

# DOE Multi-Laboratory Response to DARPA RFI

## “MICROELECTRONICS R&D FACILITY CAPABILITIES FOR PROTOTYPING”

### Response to DARPA-SN-21-06

#### Technical Points of Contact:



**Institution:**

Pacific Northwest National Laboratory  
PO Box 999, MSIN: K9-80  
Richland, Washington 99352




**Institution:**

Lawrence Berkeley National Laboratory  
1 Cyclotron Road  
Berkeley, California 94720

*The Department of Energy (DOE) National Laboratories offer a wide variety of powerful national user facilities and unique capabilities that can accelerate innovative R&D and prototyping for next generation microelectronic devices.*

-----  
**Additional Points of Contact**



Pacific Northwest National Laboratory  
Silicon Catalyst  
Lawrence Berkeley National Laboratory  
Pacific Northwest National Laboratory  
Brookhaven National Laboratory  
Battelle Memorial Institute  
Argonne National Laboratory  
Pacific Northwest National Laboratory  
Lawrence Livermore National Laboratory  
Argonne National Laboratory / U. Chicago  
Sandia National Laboratories  
Argonne National Laboratory  
Brookhaven National Laboratory  
Lawrence Berkeley National Laboratory  
Los Alamos National Laboratory  
Argonne National Laboratory  
Pacific Northwest National Laboratory  
Oak Ridge National Laboratory

# A. Background

There is no existing coordination mechanism in industry, academia, or at a consortium for the introduction of disruptive technologies that require optimization across the computing technology stack. This is necessarily the proper role for government, namely establishing crosscutting scientific fundamentals and addressing first order computing stack (2<sup>nd</sup> column in Fig. 1) alignment and scale-up challenges. By attacking these elements simultaneously and with sufficient resources, a dramatic reduction can be achieved in the time it takes to move discoveries to the point where industry will make the investments to bring the new technologies to market. Historically, that time lag has ranged from 12 to 20 years even for “drop-in” technologies such as lithography transitions and transistor improvements in CMOS semiconductor manufacturing. Advances affecting several layers of the computing stack could very well take much longer unless a more coordinated and integrated approach is taken.

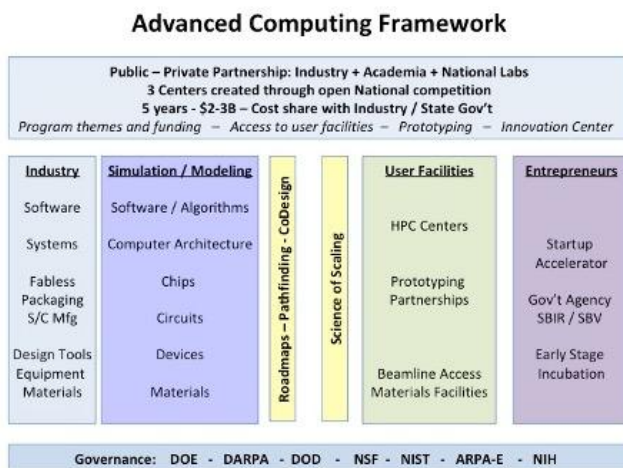


Figure 1

Combined with collective US experience in industry and academia, the DOE National Labs have the necessary scale, are fundamentally designed to address substantial interdisciplinary challenges, have the fundamental science and user facility infrastructure, and have capabilities across all hierarchical levels that need to be optimized. The National Labs are also widely viewed as objective and neutral in a way that individual competitors (or even academia) are not.

The challenges in successfully executing a collaborative program that includes such a diversity of participants are considerable. The purpose of this integrated DOE National Laboratory response is to propose solutions and facilitate a constructive dialogue, leading to a consensus on a path forward. The ultimate goal is to establish a neutral and execution-focused national initiative that can effectively down-select to the most promising emerging research innovations, attack the most critical gaps and unknowns with sufficient resources, and ultimately accelerate the technology development pipeline.

Since advances will necessarily be disruptive to current business models, there is a fundamental need for unbiased guidance and coordination. The involvement of the DOE National Labs, including their strategic assets in basic science, user facilities, and codesign capabilities can fulfill that objective. Their engagement will also assure that their unique concentration of scientists and application exploration at advanced user facilities are readily accessible, as these capabilities are not duplicated in the private sector.

Leveraging advances in electronics and computing are fundamental to the missions of several key government agencies. They share a common goal to continue US leadership and achieve asymmetric advantage by catalyzing innovations and incorporating them first. The capabilities that are described in this RFI response will be more effective in a single integrated *DOE Microelectronics Commons* than if pursued as loosely coupled Lab-centric agendas. Dispersing responsibility will not achieve the critical mass of resources necessary to invest at the scale required to tackle one of the greatest technology challenges of our generation. Moreover, the continuing advances in computing and its foundational technologies will be directly exploited by other grand challenges in science and technology where differentiation is increasingly found by applying artificial intelligence techniques (built upon advanced semiconductors, architectures and algorithms) to massive data streams to benefit our economy, defense and security.

**Success Criteria:** Implementing a private-public partnership is inherently complex, especially given the stakeholders' breadth: multiple government agencies, horizontally integrated industry participants within a vertical technology stack, academia from leading institutions, and the National Labs. This RFI response is intended to lead to follow up interactions for how DOE National Lab capabilities can be integrated into, and contribute to a multi-agency National Initiative in microelectronics prototyping, research and development.

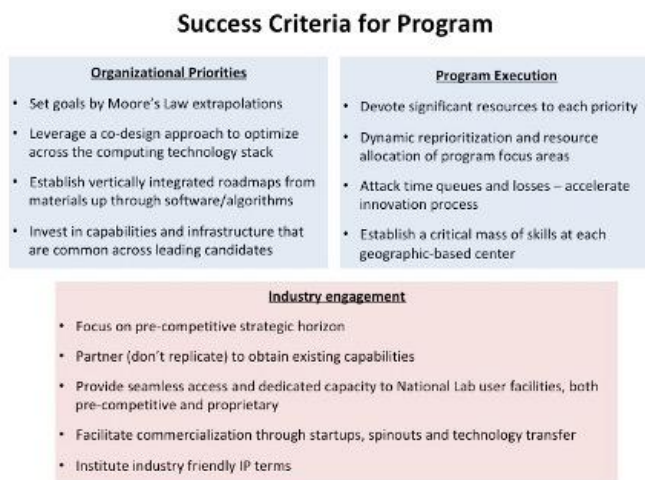


Figure 2

Hard won experiences from past collaborative efforts in basic science, technology development and manufacturing inform the proposed solutions. Of prime importance is governance, organization design, and advisory bodies' roles, and the ability to dynamically allocate resources to evolving program priorities.

The DOE National Labs have established guiding principles and success criteria (Fig. 2) for setting goals, establishing priorities, executing projects, and engaging

industry. Initially, staffing private-public partnerships with the right leadership and technical skills will be the key hurdle. Attention then must shift to developing a repeatable process to evaluate candidate technologies against a clear set of goals and metrics that sustain historical scaling trends. There needs to be rapid feedback loops using models and simulations that connect and optimize across the technology hierarchy to provide guidance in technology selection and identification of key gaps to basic understanding and scaling. These in turn must be translated into projects that are prioritized with appropriate budgets and resources.

## B. Facilities Overview

The DOE offers groundbreaking microelectronics capabilities and facilities in the following areas.

***Fundamental Materials Science Discovery, Characterization and Synthesis:*** DOE basic materials science research, has a system of user facilities that accelerate our ability to discover, and design elemental excitations in condensed matter at the atomic scale. Those include systems where the principles of quantum mechanical and competing spin-orbital and lattice interactions will enable new states of correlated matter and ordered phases. This will allow an exquisitely fine energy scale control of the states, such as reducing the energy consumption in a nanoscale switch from 1pJoule to 1AttoJoule. The DOE Basic Energy Sciences “Grand Challenges” report as the guiding principle leading to the discovery of new materials, phenomena and processes that will ultimately lead us to a totally new paradigm of information processing, storage and communication.

***Light Sources for Advanced Metrology and Imaging:*** the DOE’s X-Ray light sources (synchrotrons such as the APS, NSLS-II, ALS, SSRL, X-ray Free Electron Lasers such as LCLS-II, and beam accelerators such as CEBAF) will enable researchers to identify, model, and demonstrate the new materials and devices for ultra-efficient computing. Examples include low voltage transistor concepts such as the TFET, photonic devices, spintronics, and novel memory devices. It requires advances in DOE’s material modeling capability together with advances in algorithms from applied mathematics and data processing from DOE’s CS and applied mathematics research programs.

***Material/Device Scale Modeling and simulation*** DOE ASCR/ASC can apply its expertise in advanced computing and system performance modeling/simulation to exploit new device and materials systems and packaging technologies developed in the first two thrusts. Components include accelerators, on-chip wide-bandgap devices, photonic blocks, and emerging memory devices. Additionally, services such as DOE’s *Materials Project enables HPC high-throughput search for new materials to increase throughput for discovering new electronic materials and devices by a factor of 1000x* over current methods.

***Architecture/Systems Modeling and Simulation:*** The value of more-than-Moore, and emerging post-CMOS materials or device technologies cannot be understood without projecting their impact on the performance and efficiency of the ultimate computer systems and architectures that they ultimately target. DOE’s extensive portfolio of circuit, architecture, and system-level modeling tools enable better understanding of the performance impact of these emerging approaches on target applications and enable early exploration of new software systems that would make these new architectures useful and programmable. The goal of these architectures is to remove overheads in current designs, as well as offer hardware and thus more efficient support for important functionality for security and resiliency.

***Advanced Manufacturing, Prototyping and Integration:*** This leverages DOE’s expertise in EUV lithography e.g., the DOE led and industry supported EUREKA program. Other examples include the Sandia MESA fabrication center that can provide silicon photonics and compound semiconductor

research multi-project wafer (MPW) runs for collaborative R&D activities, as well as heterogeneous integration capabilities for making hybrid lasers and modulators on silicon photonics. Argonne offers extreme scale device processing at the Center for Nanoscale Materials cleanroom. Together, these facilities use advanced materials to develop novel nanomanufacturing methods, including EUV lithography, heterogeneous integration of advanced photonics and wide bandgap devices, and scalable engineered multi-layer thin films to fabricate photonics nanowire heterostructures. 3D stacking-is increasing density enabling photonic and memory layers on top of logic layers, and even multiple memory and logic layers interleaved. This radical change challenges assumptions embedded in current architectures, and would provide a new dimension to extend Moore's Law scaling.

## B.1 Fundamental Materials Science Discovery, Characterization and Synthesis

The DOE labs contain the largest collection of nanoscale materials synthesis and characterization capabilities in the US at their user facilities. The Nanoscale Science Research Centers (NSRCs) offer extensive extreme scale nanofabrication tools and are DOE's premier user facilities for interdisciplinary research to understand and control matter at the nanoscale, serving as the basis for a national program that encompasses new science, new tools, and new computing capabilities.

The NSRCs are a set of DOE Office of Science sponsored research user facilities available for use by the international science community to advance scientific and technical knowledge in the areas of nanoscale science. The NSRC Program was established as a major component of the DOE contribution to the U.S. Government National Nanotechnology Initiative (NNI). NNI involves twenty departments and agencies that collaborate toward *"a future in which the ability to understand and control matter at the nanoscale leads to a revolution in technology and industry that benefits society."* The nanoscience centers are co-located with other major nanoscience-related user facilities such as neutron or synchrotron light sources.

The mission of the NSRCs is twofold: to enable the external scientific community to carry out high-impact nanoscience projects through an open, peer-reviewed user program, and to conduct in-house research to discover, understand, and exploit functional nanomaterials for society's benefit. To fulfill this mission, the NSRCs house the most advanced facilities for nanoscience research and employ world-class scientists who are experts in nanoscience and enjoy working with external users. The NSRCs complement each other with their instrumentation and capabilities, the different thrusts of their in-house research programs, and the technical expertise of their staff.

### The NSRC Program

- Operates a national network of geographically distributed Facilities that leverage other facilities and expertise at DOE National Laboratories.
- Has world-leading capabilities and scientific expertise to create, characterize and understand novel nano-structured materials.
- Provides state-of-the-art nanoscience tools and expertise for research by non-profit or business organizations, whether small or large, for use-inspired research.

- Is available free-of-charge for non-proprietary work if the user intends to publish the research results in open literature.
- Proprietary work can also be supported on a full cost recovery basis.
- Serves users from all U.S. states and many countries around the world.
- Enables thousands of scientists to perform cutting edge nanoscience research each year.
- Contributes to the success of America's current and future research leaders.

The five NSRCs are:

- Center for Functional Nanomaterials (CFN) at Brookhaven National Laboratory
- Center for Integrated Nanotechnologies (CINT) at Los Alamos National Laboratory and Sandia National Laboratories
- Center for Nanophase Materials Sciences (CNMS) at Oak Ridge National Laboratory
- Center for Nanoscale Materials (CNM) at Argonne National Laboratory
- The Molecular Foundry (TMF) at Lawrence Berkeley National Laboratory

Across the National Laboratories there are a broad-array of microelectronics-relevant capabilities. For example, the CFN nanofabrication facility is focused on preparation of nanostructures and prototype devices made from a wide variety of materials. This research-focused cleanroom houses tools for the full fabrication workflow, including optical (2D and 3D) and electron-beam lithography, deposition of a wide variety of materials, and etching for pattern transfer. This facility is strongly focused on demonstration of novel nanomaterials, and not for fabrication of full CMOS microelectronics. The CFN nanofab thus provides crucial access to tools for the earliest phases of R&D into novel concepts that go beyond the current assumptions of commercial foundries.

Another example is the Molecular Foundry's nanofabrication facility, which focuses on understanding and applying advanced lithographies, thin film deposition, and characterization, with emphasis on the integration of inorganic, organic, and biological nanosystems that have the potential for nanoelectronic, nanophotonic, and energy applications. One of the Foundry's central research themes is on single-digit nanofabrication and assembly that aims to organize and structure material with critical features of dimensions at or below 10 nm, i.e., on the single-digit nanometer and atomic scales, to create nanoscale devices and architectures in inorganic, biological, or hybrid systems. Work in this theme is accomplished by developing protocols to visualize, understand, and implement methods of self-assembly and lithography in a variety of systems. Nanofabrication activities are facilitated in a roughly 4,850 square foot (450 m<sup>2</sup>) cleanroom, mainly Class 1000, which also includes Class 100 and Class 10 areas for nanofabrication/lithography, clean measurements and electron beam lithography.

## B.2 Light Sources for Advanced Imaging and Metrology

The DOE operates world-class light sources that enable the scientific study of materials with nanoscale resolution and exquisite sensitivity. These facilities have cutting-edge capabilities for x-ray imaging and high-energy resolution analysis to capture atomic-level data on a wide variety of

materials, from biological molecules to semiconductor devices. X-ray characterization tools offered at synchrotron sources provide comprehensive quantitative capabilities for supporting microelectronics research and development. The tools range from matured techniques to novel methods taking advantage of the cutting-edge capabilities that are available at synchrotron X-ray light sources across the labs. Below are some examples of the capabilities of these facilities for microelectronic materials discovery and characterization.

**Hard X-Ray Sources (APS and LCLS-II):** The Advanced Photon Source (APS) offers a 100X-1000X increase in brightness and beam coherence expected by 2023. This will enable examination of 100 micron<sup>3</sup> sized material at 10s of nm resolution. Ultrafast electron microscopy for stroboscopic imaging and diffraction examination of materials at ps timescales, and dynamic imaging at sub-ns timescales. Extensive molecular beam epitaxy, atomic layer deposition and sputter capabilities. When it comes online, LCLS-II will be able to image complete chip designs in 3D at the nano-scale resolution.

**X-Ray, Ultraviolet, and infrared Sources (NSLS-II):** For example, the NSLS-II has a suite of x-ray characterization tools using soft and hard x-rays, revealing crystalline structure, atomic coordination, elemental composition, chemical states, electronic band structure, magnetic ordering, and three-dimensional imaging with nanoscale resolutions. These tools can be used under in-situ or operando conditions. Structural and spectroscopic tools can reveal the atomic-level heterogeneity and switching mechanisms in non-volatile memory (NVM) devices or material systems. X-ray scattering using coherent x-rays exhibit high sensitivity to electronic and magnetic ordering in a range of material systems that have potential application in microelectronics. NSLS-II has a unique strength in nanoscale x-ray imaging in two separate modalities: full-field imaging and scanning microscopy. Its full-field imaging capability has the world-leading measurement throughput with spatial resolutions at 30 nm, at least by a factor 10 faster than other synchrotron facilities. This method is highly effective in visualizing the 3D structure in integrated circuits (IC) for validating structural integrity. The scanning microscopy capability utilizes the world-leading x-ray beam size down to 10 nm, providing high sensitivity to crystalline strain, chemical heterogeneity, and morphology. This multimodal imaging capability finds effectiveness in quantifying strain-field in 3D gate interfaces such as nanosheets or providing chemical mapping for the buried nanoscale junctions that are inaccessible by electron-based imaging tools. NSLS-II has active collaborations with DMEA, academia, and industry, utilizing these novel capabilities for microelectronics investigations.

**Soft-X-Ray Sources (ALS):** For example, the ALS is a DOE-BES funded user facility and operates synchrotron radiation experiments for academic and industry users. The ALS together with its partners offers proprietary and nonproprietary access to its beamlines. A range of tools are available to investigate processes and materials relevant to the microelectronics community. Together with the Center for X-ray Optics (CXRO) the ALS offers capabilities to study and test EUV lithography technology and materials, e.g photoresists, photomasks and optical surfaces. The ALS also offers soft x-ray tools for the nanoscale chemical characterization of materials using x-ray scattering and imaging probes. Chemical processes, e.g. in resists or at interfaces, can be interrogated using operando characterization techniques such as ambient pressure photoemission and absorption spectroscopy. The ALS also offers a full suite of fundamental materials

characterization techniques, such as angle resolved photoemission for electronic structure studies, e.g. of quantum materials for next generation computing technologies, scattering and diffraction imaging techniques to interrogate electronic and magnetic phases, and basic structural and physical characterization tools, for example x-ray diffraction and microtomography to understand the fundamental operation of microelectronic devices at the nano-scale.

## B.3 Material & Device Modeling and Simulation

Synthesizing a new material and implementing a new device heterostructure at atomic scale can be extremely time consuming and costly. To accelerate the discovery process, the DOE has created scalable computational models of materials and devices that make use of DOE's supercomputers to perform high-throughput simulation of materials and microelectronic devices *before* they are built. This greatly accelerates the evaluation of a wide-variety of options to focus efforts on the subset that offers the most promise.

**The Materials Project for Automated Discovery of Microelectronic Materials:** The Materials Project, is pushing the bounds of what is possible for microelectronic materials by design. MP can target data generation and uncovering design rules for novel materials relevant to domestic prototyping, including discovering and tuning multiferroics, wide bandgap semiconductors, and spin filters. This will significantly reduce the financial burden associated with failed integration of category materials because of theoretical limits on performance or synthesizability. The Materials Project is one the most visible databases of computed materials properties -- containing detailed properties of over 135,000 materials, including fundamental thermodynamics, electronic structure, elastic, piezoelectric, and dielectric tensors. All of this data can be accessed via easy-to-use interfaces through the web interface or the REST API. There is also the ability to contribute data back to the Materials Project via MPContribs, which provides the infrastructure for searchable data sets with full APIs and visibility on the Materials Details pages. This provides a rapid proliferation of data to the over 165,000 registered users for the Materials project.

As a core program of the LBL Materials Sciences Division, access to MP tools and data are free of charge. Access to staff time for customization, exploration of new scientific challenges, or generating new data sets require direct funding. For funded projects, MP has a sandbox model that enables controlled access to data available via the website. Select users have access to the core MP data and data produced for these sandboxes via the same familiar interfaces. MP currently remains a core program with a scientific focus rather than a user facility with the funding to support infrastructure upkeep and active user engagement at no or little cost. MP already serves a broad spectrum of the community from industrial to academics via community outreach: user forum, workshops, open-source codes, and significant email engagement. MP continues to grow, in-particular as MP expands to serve the large semiconductor community.

**Materials Science and Electronic Structure Simulations** – relevant to microelectronics, optics, quantum materials and device physics BNL lab leads a multi-lab team ExaLearn – which is a



codesign center for Exascale Machine Learning Technologies – we can leverage ExaLearn’s “Design” pillar to integrate AI/ML methods into eCAD tools, and to support the National Labs investments in open source hardware design and simulation tools – although this is beyond the remit of the current project. ExaLearn coordination and integration could act as a model for future efforts in codesign for microelectronics hardware and architectural improvements.

**Simulation and Modeling of Microelectronic Devices on Leadership Class Systems** Applied mathematicians and computer scientists at LBNL, under funding from the DOE Exascale Computing Project, have created scalable modeling and simulation tools to analyze emerging post-CMOS microelectronic devices (electronic, spintronic, nanomagnetic and nanomechanical). AMReX (<https://amrex-codes.github.io/>) software framework to provide a tool for multiscale physical modeling of electromagnetic fields and related physics involved in microelectronic circuitry -- enabling device-scale simulations that are far beyond the reach of commercial tools. The resulting open-source code, ARTEMIS (Adaptive mesh Refinement Time-domain ElectroDynaMics Solver), contains support for dispersive material properties, user-defined excitation and boundary conditions, and heterogeneous physical coupling present in next-generation microelectronics. Future plans include the incorporation of multi-level adaptive mesh refinement, support for more complex geometries, additional physical modules such as superconducting and solid mechanics, and a design feedback workflow in collaboration with researchers performing atomistic-level simulations to design new materials for microelectronics applications. All of the code is open source and therefore universally available.

**LAMMPS Molecular Dynamics Simulator:** LAMMPS is a classical molecular dynamics code with a focus on materials modeling. The code has potentials for solid-state materials (metals, semiconductors) and soft matter (biomolecules, polymers) and coarse-grained or mesoscopic systems. It can be used to model atoms or, more generically, as a parallel particle simulator at the atomic, meso, or continuum scale. Source code is released for public use under the GNU General Public License. Information and source code releases are available at: <https://lammps.sandia.gov>

There are many more such open source materials and chemistry modeling codes available supported by the CMS initiative

<https://science.osti.gov/bes/Research/Computational-Materials-and-Chemical-Sciences-CMS-CCS>

## B.4 Circuit, Architecture and System-level Modeling and Simulation

Next generation microelectronics will leverage extreme domain specialization at all scales, from high-performance computing devices, to autonomous systems, to distributed arrays of sensors, for a variety of applications. Areas such as machine learning and data analytics, which can greatly benefit from domain-specialized accelerators, are still quickly evolving the algorithmic methods, making it extremely difficult to design such accelerators by hand. Novel hardware design automation tools are required that can translate a high-level language formulation of an algorithm, to a variety of accelerator implementations ready to be fabricated and evaluated along different metrics and

constraints. Such design automation tools today are a critical part of the hardware prototyping and fabrication pipeline, providing a critical frontend element to microelectronics evaluation and fabrication activities. Tools to perform high-level and logic synthesis, technology mapping, and floorplanning, together logic cells represent critical elements of the design flow that so far have presented significant limits in availability and integration.

The DOE laboratories are developing a new generation of open source generators and synthesis tools to enable a quick design cycle from high-level specification to fabrication and rapid prototyping.

**Circuit Simulation Capabilities:** Xyce - is an open source, SPICE-compatible, high-performance analog circuit simulator, capable of solving extremely large circuit problems (10s to 100s of times larger and more complex than commercial SPICE-simulators can handle), using a differential-algebraic equation formulation for execution on large-scale parallel computing platforms. Xyce was developed internally at Sandia National Laboratories and funded by the National Nuclear Security Administration's [Advanced Simulation and Computing](#) (ASC) program. It also supports serial execution on all common desktop platforms, and small-scale parallel runs on Unix-like systems. The code is open source and released under the GNU General Public License. Information and source code is available at, <https://xyce.sandia.gov>

**Computer Architecture Simulation:** The Sandia Structural Simulation Toolkit (SST) was developed to explore innovations in highly concurrent systems where the ISA, microarchitecture, and memory interact with the programming model and communications system. The package provides two novel capabilities. The first is a fully modular design that enables extensive exploration of an individual system parameter without the need for intrusive changes to the simulator. The second is a parallel simulation environment based on MPI. This provides a high level of performance and the ability to look at large systems. The framework has been successfully used to model concepts ranging from processing in memory to conventional processors connected by conventional network interfaces and running MPI. Software repository: <https://github.com/sstsimulator/sst-core>

**PARADISE++ is a post-Moore HPC (High-Performance Computing) system simulation** framework to enable large scale simulations of post-Moore architectures built using emerging post-CMOS devices and technologies such as nanomagnetic, carbon nanotube, MESO, and is extensible to other emerging device technologies. Efficient simulation methods and algorithms are needed to successfully employ large scale parallel simulation of the future Beyond Moore HPC systems. PARADISE extends state of the art simulators like the SST simulator with a view of making post-CMOS system simulation flexible, scalable and extensible.

**Logic in Memory Emulator (LiME)** is an FPGA-based hardware memory subsystem emulator. The emulator can be configured to capture every memory access issued to the memory subsystem without perturbing the running software application. Memory traces for graph traversal (BFS) and DOE simulation benchmarks have been collected and made available to the DOE community. LiME can be configured to emulate a wide range of memory read and write latencies. The emulator is also

designed for easy insertion of custom logic to emulate near memory processing. Two accelerators, a programmable gather/scatter unit (Data Rearrangement Engine) and a key/value lookup engine have been evaluated through LiME and are available in GitHub. LiME is particularly valuable for codesign since a full software stack including application software components, OS, and drivers are emulated. Software repositories: <https://github.com/LLNL/lime> and <https://github.com/LLNL/lime-apps>

**High Level Synthesis for Rapid Prototyping:** PandA is an open source framework that includes methodologies supporting the research on high-level synthesis of hardware accelerators, on parallelism extraction for embedded systems, on hardware/software partitioning and mapping, on metrics for performance estimation of embedded software applications and on dynamic reconfigurable devices. The high-level synthesis tool, named Bambu, is able to generate custom hardware accelerators in hardware description languages (Verilog and VHDL), starting from high-level languages (such as C/C++ and others). Bambu interfaces with GCC and LLVM. Maintained at Politecnico di Milano, the toolchain includes significant contributions from PNNL, with several advanced synthesis methodologies that target parallel applications and irregular workloads (such as graph analytics). The tool can generate accelerators quickly instantiable on Field Programmable Gate Arrays (FPGAs), and is also able to interface to logic synthesis tools for Applications Specific Integrated Circuits (ASICs), including the opensource flows developed under the DARPA IDEA program (OpenRoad, and LS Oracle), PandA is available at: <https://panda.dei.polimi.it>

**SystemC-clang is an open source SystemC to Hardware Description Language translator** in active development through collaboration between LLNL and U of Waterloo. A floating point compression pipeline optimized to scientific data (1D, 2D, 3D arrays) has been implemented in SystemC and translated to FPGA. SystemC was particularly valuable for codesign since a scientific C++ application could drive the hardware accelerator, and iterative refinement from behavioral to cycle accurate descriptions could be evaluated in a single test environment. Software repository: <https://github.com/anikau31/systemc-clang>

**OpenCGRA:** OpenCGRA is a parameterizable and powerful open source CGRA (Coarse-Grained Reconfigurable Arrays) generator to generate synthesizable Verilog for different CGRAs based on user-specified configurations (e.g., CGRA size, type of the computing units in each tile, communication connection, etc.). OpenCGRA uses modular design and standardized interfaces between modules. The configurability and extensibility are maximized by its parametrization system to fit in various research and industrial needs. implemented leveraging PyMTL, OpenCGRA has been also exercised with heuristic design space algorithms. OpenCGRA is available at: <https://github.com/pnnl/OpenCGRA>

**Minos Computing Library (MCL):** MCL is a modern task-based, asynchronous programming model for extremely heterogeneous systems. MCL abstracts the low-level hardware details of a system, supporting the execution of complex workflows that consists of multiple, independent applications. This capability can also drive the execution of computational tasks on custom hardware designs implemented using hardware simulators or emulators, such as SST or FPGA accelerators. Importantly, MCL facilitates the execution of current applications on emerging data flow

accelerators (SambaNova, Cerebras, Xilinx Versal). MCL is available at::

<https://minos-computing.github.io/>

**SODA Synthesizer:** The Software Defined Architectures (SODA) Synthesizer is a new modular open source synthesis infrastructure. Leveraging several community efforts, the synthesizer is able to consider a variety of inputs that can interface with the Multi-Level Intermediate Representation (MLIR) compiler infrastructure, implements a synthesis backend fully integrated within the LLVM framework, and is able to generate a circuit representation in FIRRTL (the Flexible Intermediate Representation for RTL), a circuit-level IR that further simplifies the retargetability and integrability of the generated accelerators to different types of devices (FPGAs, ASICs, and more). Initially focused on Machine Learning frameworks in support of the DARPA Real Time Machine Learning Program (RTML), the retargetable frontend will be able to accept several general purpose and domain specific languages that can interface with MLIR. The backend will be able to interface with OpenCGRA, or generate RTL for FPGAs and ASICs, supporting both open source and commercial toolchains. The whole design flow will be able to perform optimization and design space exploration both at the front-end and middle-end level.

**Microelectronic testing in harsh environments** Brookhaven hosts the NASA Space radiation laboratory (NSRL) which provides users with beams of all ions from protons to thorium, ranging in energy from 50 MeV to 1500 MeV (ion species dependent), and extracted from the Booster accelerator with masses and energies similar to the cosmic rays encountered in space. A 100-meter transport tunnel and beam line then deliver these beams of simulated space radiation to a 400-square foot shielded target hall. NSRL's beams of ions impinge on user-provided targets and allow researchers to determine the effects of radiation on microelectronics—As the total ionizing dose radiation tolerance of most processes increased in recent years the focus has shifted towards susceptibility for single event effects caused by radiation, specifically high energy protons and heavy ions. Irradiation studies at NSRL enable users to characterize failure modes in electronics and develop the corresponding mitigations.

## B.5 Operating Model

### B.5.a Is your facility available only to your organization and its collaborators?

The DOE operates national user facilities that can be used by any scientific organization on a no fee basis for open science research, or full cost recovery basis for proprietary research. One important reason for establishing America's National Laboratory system immediately after World War II was to provide a home for large-scale, costly scientific facilities that universities could not afford. Such facilities were believed to be essential to sustaining America's leadership in science. The construction and operation of large-scale scientific user facilities have been integral to the mission of the DOE Office of Science from the earliest days.

Today, the DOE Office of Science maintains and operates 28 user facilities at DOE National Laboratories across the country as shared resources for the scientific community, with no fee

access determined on a competitive basis using peer review. Tens of thousands of researchers make use of these facilities each year.

These facilities—including advanced supercomputers, particle accelerators, large x-ray light sources, neutron scattering sources, specialized facilities for nanoscience, and others—have become increasingly vital tools of scientific discovery. They have also become an important component of national economic competitiveness.

### B.5.b Is your facility available to external groups from academia, the commercial sector, and the start-up community for prototyping?

DOE User Facilities are open to the global scientific community and generally follow a simple proposal based system. Proposals, if selected, enable use at no cost to the user - in return for publishing or sharing of scientific data using FAIR data principles. These facilities include the DOE BES Nanoscience Research Centers, the Advanced Computing Centers and the Synchrotron light and neutron source centers. The facilities noted above are available to external groups both from academia and the commercial sector.

**Broadly Accessible Open Source Tools for Codesign:** The DOE National Laboratories create open source hardware/software codesign tools to assist with architectural analysis and hardware/software codesign. For example, DOE advanced microelectronics simulation codes such as Xyce, AMReX, and PARADISE++ are all broadly available via open source licenses so that they can be used by scientists and engineers in industry or academia. DOE and DARPA applications and workflows are complex and heterogeneous. They execute multiple computationally different methods in parallel on highly-customized processors and require accuracy, response time, and low-power consumption that might not be necessary in other domains. The time-to-market for a new hardware design is also critical in many domains and could be as tight as 12-18 months. Evaluating new architectural designs in this context is challenging: Methodologies and tools to perform HW/SW codesign studies and evaluate novel technologies (especially if disruptive) need to be agile and flexible but, at the same time, accurate and reliable. These methodologies should enable researchers to 1) focus on the relevant part of the application that may benefit from the new design or technologies, 2) understand how the new hardware design relates to the rest of the application (performance and power improvements, data movement across devices), 3) provide quick turn-around information to discard architectural designs that will not be efficient, 4) easily combine the novel design with existing compilers, programming models, and workflow managers. In other terms, it is key to be able to take an existing application and replace computing devices with novel architectural designs (drop-in) without excessive disruptive modifications to the original code and while maintaining a realistic execution environment for all the other application components and methods.

**Access to Leadership HPC Facilities:** The DOE National Labs support high performance computing via the Leadership Computing Facilities at ANL and ORNL, and the National Energy Research Scientific Computing center at LBNL. The national impact of these computing facilities is realized via DOE's broad portfolio of scientific and engineering computational applications.

Many of these applications are open source and can be strategic elements of the infrastructure to support domestic microelectronics prototyping. Under the goals of the DOE Exascale Computing Initiative, a broad spectrum of these applications were developed to support codesign collaborations among National Lab computer and computational scientists and US computer industry counterparts that developed advanced architecture concepts for exascale computing systems. Both companies and academic researchers can apply for free access to the leadership-class supercomputing facilities (OLCF, ALCF, and NERSC) using the DOE INCITE program (<https://www.doeleadershipcomputing.org/proposal/call-for-proposals/>)

**Additional Services for Business Development Support:** All of the DOE National Labs have business development centers and business incubators with a successful track record. These include Argonne's Chain Reaction Innovations, LBNL's Cyclotron Road, SNL's MESA, SNL and LANL's CINT, LANL's National High Magnetic Field Laboratory, Laboratory for Ultrafast Materials and Optical Science, ORNL's Spallation Neutron Source.

## C. Domestic Prototyping Infrastructure

### C.1 Advanced Manufacturing and Prototyping

DOE's manufacturing and prototyping capabilities use advanced materials to develop novel nanomanufacturing methods, including EUV lithography, heterogeneous integration of advanced photonics and wide bandgap devices, and 3D stacking, are increasing density enabling memory layers on top of logic layers, and even multiple memory and logic layers interleaved. This leverages DOE's expertise in EUV lithography (the DOE led and industry supported EUREKA program). This radical change challenges assumptions embedded in current architectures, and would provide a new dimension to extend Moore's Law scaling.

**EUREKA EUV Program:** As an example of the high-impact microelectronics-relevant technology-translation enabled by DOE light-sources, the Center for X-Ray Optics (CXRO) at the ALS offers significant prototyping and design services for EUV for the microelectronics industry. The industry-funded EUREKA Program at Berkeley Lab is the current incarnation of the 25 year EUV lithography program at Berkeley Lab originally started in 1997 in partnership with the EUV LLC industry consortium and then taken over by the SEMATECH consortium in 2001 and finally transitioned to the EUREKA partnership in 2015. With the successful commercialization of EUV lithography, which the Berkeley Lab program played a key role in, EUREKA is now focused on the long-term extension of EUV lithography to sub-3-nm nodes. The EUREKA facilities at Berkeley Lab include several unique resources critical to developing the manufacturing materials of the future including both photoresists and masks. On the photoresist front, the facilities include the 0.5-NA EUV micro-exposure tool (MET5) which is the world's highest resolution projection lithography tool as well as a variety of radiation chemistry diagnostics tools enabling the detailed understanding of the interaction of EUV photons with patterning materials. On the EUV mask front, the facilities include the world's highest resolution EUV microscope for the study of EUV mask defects and next generation mask materials as well as an EUV scatterometry tool enabling the measurement of optical properties of materials in the EUV regime as well as characterizing EUV mask phase shift materials which is seen as a critical enabling technology for scaling EUV lithography to future technology nodes. This unique instrument also enables the characterization of EUV mask phase shift materials which are seen as a critical enabling technology for scaling EUV lithography to future technology nodes.

**MESA MPW Fabrication Center:** The Sandia MESA fabrication center that can provide silicon photonics and compound semiconductor research multi-project wafer (MPW) runs for collaborative R&D activities, as well as heterogeneous integration capabilities for making hybrid lasers and modulators on silicon photonics.

**Materials Characterization and Synthesis for Microelectronics:** The Brookhaven CFN nanofabrication facility is focused on preparation of nanostructures and prototype devices made from a wide variety of materials. This research-focused cleanroom houses tools for the full fabrication workflow, including optical (2D and 3D) and electron-beam lithography, deposition of a wide variety of materials, and etching for pattern transfer. This facility is strongly focused on

demonstration of novel nanomaterials, and not for fabrication of full CMOS microelectronics. The CFN nanofab thus provides crucial access to tools for the earliest phases of R&D into novel concepts that go beyond the current assumptions of commercial foundries.

**Center for Nanoscale Sciences:** Argonne offers extreme scale device processing at the Center for Nanoscale Materials cleanroom. The JEOL 8100 electron beam lithography with consistent, multilayer alignment at <10 nm. Ultrathin dielectrics including 2D ferroelectric doped oxides, conformal deposition tools.

**Integration of synthesis, fabrication, characterization and modeling:** All National Labs have the multi-disciplinary teams that are able to provide vertical integration of manufacturing capabilities with novel materials synthesis, characterization diagnostics, unique fabrication technologies, and associated modeling and simulation capabilities. For example, PNNL is pursuing research to understand new quantum phenomena with the integration of 2D magnetic or superconducting materials and highly doped oxides. The research team is using a unique pulsed laser deposition system to fabricate scalable engineered multi-layer chalcogenide thin films into quantum photonics nanowire heterostructures.

## C.2 Gaps & Opportunities

Research into next-generation concepts for computing—including beyond-CMOS designs, co-located memory/computation, and neuromorphic circuits—critically requires testing novel materials and nano-devices in realistic contexts; that is, integrated into a semiconductor device. Currently there is a critical gap between the traditional foundries—which can produce full integrated circuits but are necessarily resistant to admitting novel materials that would contaminate their stringent workflow—and research nanofabs—which can produce novel structures but cannot typically produce full circuits. Resources are needed to bridge this gap, and to enable research and design workflows where novel materials and designs for circuit elements (transistors, memory cells, neuromorphic switches) and integration concepts (e.g. 3D and heterogeneous integration) can be tested in integrated circuits. Such a capability could be enabled by upgrading existing cleanrooms to provide the needed capability, or by streamlining passage of substrates between different nanofabrication facilities to enable coordination of their respective capabilities. If this critical gap is filled, progress in the design, testing, and validation of novel concepts for computing elements would dramatically accelerate. Understanding the physics of devices under extreme conditions (e.g. temperature or radiation) and developing models for dedicated design and processing is another gap that access to facilities which adequate characterization-modeling-implementation cycle flow would help to fill.