DEVCOM ARMY RESEARCH LABORATORY

High Performance Computing Review

2022-2023

CONTENTS

INTRODUCTION & INFORMATION

FACILITIES

SERVICES

PARTNERSHIPS

COMPUTATIONAL SYSTEMS

HPC SUCCESS STORIES

WORKFORCE OF THE FUTURE

MESSAGE FROM THE DIRECTOR

Ms. Cynthia Bedell, SES

Director,
DEVCOM Army Research Laboratory
Army Research Directorate

In the last HPC Review I wrote of the reorganization within DEVCOM Army Research Laboratory (ARL) to align competencies, more rapidly and efficiently bringing our technical capabilities to bear for Army Futures Command challenges. The intervening two years brought us further realignment and efficiencies, yielding a streamlined ARL consisting of the Army Research Directorate, which I lead, the Research Business Directorate led by Dr. Jeffrey Zabinski, and the Army Research Office Directorate led by Dr. Barton Halpern. This new structure fully integrates our internal and external foundational research efforts with simplified lines of communication and expanded authorities to meet current and future Army requirements.

To recognize the value and potential of computational science for all areas of the Army Research Directorate, I have elevated the DEVCOM ARL Department of Defense (DOD) Supercomputing Resource Center (DSRC) from the Branch level to a staff function at the Directorate level, providing it more visibility within my organization. I am also encouraging my Division Chiefs to engage with Matt Goss and his staff to learn what capabilities are available and how they can amplify our basic and applied research.

Our Army Artificial Intelligence Innovation Institute (A2I2) is an ecosystem of networks, data repositories, laboratories, tools, computing resources, and expertise. It continues to provide advanced AI capabilities for autonomous maneuver to support multi-domain operations by serving as a focal point for Army AI/ML research and is more fully engaging HPC resources. The A2I2 Computational Facility, designed to support AI/ML, modeling and simulation, and storage, is building out the first two of four phases, providing 24,000 ft^2 of floor space for future unclassified and classified compute requirements. Phases three and four will add another 24,000 ft^2 of compute space, 40 MW of combined utility and IT power, and about 2,100 tons of cooling. This is the future home of ARL advanced computing capabilities.

Our mission to operationalize science for transformational over-match as part of persistent Army modernization has not changed. The Army's ability to operationalize science is an enduring strategic competitive advantage against any adversary. Computational science is key to realizing that advantage and for meeting the core challenges of the future fight: speed and scale. To this end, the DEVCOM ARL DSRC was awarded $99M in Army MILCON funds for a new AI/ML Computational & Informational Sciences facility. This facility will provide state-of-the-art infrastructure for increased capability to develop new algorithms, tools, and methods that support AI/ML foundational research efforts to inform force formations for 2040 and beyond at the speed of relevance.

MESSAGE FROM THE CENTER DIRECTOR

Mr. Matt Goss

Director,
DEVCOM Army Research Laboratory
Department of Defense
Supercomputing Resource Center

We have finally emerged from the shadow of COVID-19 and are defining the new normal as we proceed. Continual change is our status quo; however, I am not sure we could have forecast the changes that have emerged. TI-20 and TI-21 really challenged all the DSRCs, seeing procurement from two new providers, Liqid (DEVCOM ARL and ERDC), and Penguin Computing, Inc. (Navy and AFRL). The Navy brought heir systems online and the Air Force is not far behind. The Liqid systems presented greater hurdles which we are finally overcoming, but not without an unprecedented step.

The Army formally and contractually released Liqid from further involvement in TI-20, and accepted the systems as is. We are providing our own system administration team and have assumed responsibly for ongoing operations, support, and sustainment of all related systems. This is new territory for the DEVCOM ARL DSRC but it provides us with new flexibility. The DSRCs are the experts in their own facilities and resources. No one is better qualified to integrate new resources. Having authority and control over system administration opens the shortest path for getting these, and future, resources allocated and in the hands of the users. We cannot promise an uneventful transition, but we do promise to work with users to get you productive on Jean (unclassified, named for Jean Jennings Bartik) and Kay (classified, named for Kathleen 'Kay' McNulty Mauchly).

TI-22 acquisitions were announced in the fall of 2022 and began arriving in mid-FY23. Ruth, an unclassified system named for Ruth Lichterman, is an HPE EX-4000 having 127,488 AMD EPYC Genoa compute cores and 64 AMD MI-250 General-Purpose Graphics Processing Units, supported by 335 terabytes of memory and 17 petabytes of storage, including 2 petabytes of NVMe-based solid state storage. Marlyn, a classified system named for Marlyn Wescoff, is an HPE EX-2500 system with 57,600 AMD EPYC Genoa compute cores supported by 135 terabytes of memory and 13 petabytes of storage, including 2 petabytes of NVMe-based solid state storage.

Betty, our TI-18 classified investment named for Elizabeth "Betty" (Snyder) Holberton, is exhibiting greatly improved availability. HPE has supplied ARL with a "consideration" system, i.e., in consideration of the extensive delays experienced in getting Betty stable. It is named Snyder, Betty's last name, and is an HPE EX-4000 like Ruth. Since TI-18, we have named our new supercomputers to celebrate the remarkable achievements and enduring legacy of the original team of ENIAC programmers.

Our newest developmental platform, Chessie, is named after the Chesapeake Bay Monster. Users can quickly explore tradeoffs in algorithm performance and programming requirements for a range of processor combinations, and mix and match x86 and ARM processors, and GPU accelerators, all supported by a WEKA file system. It should prove especially productive for those wanting to evaluate models for edge and small form factor deployment on ARM-based devices. We plan to add Nvidia Grace nodes, Nvidia's new ARM-based processors, for additional user options.

OUR VISION, OUR MISSION

Operationalizing Science And Technology

The mission of the DEVCOM ARL is to operationalize science for transformational over-match. The DEVCOM ARL is strategically placed under the Army Futures Command as the Army's sole foundational research laboratory focused on cutting-edge scientific discovery, technological innovation, and transition of knowledge products that offer incredible potential to improve the Army's chances of surviving and winning any future conflicts.

DEVCOM ARL's unique facilities and diverse, preeminent workforce comprise the largest source of world-class integrated research in the Army. The Laboratory's facilities and regional locations focus on technology areas critical to performing disruptive foundational research to answer the hardest S&T questions for Army modernization and future Army capabilities. ARL operates laboratories and experimental facilities, offices, and many one-of-a-kind facilities in several prominent locations around the United States. In many cases, the laboratory's collaborations with other nations, laboratories, academia, and industry span the globe.

DEVCOM ARL's research portfolio addresses critical Army knowledge gaps from the near-term to the far-term through disruptive essential research. This research informs and is informed by the global academic body of knowledge with support from our regional campuses. In addition, a disciplined assessment process is used to accurately forecast the impact that scientific advancements can have across the spectrum of Multi-Domain Operations. ARL has built upon its core technical competencies to form Essential Research Programs (ERPs) that drive research questions to closure. An ERP is an integrated program designed to proactively close gaps in knowledge that could not be achieved by subsets of the individual efforts alone. ERPs have long-term objectives consistent with DEVCOM ARL's corporate research laboratory mission. However, the research outcomes from the programs should influence the Army Cross-Functional Team investments, which are aligned with the Army Modernization Priorities.

We are also investing in the future workforce, for example, through the Army Educational Outreach Program (AEOP) and the High Performance Computing Modernization Program (HPCMP) Internship Program (HIP). The AEOP offers our nation's youth and teachers a portfolio of opportunities that effectively engage future workforce generations in meaningful, real world Science, Technology, Engineering and Mathematics experiences, competitions, and paid internships. The HIP provides experience to prospective DoD employees in defense related research and development and cultivates the next-generation workforce in an active and collaborative effort among the HIP Mentors/ Supervisors, HIP Interns, and the HIP Program Manager. The demand for qualified scientists and engineers, and indeed the full array of skills required to execute the mission of the DEVCOM Army Research Laboratory, is greater than ever and the Army is actively pursuing innovative approaches for attracting and retaining the future workforce.

WHY ARMY RESEARCH?

New Discoveries, Cost Cutting Innovation

The DEVCOM ARL is the Army's research laboratory strategically placed under the Army Futures Command. It is the Army's sole foundational research laboratory focused on cutting-edge scientific discovery, technological innovation, and transition of knowledge products that offer incredible potential to improve the Army's chances of surviving and winning any future conflicts. The DEVCOM ARL supports persistent modernization and is focused on disruptive science and technology for the long term, performing research to answer the hardest science and technology questions for future Army capabilities.

The DEVCOM ARL Army Research Directorate focuses on exploiting concept development, discovery, technology development, and transition of the most promising disruptive science and technology to deliver to the Army fundamentally advantageous science-based capabilities through the laboratory's 11 research competencies. This intramural research directorate also manages the laboratory's Essential Research Programs, which are flagship research efforts focused on delivering defined outcomes.

The DEVCOM ARL Army Research Office Directorate (ARO), founded in 1951 as the Army Research Office and based in Research Triangle Park, N.C., has more than 100 scientists, engineers, and support staff who manage the Army's extramural research program. ARO drives cutting-edge and disruptive scientific discoveries that will enable crucial future Army technologies and capabilities through high-risk, high pay-off research opportunities. The DEVCOM ARL Research Business Directorate (RBD) centralizes the laboratory's business operations, such as laboratory operations, strategic partnerships and plans, programs and budget synchronization, and focuses on facilitating strategic decision-making among cross-disciplinary internal and external teams.

We are integrating science towards 2040 outcomes by creating cross-competency programs designed to operationalize science, leading to transformational over-match at the speed of relevance, and leveraging tailored engagements across the ecosystem. Winning requires this 21st century engagement model.

FACILITIES

DEVCOM Army Research Laboratory DOD Supercomputing Resource Center

The DEVCOM ARL DSRC is one of five HPC centers provided by the HPCMP and hosts unclassified and classified computing platforms and storage. Combined with HPC software, secure broadband networks, and subject matter expertise, it is a powerful tool for research, discovery, innovation, problem solving, and creation and sustainment of future weapon systems. The confluence of basic and applied research expertise and facilities is the key to unlocking basic physical phenomena and harnessing the potential for defeating future threats and protecting U.S. personnel and property.

The DEVCOM ARL DSRC continues to experience growing demand for services beyond those provided by the HPCMP. We host customer HPCs, providing our experience in systems administration along with available space, power, cooling, and backup power. We provide expertise in data science, machine learning, data visualization, and tailor ML tools for customer-specific needs. We help containerize applications and build applications in a Persistent Services Framework, create user interfaces and dashboards to simplify user interaction with applications, and have developed cloud instances of customer applications. We re-purposed an HPC platform, FOB, and obtained accreditation for processing PII, an area where we see growing requirements. Finally, we stood up a Liqid developmental platform that contains x86, ARM, and GPU processors as well as a WEKA IO file system.

Defense Research and Engineering Network (DREN)

The Defense Research and Engineering Network (DREN), and its classified counterpart, the SECRET DREN (SDREN), are the HPCMP's nationwide high performance networks supporting the unique mission and wide area performance requirements of DOD's Science and Technology, Test and Evaluation, and Acquisition Engineering communities. Enabled for complex protocols, distributed engineering and testing, and high throughput transport, DREN interconnects HPC assets, DOD laboratories, and T&E ranges for DoD, industry, and partnered universityusers.

DREN has been at the forefront of IPv6 transition since the stand up of an IPv6 testbed in the late 1990s with full dual stack enablement in production operations beginning in 2002. In 2003 the DOD formally stamped DREN as the official DOD IPv6 Pilot network. With two decades of operational experience, the DREN team continues to lead the Department, the Federal Government, and Industry in IPv6 implementation with critical inputs to policy, architecture, operations, security, lessons learned, and best practice. These efforts have served as an example for other U.S government networks, paving the way for enterprise modernization based on the Internet's latest generation of core protocols.

Root Name Server

The DEVCOM ARL operates one of the 13 global root name servers, a critical element of Internet infrastructure. Root name servers are the first stop in translating (resolving) human readable host names into machine readable Internet protocol (IP) addresses used in communications among hosts and devices. Along with 11 other organizations in the US and around the world, DEVCOM ARL works to improve and evolve the root server system (the collection of all 13 root servers), focusing on the security, stability, and resilience of the collective service.

The DEVCOM ARL root name server legacy dates back to 1985 and has been in continual operation since. Starting in 2019, DEVCOM ARL expanded its infrastructure from two US network locations to 12 global locations, providing a significant increase in capacity as well as reducing latency for users around the world. Current locations for ARL root servers are: Aberdeen Proving Ground, MD, San Diego, CA, Denver, CO, Miami, FL, Frankfurt, Germany, London, England, Dubai, UAE, Johannesburg, South Africa, Hong Kong, China, Sydney, Australia, Tokyo, Japan, and Sao Paulo, Brazil.

Persistent Services Framework

The Persistent Services Framework (PSF) at the DEVCOM ARL DSRC provides functionality to bridge gaps users experience as part of their HPC workflows. Traditional HPC jobs “tear down” services and connections established during execution. Persistent services allow continually accessible applications, such as user-owned databases, a container image library, and ML toolboxes. The capability is targeted for those users who need to in incorporate data science, for example, into their existing simulations, and new users interested in leveraging the compute power offered by HPC systems to perform large-scale analysis and machine learning on their existing data sets.

The DEVCOM ARL DSRC TI-20 system, Jean, has dedicated PSF resources (20 compute nodes with 96 cores per node, and 4 storage nodes) and is available as a limited-access capability, with access and resources allocated to HPC projects/applications having bona fide PSF requirements.

Data Science Team

Since inception of the HPCMP, the PET initiative has served as the primary source of subject matter expertise in HPC as well as in the various technical domains previously organized as computational technology areas but today known as high-end computing (HEC) relevant disciplines. With the rapidly growing demand in data science (ingesting and working with large, diverse data sets and applying statistical methods to the data); training, testing, and applying ML tools; visualizing in conventional and novel ways; and containerizing workflows, the DEVCOM ARL DSRC Data Science Team provides these services to customers. Our Army civilian and contractor team assists customers in transitioning to HPC platforms, brokers partnerships between organizations, develops workflows and applications, and integrates with PET. The DEVCOM ARL DSRC capabilities, combined with foundational research in ARL and decades of PET expertise, create a powerful resource for new and established HPC users, and are positioned to address the next big challenges.

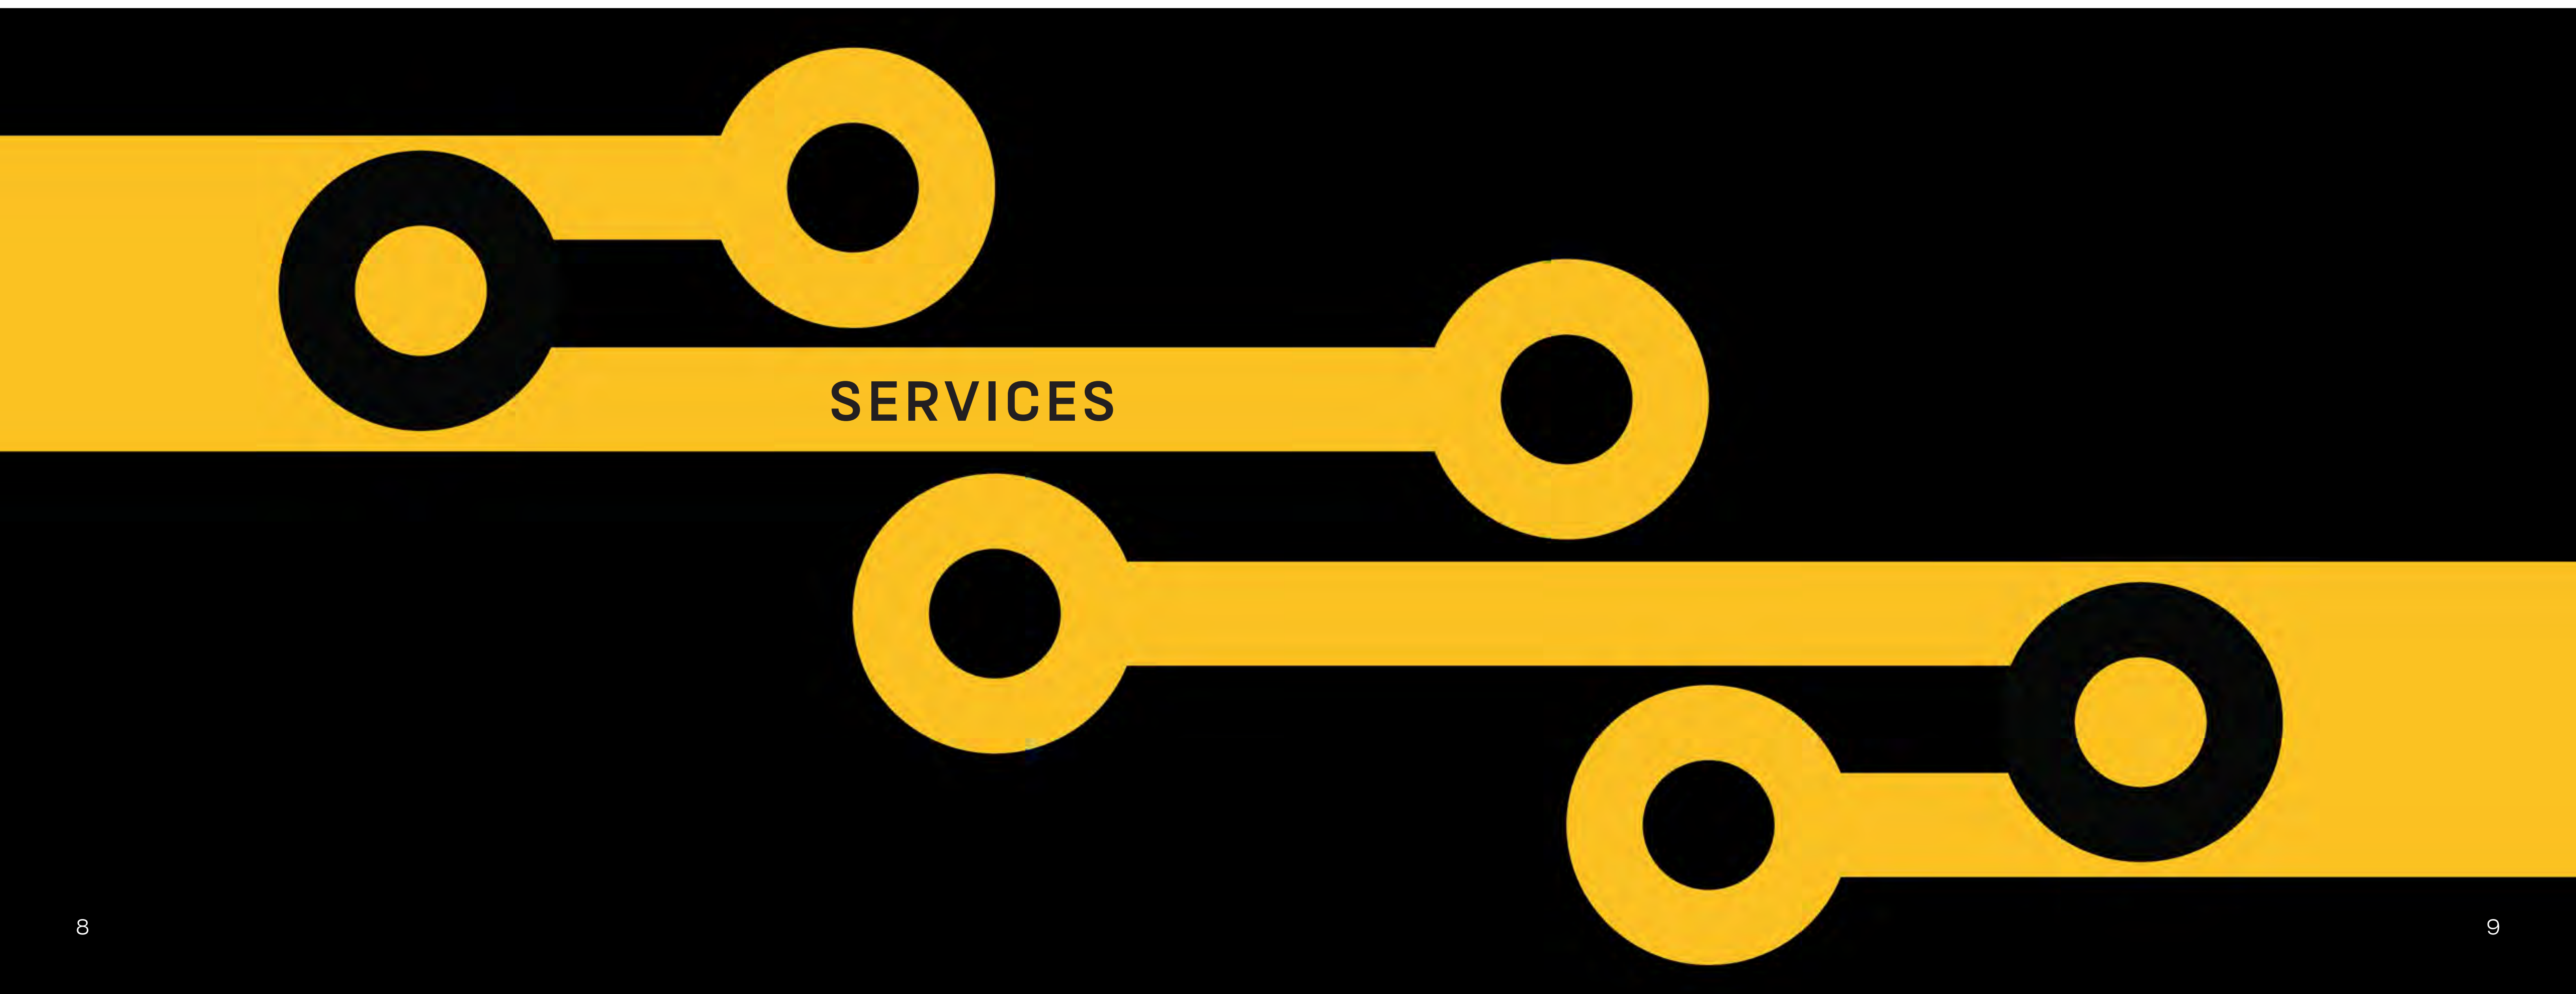

DOD High Performance Computing Modernization Program

The HPCMP was created when President George H. W. Bush signed into law the National Defense Authorization Act for Fiscal Years 1992 and 1993, on December 5, 1991. It was established under the Office of the Director, Defense Research and Engineering and in 2011 was transferred to the U. S. Army, which assigned it to the Engineer Research and Development Center in Vicksburg, Mississippi.

After more than 30 years, the Program still executes its mission to modernize the supercomputer capability of Department of Defense laboratories and test centers with large-scale computing resources; high bandwidth, low latency networks; commercial and government owned software and open source tools; and resident subject matter expertise and training. The Program was born amid rapid technological change internationally, which we continue to witness today. It has adapted to changing requirements, responded to national emergencies, and enables current and future generations of DOD scientists and engineers.

Originally focused on support for science and technology requirements, the HPCMP expanded to support DOD test and evaluation, and more recently, to the acquisition engineering community. The ever-growing set of capabilities and services include digital engineering, artificial intelligence and machine learning, virtual prototyping, virtual testing, and large-scale data analytics.

PARTNERSHIPS

Army Artificial Intelligence Innovation Institute (A2I2)

In 2019, DEVCOM ARL established the Army Artificial Intelligence Innovation Institute (A2I2) as an ecosystem of networks, data repositories, laboratories, tools, computing resources, and expertise. Its goals are to advance AI capabilities for autonomous maneuver for multi domain operations (MDO); host challenge problems addressing MDO; facilitate collaborative research with academia, industry, and government; and establish a repository of proven AI algorithms and labeled data. It addresses complex issues, such as how do intelligent systems reason about, interact with, and manipulate the environment to achieve complex military relevant actions for the future multi domain operations, and how do diverse, embodied agents collectively sense, infer, reason, plan, and execute in collaboration with Soldiers and in the face of a peer adversary?

The A2I2 Computational Facility is designed to support artificial intelligence and machine learning modeling and simulation, and storage. With construction already in progress for phases 1 and 2, the DEVCOM ARL DSRC will provide 24,000 ft^2 of floor space for future compute requirements, with TI-24 being the first production HPC resources planned to occupy the space. Phases 3 & 4 will add another 24,000 ft^2 of compute space, 40 MW of combined utility and IT power, and about 2,100 tons of cooling.

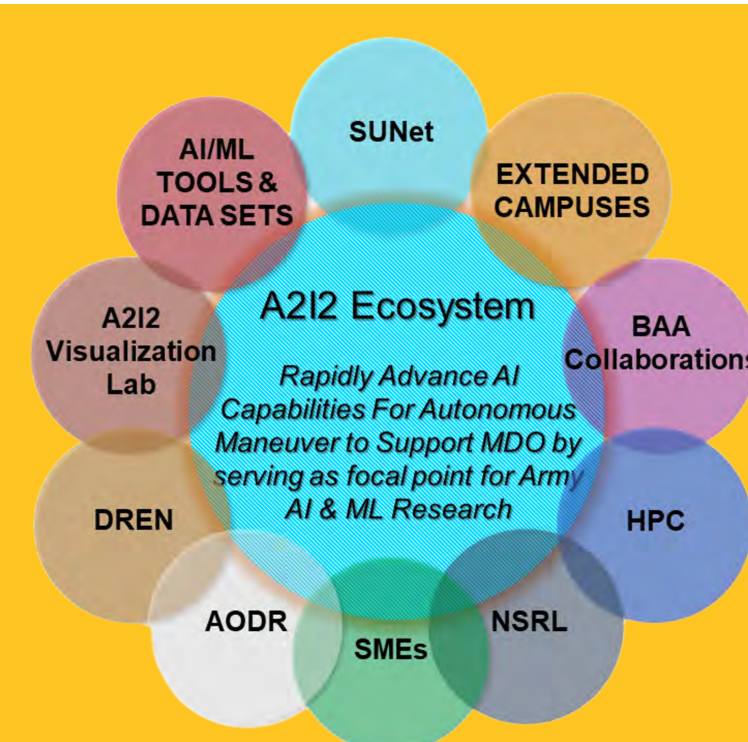

COMPUTATIONAL SYSTEMS

GPUs (Graphics Processing Units)
NVMe (Non-Volatile Memory express)
Gbit/s (Gigabit per second)

BETTY

Cray CS500
102,400 2.9 GHz AMD (EPYC) Rome Cores
446 TB System Memory
350 TB NVMe Solid State Drive
15 PB RAID Storage
GPUs: 292 NVIDIA Volta V100
Interconnect: EDR InfiniBand
Operating System: Red Hat Enterprise Linux

SCOUT

IBM Power9
22 Training Nodes
132 NVIDIA Volta V100 GPUs
11 TB Memory
330 TB Solid State Storage
128 Inference Nodes
512 NVIDIA T4 GPUs
32 TB Memory
512 TB Solid State Storage
Interconnect: EDR InfiniBand
Operating System: Red Hat Enterprise Linux

JEAN

Liqid Computing
8.8 petaFLOPS Theoretical Peak Capacity
55,392 Intel XEON Cascade Lake Advanced Performance compute cores
308 TB of memory
12.5 PB NVMe Solid State Disk
GPUs: 294 NVIDIA Ampere A100
Interconnect: 200 Gbit/s Mellanox InfiniBand
Operating System: Red Hat Enterprise Linux

KAY

Liqid Computing
4.9 petaFLOPS Theoretical Peak Capacity
46,176 Intel XEON Cascade Lake Advanced Performance compute cores
215 TB of memory
10.3 PB NVMe Solid State Disk
GPUs: 86 NVIDIA A100 Ampere
Interconnect: 200 Gbit/s Mellanox InfiniBand
Operating System: Red Hat Enterprise Linux

New In 2023

RUTH

HPE EX-4000
127,488 AMD (EPYC) Genoa compute cores
335 TB System Memory
17 TB Storage, including 2 PB of NVMe-based solid state storage
GPUs: 64 AMD MI-250
Interconnect: Slingshot
Operating System: SLES - SUSE Enterprise Server 15

MARLYN

HPE EX-4000
57600 AMD (EPYC) Genoa compute cores
135 TB System Memory
13 TB Storage, including 2 PB of NVMe-based solid state storage
GPUs: 64 AMD MI-250
Interconnect: Slingshot
Operating System: SLES - SUSE Enterprise Server 15

WHAT IS HIGH-PERFORMANCE COMPUTING?

The following pages contain HPC Success Stories from across DOD Services and Agencies. They represent workdone by DEVCOM ARL researchers at any of the HPCMP DSRCs, and work done by all users at the DEVCOM ARL DSRC.

The HPCMP has deployed the full range of supercomputers over its 30+ years of service, from vector computers and massively parallel platforms to Linux clusters and reconfigurable architectures; and from custom chips to commodity x86 multi-core processors and GPU accelerators to the recent ARM processors. The growth of the Program was contemporary with the TOP500 (top500.org), which was first published in 1993. Recognition of the TOP500 eventually had an impact on architectures as manufacturers tailored their products for TOP500 performance. Through the end of Dennard scaling and approaching sunset on Moore's Law, peak floating point operations per second has driven the high end market and is reflected in every computer the HPCMP has deployed, and in fact, every computer that was available.

Physics-based modeling and simulation was the mainstay of requirements for the first two decades and remains the predominant consumer of cycles. The rapid emergence of requirements for machine learning, large-scale data analytics, and statistical analysis has expanded the workload. The result is High End Computing, which merges the traditional workload with the new requirements and high performance computing.

The future portends big changes as noted in "HPC Forecast: Cloudy and Uncertain" published in the February 2023 issue of Communications of the ACM. The market is driven by demand and investment. Investment in big iron systems amounts to maybe $1B every 5 to 7 years; investment in mobile and cloud services is many billions of dollars every year. Future HPC will be guided by mobile and cloud providers in that they will influence what gets designed and built, in terms of processors, networks, memory, accelerators, etc. and the TOP500 will have less and less effect on architectures. To be successful, we must start thinking about how we will design, program for, and sustain code for this new and coming environment. As we reach the end of Moore's law, cloud vendors will be able to specialize processors in a way that Governments cannot. THAT could be the tipping point.

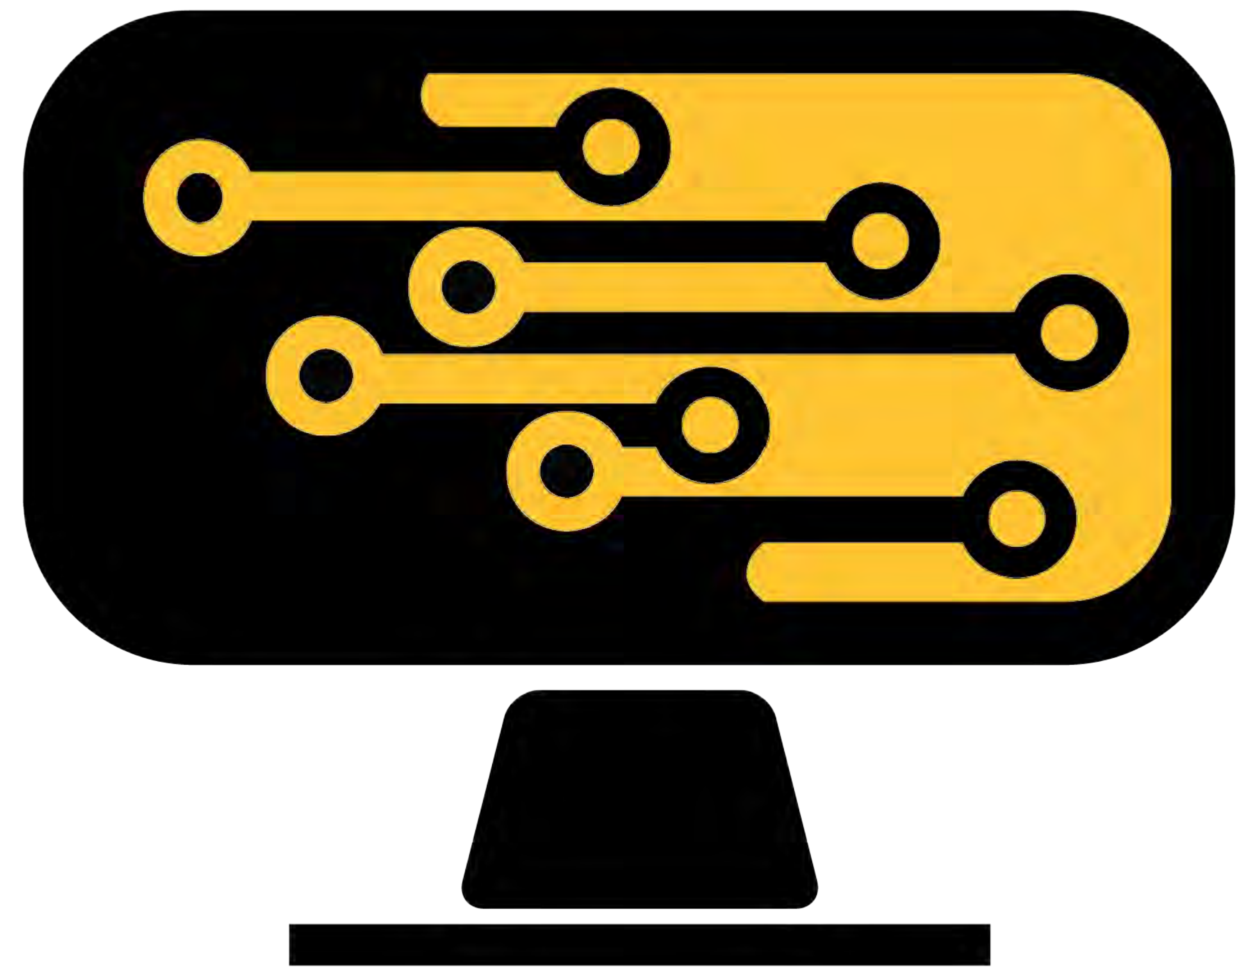

Dr. James DeSpirito

Army Research Diretorate, DEVCOM Army Research Laboratory Adelphi MD

Prediction of Vortex Interaction Effects on Missile Aerodynamics

Project Description

The control of guided missiles and projectiles for maneuverability and agility requires analysis of the complex, vortex-dominated flow field over the vehicle. Vortices are typically shed from control surfaces (e.g., canard, wing, or strake) or the surface of the vehicle itself from smooth-body separation (Figure 1). The interaction of these vortices with downstream portions of the vehicle body (e.g., tail fins) can lead to adverse aerodynamic loads that can be difficult to predict to the required accuracy at some incidence angle (σ), roll angles (λ), and Mach numbers, (M). At a minimum, these interactions can reduce the available control authority below the design requirement. At the extreme, the vehicle can attain a state in which the guidance, navigation and control system will have difficulty maintaining control of the vehicle. The goal of the current study was to assess and quantify the prediction accuracy of our high-fidelity computational fluid dynamics (CFD) tools, extend our understanding of the vortex interaction flow physics and numerical solution techniques to ultimately enhance our capability to accurately predict these flow fields. This work was performed in collaboration with partners under a NATO Science and Technology Organization (STO) Applied Vehicle Technology (AVT) Panel task group, AVT-316, “Vortex Interaction Effects Relevant to Military Air Vehicle Performance.”

Relevance of Work to DOD

The Army modernization priority Long-Range Precision Fires includes both missile and projectile systems. Increasing range for both strategic and tactical operational missions means the munition will likely be operating at times in the flight environment where vortex interaction effects will impact the maneuverability or control of the vehicle. The need to fully assess and improve the prediction capability of our high-fidelity computational analysis tools is important to design these vehicles for range and agility.

Computational Approach

The missile configurations shown in Figure 1 were analyzed using high-fidelity CFD to predict the complex, vortex-dominated flow fields and the resulting aerodynamic loads. The prediction of the roll moment coefficient value was the main metric used to evaluate the effects of mesh resolution or solver numerics on the accuracy of the prediction. The flow conditions and vehicle orientations shown in Figure 1 were chosen for study because the predicted roll moment was not as accurate as in other parts of the flight envelope. Three mesh resolutions were investigated for each missile, up to 429 million cells for the strake-tail missile and 346 million cells for the wing-tail missile. The largest mesh sizes refined the cell sizes over the strakes or wings, including the vortex trajectories past the fins. The larger mesh sizes are generally too large for standard aerodynamic characterization, but minimized vortex dissipation and provided meshes adequate to investigate turbulence scale-resolving methods, such as hybrid Reynolds-averaged Navier-Stokes (RANS)/Large-Eddy Simulation (LES).

Two modern CFD solvers were used in the investigation, Kestrel, developed under the DoD HPCMP CREATETM-AV program, and the commercial solver, CFD++, from Metacomp Technologies, Inc. Multiple turbulence models available in both solvers were investigated, including time-accurate, scale-resolving simulations. Adaptive mesh refinement (AMR) was investigated using Kestrel’s dual solver near-body, off-body capability that uses the trimmed unstructured mesh solver around the vehicle and a Cartesian solver outside this region. The largest mesh size cases used over 3000 cores and several million hours were used over the four-year investigation. Post processing was standardized among our NATO partner using Paraview batch scripts allowing direct comparison of our results.

Results

Figure 2 shows the vortex-dominated flow fields around the two missile configurations computed from Kestrel steady-state RANS, time-averaged RANS/LES, and instantaneous time-accurate RANS/LES simulations. In both cases, the steady-state RANS captures the primary vortex structures and separation locations. Comparing to the time-averaged RANS/LES simulations, the latter resolves more of the vortex structures and shows less dissipation. The instantaneous RANS/LES images show the highly unsteady nature of the flow fields. In the strake-tail missile case, the dominant effect on the roll moment prediction was found to be the tip vortex emanating from the lower starboard strake impacting near the upper starboard fin (obscured in view). The numerical dissipation of the vortex was found to be the dominant prediction error, due to the long travel of this vortex to reach the fin. In the wing-tail missile, the dominant effect on roll moment prediction was found to be the separation over the two near-horizontal fins. In both cases the turbulence model, especially over prediction of the turbulent kinetic energy, was a key factor in prediction accuracy.

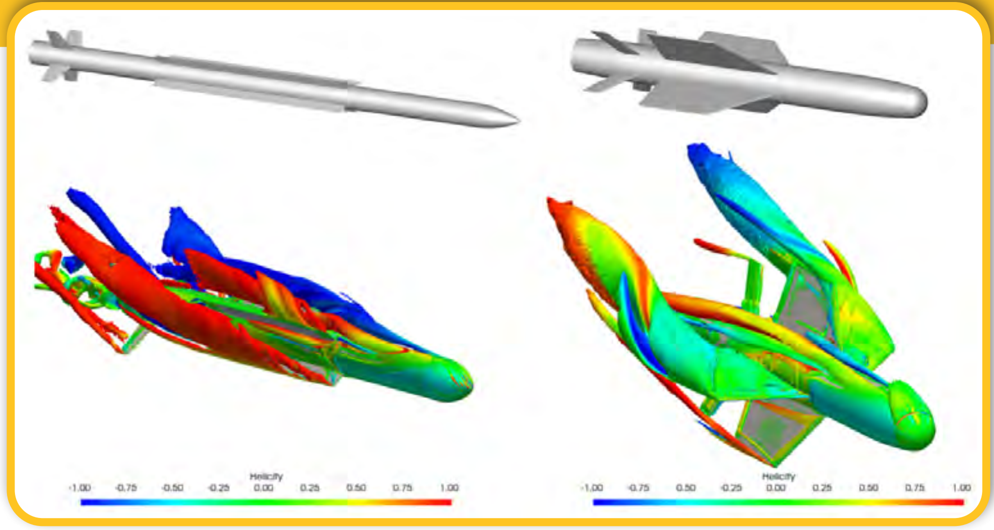

Figure 1:
(Top) Generic strake-tail missile (left) at M=1.4, σ=15°, λ=2.5°, and wing-tail missile (right) at M=0.85, σ=17.5°, and λ=45°. (Bottom) Steady-state flow fields computed using RANS CFD in fine mesh illustrating vortex-dominated flow fields, visualized via Q-criterion isosurfaces colored by helicity indicating vortex direction and strength.

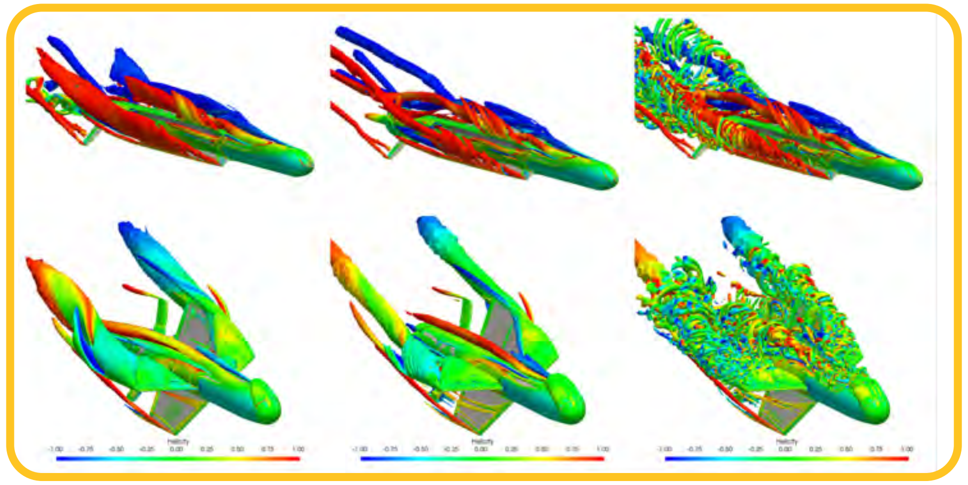

Figure 2:
Strake-tail missile (top) and wing-tail missile (bottom) flow fields from fine-mesh, steady state RANS (left), time-averaged (middle) and instantaneous (right) RANS/LES simulations.

In general, time-accurate RANS/LES provided the most accurate result but did not guarantee a close match with experimental data. These time- and compute-intensive simulations are not generally used in aerodynamic characterization of these vehicles. However, lessons learned from this study includes identifying the points of the flight envelop, such as the two shown here, where additional investigation would be warranted to confirm the predictions from standard prediction methods such as RANS simulations. These points would be identified as those where vortex structures pass very close to a wing or fin, or where intense separation is occurring.

Future

Work is continuing in a follow-on NATO task group, AVT-390, “Vortex Flow Predictions for Stability and Control of Missile Airframes.” Lessons learned will be applied to further attempt to understand the current numerical limitations and improve the prediction accuracy. Additional wind tunnel data are now available that include particle image velocimetry and surface oil flow images to allow better qualitative and quantitative comparison with experimental data. Additional flow visualization techniques will also continue to be used to extract the nature of the flow phenomena.

Co-Investigators

Partners in NATO STO AVT-316 included members from *Corvid Technologies (USA), NASA LaRC (USA), MBDA-UK, MBDA-France, DLR (DEU), FOI (SWE), TÜBİTAK SAGE (TUR), and University of Glasgow (UK).*

Publications

AVT-316 Final Report, “Vortex Interaction Effects Relevant to Military Air Vehicle Performance,” STO-TR-AVT-316, February 2023.

Taylor, N. J., “AVT-316 Missile Facet: Lessons learned concerning the prediction of Vortex Flow Interactions about Generic Missile Configurations,” AIAA SciTech 2022 Forum, AIAA-2022-0002, Jan. 2022.

Park, M. A., and DeSpirito, J., “The Influence of Adaptive Mesh Refinement on the Prediction of Vortex Interactions about a Generic Missile Airframe,” AIAA SciTech 2022 Forum, AIAA-2022-1177, Jan. 2022.

DeSpirito, J., et al., “Comparisons of Predicted and Measured Aerodynamic Characteristics of the DLR LK6E2 Missile Airframe (Scale Resolving),” AIAA SciTech 2022 Forum, AIAA-2022-2308, Jan. 2022.

Dr. Manoj K. Shukla
US Army Engineer Research and Development Center, Environmental Laboratory, Vicksburg, MS

Multiscale Modeling of Polyamide 6 Using Molecular Dynamics & Micromechanics

Project Description

Polyamide 6 (PA6) is a semi-crystalline thermoplastic used in many engineering applications due to high strength, good chemical resistance, and excellent wear/abrasion resistance. The mechanical properties of semi-crystalline PA6 are dependent on two forms of crystallinity (α/γ). The effect of these two forms of crystallinity on the mechanical properties is largely unknown as most studies in the open literature do not report quantitative data on mechanical properties and the associated crystallinity. The microstructure of semi-crystalline polymers is hierarchical and multiscale.

While molecular dynamics (MD) simulation is a reliable and efficient tool for predicting mechanical properties of bulk polymers and their composites, it is unable to simulate multiscale structures due to the computational rigor. However, the NASA Glenn Research Center-developed micromechanics code "Micromechanics Analysis Code based on the Generalized Method of Cells (MAC/GMC)" can easily model the hierarchical, multiscale microstructure of semi-crystalline polymers. In this work, we predicted mechanical properties as a function of crystallinity and crystal form of PA6 using a multiscale modeling framework.

Relevance of Work to DOD

Lightweight composite materials are needed for force-protection and force-projection in the current and future complex operational environments. There are many semi-crystalline thermoplastics with thermo-mechanical properties relevant to DoD missions and priorities. This multiscale modeling framework (Figure 1) can be used to predict mechanical properties of other semi-crystalline thermoplastics as a function of crystallinity. Comprehensive studies of the effect of crystallinity and crystal forms on the mechanical properties of this class of polymers do not yet exist. The stiffness of semi-crystalline polymers is dependent on the crystallinity: as the crystallinity increases, the stiffness typically increases, but the degree of increase is largely unknown.

Quantifying this change is crucial when designing components using semi-crystalline polymers, since the mechanical properties of the material can be tuned for the specific application. These semi-crystalline polymers can be further improved by nanofillers such as carbon, boron nitride, and other 1D and 2D nanomaterials. These semi-crystalline polymer composites have stiffness-to-weight ratios greater than those of state-of-the-art metal alloys and thus enable significant weight savings in numerous DoD applications such as airplanes, automobiles, naval vessels, drones, personal protective equipment, etc. Additionally, the addition of certain nanofillers such as transition metal dichalcogenides to semi-crystalline polymers can create multifunctional materials such as structural components that can also perform as sensors.

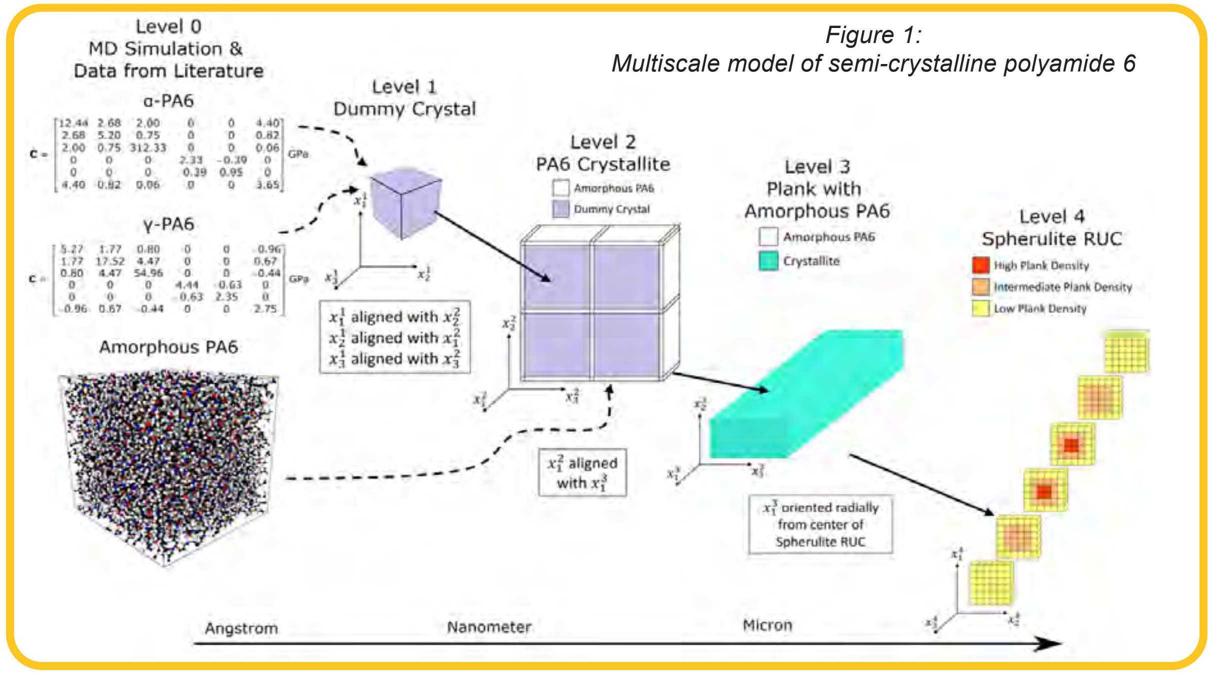

Figure 1:
Multiscale model of semi-crystalline polyamide 6

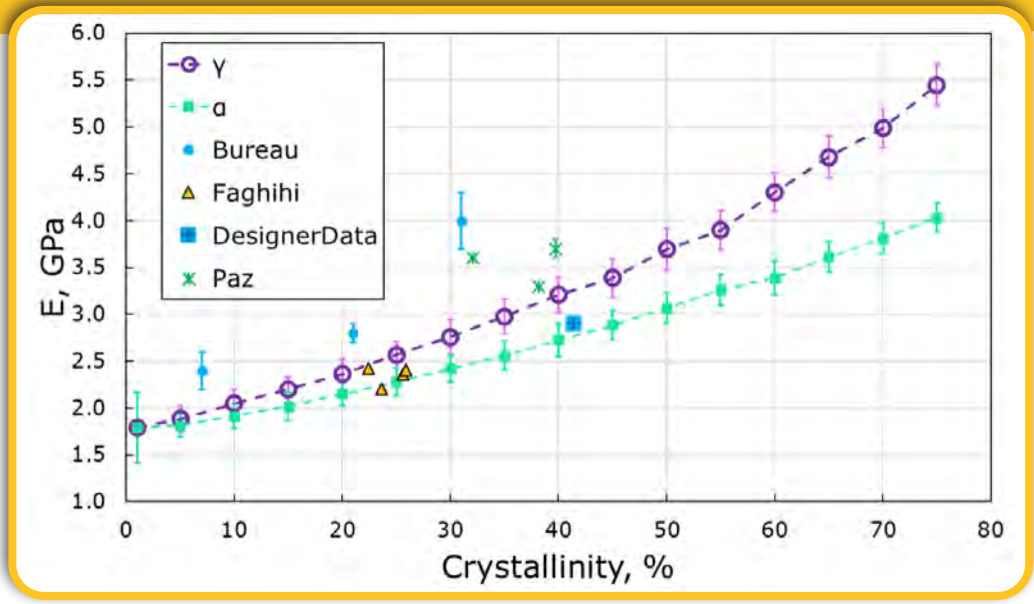

Figure 2:
Elastic modulus of semi-crystalline PA6 as a function of crystallinity and crystal form
(taken from ERDC TR-23-4; https://doi.org/10.21079/11681/46713).

Computational Approach

A semi-crystalline thermoplastic is ultimately composed of two phases: an amorphous phase and a crystalline phase. Amorphous polymer can be likened to spaghetti: chains of polymer crisscrossing randomly. Crystalline polymer is composed of chains of polymer folded in a repeating, ordered fashion. These two phases combine to form hierarchical, multiscale microstructures. This microstructure from smallest to largest consists of granular substructures, lamellae stacks, and spherulites. These real-life structures correspond to the PA6 crystallite, plank with amorphous PA6, and spherulite RUC model structures in Figure 1, respectively.
In this work, the amorphous PA6 was simulated using MD simulations with DoD HPC resources while the crystalline PA6 properties were taken from literature. DoD HPC resources enabled the use of larger model sizes, more statistical samples, and reduced the total wall-clock time. The DoD HPC resources Excalibur and Onyx were used. The mechanical properties predicted by the amorphous PA6 MD simulations include elastic modulus, shear modulus, and Poisson's ratio. These three properties along with the stiffness matrices of the two crystal forms from literature were used as input into the multiscale model.

Results

Figure 2 shows how the elastic modulus of semi-crystalline PA6 is affected by crystallinity and crystal form. The overlaid experimental data show that the multiscale model results are accurate, thus validating our modeling protocol. Figure 2 shows that PA6 materials with more γ crystals will have greater stiffness than if there were more α crystals. Crystal form and crystallinity can be controlled via manufacturing methodology. Therefore, knowledge in Figure 2 can be used to manufacture PA6 materials at certain crystallinities with a certain crystal form for specific DoD applications.

Future

From here, this multiscale modeling framework can be applied to other semi-crystalline thermoplastics to generate comprehensive studies on the effect of crystallinity and crystal structure on the mechanical properties of the semi-crystalline polymer. These studies will enable the manufacture of semi-crystalline materials with precise properties for specific applications. Generally, very little data is available on the mechanical properties of the crystalline phase of these polymers, and, depending on the polymer, the same may be true for the amorphous phase. Therefore, DoD HPC resources will be instrumental to predict those missing properties.

Co-Investigators

Dr. William A. Pisani *(Oak Ridge Institute for Science and Education, Oak Ridge, TN)*
Dr. John K. Newman *(CUS Army Engineer Research and Development Center, Geotechnical and Structures Laboratory, Vicksburg, MS)*

Publications

Pisani, W., Wedgeworth, D., Roth, M., Newman, J., & Shukla, M. (2023). Exploration of two polymer nanocomposite structure-property relationships facilitated by molecular dynamics simulation and multiscale modeling. US Army ERDC Technical report. *https://doi.org/10.21079/11681/46713*

Pisani, W. A., Newman, J. K., & Shukla, M. K. (2021). Multiscale Modeling of Polyamide 6 Using Molecular Dynamics and Micromechanics. Industrial & Engineering Chemistry Research, 60(37), 13604–13613. *https://doi.org/10.1021/acs.iecr.1c02440*

Dr. Manoj K. Shukla
US Army Engineer Research and Development Center, Environmental Laboratory, Vicksburg, MS

Computational Modeling of Per-& Polyfluoroalkyl Substances

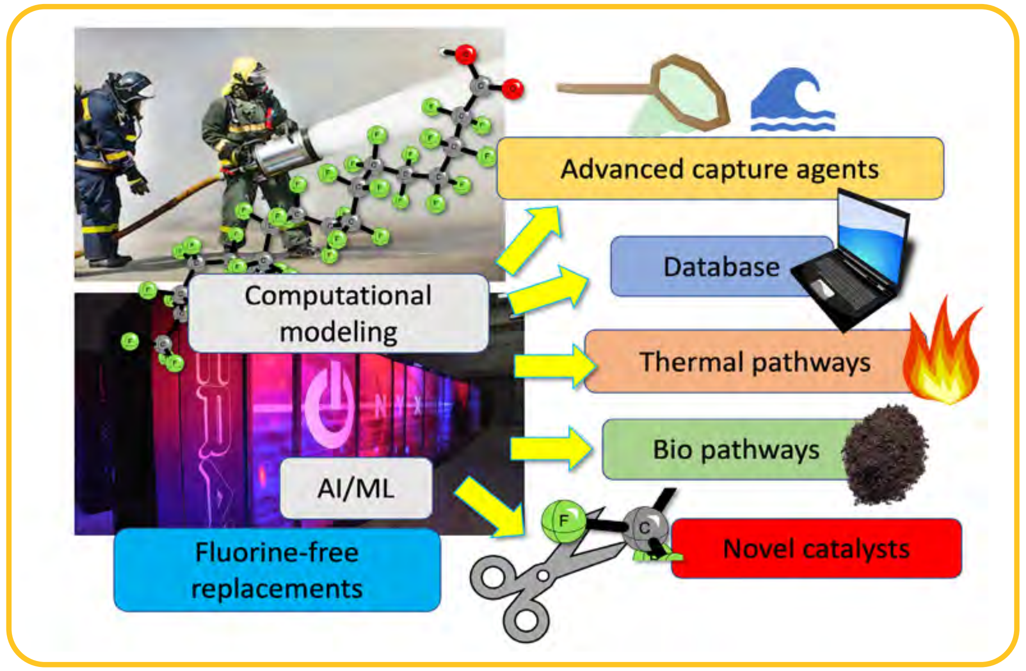

Figure 1: Overview of computational approach to PFA

Project Description

Per- & polyfluoroalkyl substances (PFAS) constitue a large family of manufactured industrial chemicals developed in the 1940s with diverse applications in almost every aspect of our life. These chemicals, classified as emerging contaminants, are detected around the globe, and nicknamed "forever chemicals" due to their resistance to degradation. As per Environmental Working Group (www.ewg.org; access on April 24, 2023), as of June 2022, the 2,858 locations in 50 states and 2 territories in the USA were known to be contaminated with PFAS. These compounds are potentially toxic to living organisms and bind to the transport protein serum albumin and thus accumulate in liver, blood, and other organs.

Recently, the Biden-Harris Administration and U.S. Environmental Protection Agency (EPA) proposed a national drinking water standard for perfluorooctanoic acid (PFOA) and perfluorooctane sulfonic acid (PFOS) as well as mixtures of several PFAS. According to the US National Health and Nutrition Examination 2011-2012 survey, detectable serum PFAS concentrations were revealed in 97% of tested individuals. A viable technology for PFAS remediation at the commercial level is not available yet, though several technologies are at various stages of development and testing. We combine computational chemistry and artificial intelligence (AI)/machine learning (ML) techniques on the DoD High Performance Computing (HPC) resource Onyx to study degradation, destruction, isolation, and removal of PFAS from the contaminated environments.

Relevance of Work to DOD

PFAS have been used heavily in the military within aqueous film forming foams for fire training and emergency response purposes. Several military sites are contaminated with PFAS. The DoD has a potential liability of several billion dollars associated with PFAS contaminated sites. There are several programs within DoD for PFAS regarding the development of technologies for contaminated site assessment, remediation, and development of fluorine-free fire-fighting foams.

Computational Approach

The regular and periodic Density Functional Theory (DFT) approaches are used to compute the potential degradation of select PFAS compounds by reactive species, unraveling potential degradation and destruction mechanisms. Molecular dynamics (MD) simulations are used to study the adsorption of select PFAS molecules on some clay surfaces and different carbon-based nanomaterials, as well as their propagation through lipid bilayers. Moreover, an AI/ML approach has been used to classify several thousands of PFAS and related compounds to establish structure-property relationships that could be used for remediation purposes.

Results

Our rigorous DFT calculations on the catalytic mechanism of silylium-carborane chemistry for PFAS degradation revealed that PFAS molecules have a low-lying anti-bonding orbital. This orbital is occupied in the presence of silylium and the addition of carborane enhances this occupation, which leads to a weakening of the C—F bond and subsequent degradation.

Further, aqueous redox induced chemical degradation reaction mechanisms for PFAS at the DFT level were found to depend upon oxidation or reduction environments and the chemical structure of PFAS under investigation. Using DFT, we have also investigated the gas-phase and in water solution degradation of trisiloxane surfactants, PFAS alternatives in firefighting foams. Our computation revealed that trisiloxanes would readily decompose through conventional atmospheric oxidation, thermal treatment, and acidic/basic hydrolysis highlighting its ability to degrade in the environment. Using MD simulation, we probed the intra-membrane interactions of lipids with PFAS and other

Future

Technologies for the degradation, destruction, detection, isolation, removal, and sensing of PFAS are at different stages of active development. Despite the current research activities on PFAS, we still have very limited information that can help in developing a robust remediation technology that could be applied across a large group of PFAS compounds. Moreover, given the large number of PFAS and PFAS-related compounds, it would not be possible to general physico-chemical property data using experimental approaches. Therefore, computational chemistry approaches along with DoD's HPC resources offer a powerful alternative to generate such data that could be used in developing treatment technologies for PFAS.

Co-Investigators

Dr. Glen Jenness, Dr. Harley McAlexander, Dr. Timothy Schutt, and Dr. Caitlin Bresnahan *(Environmental Laboratory, Vicksburg, MS)*
Dr. Robert Lamb and Dr. Brain Etz *(Oak Ridge Institute for Science & Education, Oak Ridge, TN)*
Dr. Ashlyn Koval *(Simetri, Inc., Winter Park, FL)*

Publications

PFAS chemical degradation strategies: insights into the underlying reaction mechanisms, Brian D. Etz and Manoj K. Shukla, Current Opinion in Chemical Engineering (Submitted)

High-Temperature Decomposition Chemistry of Trimethylsiloxane Surfactants, a Potential Fluorine–Free Replacement for Fire Suppression, Brian D. Etz, Maleigh Mifkovic, Shubham Vyas and Manoj K. Shukla, Chemosphere 308 (2022) 136351. *https://doi.org/10.1016/j.chemosphere.2022.136351*

Atomic Insights into the Hydrodefluorination of PFAS Using Silylium Catalysts, Glen R. Jenness, Ashlyn M. Koval, Brian D. Etz, and Manoj K. Shukla, Environmental Science: Processes & Impacts 24 (2022) 2085-2099. *https://doi.org/10.1039/D2EM00291D*

Properties and Mechanisms for PFAS Adsorption to Aqueous Clay and Humic Soil Components, Charles M. Luft, Timothy C. Schutt and Manoj K. Shukla, Environmental Science and Technology 56 (2022) 10053-10061; *https://doi.org/10.1021/acs.est.2c00499*

*Im*pact of Submicellar Aggregation on Reduction Kinetics of Perfluorooctanoate by the Hydrated Electron, William A. Maza, Brian D. Etz, Timothy C. Schutt, Brian L. Chaloux, Vanessa M. Breslin, Bradford B. Pate, Manoj K. Shukla, Jeffrey C. Owrutsky, and Albert Epshteyn, Environ. Sci. Technol. Lett. (2022), 9, 226–232. *https://doi.org/10.1021/acs.estlett.1c01020*

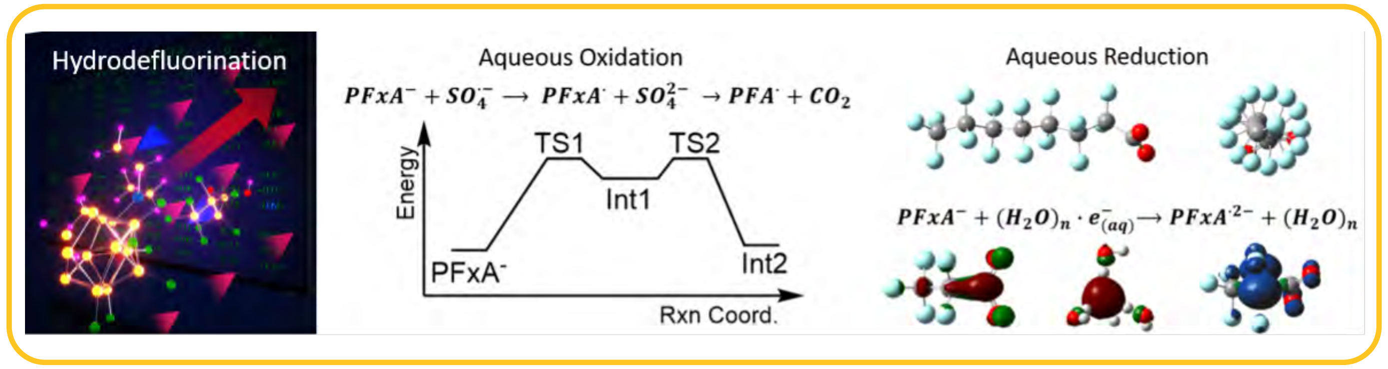

Figure 2: DFT methods were employed to investigate PFAS chemical degradation approaches involving hydrodefluorination, oxidation, and reduction.

Robert Oshiro
Parson Corporation

High Performance Computing to Improve Biometric Evaluations

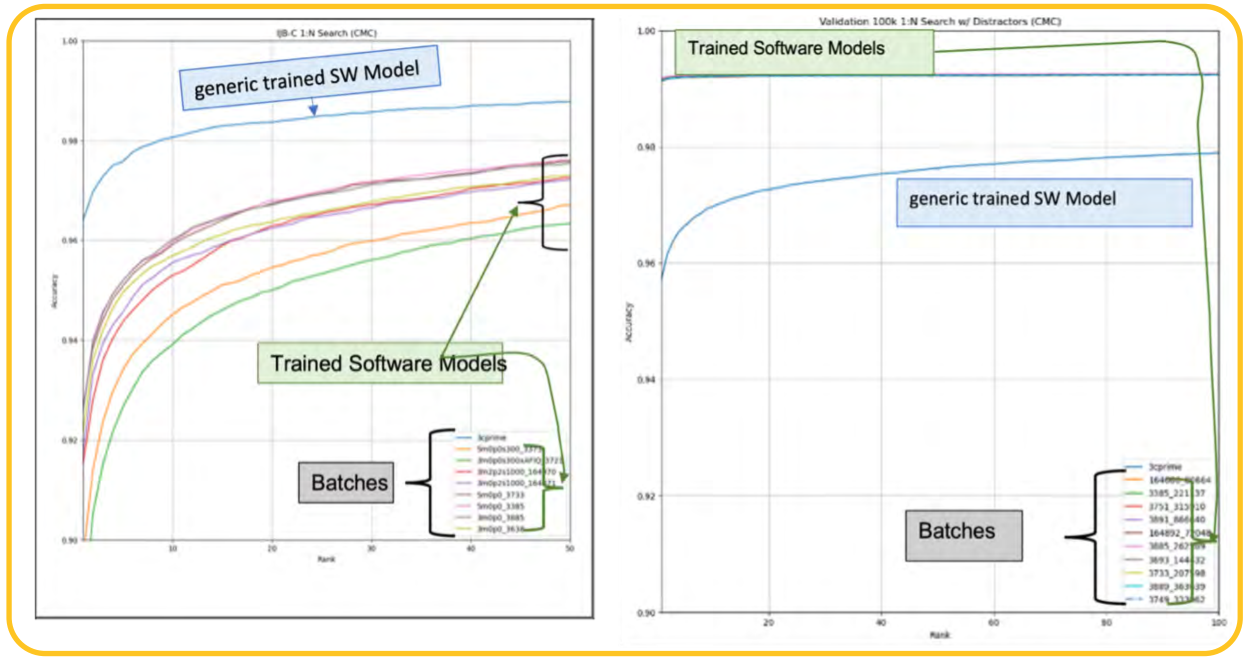

Figure 1:
(Left) Using generic image data, "Trained models" vs generic trained Model;
(Right) Using collection data, shows "Trained models" improvement vs generic trained Model

Project Description

The high performance computing (HPC) image processing project selected biometric facial image identification to assess HPC enhancements. Most biometric collection databases consist of biometric images and demographic information originally stored as binary files. The image processing system correlates biometric images, usually without regard to demographic information. When a biometric collection occurs, a binary file is built and sent to the biometric evaluation system, where images and demographic information are extracted and compared to the biometric database. Facial images are processed through automated recognition and Identification software (Identification matches are sometimes manually verified).

The goals of the biometric image processing project were to evaluate how HPC environments can contribute to training image software processing models, improve the image count limitations found during previous machine learning testing and determine HPC limitations. The image processing software used was produced under the IARPA Janus project. The Janus software provided the machine learning (ML) image processing baseline and data for facial recognition.

Relevance of Work to DOD

The Federal Government relies heavily on image processing to secure and defend our homeland.
Object recognition and Identification are used by most overwatch systems, including vehicle (ground and air) protection and overwatch surveillance (local and space) for personnel identification. Many systems today use large image and attribute databases to distinguish and filter objects. Moving objects detected by protection systems use speed detected and object size to select how to best protect the platform, and personnel detected can be voluntary and non-voluntary. Voluntary biometric collections can contain more extensive demographic information. In all cases, speed and accuracy are essential and constantly being improved.

Computational Approach

Shaping complex image processing software and then evaluating improvements is a complex task that requires enormous CPU, GPU, and storage capacity. In addition, diverse pose and demographic attributes further complicate the processing. After the raw binary files were received, we utilized three HPC systems to extract data from binary files, train image processing software using machine learning, and formulate improvements.

Of the three HPC systems utilized, Centennial and SCOUT are shared systems within the DoD High Performance Computing Modernization Program, while Sofia was dedicated to the project. SCOUT and Sofia are IBM Power9 platforms designed for ML workloads and contain specific nodes for training and inferencing. A strategy using these systems was designed to conduct training and assess performance of the Janus facial recognition code developed through previous research and development efforts.

All data were stored on the center wide file system (CWFS). Final Janus trained models was configured to run in a Singularity container in the HPC environment. Centennial provided a CPU integration platform for the Janus software while SCOUT and Sofia were used for CPU/GPU configurations. Use of Centennial was discontinued early in the effort due to the superior performance of the Power9 architectures for this application. GPU nodes were used during training, while inference nodes were used for verification and testing of the trained software. The CWFS provided 3.3 PB bulk storage accessible from all ARL HPC systems. A benefit of SCOUT over Sofia is the direct link from compute nodes to the CWFS, whereas with Sofia, data had to be staged prior to computations. The HPC systems can accommodate many simultaneous tasks. Our evaluation processed 13.5 million biometric binary files containing 22 million facial images of various poses. To prepare image data for training and testing, Bash scripts were configured to form a workflow pipeline. This pipeline was re-configured by the HPC scheduling software, running 270 instances concurrently (assigning HPC resources as they came available). The final output consisted of an image database aligned to the identity, pose, and demographic data.

Results

Machine learning models were built by training the image software using a wide variety of images and demographics, to render the models as insensitive as possible to variations in new input images. After setting up the biometric collection database, training was run on image data sets of various sizes (0.3M, 1.3 M, and 2.3 M identities). The training auto- evaluates the results and adjusts parameters. Training was conducted iteratively greater than 20 times, until target accuracy was achieved.

After a training was complete, an evaluation of each trained model was run on the original collection of images and on generic image data.

Our project answered three main questions regarding using high performance computing for biometric evaluations:

- Can HPC be used to train machine learning models faster than using desktop systems and clusters?
- Can it be used to train models on very large datasets, even the complete image set?
- Are the resulting models trained using the complete image set more accurate than models trained on subsets of the images?

The unequivocal answer to all three is yes. We trained on the full set of images, whereas previous work was limited to 300K images. Even using the full image set, HPCs reduced the training time by an order of magnitude, from 2 – 3 weeks to 2 days. The resulting models were more accurate than those trained with 300K images. High performance computing is a valuable tool for biometric evaluation and will only grow in importance as the volume of data grows, and as multi-mode identification is required.

Future

Biometric data collection is not limited to image data; it also includes other modalities, such as gait, fingerprint, and iris. The amount of data grows exponentially and more complex algorithms require even more computational resources. In addition, considerably more work needs to be done with demographic analysis of biometric data and the limitations in currently available datasets.

Co-Investigators

Patrick Rauss *(Army Research Laboratory);*
Jakob Adams, James Gabberty, Jamie Johnson *(Parsons Corporation);*
Dr. David Richie *(Brown Deer Scientific);*
Prof. Rama Chellappa & Joshua Gleason *(John Hopkins University).*

WORKFORCE OF THE FUTURE

The High Performance Computing Modernization Program sponsors the annual HPCMP Internship Program as part of workforce development, supporting undergraduate and graduate students. It is a mentor driven process where mentors submit proposals to the HIP team and successful mentors then select interns. Interns are expected to write a paper and prepare and present a brief at the end of the internship. The DOD and its contractors have succeeded in hiring greater than 20% of the interns. ARL submits proposals every year and has been quite successful, supporting 66 students over the last 5 summers.

Candidate HIP interns are US citizens exploring science, technology, engineering, and mathematics (STEM)-related career fields. However, the pool of STEM students who are US citizens and are interested in pursuing federal employment continues to dwindle due to competing opportunities offered by industry. Additionally, the country continues to see a relative decline in the number of US citizens enrolled in science and engineering undergraduate and graduate schools. It is essential to develop interest and computational skills earlier in the student pipeline, in high school, before students decide on an academic career path. In 2020, the ARL DSRC submitted a proposal to the HIP team to extend the STEM workforce pipeline. In addition to working with local schools to recruit high school interns, the project included engaging and training teachers to work with computer scientists who will impart their new knowledge and experience in their own classrooms. Although this proposal was not a typical HIP proposal, which encompasses computational technologies and high performance computing, Dr. Kevin Newmeyer, Deputy Director of the HPCMP, funded the proposal. Three teachers and four high school students from the surrounding counties were supported by that non-traditional internship.

We are pleased to feature the work of four of the interns from the summer of 2022. The students are also co-authors on articles that appeared in the March and June 2023 issues of the ITEA Journal.

Richard Healy
Rensselaer Polytechnic Institute

Aerodynamic Simulation for Reconfigurable UAS Performance Improvement

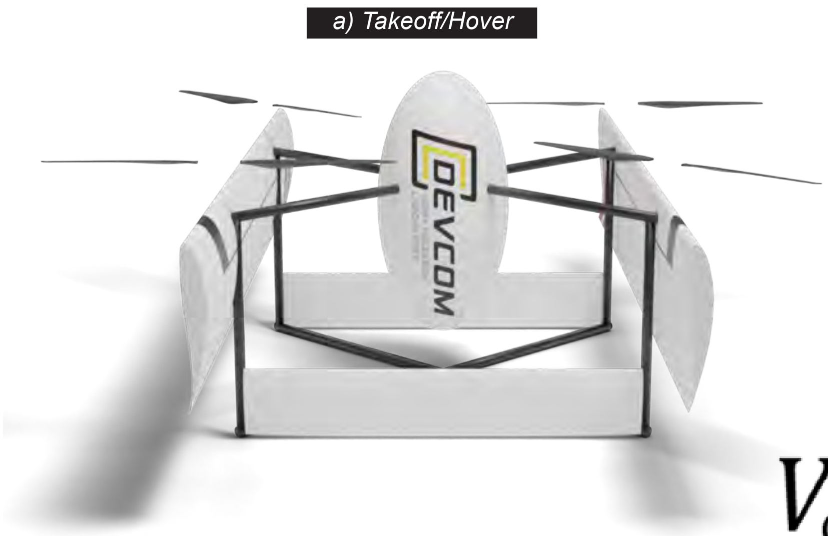

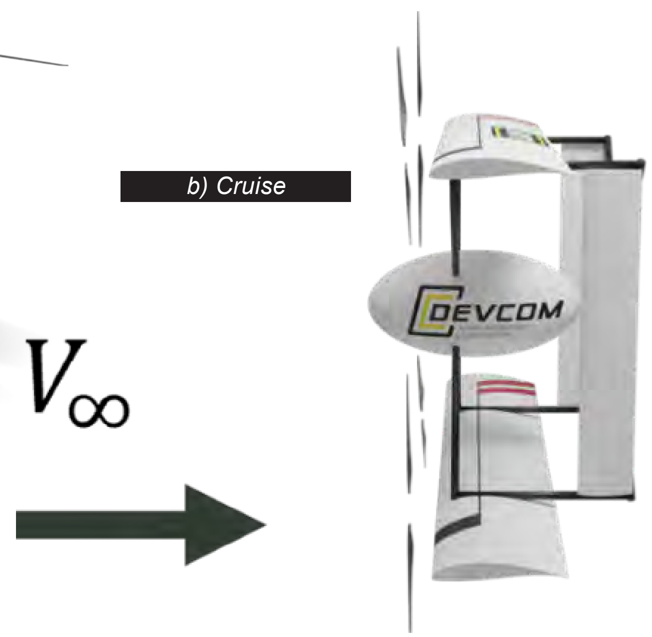

Figure 1: CRC-20 Flight Modes. The hover (left) and cruise (below) orientations of the CRC-20

Project Description

In recent years, a range of Unmanned Aerial System (UAS) technologies have been developed to fulfill the goals of the Army's multi-domain operation concept. In this study, several modifications are considered for the CRC-20, Figure 1, a quadrotor biplane tailsitter UAS. As an alternative to variable blade pitch, rotor tilt is investigated as a method for improving rotor performance in cruise. Mid-fidelity aerodynamic models of the baseline and modified CRC-20s are developed using the DoD's RCAS software. These models are run in parallel using high performance computing resources to quickly evaluate vehicle performance across a range of flight conditions. With the merit of a tilted rotor configuration established, a CFD model of the vehicle is developed using the DoD's HELIOS software.

Relevance of Work to DOD

Modern advancements in electronics and battery technology have made UAS highly capable for a range of missions including surveillance, reconnaissance, and resupply. Their flexibility, among other qualities, has led researchers at the Army Research Laboratory to develop a range of UAS technologies to fulfill the goals of the Army's multi-domain operation concept. Electrification has also enabled the development of reconfigurable UAS that can adapt how they operate to maximize their effectiveness in different mission roles or flight conditions. One configuration enables vertical takeoff and landing capabilities with significantly less mechanical complexity than more traditional tiltrotor configurations.

Computational Approach

Mid-fidelity simulation models are developed for the baseline and tilted-rotor CRC-20s using the DoD's RCAS software. For a given flight speed and aircraft design, the RCAS model solves for the rotor thrust and vehicle orientation needed to achieve zero vehicle acceleration (trim the vehicle). The baseline CRC-20 is simulated either varying rotor RPM at a fixed blade pitch angle, or varying blade pitch angle at a fixed RPM. All RCAS simulations were performed using DoD High Performance Computing (HPC) Modernization Program resources. Trim sweeps were performed with each flight speed / vehicle configuration running in parallel on a single node with 128 processors.

A computational fluid dynamics (CFD) model of the CRC-20 was developed using the DoD HELIOS software. HELIOS uses an overset grid where near-body flow is solved using a separate solver from the off-body flow. Near-body solvers used included FUN3D and mStrand. An isolated rotor in propeller mode is simulated both by fully resolving the blade with FUN3D

Figure 2. Tilted rotor configuration power. Power required by the tilted-rotor configuration at varying rotor tilt angles and flight speed

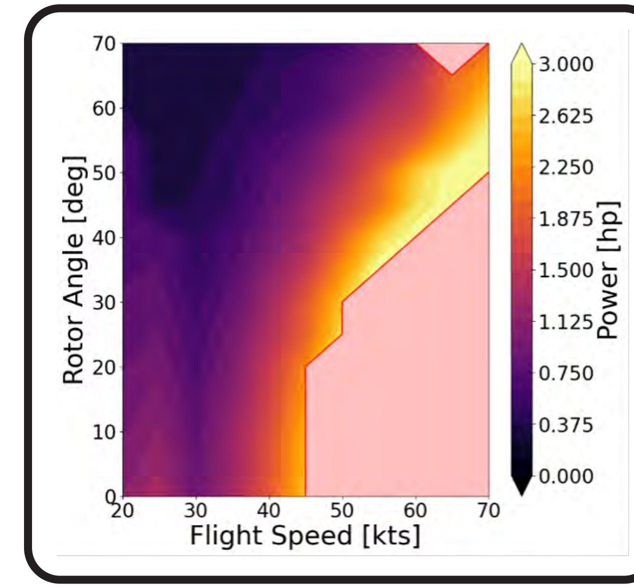

Results

The baseline CRC-20 is simulated in RCAS at flight speeds ranging from 20 kts – 70 kts. At each flight speed, the vehicle is trimmed and the associated rotor power requirement is evaluated. The baseline CRC-20 with variable RPM shows nonlinearly increasing power requirements as the flight speed increases. This is because as flight speed increases, the rotor RPM must be increased to produce sufficient thrust, thereby increasing both induced and profile power. The maximum achievable speed with variable RPM is also constrained to 50 kts, with the rotors unable to generate sufficient thrust at higher flight speeds. If variable blade pitch is used instead, the power required at speeds above 30 kts is reduced by up to 61%, and the maximum achievable speed is extended beyond 70 kts. Alternatively, if the tilted rotor configuration is used (even with fixed blade pitch), performance improvements are like the variable pitch configuration. These results suggest that employing tilted rotors on the CRC platform is a viable alternative to a mechanically complex variable blade pitch mechanism.

To test how the rotor tilt angle affects performance, the tilted rotor configuration is trimmed at flight speeds ranging from 20 kts – 70 kts, and with rotor angles ranging from axial (α= 0°) to α= 70° from the vertical. Figure 2 plots the power required for each combination of flight speed and rotor tilt angle, with light red areas denoting unachievable conditions. At 0° tilt angle (rotors in propeller mode), the maximum endurance flight speed is 30 kts. At 30 kts, as rotor angle increases (and it is transitioned from axial flow to edgewise flow), the power required drops by up to 72% at 70°. Alternatively, α= 70° rotors can achieve 55 kts at the same power required by untilted rotors at 30 kts. Rotor tilt also extends the flight envelope, which is constrained below 45 kts at 0° but extends beyond 70 kts at 60°. Introducing rotor tilt improves vehicle performance in terms of power required and top speed at all flight conditions tested.

The tilted rotor configuration is simulated using HELIOS CFD. An isolated rotor and isolated wing are also simulated independently to give points of reference for aerodynamic performance without aerodynamic interaction. All rotors are simulated using FUN3D as the near-body solver and all wings are simulated using the mStrand solver. Wing-induced downwash on the rotors is found to reduce total system lift by 3.91%. However, this downwash also reduces the rotors' induced drag by up to 11.5%, leading to an overall performance improvement of 16.8% (based on total L/D_e).

Future

A modified CRC-20 UAS with tilted rotors is investigated to improve reconfigurable UAS performance. Despite a rotor thrust deficit, decreases in rotor induced drag led to an overall 16.8% performance increase. These findings suggest that tilted rotors, even with their interactional aerodynamics, have the potential to offer substantial performance improvements over the baseline CRC-20.

HIP Mentors

Dr. Matthew Floros and Dr. Phuriwat Anusonti-Inthra, *(ARL Army Research Directorate)*

Publications

R. Healy, P. Anusonti-Inthra, M. Floros and H. Kang, "Air-Launched Effect Trajectory Prediction using Computational Fluid Dynamics," ITEA Journal, June 2022.

M. Misiorowski, F. Gandhi and P. Anusonti-Inthra, "Computational Analysis of Rotor-Blown-Wing for Electric Rotorcraft Applications," AIAA Journal, vol. 58, p. 2921–2932, July 2020.

Brendan Smith
Rensselaer Polytechnic Institute

Rotor Broadband Noise Modeling and Propeller Wing Interaction

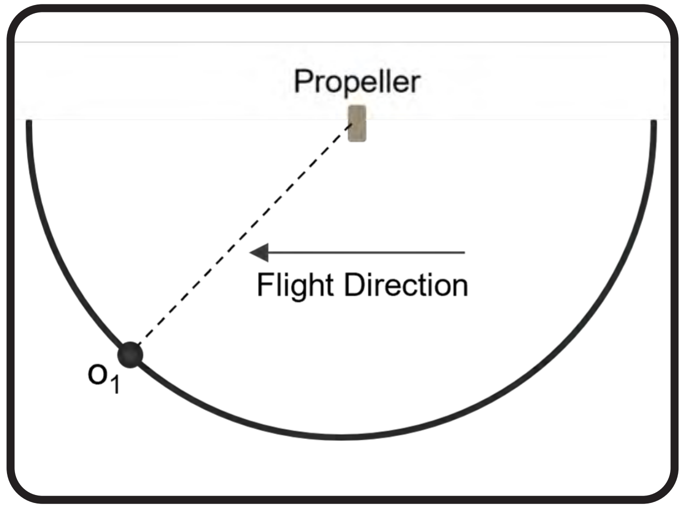

Figure 1: Side view of propeller, showing single observer and observer hemisphere locations

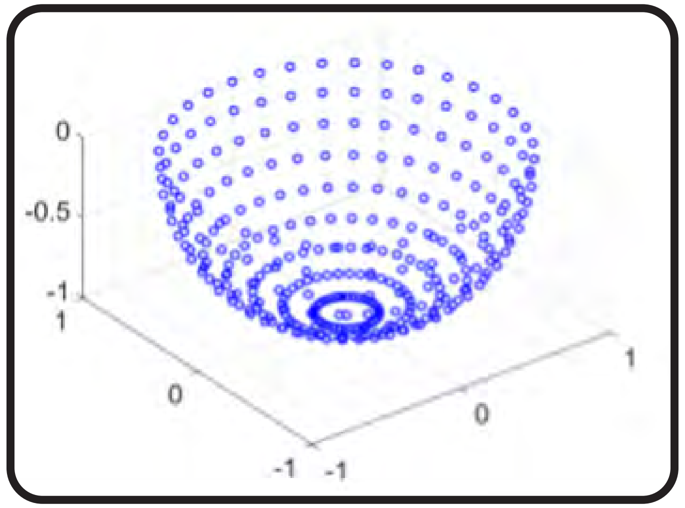

Figure 2: Hemisphere of observers placed in 10° increments in azimuth

Project Description

The goal of this project was to integrate NASA's Aircraft Noise Prediction Program (ANOPP2) into the RCAS2WOPWOP code suite. The implementation was made to be generalized for an arbitrary number of rotors. Individual rotor noise predictions are combined into a total noise prediction. Another second goal was to use the updated code suite to predict noise for a propeller-wing case, focusing specifically on the influence of interactional aerodynamics on the predicted noise. A case with an isolated propeller was run, then a propeller-wing case, and the results were compared to highlight where interactional aerodynamics altered the noise produced.

Relevance of Work to DOD

Improvements to large scale electric batteries and motors has led to increased interest in electric takeoff and landing (eVTOL) vehicles. The designs of these vehicles differ when compared to traditional rotorcraft due to benefits gained by distributed electric propulsion. eVTOL are being considered for many different missions, with the civilian side seeing missions like delivery of goods and people while military applications like surveillance and cargo transportation are being considered.

Computational Approach

A key technical challenge is understanding the noise generated by these new eVTOL configurations. The Army Research Laboratory (ARL) has a robust code suite pairing the aerodynamic tool, Rotorcraft Comprehensive Analysis System (RCAS), with the acoustic prediction code PSU-WOPWOP. This code suite handles tonal and broadband noise predictions for arbitrary vehicle designs, but previous work has shown NASA's ANOPP2 to be better at predicting broadband noise when comparing to experimental data. The ARL code suite is updated to include ANOPP2.

The second part of this work investigates a propeller-wing case common in some eVTOL designs that look towards wing borne cruise. The conditions found in a propeller-wing case are unlike those found on traditional helicopter platforms and need to be investigated to develop more understanding for the designs. The prop-wing case is run using the updated code suite and then compared to a case with just the isolated propeller. This allows for understanding of how the aerodynamic interactions present in the prop-wing case effect the aeroacoustics, with noise investigated for a single observer as well as an observer hemisphere. The simulations are run with RCAS settings using the viscous vortex particle method to allow for aerodynamic interactions between the propeller and wing to be captured. Cases are run with and without the wing present to capture the effects the aerodynamic interactions have on the noise produced by the propeller. Noise is predicted at a single observer placed at 45° elevation angle, as well as a hemisphere of observers as shown in Figures 1 and 2

Results

For the single observer the no wing case has much larger tonal noise. But, when looking specifically at multiples of the 8th blade passing frequency, the with-wing case has large tonal noise spikes that far exceed the peaks in the no-wing case. These spikes can be directly attributed to the aerodynamic interactions caused by the presence of the wing. Broadband noise is comparable between the two cases. Examining the overall sound pressure level (OASPL) shows that even though the lower blade passing frequency peaks are higher for the no-wing case, the with-wing case has a higher tonal noise value due to those large spikes exhibited. But when broadband noise is added the overall noise values become much closer.

The OASPL for the no-wing case and the prop-wing case are studied. The noise in the two cases has the same shape across the hemisphere, with higher noise regions concentrated near to in plane with the propeller, with slightly higher noise concentrated behind the propeller. The low noise regions are directly in plane with the propeller as well as directly in front of and behind the propeller. The differences easily identified when the no-wing case noise is subtracted from the prop-wing case, Figure 3. Here we see that in general, the prop-wing case is quieter across the hemisphere, with only small pockets of higher noise. The cause of this remains to be investigated in future research.

Future

Prop-wing results are encouraging but more work must be done to understand where the differences arise. A good first step includes looking into how the loading across the propeller changes with and without the wing and see if the reduction in tonal noise comes from a reduction in the overall lift on the propeller, or merely a shifting of the high lift regions. There is also a need to investigate the pressure signal from the propeller and see how the signal changes, and what sort of implications that might have on the noise of the system.

HIP Mentors

Dr. George Jacobellis, Army Research Directorate, US Army Research Laboratory

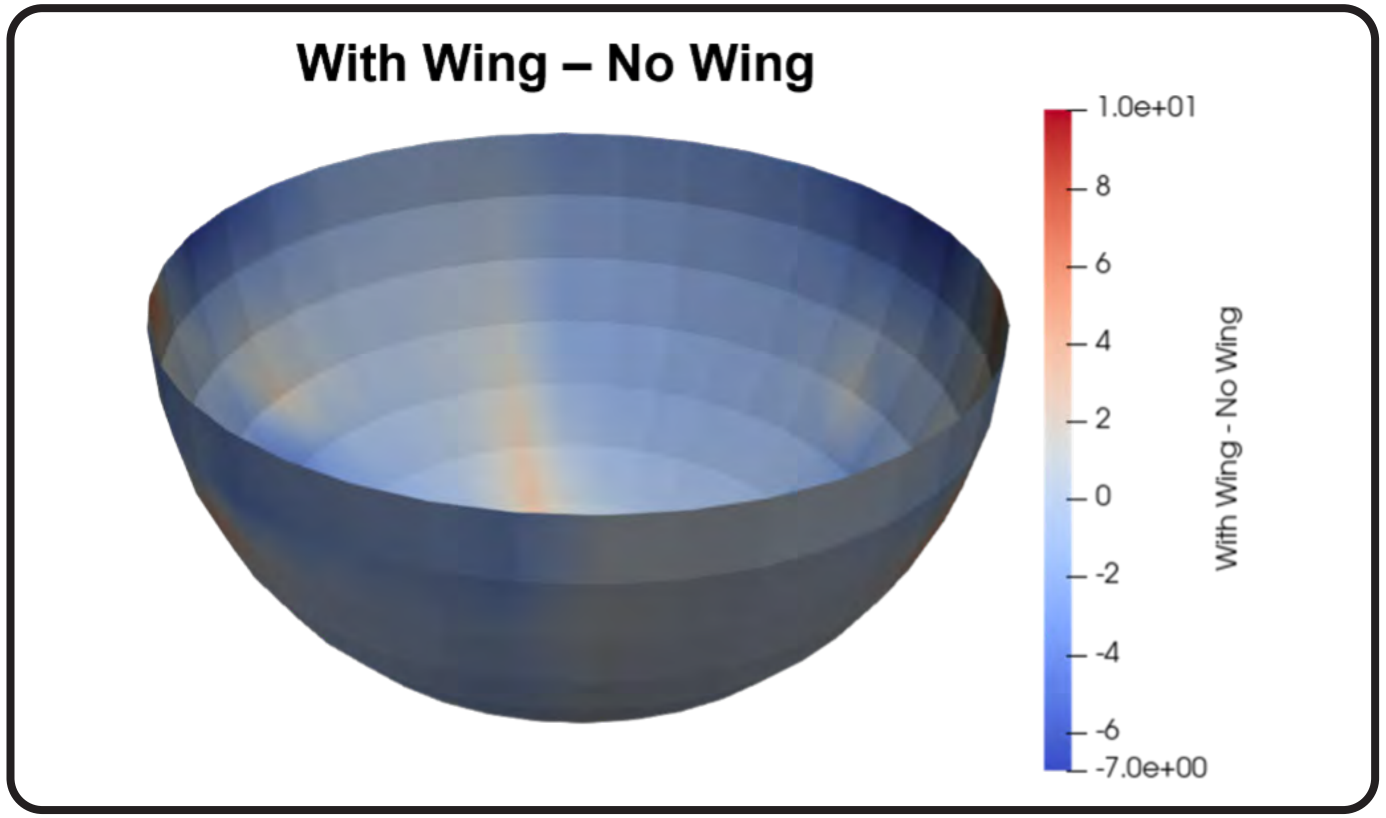

Figure 3:Difference hemisphere of (prop-wing) – (no wing)

Nikos P. Trembois
University of California, Davis

Broadband Noise Prediction from CFD Wall Spectrum

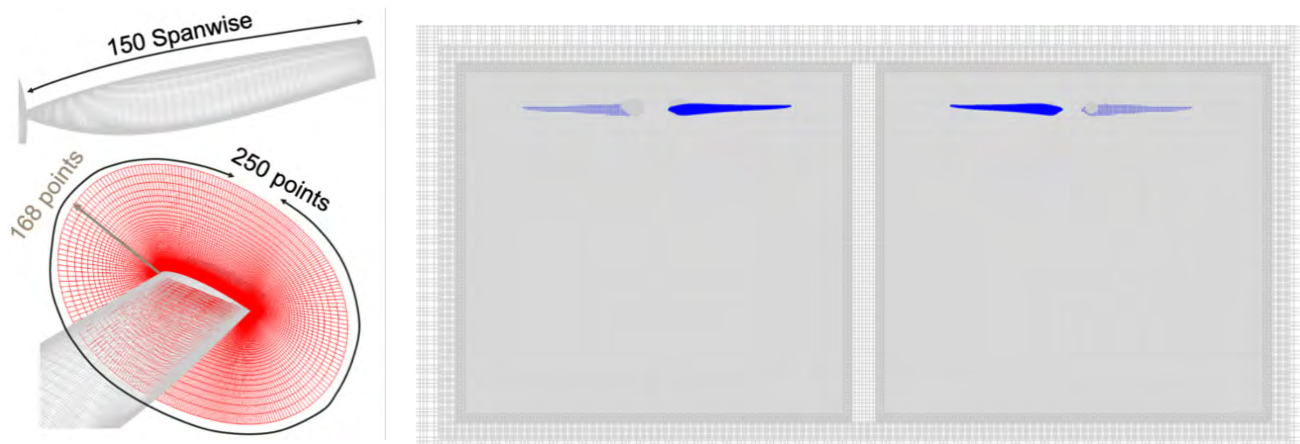

(a) (b)

Figure 1:
(a) Near body grids and (b) off-body grids..

Project Description

The objective of this work is to assess the ability to predict noise generated by incident turbulence on aerodynamic forces. New battery technology has made multi rotor vehicle designs feasible that have disparate sound signatures compared to conventional rotorcraft (i.e., helicopters). Leading edge noise created by incident turbulence is required for a comprehensive understanding of rotorcraft noise. This research uses high-fidelity computational fluid dynamics (CFD) simulation to evaluate the turbulence near the leading edge of rotor blades. High fidelity grids are required to resolve the blade tip vortices and the surrounding turbulence in the blade wakes. From the CFD solutions, the turbulence kinetic energy is extracted and the turbulence intensity and turbulence integral length scale are calculated. The broadband noise is then calculated from the turbulence quantities using Amiet's leading edge noise formula.

Relevance of Work to DOD

Urban Air Mobility is an emerging technology that promises safe, affordable, and environmentally friendly air transport pursued by automotive companies, commercial aircraft industry, and startups. Smaller scale multicopters see use in remote location to deliver medicine and other resources. The agriculture industry has adopted imaging drones to improve land plotting. This includes tracking invasive species and improving planting processes. Military surveillance techniques are not only aided by the unique aviation capabilities of small drones, but also the ability to deploy multiple small-scale rotorcraft that communicate with each other. Further, the Future Vertical Lift aircraft will incorporate new designs providing extended range potentially changing military planning for base locations

Computational Approach

Leading-edge noise prediction requires accurate simulation of flow around rotor blades, including the effect of flow interactions.Capturing flow structures like tip vortices and blade wakes is essential for modelling leading-edge noise. The CFD software Helios was used to simulate the flow and Helios Input Generator was used to set up the parameters. Helios accommodates and facilitates communication between a variety of near-body and off-body solvers. The structured, overset solver OVERFLOW is used for the near-body grids. The blades are modelled with three separate grids, shown in Figure 1.

The chordwise spacing uses 0.02% and 0.05% of the reference chord length as spacing at the trialing edge and leading-edge respectively. Since turbulence values are desired near the leading edge, it follows that fine spacing to model the impacting flow at precise locations is important for accurate acoustic predictions. It is equally important to model any flow separation that occurs at the trailing edge.

The off-body grid uses 8 levels of refinement, each refinement level doubles in spacing, extending to 20 times the rotor radii, which is necessary to ensure the freestream boundary conditions are met. A 2nd order dual time-stepping scheme is used to advance the simulation in time yet reduce residuals on each physical time step. The spatial scheme is a 5th order central scheme with rotational correction and delayed detached eddy simulation.

The k-omega shear stress transport Reynolds Averaged Navier-Stokes (RANS) turbulence model is used to model the flow turbulence that occurs at scales smaller than the grid resolution. The turbulence integral length scale and turbulence intensity are calculated from the turbulence kinetic energy and dissipation rate. After obtaining appropriate leading-edge values, the implementation of Amiet's model in UCD-QuietFly was used to predict interactional noise. In Amiet's model, the power spectral density of the sound is frequency dependent and can be simplified by assuming the wavelength is small compared to the rotor radius. The turbulence sound contribution is primarily modelled in the vertical velocity fluctuations. The vertical velocity fluctuations attenuate the pressure spectrum along with airfoil lift response. The pressure spectrum scales linearly with the semi-span and velocity and increases quadratically with the frequency, freestream density, and chord, but decays with the speed of sound and distance.

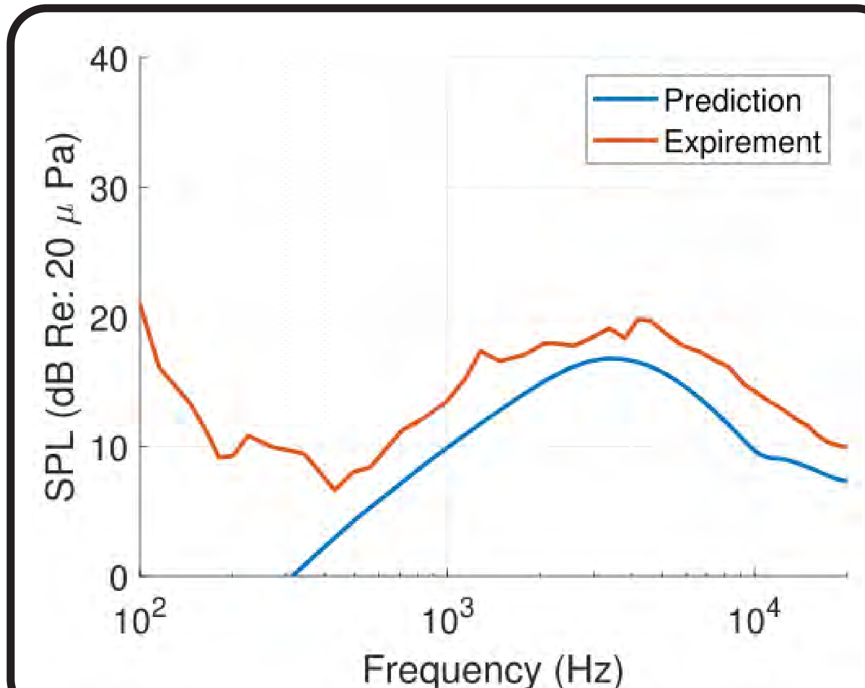

Figure 2:
Acoustic prediction comparison with experiment.

Results

Acoustic predictions are carried out on a one-third octave band scale that ranges from 100 to 10,000 Hz, covering a significant portion of human hearing and effective range of the acoustic models. Figure 2 compares the experimental and predicted sound pressure level (SPL). The prediction model only includes broadband noise; thus the noise is underpredicted at lower frequencies where tonal noise dominates. Overall, the shape of the broadband noise signature is captured by the prediction model, but underpredicts.

Future

Future work will quantify how much of rotorcraft broadband noise is created from self-noise and how much of the discrepancy between experiment and prediction is made up of leading-edge turbulence noise. These same prediction methods will be applied to a variety of geometries including a tilt-wing with closely stacked propellers. This will help evaluate the contributions from interaction noise and noise generated from aerodynamic bodies other than blades, like a downstream wing.

HIP Mentors

George Jacobellis, PhD *(Army Research Directorate, US Army Research Laboratory)*

Publications

Nikos P. Trembois and George Jacobellis, "Broadband Noise Prediction from CFD Wall Spectrum, "The ITEA Journal of Test and Evaluation 2022; 43 (1): 19-23

How to Request an Account

https://centers.hpc.mil/users/index.html#accounts

Who Can Use DoD HPC?

- DOD researchers and engineers
- DOD contractors
- University staff working on a DOD Research Grant. (Non-US Citizens may not have access to resources)
- Accounts are maintained through the Portal to the Information Environment (pIE) and managed by the Service/ Agency Approval Authorities (S/AAAs)

To Get Started:

(You will need to know the following items to complete your pIE account application)

- Citizenship Status
- Preferred Kerberos Realm
- Organizational ID
- Name, Title and Position
- Mailing address (no PO Box), Company Name, Phone, Fax and Email address
- Preferred User name
- Government Employee Status
- Who is your S/AAA?

High Performance Computing Modernization Program

Website: https://centers.hpc.mil/users/
HPCMP Help Desk Email: help@helpdesk.hpc.mil
HPCMP Help Desk Phone: 1.877.222.2039
Accounts Email: accounts@helpdesk.hpc.mil

Contact Information

DEVCOM Army Research Laboratory (DEVCOM ARL),
DEVCOM Army Research Laboratory DOD Supercomputing Resource Center (ARL DSRC)

Website: https://www.arl.hpc.mil/
Email: outreach@arl.hpc.mil
User Services Phone: 1.800.275.1552