarily those in the 5+ or 6+ oxidation state, which can be difficult to crystallize. Additionally, this technique could be used to crystallize actinides from low-concentration solutions, given the difficulty in obtaining large quantities of many radioactive elements.

Beyond synthesis of f-element oxalates, subsequent characterization of the solids was performed using thermogravimetric analysis and microscopy. A result within this portion of the scope was the progressive method development to enable analysis of neptunium(IV) oxalate, which has the highest radioactivity per mass quantity of the all the elements studied in this project and requires highly specialized work controls. Through progressive analysis of cerium, a nonradioactive f-element, followed by thorium and uranium, which have significantly lower radioactivity per mass quantity, sample preparation, transfer, and analysis methods were refined to prepare for neptunium analysis. Analysis of neptunium(IV) oxalate is limited in the literature, and the measurements performed in this work represent the first revisitation of this material in decades. Further characterization of materials, including imaging using digital and electron microscopy, added additional insights into physical properties of the series of materials studied throughout this project.

RADIOISOTOPE SCIENCE AND TECHNOLOGY DIVISION

11347: Bringing Promethium Chemistry Up to Date

F. White

Project Description

Promethium is the only radioactive lanthanide. Because of its radioactive nature, it is the least-studied lanthanide. This lack of research has left a significant knowledge gap in promethium chemistry, especially in the solid state. This work aimed to increase knowledge of promethium by examining its structure, bonding, and spectroscopic nature through cyanometallate structures. The project prepared a promethium cyanometallate structure, which was one of the first examples of a promethium single crystal structure, and provided one of the first looks at promethium emission in the solid state.

Mission Relevance

Promethium, particularly promethium-147, is one of the few beta-emitters that can be used for lightweight, nuclear betavoltaic batteries, which are relevant to DOE's mission to address America's energy challenges. However, chemical knowledge of promethium is scarce, thus inhibiting growth in this area of technology. ORNL is currently the sole producer of weighable quantities of promethium-147, and this work elucidated properties of promethium in the solid state, which is the state in which promethium is primarily used.

Results and Accomplishments

Very little information about promethium is known compared with many other elements, particularly other lanthanides, because promethium is the only lanthanide that has no stable isotopes. Moreover, the chemistry of promethium in the solid state is even less defined, and the solid-state chemistry of promethium is important because most applications using promethium are in the solid state. The primary objective of this work was to develop the chemistry of promethium in the solid state to better understand the bonding and properties of this elusive element.

This project established reaction parameters for promethium using the nonradioactive lanthanide surrogates neodymium and samarium. Because promethium is extremely rare and valuable, an effort was made to meticulously determine suitable reaction conditions using the surrogates. Synthetic efforts to make a platinum cyanometallate compound transpired using neodymium and samarium. Interestingly, it was observed that under these identical reaction conditions, neodymium and samarium repeatedly formed two different structure types, although their charge, ionic radii, and chemistry were extremely similar. The structural isomers differed in their coordination environment around the lanthanide metal center. The neodymium compound crystallized in the triclinic space *P1*, and the samarium compound crystallized in the orthorhombic space group *Pbcn*. The coordination of the neodymium metal center with the $[Pt(CN)_4]^{2-}$ group can be best described as *cis*-bridging and, likewise, *trans*-bridging in the samarium structure.

The reaction conditions used with promethium were scaled down to where only 200 µg of promethium-147 was needed for the reaction. Because promethium lies between neodymium and samarium on the periodic table, it was hypothesized that promethium would form either the neodymium or samarium structure type. Surprisingly, promethium formed neither structure type and instead formed a different structure type in which a nitrate anion was reduced to a nitrite. This anion formed a bridging structure between promethium metal centers. Preliminary results attributed this to promethium-147 radiolytic damage. Regardless, this work provided one of the first promethium single-crystal structures. This change in structure type could potentially provide insight into ways to more efficiently process promethium from the irradiated waste stream at ORNL.

Next, the spectroscopic properties of promethium were probed in the solid state for this crystal system. This work provided some of the first solid-state absorbance and emission measurements of promethium. Although cyanometallates were selected because of known spectroscopic characterization with lanthanides, the complexity and richness in the near-infrared emission spectra of promethium was unexpected. The results of the emission spectrum were validated by computational techniques.

SUMMARIES OF PROJECTS SUPPORTED BY THE UT-ORII FELLOWSHIPS

12041: Online Metrology for Convergent Manufacturing

M. Borish, W. Carter

Project Description

The overarching goal of this project was to improve the efficacy and efficiency of convergent manufacturing processes by creating and demonstrating online metrology capabilities. This project aimed to accelerate the growth of national manufacturing capacity by demonstrating desired manufacturing outcomes through two specific aims: (1) develop simultaneous localization- and mapping-based metrology algorithms and a software toolset and (2) demonstrate desired manufacturing outcomes with sensors and equipment at ORNL's Manufacturing Demonstration Facility. Specific aim 1 adapted real-time 3D reconstruction algorithms from simultaneous localization and mapping to deliver metrology capabilities with a flexible, reusable software toolset. Specific aim 2 deployed these new capabilities at the Manufacturing Demonstration Facility and demonstrated desired manufacturing outcomes that showed improved efficacy and efficiency consistent with the project's overarching goal. This project allowed for selecting and integrating new sensor capabilities as well as novel inspection algorithms for in situ object reconstruction. Example objects were characterized, which demonstrated the efficacy of the approach. The project was also able to complete an initial investigation into defect identification and analysis.

Mission Relevance

This work served the national need to rapidly increase manufacturing capacity and aligned closely with DOE key mission areas. Establishing this new capability at the Manufacturing Demonstration Facility will improve manufacturing outcomes and speed up multiple avenues of manufacturing research.

Results and Accomplishments

This project investigated appropriate sensors for digital reconstruction of 3D objects, settling on structured light because of its robust combination of accurate measurement, speed, and flexibility. Photoneo was chosen, and a MotionCam-3D was purchased. This sensor provided dense data collection at real-time speeds of a few hundred milliseconds. Additionally, the sensor was wrapped inside established software libraries that will allow interoperability and reuse with additional platforms in the future. The sensor was then attached to a Universal Robotics platform.

With initial setup complete, the project team began software development on the autonomous scanning methodology. Several insights were identified during code development. At a high level, this code assists the user in choosing a view of an object to maximize how much of the object is seen. Because the object may not be known ahead of time, this problem is contained in a class of computer science problems known as NP-hard, meaning no known algorithmic solution guarantees an optimal result other than trying every possible input combination. This work involved trying every possible view around an object. To address this problem, the project team developed a heuristic scoring mechanism to define a good view and balance competing concerns during the selection process. A good view was defined as one that maintained a minimum amount of overlap with already existing views while attempting to maximize the amount of unseen object. Essentially, the program enforced locality on the next view while trying to view as much as possible of the unseen object. This approach produced a balanced result of good coverage at fast speeds because the robotic arm did not have to move far to reach a new view, and each new view scanned as much as possible of the unseen portion of an object.

In the process of developing the heuristic, it became apparent that this approach could not be applied to the collected data without downsampling, which was required because of the volume of data collected from the sensor combined with the goal of providing an in situ reconstruction. To accomplish downsampling, the point cloud data were triangulated into a 3D mesh. Then, this mesh was overlaid with a voxel grid. The voxel grid was of arbitrary resolution, but this project was able to achieve a voxel size

of 2 mm, allowing an object the size of a coffee mug to have its candidate views computed in approximately 3 s. Notably, because of the NP-hardness of the problem, the candidate view selection was constrained. In this case, the computation time was for 100 candidate views. In total, the entire algorithmic pipeline for view selection took approximately 7.5 s and included data collection, data processing, and view selection.

The final set of insights came from developing code for defect identification. The heuristic development mentioned previously did not consider the correctness of an object scan but instead simply whether the necessary coverage had been achieved and at what locations coverage should occur. To assist in the convergent manufacturing pipeline, the code library must also be able to identify defects for further scanning in situ so that corrective action can be taken if necessary. This project found two separate use cases determined by whether a reference mesh or CAD existed. For each of these use cases, a unique implementation was developed.

For the case in which no reference mesh existed, eigenvalue-based curvature analysis was investigated. Eigenvalues represent how much variance is present in a local area of the surface. In this case, this variance was a curvature analysis. Objects that were successfully built should have relatively smooth surfaces and gradual transitions among features on the surface. Thus, high eigenvalues indicated abrupt or disjoint surfaces that could indicate defects. This eigenvalue-based approach proved effective in identifying disjoint and abrupt surface transitions in example scans that were taken with defects purposely attached to the object. This work indicated that the approach could be used to supplement the original view computation code to override location regions of interest for additional scanning or mitigation. Similar to the view computation, this approach required only a few seconds to compute the eigenvalues and analyze the surface.

The second case was one in which a reference mesh existed. In this case, the object was voxelized and downsampled just like the simple view computation because of the sheer amount of data. In each voxel, a truncated signed distance field was computed. The truncated signed distance field held a value that described the distance from the voxel to the nearest surface and could be positive or negative. A positive value indicated the voxel was outside the object, and negative indicated the voxel was on the interior of the surface. This representation could then be compared with a reference surface to determine deviations. In this case, the deviations represented defects, and the ability of the software to identify those defects was determined by the resolution of the voxel overlay. This project was able to achieve a reconstruction and comparison in slightly less than 1 min using a resolution of 5 mm. Unlike the eigenvalue approach, the voxel overlay with a truncated signed distance field can be much more precise at the cost of additional computation time.

INDEX OF PROJECT NUMBERS