DISCLAIMER

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

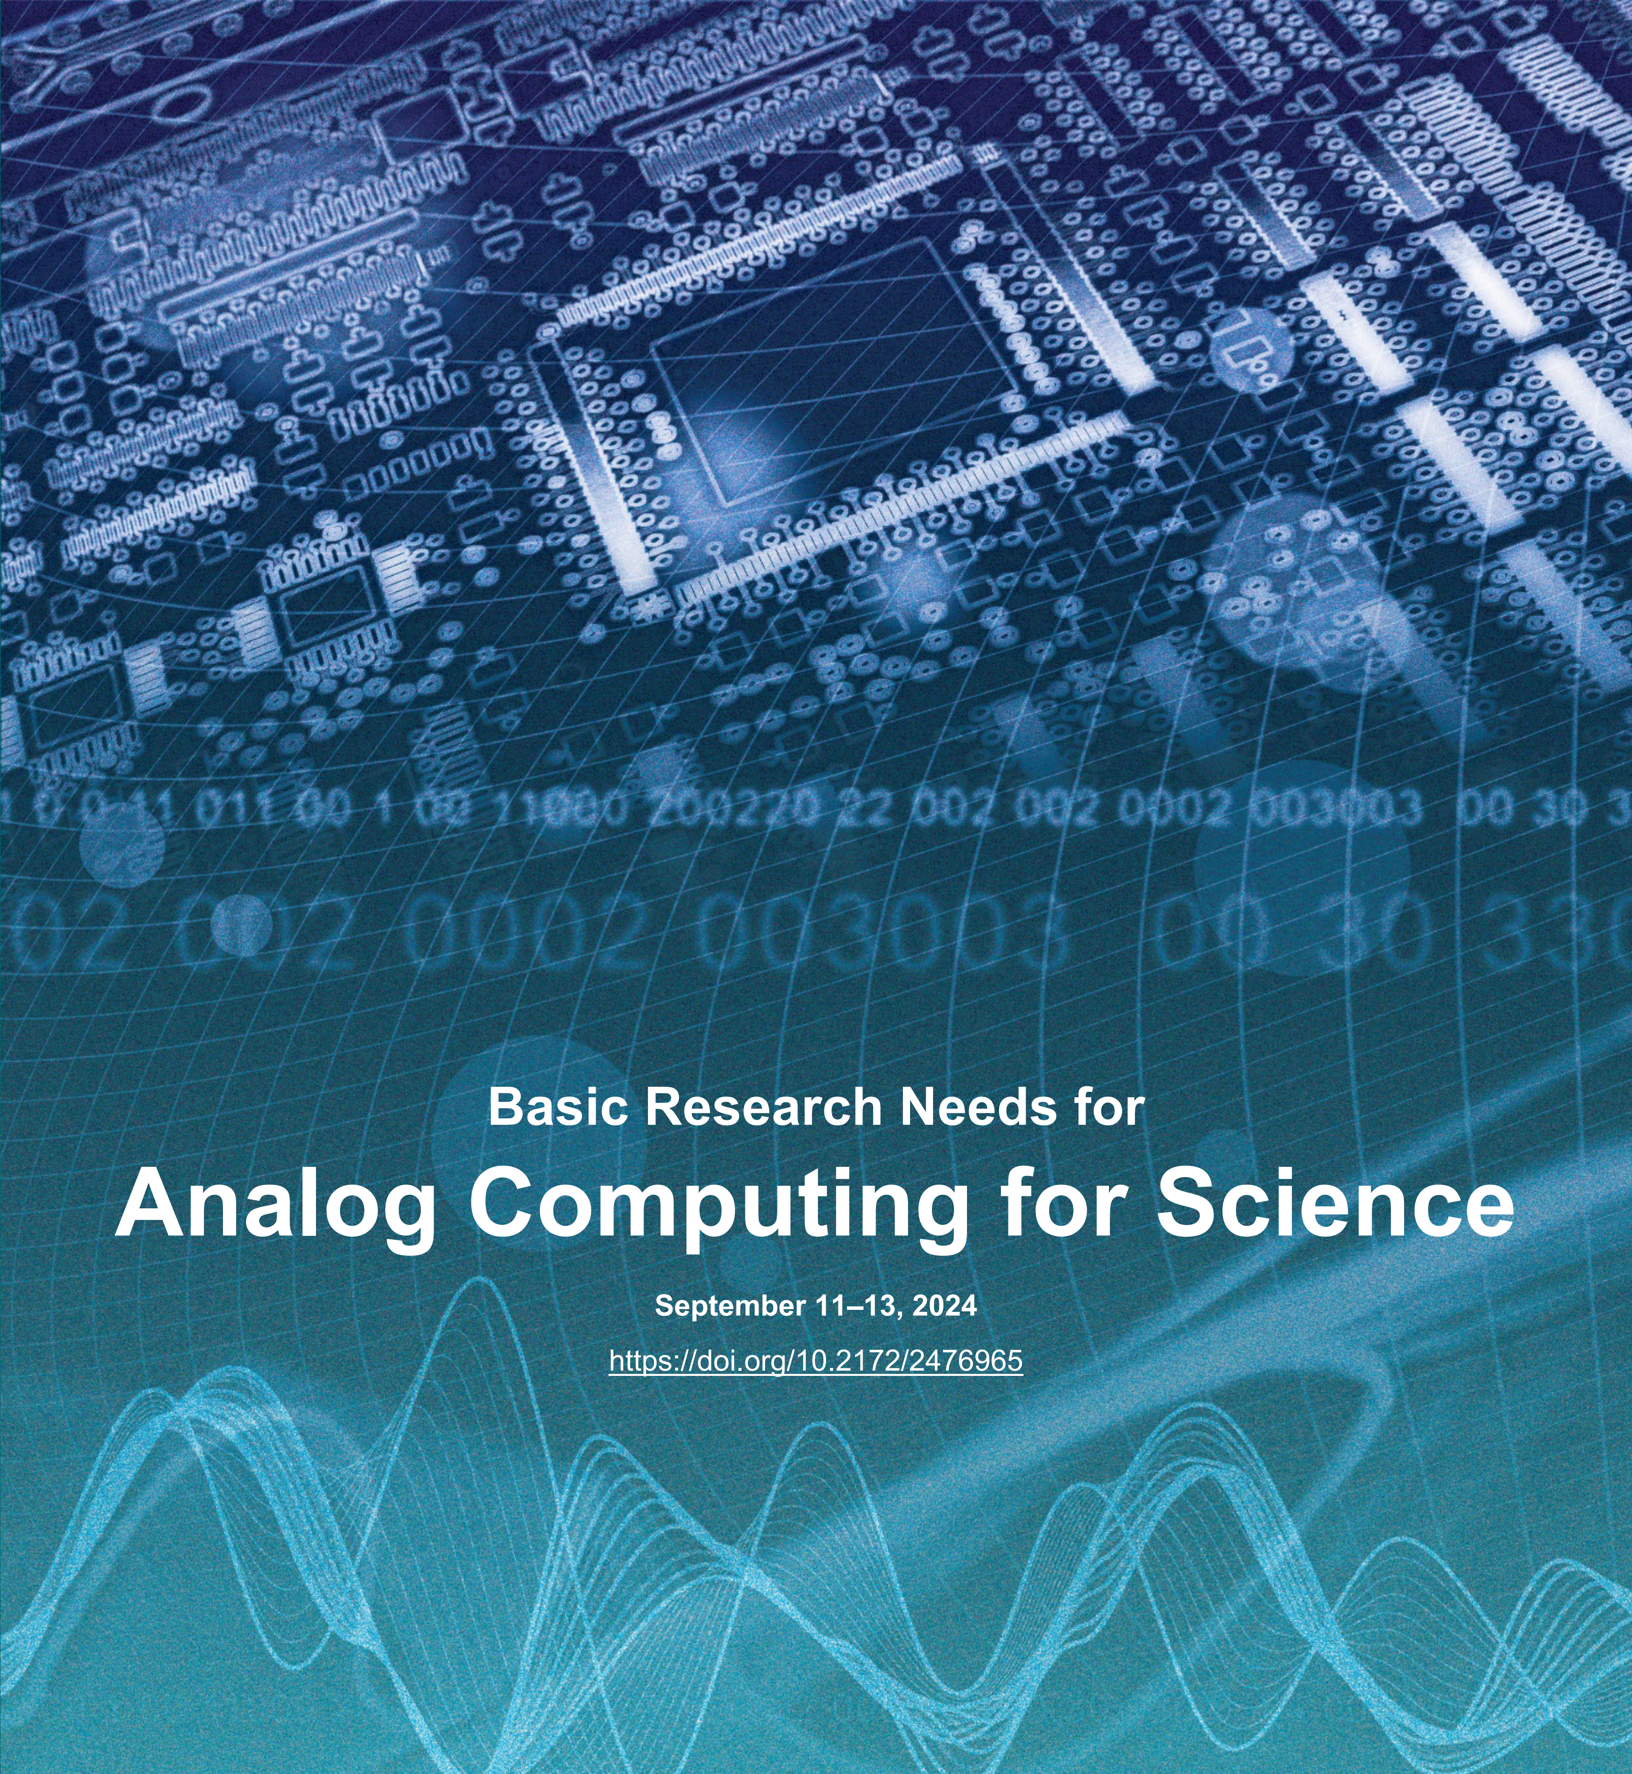

Basic Research Needs for

Analog Computing for Science

September 11–13, 2024

https://doi.org/10.2172/2476965

U.S. DEPARTMENT of ENERGY | Office of Science

Basic Research Needs for Analog Computing for Science

March 20, 2026

Workshop Organizing Committee

David Soloveichik
University of Texas at Austin *(Co-Chair)*

Antonino Tumeo
Pacific Northwest National Laboratory *(Co-Chair)*

Sara Achour
Stanford University

Natalia Berloff
University of Cambridge

Suma George Cardwell
Sandia National Laboratories

Shantanu Chakrabartty
Washington University in St. Louis

Jennifer Hasler
Georgia Institute of Technology

Siddharth Joshi
University of Notre Dame

Jaijeet Roychowdhury
University of California, Berkeley

Abstract

Conventional digital computing faces fundamental physical limits: large scale computing systems already consume tens of Megawatts of power, Dennard scaling has ended, and data movement costs dominate application performance. Next generation experimental facilities generate data at rates that overwhelm conventional processing and demand real-time analysis at the source. Analog computing, which exploits the continuous dynamics of physical systems to perform computation, promises a transformative path toward orders-of-magnitude gains in energy efficiency and time-to-solution for scientific workloads.

In September 2024, the US Department of Energy (DOE)'s Advanced Scientific Computing Research (ASCR) program convened a Workshop on Analog Computing for Science to address the critical research challenges and opportunities in this field, bringing together experts in applied mathematics, computer science, device physics, and applications domains from academia, government, and industry. The participants identified six interconnected priority research directions (PRDs): (1) developing a rigorous mathematical foundation for analog computation, (2) designing high-performance analog computing architectures, (3) establishing reliable device primitives, (4) enabling edge computing for real-time analysis, (5) exploring natural computing substrates, and (6) creating co-design methodologies that integrate software and hardware. The workshop's findings also identified several cross-cutting challenges underpinning all six PRDs, including a mathematical theory of continuous computation under noise, programming abstractions and compiler toolchains, and community benchmarking infrastructure, emphasizing the needs for a coordinated, multi-faceted effort to enable rapid progress in the area. The research directions outlined in the workshop report aim to guide the development of energy-efficient analog computing technologies that can support future scientific discoveries and address the growing energy demands of scientific computing and artificial intelligence.

CONTENTS

EXECUTIVE SUMMARY

The Department of Energy (DOE) Advanced Scientific Computing Research (ASCR) program advances applied mathematics and computer science to enable scientific discovery across DOE mission areas. Decades of scientific progress built on Moore's Law scaling have reached a physical wall: exascale systems already consume on the order of 20–30 megawatt (MW), Dennard scaling has ended, and data movement costs dominate application performance. Simultaneously, next-generation experimental facilities generate data at rates that overwhelm conventional processing and demand real-time analysis at the source. These converging pressures threaten DOE's mission across fusion energy, high-energy physics, materials science, geoscience, and other critical domains such as artificial intelligence (AI) for science.

Analog computing offers a transformative path forward by exploiting the continuous dynamics of physical systems—electromagnetic, chemical, quantum, and biological—rather than discretizing all operations into bits. Recent advances in materials, fabrication, and hybrid architectures have matured analog approaches from laboratory curiosities to practical technologies, with reported orders-of-magnitude energy-efficiency gains for differential equations, optimization, and linear algebra. Hybrid analog-digital architectures that combine analog efficiency with digital programmability and control offer the most practical near-term integration path. However, realizing these gains reliably and at scale requires resolving deep open problems in mathematics, programming models, device science, and system design.

This report addresses the **Grand Challenge:** *Can we develop the scientific foundations, algorithmic principles, and system architectures to harness analog and analog-digital computation for 100–1000× improvements in energy efficiency and 10–100× improvements in time-to-solution for DOE scientific computing workloads, while establishing the precision, programmability, and integration pathways needed to make these systems practical for production science?*

To chart a path forward, ASCR convened a workshop in 2024 bringing together experts in analog computing, applied mathematics, computer science, device physics, and DOE application domains. The workshop identified six priority research directions (PRDs) addressing the most critical research gaps: **(1) Mathematical Foundations**—a rigorous theory of analog computation, including completeness proofs, noise-resilient algorithms, and formal precision bounds, does not yet exist; **(2) High-Performance Analog Computing**—scalable architectures for high-performance computing (HPC) workloads remain undemonstrated beyond narrow kernels; **(3) Device Primitives**—reproducible fabrication, variability models, and validation testbeds are absent, blocking systematic co-design; **(4) Edge Computing**—principled frameworks for real-time analog processing at experimental instrumentation are lacking; **(5) Natural Computing**—theoretical and algorithmic foundations for diverse substrates (physics-based, biological, chemical, and brain-inspired) at scale are in their infancy; **(6) Co-design Methodologies**—integrated software stacks, compilers, and design-space-exploration methods spanning devices through applications do not yet exist.

Furthermore, several research challenges cut across all six areas. A mathematical theory of continuous computation under noise and uncertainty—encompassing error propagation, algorithm correctness, and formal precision limits—is needed to underpin every other advance. Programming abstractions and compiler toolchains that make diverse analog hardware accessible without bespoke low-level expertise are a prerequisite for broad adoption and portability. Community benchmarking infrastructure with standardized metrics and reproducible evaluation protocols is essential for measuring progress across all PRDs and guiding research priorities.

Addressing these needs would enable the DOE Office of Science (SC) to sustain computational performance growth beyond the limits of digital scaling, optimize energy costs at HPC facilities, accelerate time-to-solution for mission-critical simulations, transform experimental capabilities through real-time analog processing, and position the U.S. at the forefront of post-Moore computing technologies. The research agenda outlined in this report represents a coordinated fundamental research effort—spanning mathematics, computer science, device physics, and application domains—that no single discipline or institution can accomplish alone.

1 INTRODUCTION

1.1 The Scaling Crisis in Scientific Computing

For over five decades, exponential improvements in computing performance have enabled transformative advances in computational science, from large-scale multi-physics simulation and materials discovery to fusion energy research and quantum chemistry. This progress, historically driven by Moore's Law scaling of transistor density and clock speeds, has reached fundamental physical limits. Modern digital processors face a *power wall*: exascale computing systems consume on the order of 20–30 MW, approaching the practical limits of power delivery and cooling infrastructure. Further scaling of conventional complementary metal-oxide-semiconductor (CMOS) technology offers diminishing returns as transistors approach atomic dimensions and leakage currents dominate energy consumption.

Simultaneously, DOE experimental facilities face a parallel challenge: next-generation scientific instruments, from high-energy physics detectors and X-ray light sources to electron microscopes and environmental sensor networks, generate data at rates that overwhelm conventional processing architectures. These instruments require real-time analysis and decision-making at the source, yet transmitting raw data to remote computing facilities is increasingly impractical due to bandwidth limitations. The data deluge from scientific instrumentation, combined with the need for closed-loop control and autonomous operation, demands computing capabilities that are efficient, fast, and physically co-located with sensors.

The DOE SC mission to deliver transformative science and technology solutions requires computational capabilities that continue to grow exponentially—both for large-scale simulation and for real-time experimental data processing. The ASCR program, which advances applied mathematics and computer science to enable scientific discovery across DOE mission areas, faces a critical question: *how can we sustain performance improvements when conventional digital scaling has ended?*

This challenge is particularly acute for the differential equations, optimization problems, and linear algebra operations that dominate scientific computing workloads. These problems often involve continuous variables and smooth functions, yet conventional digital computers discretize all operations, expending vast energy to represent continuous quantities with discrete bits. This fundamental mismatch between problem structure and computing substrate suggests an opportunity: *can we exploit the continuous nature of physical systems themselves to perform computation more efficiently?*

Analog computing—leveraging the continuous dynamics of physical systems to perform computation—offers a path to address these converging challenges. This report, which collects and systematically organizes the discussions of the 2024 Workshop on Analog Computing for Science, examines the scientific foundations, technological readiness, and research directions needed to realize this vision, organized around the following grand challenge:

Grand Challenge: *Can we develop the scientific foundations, algorithmic principles, and system architectures to harness analog and analog-digital computation for 100–1000× improvements in energy efficiency and 10–100× improvements in time-to-solution for DOE scientific computing workloads, while establishing the precision, programmability, and integration pathways needed to make these systems practical for production science?*

1.2 Physical Computation and Analog Computing Fundamentals

Analog computing refers to computational systems that represent and process information using continuous physical quantities such as voltages, currents, optical intensities, chemical concentrations, mechanical displacements, rather than discrete digital states. Unlike digital computers that encode all information as binary sequences and perform computation through logic gates, analog computers harness the governing dynamics of physical systems to solve problems.

The distinction between *digital* and *analog* computation is fundamentally about representation: digital systems discretize variables into finite sets of values (typically binary), while analog systems operate on continuous quantities. This distinction leads to essential differences in how computation is performed:

- **Digital computation:** discretizes time, space, and amplitude; represents all variables as bit sequences; performs operations through combinational and sequential logic; achieves precision through bit width; and can be reprogrammed arbitrarily through software.
- **Analog computation:** operates on continuous signals; represents variables directly as physical quantities; per-

forms operations through physical dynamics (ohmic flow, wave propagation, chemical kinetics); achieves speed and efficiency by exploiting parallelism inherent in physical laws; but often requires hardware reconfiguration to change the computed function.

Throughout this report, we use the term *physical computation* to encompass the broader concept of harnessing physical dynamics for information processing. This includes: (1) *physical dynamics*—exploiting natural evolution of physical systems (electromagnetic, mechanical, chemical, quantum) to perform computation; (2) *physical substrates*—the materials and devices that embody computational primitives; and (3) *physical constraints*—the fundamental limits (noise, precision, bandwidth) that govern what is computable and at what cost.

Modern analog computing synthesizes emerging materials, nanoscale devices, and hybrid analog-digital architectures to combine analog efficiency with digital programmability. Representative substrates include CMOS analog circuits, memristive arrays, photonic integrated circuits, superconducting electronics, micro-electro-mechanical systems (MEMS), and molecular/biological systems. For readers seeking additional background on analog computing implementations and applications, a recent survey[1] provides comprehensive coverage of in-memory computing, Ising machines, analog solvers, neuromorphic processors, and machine learning accelerators.

Patch-cord analog machines, such as the EAI TR-10[2] and THE ANALOG THING[3], were early analog computers that performed calculations by manually connecting components with physical wires (*patch-cords*). This approach enabled flexible, real-real time configuration of mathematical operations, but offered limited repeatability and automation, making them important precursors to modern programmable analog systems. Our scope considers a much broader class of approaches, including electronically configured matrix- and tensor-centric systems based on dense in-memory vector-matrix multiplication (VMM) and photonic primitives, as well as alternative paradigms in which physical processes directly implement mathematical models, including biological and chemical substrates. When combined with digital orchestration and calibration to enable compiled, repeatable execution, these approaches open opportunities for analog and hybrid systems that were not possible in the patch-cord era. Crucially, enabling capabilities differ substantially from earlier eras. Key ingredients that were unavailable or immature until recently include: (i) reliable non-volatile analog memories and foundry-grade photonic integrated circuits that store and manipulate weights *in situ*; (ii) heterogeneous packaging and 3D integration that provide the bandwidth for tight analog-digital coupling; (iii) hardware-aware algorithms and control that tolerate noise/variation and support in-the-loop calibration and training; (iv) high-speed, energy-efficient analog-to-digital converter (ADC)/digital-to-analog converter (DAC) interfaces for closed-loop operation; and (v) software toolchains for co-design, compilation, and verification. Combined with workloads dominated by data movement and linear algebra in HPC and AI, these advances create an intrinsic opportunity for analog and hybrid systems distinct from the patch-cord era[4].

1.3 The Analog Advantage for Scientific Applications

Recent research has demonstrated that analog computing can achieve dramatic energy and speed efficiency improvements for problem classes that align with natural physical dynamics. Experimentally validated results include:

- **Differential equations:** Analog solvers for partial differential equations have reported 100–1000× improvements in energy efficiency and 10–100× speedups compared to digital implementations, particularly for systems with sparse spatial coupling[4;5].
- **Optimization:** Physical annealing systems (optical, memristive, quantum-inspired) have shown 10–1000× efficiency gains for combinatorial optimization tasks relevant to DOE missions, such as partitioning for parallel computing, quantum circuit optimization, materials configuration space exploration, resource allocation in complex systems, sensor placement optimization, and facilities operations optimization. Continuous optimization problems—including partial differential equation (PDE)-constrained optimization for inverse problems and real time control systems—are addressed through analog differential equation solvers (see Differential equations above). Best results are achieved for problems mappable to physical energy landscapes or exploiting continuous-time gradient dynamics.
- **Linear algebra:** Photonic and analog electronic matrix-vector multipliers have demonstrated 100–1000× energy-efficiency improvements for large-scale linear operations in targeted settings, with demonstrations at scale for neural network inference. Recent work[6] demonstrated precise matrix equation solving using resis-

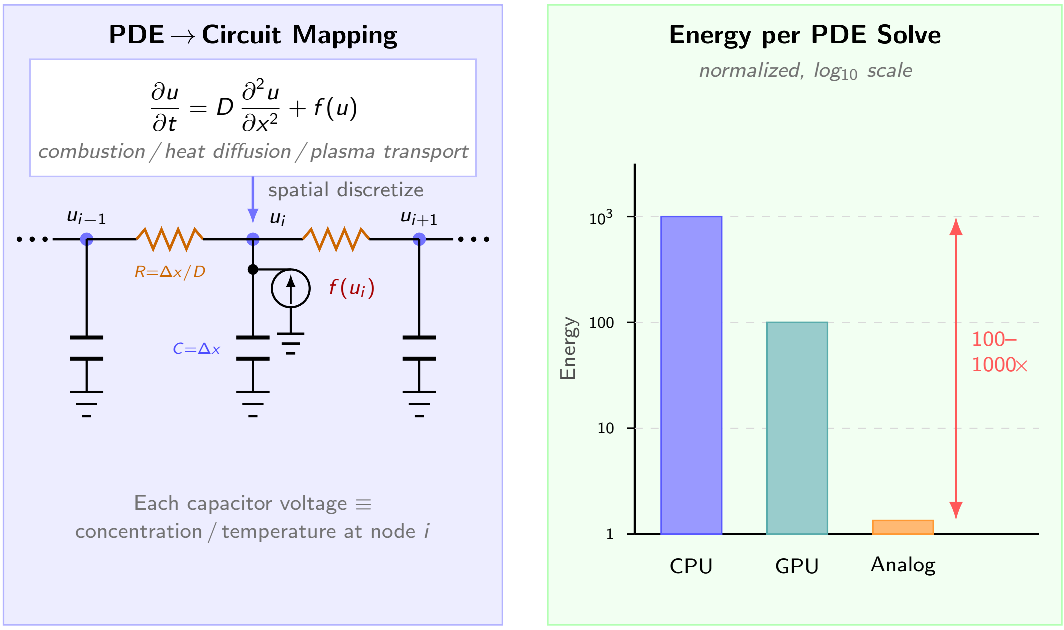

Figure 1: Analog computing advantage example. *Left:* A 1-D reaction-diffusion PDE ($\partial u/\partial t = D\,\partial^2 u/\partial x^2 + f(u)$, governing combustion, heat diffusion, and plasma transport) maps onto an RC-ladder circuit. Resistors ($R = \Delta x/D$) implement diffusion coupling; capacitors ($C = \Delta x$) store nodal state; current sources $f(u_i)$ encode local reaction rates. *Right:* Projected energy per PDE solve on a logarithmic scale; analog implementations are estimated at 100–1000× lower energy than central processing unit (CPU) or graphics processing unit (GPU) digital solvers[4;5].

tive random access memory (ReRAM), reporting 24-bit precision with 1000× throughput and 100× energy-efficiency gains.

- **Signal processing:** Analog radio frequency (RF) front-ends and analog-domain feature extraction reduce energy consumption by 10–100× for edge sensing applications by avoiding costly analog-to-digital conversion.

Figure 1 illustrates the differential-equation case in detail: a reaction-diffusion PDE governing combustion, heat diffusion, and plasma transport maps directly onto an RC-ladder circuit, and projected energy consumption places analog implementations three orders of magnitude below digital alternatives.

These problem classes are precisely those that dominate DOE scientific computing: large-scale multi-physics PDE simulations (fluids, plasmas, materials), quantum simulations (e.g., density functional theory (DFT), quantum Monte Carlo, quantum chemistry), molecular dynamics, fusion energy (e.g., magnetohydrodynamics (MHD) simulations, gyrokinetic turbulence, particle-in-cell methods), high-energy physics (e.g., track finding, event reconstruction, detector simulation), and data analysis across experimental facilities. The potential impact extends beyond performance to enabling new science: ultra-low-power analog systems could enable dense sensor networks for environmental monitoring; analog accelerators could enable interactive design cycles for materials optimization; analog-digital hybrid systems could make previously intractable simulations routine.

Analog Computing and the AI Revolution. The rapid progress in AI creates both a challenge and an opportunity for analog computing. Training and deploying large AI models drives unprecedented energy consumption, threatening the sustainability of data-center-scale computing. Yet the core operations of neural network inference—matrix-vector multiplication, nonlinear activation, attention—map naturally onto analog hardware, where they can be performed directly without the energy cost of repeated digital conversions. This advantage operates at two scales. At the data-center scale, analog platforms have achieved projected efficiencies exceeding 500 tera operations per second (TOPS)/W for AI inference, more than 100× that of leading GPUs[7], and analog in-memory architectures have demonstrated substantial reductions in large language model (LLM) inference costs by addressing parameter-fetching bottlenecks[8]. At the edge, analog computing enables AI inference directly at the sensor, reducing costly ADC/DAC conversions and making ultra-low-power, real-time processing feasible in environments where digital solutions are energy-prohibitive. A broader survey of state-of-the-art demonstrations appears in Section 2.6. These results underscore a central theme of this report: realizing analog computing's full impact on AI requires co-designing algorithms and hardware to exploit analog's unique properties.

Critically, analog computing is not a replacement for digital computation, but rather a complement. *Hybrid analog-digital architectures* that combine analog efficiency for specific operations with digital flexibility for control, precision where needed, and integration with existing software ecosystems offer the most practical path forward.

The alignment between analog computing capabilities and DOE mission requirements is particularly strong. DOE's scientific computing needs—extreme-scale simulations, operation in harsh environments (temperature extremes, radiation), stringent precision requirements, and mandatory uncertainty quantification—present challenges distinct from commodity and/or enterprise computing[9–11]. Analog hardware deployed in these contexts must also meet reproducibility, provenance, long-term stability, and verification & validation standards that consumer-grade accelerators are not designed to satisfy. DOE mission applications in geoscience, materials, fusion, and high-energy physics demand error characterization and uncertainty quantification beyond commercial standards. These unique requirements position DOE to drive fundamental advances in analog computing that may not emerge from commercially-motivated research. The benchmarking infrastructure, evaluation methodologies, and open-source toolchains developed through DOE investment can accelerate progress across the broader research ecosystem.

1.4 Fundamental Challenges and Research Opportunities

Despite demonstrated advantages, analog computing faces fundamental scientific challenges that require sustained research investment:

- **Noise and precision:** Physical systems are inherently noisy, with thermal fluctuations, device mismatch, and environmental variations limiting achievable precision. Understanding error propagation, developing noise mitigation strategies, and designing algorithms robust to inexact computation are central mathematical challenges.
- **Interfaces and integration:** Analog-digital interfaces consume significant power and introduce latency, often negating analog efficiency gains. Novel interface designs, hybrid partitioning strategies, and system-level co-optimization are needed.
- **Programmability and toolchains:** Most analog systems require hardware reconfiguration to change functionality. Developing programming models, compilation strategies, automated design tools, and verification methodologies for analog and hybrid systems is essential for practical deployment.
- **Scalability and manufacturability:** Scaling analog systems to problem sizes relevant for scientific computing while maintaining precision and managing device variability requires advances in architecture, calibration, and fault tolerance.
- **Application mapping:** Identifying which computational kernels benefit from analog acceleration, developing efficient mappings from algorithms to hardware, and quantifying accuracy-performance trade-offs requires co-design across applications, algorithms, and devices.

Importantly, these challenges are *addressable through fundamental research*; they are scientific opportunities rather than fundamental showstoppers. The PRDs outlined in this report define coordinated research directions to advance the foundations, architectures, devices, algorithms, and methodologies needed to make analog computing practical for DOE mission applications.

1.5 The 2024 ASCR Workshop on Analog Computing

To define a comprehensive research agenda for analog computing, the ASCR program convened a workshop in September 2024, bringing together leading researchers in analog computing technologies, applied mathematics, computer science, device physics, and computational science. The workshop aimed to identify key challenges, research opportunities, and promising directions that could enable analog computing to complement and enhance digital systems for scientific computing workloads. To better frame the discussions, the organizing committee circulated a set of notional questions (see Appendix 9.3) and issued an open call for contributed white papers; the accepted white papers are archived on OSTI.gov (`https://doi.org/10.2172/2506701`).

Over two and a half days, participants engaged in plenary sessions establishing state-of-the-art, breakout discussions organized around emerging research themes, and collaborative synthesis sessions. Working groups refined problem definitions, assessed feasibility and impact, and drafted preliminary research directions. A final plenary session identified cross-cutting challenges that span multiple research areas, particularly around noise management, analog-digital interfaces, and co-design methodologies.

The workshop structure balanced broad community input with focused technical discussion, ensuring that the resulting research directions reflect both fundamental scientific opportunities and practical pathways toward impact for DOE mission applications. A strong recommendation from the workshop is to regularly convene the analog computing community in focused gatherings, as the event clearly demonstrated tremendous value through positive and enthusiastic cross-disciplinary discussions.

It should be noted that this report focuses specifically on analog computing as defined above. Quantum computing and neuromorphic computing, while related to physical computation paradigms, were discussed in parallel workshops and are covered in separate reports.

1.6 Organization of This Report

This report synthesizes the outcomes of the workshop, presenting six PRDs that collectively define a comprehensive research agenda for analog computing:

- **PRD 1: Mathematical Foundations of Analog Computing (Section 3):** Develops robust mathematical frameworks including models of computation, complexity theory, and rigorous analysis of analog advantage. Essential for understanding capabilities, limitations, and optimal applications of analog computing.
- **PRD 2: Building High-Performance Analog Supercomputing Systems (Section 4):** Explores analog accelerators for large-scale scientific computing with order-of-magnitude improvements in speed and energy efficiency. Addresses challenges in scaling, precision, and integration with digital infrastructure.
- **PRD 3: Identifying Novel Device Primitives for Next-Generation Analog Systems (Section 5):** Focuses on emergent device technologies that promise scalable analog solutions with inherent efficiency. Emphasizes predictable performance at scale and compatibility with manufacturing.
- **PRD 4: Redefining the Edge: Integrating Analog Compute with Sensing and Transduction (Section 6):** Transforms edge computing by integrating analog computation directly with sensing systems, eliminating costly conversions and enabling real-time, ultra-efficient processing.
- **PRD 5: Harnessing, Enhancing, and Inspiring Efficient Computation through Natural Systems (Section 7):** Explores biological, chemical, and physical substrates for computation, leveraging natural principles of adaptation, efficiency, and robustness.
- **PRD 6: Cross-Cutting Multi-Granularity Co-Design Methods (Section 8):** Provides unified tools and methodologies spanning materials to applications, enabling holistic optimization across the full stack for both analog and hybrid systems.

Section 2 establishes the common context, challenges, and methodologies that underpin all six PRDs. Subsequent sections (Section 3 through Section 8) detail the specific research directions, gaps, and opportunities in each area. Appendices provide additional technical details, glossary of terms, and lists of workshop participants and contributors.

These PRDs are deeply interconnected and mutually reinforcing. Success across this portfolio requires coordinated effort across theoretical foundations, enabling technologies, application demonstrations, and design methodologies. Collectively, they define a research program that balances fundamental scientific questions with practical pathways toward deploying analog computing for DOE SC mission applications, with the ultimate goal of sustaining computational performance growth beyond the limits of conventional digital scaling.

2 FOUNDATIONS AND CROSS-CUTTING THEMES

2.1 Motivations: The Analog Computing Opportunity

2.1.1 The Analog Computing Paradigm

Analog computation represents a paradigm fundamentally different from its digital counterpart, operating on continuous rather than discrete values. While digital computation has dominated computing for decades, recent advances in analog neural networks, programmable analog circuits, as well as other analog devices and systems suggest that analog approaches may offer critical advantages for specific tasks. By representing data with continuous physical quantities (e.g., voltages, probabilities, chemical concentrations, or light intensities) instead of discrete binary states, analog systems can inherently solve mathematical problems through their physical behavior, and offer distinct advantages in scenarios where continuous operations are more effective than Boolean logic.

Compared to digital approaches, analog encoding and processing of information are fundamentally different in key ways. First, digital systems enforce discrete states using design margins and error correction, while analog systems employ quasi-continuous states defined by non-equilibrium thermodynamics and kinetics, with much smaller effective barriers in many implementations[12]. Second, while conventional digital logic encodes information primarily through charge, analog systems can additionally exploit electrical, thermal, and electrochemical gradients in heterogeneously integrated materials to move electrons, ions, and domains[13]. Understanding the scientific basis of these complex, frequently coupled mechanisms is essential for exploiting the computational advantages of analog systems.

2.1.2 The End of Digital Scaling and the Need for New Paradigms

Digital computing faces fundamental power and performance bottlenecks that increasingly impede progress on DOE scientific priorities. The completion of the DOE's Exascale Computing Project (ECP) has brought these limitations into sharp relief: scaling to the next 1000× is stymied by the effective end of Dennard scaling and the pronounced slowdown of Moore's Law; the energy cost of data movement (von Neumann bottleneck); and the diminishing returns of weak scaling for many complex scientific codes[11]. Future performance gains will come at disproportionate power costs, impeding DOE's mission across critical domains including materials, fusion, physics, and biological sciences[10].

Radically improving computers will almost certainly require embracing novel physical platforms and computing paradigms that can reach beyond the physical limits of current CMOS digital circuits. These limits include quantum tunneling, short channel effects, thermal limits, mobility degradation, variability, interconnect delay, and overall thermodynamic limits which make it asymptotically harder to reduce energy per operation. Overcoming these hurdles requires revolutionary advances that deliver orders-of-magnitude gains in energy efficiency.

2.1.3 Analog Advantages

Analog computing offers a fundamentally different paradigm that represents a promising path forward, with demonstrated advantages in multiple dimensions:

- **Energy efficiency**: Analog computation shows exceptional energy efficiency demonstrated both in theory[14] and experiment[15], with higher computational energy efficiency (computations for a given amount of energy) compared with digital computation.
- **Area efficiency**: Analog systems achieve higher area efficiency (computations per unit area) through dense, parallel architectures.
- **Latency**: Very low computational latency for operations that map naturally to physical dynamics.
- **Continuous-time operation**: Direct solution of differential equations and dynamical systems without discretization.
- **Natural problem matching**: Physical dynamics can directly implement computational primitives, avoiding inefficient mappings to Boolean logic.

These advantages apply to the DOE mission workloads that motivate this report: differential equations (e.g., multi-physics simulations, dynamical systems), linear algebra (e.g., quantum chemistry, scientific AI), continuous optimization (e.g., gradient-based methods for inverse problems and control systems), combinatorial optimization (e.g., energy landscape methods for graph partitioning and resource allocation), and signal processing (e.g., analog RF front-ends to minimize conversion overhead). They manifest across scales—from ultra-low-power edge systems co-located with sensors, to large-scale HPC accelerators with potential order-of-magnitude energy efficiency gains for specific computational kernels, to physics-based substrates operating near thermodynamic limits.

However, realizing this potential requires overcoming several cross-cutting challenges that appear throughout all Priority Research Directions:

- **Theoretical foundations**: A comprehensive mathematical framework for analog computation is needed to rigorously understand its capabilities, limitations, and optimal applications.
- **Device and materials innovation**: Novel physical substrates must be developed that can overcome the limitations of CMOS while maintaining manufacturability and predictability at scale.
- **System integration**: Efficient interfaces between analog components and conventional digital systems are critical for building practical, heterogeneous computing platforms.
- **Robustness**: Device mismatch, noise, parameter variability, and environmental fluctuations must be managed to ensure accuracy, reliability, and scalability.
- **Design methodologies**: A holistic, full-stack co-design approach is essential, fundamentally different from traditional digital design.

Realizing these advantages could dramatically reduce the energy footprint of DOE HPC facilities and accelerate the design-to-discovery pipeline for energy materials and technologies.

2.2 Workshop Organization and Breakout Topics

To promote focused discussion and community-driven insights, the 2024 Analog Computing for Science Workshop featured eight thematic breakout sessions spanning the breadth of analog computing research and applications in scientific computing. Each session convened experts to identify challenges, opportunities, and research priorities within a specific domain of analog computing.

The breakout topics were as follows:

- **Biological / Chemical Computing**: Exploring the use of biochemical and chemical systems, including synthetic biology and reaction networks, as substrates for analog computation in scientific modeling.
- **Hybrid Systems / Sensor+Compute Integration**: Investigating the fusion of sensing and analog computation at the edge, enabling real-time signal processing and efficient front-end data handling.
- **Devices / Photonics**: Examining the role of emerging devices and photonic technologies in enabling scalable, high-throughput, and low-power analog computing platforms.
- **Analog Electrical and Computational Memory**: Focusing on electrical analog circuits and memory-centric computing paradigms, such as those based on memristors or resistive memory arrays, to perform in-memory computation.
- **Probabilistic / Ising**: Discussing analog approaches to probabilistic computing, including physical implementations of Ising machines and stochastic models relevant to scientific computing.
- **Co-design / Software+Hardware+Physics**: Highlighting the importance of integrated co-design approaches that align physical substrates, hardware architectures, and software frameworks to support analog computing workflows.
- **Mathematical and Theoretical Foundations**: Addressing the formal models, theoretical frameworks, and algorithmic principles that underpin analog computation across different physical domains.
- **Benchmarking**: Developing rigorous benchmarks, metrics, and evaluation methodologies to assess the capabilities, performance, and scientific relevance of analog computing systems.

From Breakouts to Priority Research Directions. These eight breakout sessions were synthesized into six cohesive PRDs that define a comprehensive research agenda. The consolidation process recognized natural groupings, cross-cutting themes, and complementary expertise, organizing the PRDs into three categories:

- **Foundational PRDs:** provide theoretical underpinnings and enabling technologies that support all analog computing approaches. *PRD 1 (Mathematical Foundations)* emerged from the *Mathematical and Theoretical Foundations* breakout, establishing rigorous frameworks for analog computation capabilities, complexity theory, and error analysis that underpin all other PRDs. *PRD 3 (Device Primitives)* combined insights from *Devices/Photonics* and *Analog Electrical and Computational Memory*, with input from *Co-design* on device-circuit integration, recognizing that predictable, scalable devices enable all application domains.
- **Application-Domain PRDs:** address distinct scales and contexts where analog computing offers transformative advantages. *PRD 2 (High-Performance Computing)* synthesized contributions from *Analog Electrical and Computational Memory*, *Devices/Photonics*, *Probabilistic/Ising*, *Mathematical Foundations*, and *Benchmarking*, reflecting the diverse technologies needed for large-scale scientific computing. *PRD 4 (Edge Computing)* emerged from *Hybrid Systems/Sensor+Compute Integration*, with contributions from *Analog Electrical*, *Devices*, and *Co-design*, addressing ultra-efficient computation integrated with sensing systems. *PRD 5 (Natural Computing)* combined *Biological/Chemical Computing* and *Probabilistic/Ising*, with input from *Mathematical Foundations* and *Devices*, exploring physics-based substrates operating near thermodynamic limits.
- **Cross-Cutting Methodologies:** span all technology areas and application domains. *PRD 6 (Co-Design)* merged *Co-design/Software+Hardware+Physics* and *Benchmarking*, with all breakouts contributing cross-cutting themes emphasizing integrated design methodologies spanning materials through applications.

Key Interconnections. As illustrated in Figure 2, these PRDs form an integrated portfolio. PRD 1 provides theoretical frameworks and fundamental limits informing all applications. PRD 3 supplies physical substrates while receiving requirements feedback from application domains. PRD 6 integrates vertically across abstraction layers and horizontally across PRDs through unified toolchains and benchmarks. Application domains cross-fertilize: HPC precision techniques inform edge deployments; edge low-power methods benefit HPC at scale; natural computing optimization algorithms apply to HPC and edge problems. Shared challenges—noise and variability (Section 2.3.1), analog-digital interfaces (Section 2.3.2), scalability (Section 2.3.3)—require coordinated research where solutions in one PRD accelerate progress across others. Success depends on viewing the six PRDs as mutually reinforcing rather than independent, with advances in each enabling progress in others.

2.3 Scientific Challenges

Several fundamental scientific challenges recur across all analog computing approaches. Understanding and addressing these challenges is essential for the field's advancement. This section consolidates the theoretical and practical treatment of these cross-cutting issues, which are referenced throughout the subsequent PRDs.

2.3.1 Noise, Variability, and Environmental Effects

Variations in physical and electrical parameters of devices resulting from fabrication processes ("parameter variability"), environmental changes (such as temperature, pressure, *etc.*), and noise in devices all affect the functioning of analog circuits. Variability and noise are often discussed together because they are both modeled probabilistically, but it is useful to consider them separately: parameters subject to variability are typically constant with respect to time, hence can be modeled mathematically as random variables; whereas noise changes with time, hence requires the more complicated mathematical machinery of stochastic processes. The physical mechanisms underlying parameter variability and device noise are also typically quite different.

Device Mismatch and Parameter Variability. In digital systems, voltages representing Boolean logic values are typically widely separated, with values in between being (in principle) illegal. A corrupted digital voltage can be thresholded to the nearest legal value; this is a powerful correction mechanism that is not usually available in analog systems, where variables take a continuous range of values that are all legal. As a result, analog circuits are generally assumed to be more sensitive to the effects of parameter variability, environmental variations, and noise than digital circuits. However, it is possible for a carefully designed analog system to be robust to these effects. Detailed low-level analysis of these effects, properly grounded in mathematical theory, is critical for assessment.

As Mead wrote, "not all transistors are created equal"; as a result, identically drawn devices are not identical, but have variations between them—referred to as *mismatch*. Parameter variability and mismatch arise from inherent randomness in manufacturing processes; any fabricated device's physical and electrical parameters (such

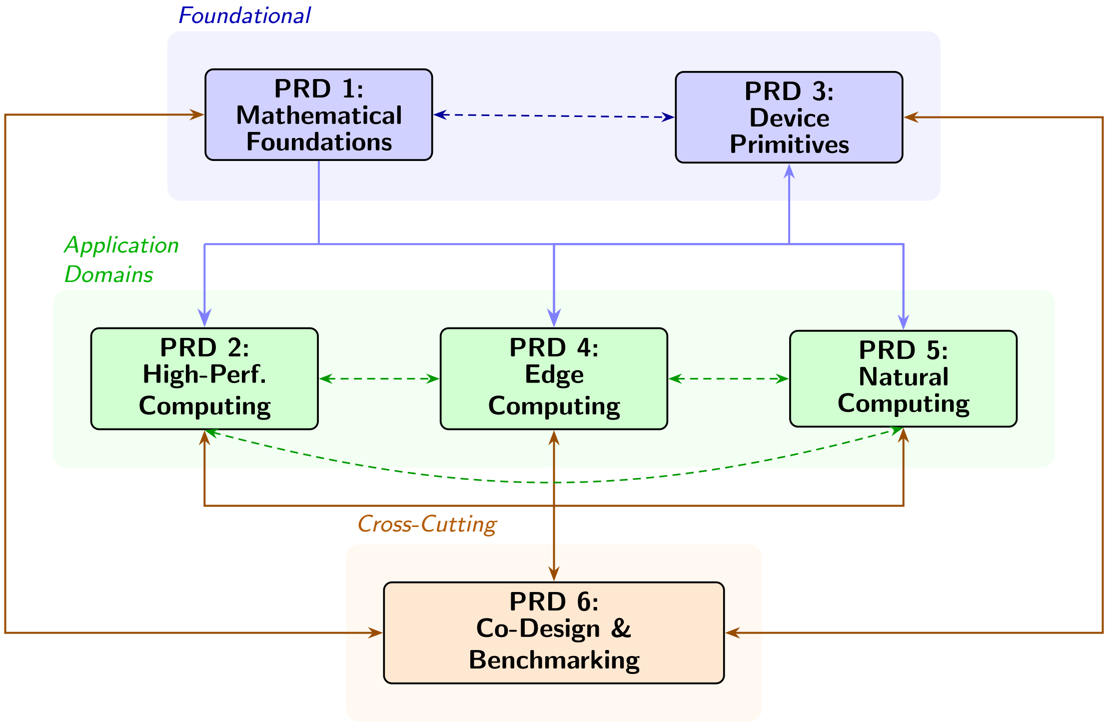

Figure 2: Interconnections among the six Priority Research Directions. Foundational PRDs (PRD 1: Mathematical Foundations, PRD 3: Device Primitives) provide theoretical frameworks and physical substrates supporting application-domain PRDs (PRD 2: High-Performance Computing, PRD 4: Edge Computing, PRD 5: Natural Computing). Cross-cutting methodologies (PRD 6: Co-Design) integrate vertically across abstraction layers and horizontally across all PRDs. Solid arrows indicate primary dependencies and information flow, bidirectional arrows show mutual reinforcement and requirements feedback, and dashed arrows represent cross-fertilization among application domains. Background shading groups PRDs by tier.

as device lengths/widths, oxide thicknesses, doping strengths, threshold voltages, *etc.*.) will generally be different from nominal values specified during design. The relative impact of parameter variations increases as devices become physically smaller, hence assessing and mitigating the impact of variability is especially important for novel nano-devices.

Mathematically, parameters subject to variability are usually modeled as random variables, with their probability distributions typically obtained empirically from measurements after fabrication. Techniques to mitigate the impact of variability can be at the process/fabrication/layout level, or at the circuit/system design level. Examples of the former include averaging over multiple devices to reduce variance, selecting the best candidate from multiple redundant devices, routing to target minimizing timing variations, and device level calibration using, *e.g.*, floating-gate (FG) devices (which feature a mechanism to change threshold voltages at run time)[16;17].

Clever circuit designs —for example bandgap reference circuits, proportional to absolute temperature (PTAT) compensation circuits—have long been used, without the need for device- or layout-level calibration, to achieve robust analog performance in the face of variability and environmental changes. Such techniques have been based on *simple analytical models* of device dependence on parameters; for example, the exponential dependence of current through a PN junction on temperature. Developing similar simple models for (*e.g.*) novel nanodevices, and using them to devise new circuit structures that are robust to variability/environmental changes is a necessary direction for future efforts.

Another important direction for future investigation is the problem of *probabilistic characterization* of circuit or system performances. Such characterization is crucial for assessing whether a circuit or system is capable of operating robustly. For virtually all interesting or important analog circuits, the mapping from device parameters to outputs/performances is strongly nonlinear, as well as dynamical. Currently, probabilistic analysis of linearized approximations[18] and brute-force Monte-Carlo methods[19] constitute the mainstay of practically-useful methods. Both approaches have significant shortcomings—*e.g.*, in terms of accuracy/relevance for linearized analyses, and computational cost for Monte-Carlo. Developing "smart" and effective analytical/computational methods for propagating probability distributions through nonlinear systems (defined by algebraic or differential equations) can make a big improvement to the *status quo*.

Noise in Analog Systems. Noise, unlike variability, stems intrinsically from unpredictable fluctuations at the atomic and molecular levels. One important intrinsic noise generation mechanism is *thermal noise* in conductors, arising from thermodynamic fluctuations in the occupation of electronic states at any non-zero temperature. Other important types of intrinsic noise include *shot noise* and *flicker noise*. Noise propagates through circuits and systems and affects the fidelity of desired signals; in analog systems, the signal-to-noise ratio (SNR) is a metric often used to assess fidelity.

As mentioned earlier, noise is modeled mathematically as stochastic processes, *i.e.*, random variables that are functions of time. Almost all intrinsic noise in circuit and system design is currently modeled as *stationary* stochastic processes[20], wherein noise statistics are invariant with respect to time shift. But if nonlinearities are present, stationary noise models are valid only if the circuit or system's deterministic operation (*i.e.* no voltages or currents) does not change with time. However, analog systems in dynamical operation, in particular ones performing computation, typically feature large changes with respect to time.

Unfortunately, (and in spite of thermal, shot, and flicker noise models having deep thermodynamic and quantum-mechanical underpinnings[21]) existing macroscopic models tend to be valid only for the simple case of no time-variations in the operation of the circuits and systems involved. Theory for understanding and modeling noise generation in the presence of large-signal changes is lacking—particularly with respect to flicker noise and nonlinear circuit effects.

Accounting for correlations between noise sources that are physically proximate at atomistic scales is also an important aspect of noise generation models that needs more attention. To devise rigorous non-stationary noise models, abstracting particle level approaches is one promising direction. For example, the Gillespie molecular-level model of chemical reactions[22] and its abstraction to macroscopic noise models in the Langevin equation is a successful example of noise modeling and analysis that connects atomistic models of noise to system-level ones. Such first-principles, bottom-up approaches hold promise for developing comprehensive, theoretically grounded models of large signal noise suitable for system level analysis.

Developing first-principles, theory-based, noise models hand in hand with experimental measurements would advance analog design significantly. Also, while stochastic analysis techniques exist for some special cases of time-varying and nonlinear systems (periodic operation, cyclostationary[23]; oscillator phase noise[24]), the general problem of stochastic characterization of noise propagation through nonlinear dynamical circuits and systems is largely open theoretically, computationally, and practically. New mathematical techniques to quantify computational complexity under physical noise would be an important component of any unified mathematical framework for analog computing.

Noise as a Computational Resource. Rather than merely combating noise in analog systems, emerging research explores harnessing it as a computational resource, recognizing that noise can provide valuable functionality in certain computational contexts.

Most conventional microelectronics is engineered to minimize noise at all scales, which is essential for deterministic calculations. However, many computational applications—including stochastic optimization, machine learning, and Monte Carlo techniques—rely on sampling methods that could potentially leverage inherent physical noise. These applications currently require the artificial reintroduction of randomness into deterministic systems via pseudo-random number generators (PRNGs). While modern PRNGs are high quality, their repeated use incurs significant computational overhead in applications requiring extensive sampling.

For this reason, in recent years, there has been renewed interest in hardware-based methods to enable probabilistic computing. One such approach leverages fluctuating devices called p-bits[25;26], essentially realizing the simulated annealing algorithm[27] using noisy analog hardware that samples a range of behavior. A variation on p-bits leverages stochastic devices as a source of samples for statistical calculations, including probabilistic neuromorphic algorithms[28–30]. Such true random number generator approaches have been explored in the past, but these have now been achieved at large scales[31] to generate non-uniform random samples[32;33], producing significant efficiency by moving complexity from other parts of the calculation into complex sampling schemes. Noise sources based on specific statistics and annealing properties[34] have been shown to be useful for solving combinatorial optimization problems. Judicious amounts of noise have also been observed to lead to improved optimization in oscillator Ising machines[35]. However, mechanisms behind these improvements are poorly understood and developing a theoretical understanding how noise affects analog optimization schemes remains essential.

2.3.2 Interfacing Analog and Digital Systems

Creating efficient, bidirectional interfaces between analog components and conventional digital systems is a critical challenge that must be addressed to build practical, heterogeneous computing platforms. This challenge appears in every application domain, from high-performance computing to edge sensing to natural computing systems.

Specialized analog devices often outperform general-purpose digital systems in a single mathematical operation but are rarely capable of executing the complete set of algorithms required by an application, thus usually requiring a digital interface. Unfortunately, analog-to-digital and digital-to-analog converters (ADC, DAC) can introduce substantial energy, footprint, and latency overheads, potentially thwarting the benefits of analog approaches[36]. A key scientific challenge is how to effectively couple analog and digital signal processing with minimal loss in energy and latency.

Addressing this challenge will require new innovations in how data is represented and manipulated, such as digital transistors that are used to emulate analog hardware, analog devices which can be used in a digital fashion, and more complex analog systems which incorporate a much wider range of computational abstractions. The conversion costs between analog and digital domains must be carefully accounted for in any system-level analysis, as these costs can dominate overall performance and energy consumption.

For optimal efficiency in analog design, it is preferable to minimize data conversions whenever possible. In-domain computing approaches—such as all-optical processing or direct chemical computation—can avoid domain conversion entirely, though they introduce their own integration challenges. The trade-off between conversion overhead and the benefits of hybrid analog-digital approaches is highly application-dependent and represents an active area of research.

2.3.3 Scalability and System Integration

Scaling analog computing systems from laboratory demonstrations to practical, large-scale implementations presents fundamental challenges distinct from those encountered in digital systems. Unlike digital circuits where identical components can be replicated with predictable behavior, analog systems face compounding effects of device variability, parasitic capacitances, signal integrity issues, and interconnect complexity as system size increases.

The connectivity required between analog components scales with system size, and layout complexity can become a bottleneck. Analog systems' sensitivity to device mismatches and limited dynamic range can degrade accuracy as systems grow. Wire capacitance, resistance, and crosstalk become increasingly significant issues in large arrays. Error propagation through multiple stages of analog computation can accumulate in ways that fundamentally limit system size.

However, these constraints do not eliminate analog advantage—rather, they guide us toward problem classes where analog excels: systems with sparse or structured connectivity, well-conditioned or noise-tolerant formulations, and problems naturally expressed as continuous dynamics. Understanding and quantifying these scaling limitations is essential for identifying appropriate application domains.

To assess scaling confidently requires understanding how a novel device fits into a larger system, how its fabrication and performance depend on other parts of the hardware, and how interactions between distinct subcomponents occur during operation. In other words, for these novel computing devices to be useful, they need to be abstractable, with models that enable the design of realistic systems that perform practically relevant algorithms. This requires the design of devices with predictable performance at scale, in diverse systems, as well as robust fabrication processes compatible with mainstream chip manufacturing[37].

2.4 Cross-Cutting Methodologies

The preceding sections identified fundamental scientific challenges—noise, interfaces, programmability, and scalability—that must be addressed for analog computing to realize its potential. Beyond understanding these challenges, realizing practical analog computing systems requires methodological advances that span the entire technology stack. Three cross-cutting methodologies emerged as essential themes throughout the workshop: co-design across hardware and software layers, hardware-aware algorithm development, and standardized benchmarking frameworks. These methodologies, detailed comprehensively in PRD 6, provide the integrated approach needed to translate analog computing research into deployable systems.

2.4.1 Hardware-Software Co-Design Principles

Hardware-software co-design has been successfully deployed by the U.S. DOE to enhance software and hardware in leadership-class HPC systems[9;38]. Recent exascale computing initiatives—including HPE's Aurora, Frontier, and El Capitan—demonstrate the efficacy of co-design by enabling mission-relevant capabilities through close collaboration with system vendors. Analog systems can benefit significantly from this form of co-design, which enables specialization of the hardware and allows practitioners to pursue designs with increased resource-efficiency and fidelity, at the cost of generality.

Analog systems are fundamentally different from digital systems, requiring a new generation of co-design tools that can simultaneously optimize across all system levels, from basic materials to high-level algorithms. These methods need to be adaptable to various existing and emerging analog computing approaches while also seamlessly integrating with digital systems. The tools must allow flexible exploration of system configurations, be reusable across projects, and include safeguards to prevent improper applications of models.

The conventional approach to developing complete analog computing systems has often been highly serialized, with device designers, circuit architects, and software developers working in isolation. This can take many years, if not decades, for a device technology to eventually make its way into a commercially viable end system. Full-stack co-design requires researchers to be able to work across multiple stack layers and communicate with colleagues at distant layers of the analog computing stack.

Key co-design principles that apply across all PRDs include:

- **Simultaneous optimization**: Hardware, software, and algorithms should be designed together, not sequentially
- **Hardware-aware algorithms**: Computation should be adapted to leverage analog hardware characteristics while compensating for non-idealities
- **Abstraction and modularity**: Design layers should be connected through well-defined abstractions that enable independent evolution
- **End-to-end evaluation**: System performance should be assessed from application workload to physical implementation
- **Iterative refinement**: Co-design is an iterative process requiring tools for rapid prototyping and evaluation

These co-design principles are elaborated in detail in PRD 6, which provides comprehensive methodologies, tools, and frameworks for implementing co-design across the full stack from materials to applications.

2.4.2 Hardware-Aware Algorithm Development

Designing algorithms specifically for analog hardware can enable faster, more energy-efficient computing by leveraging continuous signals, natural parallelism, and inherent nonlinearity. Rather than imposing a digital design approach on a continuous platform, hardware-aware algorithm development integrates both hardware and software considerations to align each with the underlying physics of analog devices.

Traditional algorithms are designed with assumptions of perfect precision, deterministic behavior, and discrete state spaces—assumptions that do not hold for analog systems. Hardware-aware algorithms must:

- **Tolerate imprecision**: Account for limited precision and dynamic range
- **Exploit physical properties**: Leverage natural dynamics, parallelism, and continuous-time operation
- **Compensate for non-idealities**: Include calibration, error correction, and adaptive techniques
- **Optimize for energy**: Minimize conversions between analog and digital domains
- **Match problem structure**: Align computational structure with hardware capabilities

Examples of hardware-aware algorithms include gain compensation techniques in neural networks implemented on analog crossbars, stochastic optimization algorithms that exploit device noise, and numerical methods designed for the unique error characteristics of analog differential equation solvers.

2.4.3 Benchmarking and Evaluation Frameworks

Effective engineering relies on the ability to compare competing designs on well-defined benchmark problems. For digital systems, benchmarks such as Standard Performance Evaluation Corporation (SPEC), High-Performance Linpack (HPL), and MLPerf have helped focus research efforts and enable comparisons between disparate architectures and systems. Analog computing requires similar standardized evaluation frameworks, but with metrics appropriate to continuous computation.

Key challenges in analog benchmarking include:

- **Diverse performance metrics**: Energy efficiency, latency, throughput, accuracy, and area efficiency must all be considered
- **Application-specific trade-offs**: The optimal design point varies significantly with application
- **Lack of standard problems**: Different researchers select different benchmark problems, complicating comparisons
- **Simulation vs. hardware**: Benchmark performance often relies on specific device behaviors that may be unavailable to other researchers

Developing community-driven benchmarks and standardized evaluation methodologies is essential for assessing progress and guiding future investments in analog computing. This includes defining representative workloads, standardized metrics, reference implementations, and accessible hardware platforms for validation. The benchmarking challenge is addressed comprehensively in PRD 6, Section 8.4.5.

2.5 Manufacturing, Characterization, and Metrology

Translating promising analog devices and circuits into deployable systems depends on a coherent, cross-cutting program in manufacturing, characterization, and metrology. Today, variability and limited process portability impede confident scaling from laboratory prototypes to large systems. A unified roadmap avoids duplicated effort across PRDs by defining common protocols, shared infrastructure, and clear maturity targets that device, system, and application teams can jointly drive toward.

Scope and Objectives. This cross-PRD effort focuses on: (1) establishing reproducible, CMOS-compatible process flows and process design kits (PDKs); (2) building standardized characterization and metrology that deliver statistical device and circuit models usable by designers; (3) quantifying reliability through accelerated life testing and drift/noise characterization; and (4) defining maturity milestones aligned with system scale-up. These needs were identified across PRDs and the broader community[37;39].

Process and PDK Foundations.

- **CMOS compatibility and integration:** Define minimal-complexity add-on modules for novel devices that preserve back-end compatibility (e.g., copper interconnect), enable monolithic 3D stacking where beneficial, and articulate design rule check (DRC)/layout versus schematic (LVS) rules suitable for foundry transfer.
- **Statistical PDK content:** Provide compact models with mismatch, variability, and aging parameters; process corners; and yield-aware abstractions consumable by circuit/system tools. Short-term targets: $<10\%$ device-to-device variability; long-term: $<5\%$ for high-volume manufacturing.
- **Multi-lab reproducibility:** Standardize key steps and witness structures to enable direct replication across sites with quantified process windows.

Characterization and Metrology.

- **Standard protocols:** Define shared test methods for I-V dynamics, noise spectra, programming linearity, retention/drift, endurance, and temperature dependence; publish reporting formats and confidence intervals enabling cross-technology comparison.
- **Test structures and fixtures:** Curate a community set of on-wafer structures and packaged vehicles for high-throughput statistical testing and accelerated stress, including interfaces for hybrid analog-digital evaluation.
- **From device to circuit statistics:** Translate device distributions into circuit- and block-level statistical models (including mismatch and temporal noise) suitable for Monte Carlo and uncertainty propagation in

system simulations.

Standardized evaluation complements these activities through shared metrics and benchmarks; see Sections 2.4.3 and 2.4.1 for evaluation frameworks and co-design integration.

Reliability and Maturity Metrics.

- **Accelerated reliability:** Establish accelerated stress protocols (temperature, voltage, cycling) and physics-informed extrapolation to mission profiles relevant to HPC facilities and fielded edge systems.
- **Maturity ladder:** Define stage gates from *Lab Prototype* → *Replicable Process + PDK* → *Pilot Manufacturing with Statistics* → *high-volume manufacturing (HVM)-ready*; illustrative thresholds for the terminal stage include $< 5\%$ device-to-device variability, scalability to 10^{6+} devices per system, full-stack statistical design flows, and reliability sign-off tied to DOE deployment environments.
- **System-relevant KPIs:** Couple metrology to system KPIs (effective precision, SNR, calibration overhead, energy-per-op including interface costs) so that process decisions map to application value.

Coordination with PRDs.

- **PRD 2 (High-Performance Computing):** Addresses HPC-specific manufacturing challenges—wafer-scale economics (1000+ wafers), mask-set cost amortization through reconfigurability, monolithic 3D integration for 10–100× density improvements, system-scale yield modeling for 100,000-element systems, and calibration/runtime metrology tied to effective precision. Adopts process/PDK foundations and statistical models from this section.
- **PRD 3 (Device Primitives):** Develops device physics understanding of modality-specific fabrication challenges—physical mechanisms governing variability in memristive, phase-change, electrochemical, and spin-based devices. Provides physics-informed fabrication strategies and device-specific characterization templates that instantiate the cross-cutting protocols defined here. Owns compact modeling frameworks integrating these physics into circuit-level abstractions.
- **PRD 5 (Natural Computing):** Extends standardized characterization and benchmarking patterns to biochemical and physical substrates; develops modality-appropriate assays, biosafety protocols, and reproducibility standards analogous to those defined here while avoiding re-specification of core manufacturing frameworks.

2.6 State of the Art in Analog Computing

Analog computing has experienced renewed interest in recent years, driven by advances in device technology, improved understanding of analog computation theory, and the pressing need for energy-efficient alternatives to digital computing. Current state-of-the-art demonstrations span multiple domains:

Analog Neural Networks and Machine Learning. The intersection of analog computing and AI represents one of the most active research areas, driven by the energy crisis in AI deployment and the natural alignment between neural network operations and analog hardware. In-memory computing using analog crossbar arrays has demonstrated significant energy-efficiency improvements for neural network inference[40], with commercial efforts (e.g., Mythic AI, Analog Inference) showing practical implementations of analog machine learning (ML) accelerators. Recent work has demonstrated unified platforms addressing both AI inference and combinatorial optimization: an analog optical computer (AOC) that achieves a projected efficiency of 500 TOPS/W (over 100× more efficient than leading GPUs) by eliminating analog-digital conversions through fixed-point searches in the analog domain has been introduced[7]. For LLMs specifically, it has been demonstrated[8] that mixture-of-experts architectures deployed on 3D non-volatile memory (NVM)-based analog in-memory computing can substantially reduce inference costs, leveraging conditional computing to address parameter-fetching bottlenecks.

Reconfigurable Analog Platforms. Field-programmable analog arrays (FPAAs) provide reconfigurable platforms for rapid prototyping and deployment of analog signal processing systems. Large-scale FPAAs with hundreds of thousands of programmable elements enable sophisticated analog computations while maintaining flexibility comparable to field-programmable gate arrays (FPGAs) in the digital domain. These platforms have demonstrated applications in edge computing, sensor processing, and adaptive signal processing[41;42].

Photonic and Optical Computing. Photonic integrated circuits enable analog computation at the speed of light with exceptional energy efficiency for specific operations. Coherent optical systems have demonstrated matrix-vector multiplication and neural network inference with bandwidths exceeding electronic implementations[43]. Optical Ising machines and photonic tensor cores show promise for optimization and machine learning applications[44;45], though challenges remain in integration, reconfigurability, and interfacing with electronic systems.

Superconducting Analog Computing. Superconducting circuits operating at cryogenic temperatures offer unique advantages for analog computation, including extremely low energy dissipation and fast switching speeds. Josephson junction-based circuits demonstrate potential for high-speed signal processing and neuromorphic computing[46]. While requiring cryogenic infrastructure limits deployment, superconducting approaches may be valuable for specific high-performance computing applications where cooling overhead is acceptable.

Physics-Based and Neuromorphic Computing. Coherent Ising machines, oscillator networks, and other physics-based computing approaches have demonstrated solutions to optimization problems with thousands of variables[44;47;48]. Neuromorphic systems such as Intel's Loihi[49] and IBM's TrueNorth[50] explore brain-inspired architectures, though most current implementations use digital circuits. Analog neuromorphic approaches promise greater efficiency but face challenges in connectivity and fan-out.

Biological and Chemical Computing. Deoxyribonucleic Acid (DNA)-based computation has achieved proof-of-concept demonstrations of neural networks, oscillators, and pattern recognition[51;52]. Synthetic biology approaches have implemented analog control systems and logic gates in living cells. While these approaches remain largely at the research stage, they demonstrate the feasibility of computation in molecular substrates[53].

Hybrid Analog-Digital Systems. Commercial analog-digital hybrid systems are deployed in applications ranging from sensor interfaces to power management[13]. The trend toward heterogeneous computing platforms creates opportunities for analog accelerators integrated with conventional processors[9;38]. However, interface costs remain a significant challenge that limits the broader adoption of analog approaches[54;55].

Despite these advances, significant gaps remain between laboratory demonstrations and practical, scalable analog computing systems suitable for DOE mission applications. The following Priority Research Directions identify the key challenges and opportunities for advancing analog computing from current capabilities toward transformative impact in scientific computing.

2.7 Relative Research Maturity Across Priority Research Directions

To assess research priorities and guide strategic investment decisions, we conducted a systematic evaluation of scientific maturity across all six PRDs. The assessment examines seven fundamental capability dimensions—theoretical foundations, computational tools, hardware demonstrations, performance characterization, noise and error management, system integration, and manufacturing—using the following five-level scale:

- **Level 1 (Exploratory):** Fundamental concepts emerging with major gaps; research questions forming
- **Level 2 (Foundational):** Basic principles established; significant challenges remain
- **Level 3 (Established):** Validated models with reproducible results; community adoption beginning
- **Level 4 (Validated):** Robust DOE-scale performance; transition pathways clear
- **Level 5 (Research-Mature):** Comprehensive understanding; ready for applied research and development (R&D) transition

Each rating reflects expert judgment informed by the state-of-the-art analyses in Sections 3–8, and should be read as a relative indicator rather than an absolute technology readiness assessment.

As shown in Figure 3, the evaluation reveals coherent maturity progression aligned with PRD roles and technology readiness. The entire portfolio resides at Levels 1–3, reflecting the early-stage nature of analog computing as a field with clear pathways toward higher maturity rather than near-term deployment technology. Foundational PRDs (PRD 1: Mathematical Foundations, PRD 3: Device Primitives) establish theoretical frameworks and hardware substrates at Level 1–2, providing the basis for application domains. Application PRDs leverage these foundations with varying maturity: PRD 2 (High-Performance Computing) and PRD 4 (Edge Computing) achieve Level 2–3 by building on CMOS-compatible approaches with commercial validation; PRD 5 (Natural Computing) remains at Level 1–2, reflecting fundamental challenges in novel biological and chemical substrates. Cross-cutting PRD 6 (Co-Design) enables integration across all domains but shows tool infrastructure gaps (Level

1) that limit broader accessibility.

The matrix highlights cross-cutting patterns in the field's current state. Common capability gaps span all PRDs: computational tools infrastructure (Level 1–2) represents a critical bottleneck requiring coordinated development; system integration (Level 1–2) remains a research frontier with multi-scale challenges; manufacturing reproducibility varies widely from exploratory demonstrations (Level 1) to commercial deployment (Level 3). Conversely, bright spots demonstrate field vibrancy: hardware demonstrations reach Level 3 across five of six PRDs with both commercial products and DOE instrumentation deployed; PRD 2 achieves Level 3 theory for specific validated paradigms; noise management shows uniform foundational progress (Level 2) across all approaches. These patterns are consistent with the shared scientific challenges (Section 2.3), cross-cutting methodologies (Section 2.4), and manufacturing priorities (Section 2.5) discussed in preceding sections.

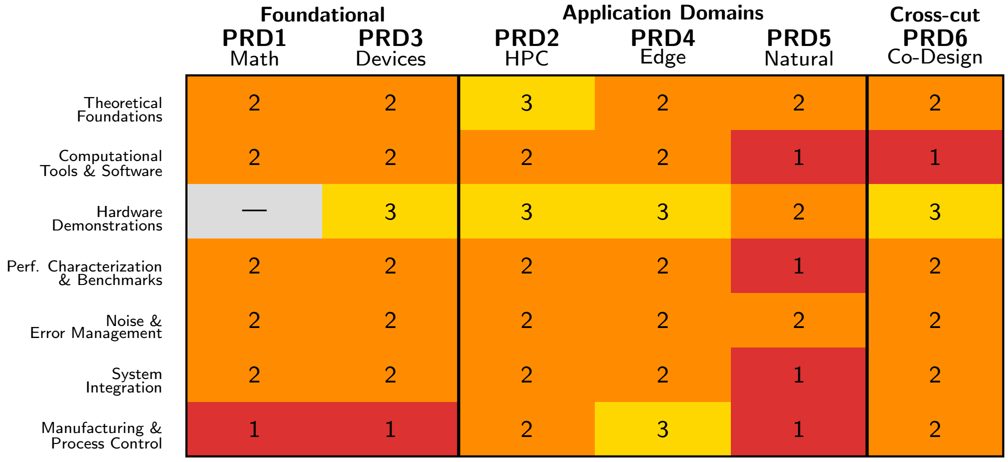

	Foundational		**Application Domains**			**Cross-cut**
	PRD1 Math	**PRD3** Devices	**PRD2** HPC	**PRD4** Edge	**PRD5** Natural	**PRD6** Co-Design
Theoretical Foundations	2	2	3	2	2	2
Computational Tools & Software	2	2	2	2	1	1
Hardware Demonstrations	—	3	3	3	2	3
Perf. Characterization & Benchmarks	2	2	2	2	1	2
Noise & Error Management	2	2	2	2	2	2
System Integration	2	2	2	2	1	2
Manufacturing & Process Control	1	1	2	3	1	2

Figure 3: Relative research maturity matrix for the six PRDs across seven capability dimensions, using the five-level scale defined above. Heatmap shows Foundational PRDs (Math, Devices) at Level 1–2 providing theoretical and hardware bases, Application PRDs (HPC, Edge, Natural) at Level 2–3 building on these foundations, and Cross-cutting Co-Design enabling integration. Key patterns: computational tools (Level 1–2) represent common bottleneck across all PRDs; hardware demonstrations (Level 2–3) are ahead of manufacturing readiness (Level 1–3); noise management shows uniform foundational progress (Level 2); system integration (Level 1–2) remains frontier challenge.

3 PRD 1: MATHEMATICAL FOUNDATIONS OF ANALOG COMPUTING

3.1 Introduction and Vision

Analog computation represents a fundamentally different computational paradigm offering critical advantages for specific tasks. While concrete physical devices (PRDs 2–6) provide crucial proof-of-principle demonstrations, these advances must be coupled with rigorous mathematical frameworks to unlock analog computation's full potential and enable predictable, scalable implementations. Mathematical theory guides hardware development, defines fundamental limits, characterizes complexity classes, and enables principled algorithm design.

This PRD establishes cross-cutting theoretical foundations that underpin the domain-specific PRDs and are concretely instantiated in each: high-performance analog computing (PRD 2), edge and in-sensor computing (PRD 4), device primitives and materials (PRD 3), co-design and toolchains (PRD 6), and natural and bio-inspired computing (PRD 5).

Key Research Questions: *What are the canonical models of analog computation, and can we establish completeness results analogous to Boolean circuit completeness? Where does analog computation provide provable advantages with rigorous characterization of problem classes showing asymptotic speedups? How do physical constraints (noise, precision, energy) affect computational power and complexity? Can continuous methods illuminate discrete problems, revealing hidden structure in combinatorial problems and enabling novel hybrid algorithms?*

Vision: Establish comprehensive mathematical foundations for analog computation through: (1) rigorous theoretical framework with completeness proofs, complexity classes, and formal models guiding hardware development (PRDs 2–6); (2) validated predictive models for noise, variability, and precision enabling reliable algorithm design; (3) principled algorithm design methodologies exploiting analog-specific advantages while managing physical constraints; (4) rigorously characterized problem classes where analog provides demonstrable advantages; and (5) widely accepted frameworks, verification methods, and community standards adopted across research institutions and industry.

3.2 Goals and Desired Outcomes

Establishing rigorous mathematical foundations is essential to transform analog computing from empirical demonstrations to principled engineering practice. Main research needs follow.

- **Canonical models and complexity theory:** Establish formal canonical models with completeness proofs for major analog computation approaches (dynamics-based ordinary differential equation (ODE)/PDE systems, optimization-based frameworks, stochastic/probabilistic methods) and develop rigorous complexity-theoretic framework characterizing computational classes naturally suited to analog hardware. Identify minimal complete primitive sets to guide hardware implementation.
- **Analog advantage characterization:** Establish rigorous asymptotic speedup theorems for problem classes where analog computation demonstrates theoretical advantage, while identifying fundamental limits clarifying where analog methods cannot compete with digital approaches.
- **Precision and noise theory:** Formalize relationships between numerical precision, noise characteristics, and algorithmic convergence for analog systems, developing analog numerical analysis that accounts for unique error modes and noise-aware optimization frameworks.
- **Theory-hardware co-development:** Establish iterative co-development cycle where theoretical predictions guide hardware experiments on analog platforms (PRDs 2–6) while hardware insights drive theoretical advances, enabling parallel progress in both domains. This bidirectional feedback accelerates both mathematical foundations and practical implementations while building research community adoption.
- **Unified Mathematical Framework:** Develop comprehensive mathematical framework unifying diverse analog computation paradigms, clarifying relationships between dynamics-based, optimization-based, and stochastic approaches while establishing bridges to digital complexity theory.
- **Algorithm Design Principles:** Establish principled approaches to algorithm design that exploit analog-specific capabilities (continuous-time operation, natural parallelism, physical dynamics) while managing constraints, enabling systematic development of noise-aware optimization algorithms that leverage stochastic dynamics as both constraint and resource.

- **Cross-Domain Validation and Adoption:** Demonstrate mathematical foundations enabling applications across scientific domains with theory-driven algorithm development validated experimentally, while establishing community-accepted standards for theoretical analysis, verification methods, and educational resources that make the field broadly accessible.

3.3 State of the Art

Fragmented Theoretical Landscape: Unlike digital computation's unified framework (Turing machines, Boolean circuits, random access memory (RAM) models with well-understood equivalences), analog computation lacks canonical models with proven completeness results. Shannon's foundational work[56] on polynomial ODEs and recent extensions[57;58] provide important starting points, but comprehensive theory remains underdeveloped compared to digital foundations from the 1960s–1970s. The contrast with digital theory is stark: Post's lattice[59] exactly enumerates the closed classes of Boolean functions, providing a complete mathematical taxonomy for digital circuits, while no similar comprehensive understanding exists for function classes relevant to analog computation. Initial efforts in analog abstraction[60] provide starting points, but developing analogous theoretical frameworks for analog function classes represents a fundamental research need. However, unlike the sequential development path that digital computing followed, modern capabilities—automated theorem proving, symbolic computation systems, machine learning-driven mathematical discovery, and validated simulation frameworks—enable parallel co-development of theory and hardware, with each informing and accelerating the other. This synergy, combined with lessons learned from 75 years of digital computing, can dramatically compress the timeline from foundational theory to practical deployment.

Incomplete Noise Theory: While fundamental noise mechanisms are understood (Section 2.3.1), mathematical theory for non-stationary operation—critical for analog computing systems with large signal variations—remains incomplete. Existing models primarily address stationary conditions, limiting predictive power for dynamical computation.

Uncharacterized Advantage Domains: Despite promising examples (chaotic systems, stiff ODEs, well-conditioned linear systems), rigorous mathematical characterization of where analog computation provides asymptotic advantages remains incomplete. The lack of a complexity-theoretic framework hinders systematic identification of suitable problem classes.

Limited Algorithm Design Principles: Algorithm development for analog systems remains largely *ad hoc*; systematic design methodologies that account for analog-specific capabilities (continuous-time, natural dynamics) and constraints (noise, variability) from first principles are absent.

Theory-Hardware Gap: Bidirectional feedback between mathematical theory and hardware implementation (PRDs 2–6) remains insufficient: theoretical advances often lack experimental validation, while hardware innovations proceed without rigorous mathematical foundations.

3.4 Research Directions

Three interdependent research directions structure this PRD, as illustrated in Figure 4. Canonical models and complexity theory establish the formal foundations; analog advantage characterization and algorithm design identify where and how those foundations translate into computational gains; and noise, precision, and physical constraints theory defines the physical envelope within which both operate—feeding back to shape the models and motivate the algorithms above.

3.4.1 Canonical Models, Completeness, and Complexity Theory

Digital computation's exponential growth from the 1970s to early 2010s was enabled by strong theoretical foundations. Turing machines, Boolean circuits, and RAM models serve as well-established canonical models with understood connections. Critically, completeness results guide hardware: since NOT-OR (NOR) gates are complete for Boolean circuits, one need only build NOR gates and compose them to realize arbitrary circuits—a principle demonstrated by the Apollo Guidance Computer's all-NOR CPU[61]. Analog computation lacks such canonical models and completeness understanding, hindering systematic hardware development.

Research objectives center on establishing canonical models for major analog computation paradigms, with rigorous completeness proofs:

Dynamics-based models: Shannon's foundational result[56] showed polynomial ODEs (adders, multipliers,

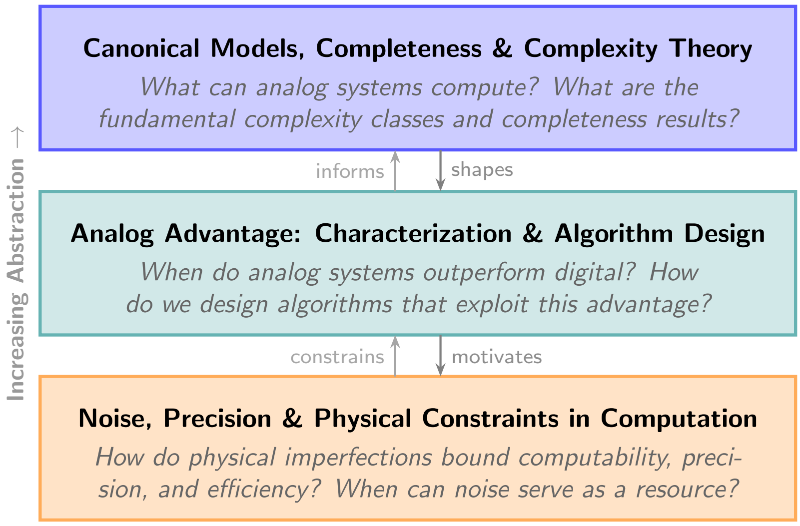

Canonical Models, Completeness & Complexity Theory

What can analog systems compute? What are the fundamental complexity classes and completeness results?

Analog Advantage: Characterization & Algorithm Design

When do analog systems outperform digital? How do we design algorithms that exploit this advantage?

Noise, Precision & Physical Constraints in Computation

How do physical imperfections bound computability, precision, and efficiency? When can noise serve as a resource?

Figure 4: Overview of the three research directions in PRD 1. Arrows indicate the primary relationships between directions: canonical models shape the space of algorithmic inquiry, algorithm design informs which models warrant formalisation, noise and precision theory motivates new algorithmic strategies, and in turn algorithmic requirements constrain what physical realisations must achieve.

integrators, constants) generate a broad function class. Systematize implementation methods across mechanics, electronics, chemical reaction networks[62], and other physical domains. Extend to differential-algebraic equations (DAEs) modeling complex circuits (Simulation Program with Integrated Circuit Emphasis (SPICE)-like systems).

Optimization-based models: For discrete problems, examples include Hopfield networks[63], ODE-based Boolean satisfiability (SAT) solvers[64], and analog Ising machines[65] with applications spanning graph-based problems, resource allocation, and configuration space exploration. For continuous problems, gradient-based analog methods address PDE-constrained optimization and optimal control challenges. Characterize when global minima appear as attractive fixed points through arithmetization and Lyapunov function methods, and establish when continuous relaxations of discrete problems enable tractable analog solutions.

Stochastic models: Develop completeness theory for probabilistic analog computation, including noise-enhanced optimization and sampling methods.

Complexity-theoretic framework: Develop complexity-theoretic framework defining analog complexity classes comparable to digital counterparts (P, nondeterministic polynomial time (NP), etc.). Build on recent work bridging analog dynamics and Turing machines[57;58], where polynomial time is preserved when measuring time as trajectory curve length. Extend these results to broader paradigms.

Semantic foundations for representation: Define formal semantics, correctness criteria, and verification frameworks that inform the design of analog intermediate representations and compilers (developed in PRD 6), ensuring that abstraction layers preserve meaning between high-level algorithms and physical hardware implementations.

Mathematical characterization of analog advantage: Establish mathematical frameworks characterizing when analog systems are fundamentally difficult to simulate digitally, including systems with unbounded derivatives, chaotic dynamics with positive Lyapunov exponents, and high-condition-number matrix problems where continuous-time evolution provides inherent advantages.

Computational capacity characterization: Characterize computational capacity including physical resources: integration time, detector response time, experimental trials/repetitions, number of computational elements, total energy consumption.

Formal model connections: Establish formal connections between analog models, demonstrating equivalences or proving separations. Clarify when different physical implementations realize equivalent computational models.

Research in this area may lead to completeness proofs identifying minimal primitive sets for major analog paradigms directly guiding hardware development (PRDs 2–6), with rigorous complexity classes and proven relationships to digital classes. These results will also include a formal framework enabling systematic comparison of analog approaches and principled design choices, as well as publications establishing foundational results accepted by the theoretical computer science community.

3.4.2 Analog Advantage: Rigorous Characterization and Algorithm Design

While promising examples suggest analog advantages (chaotic systems, stiff ODEs, well-conditioned linear systems), rigorous characterization of where analog computation provides asymptotic speedups remains incomplete. Digital numerical analysis has 50+ years of mature theory[66], yet finite precision errors remain surprisingly problematic[67], with extreme examples showing complex computation possible in rounding errors alone[68]. Analog numerical analysis—the systematic study of analog system behavior under noise, drift, variability, and finite precision—is only now emerging[67]. We must rigorously identify analog advantage domains while clarifying fundamental limits.

Research objectives address developing an analog numerical analysis framework that creates systematic theory for analog algorithm analysis, accounting for unique error modes (parameter variability, noise, drift) while exploiting advantages (continuous-time operation, absence of discretization error in accumulations). Research must formalize advantages for chaotic/stiff systems where digital methods face fundamental barriers: for chaotic systems with Lyapunov exponent λ, digital integration requires step count $N(T) \sim T \cdot (\delta_{\max}^{-1} e^{\lambda T})^{1/p}$ growing exponentially, while analog continuous-time emulators avoid discretization enabling real-time trajectory evolution. Research must establish rigorous speedup results for well-conditioned linear systems via analog solution with $O(N\kappa(\mathbf{A}))$ time scaling versus digital $O(N^{2.3})$–$O(N^3)$. Research must identify noise-tolerant applications (e.g., preconditioners for iterative solvers[69–71]) where approximate analog solutions accelerate digital convergence. Continuous-time analog formulations that reveal hidden structure in discrete problems—SAT arithmetization, Ising machines—introducing gradient information and energy landscapes absent in discrete formulations[64;65;65], must be investigated. Theory explaining when and why hybrid approaches combining periodic digitization[35] outperform pure discrete or pure continuous methods must be developed. Research must establish systematic methodologies for designing algorithms that exploit analog capabilities (natural parallelism, physical dynamics, continuous time) while managing constraints (noise, precision limits, connectivity).

Research in this area may lead to rigorous asymptotic speedup theorems for problem classes with proven analog advantages, with the analog numerical analysis framework adopted as a standard for algorithm design. These results will also include identified problem domains where analog serves as a powerful accelerator in hybrid architectures, novel algorithms with improved scaling bounds, and a mathematical understanding of the discrete-continuous interplay.

3.4.3 Noise, Precision, and Physical Constraints in Computation

Device mismatch, parameter variability, environmental effects, and noise fundamentally constrain analog computing. Existing models address stationary conditions; dynamical computation requires non-stationary theory. Paradoxically, noise can be a computational resource (p-bits, stochastic devices, probabilistic computing from Section 2.3.1), but *how* and *why* noise improves performance remains poorly understood. Rigorous mathematical frameworks are essential for predicting system behavior and enabling reliable algorithm design.

Research objectives address developing rigorous noise models for dynamical operation with large time-variations, extending understanding of flicker noise and nonlinear circuit effects in non-stationary regimes. Research must account for correlations between physically proximate noise sources at atomistic scales using first-principles bottom-up approaches (e.g., extending Gillespie-type stochastic models to analog circuits) coupled with experimental validation. Research must formalize relationships between numerical precision, noise characteristics (thermal, shot, flicker), and algorithmic convergence, developing theory predicting when analog's intermediate numerical accuracy advantages (no quantization in accumulations) outweigh susceptibility to environmental fluctuations and component mismatch (discussed in Section 2.3.1). Research must develop mathematical understanding of noise-enhanced optimization and sampling, explaining mechanisms behind observed improvements in Ising machines and optimization hardware. Rigorous frameworks guiding when to exploit noise versus when to suppress it must be created, enabling principled noise-aware algorithm design. Research must develop mathematical techniques quantifying computational complexity accounting for physical noise, energy constraints, and finite precision. Building on Section 2.3.1, research must develop probabilistic characterization methods for circuit performance, analytical models for device dependencies, and computational methods for propagating probability distributions through nonlinear dynamical systems.

Research in this area may lead to validated non-stationary noise models enabling accurate prediction of analog computing system behavior, with theory-driven design principles for noise-tolerant algorithms experimentally validated on platforms from PRDs 2–6. These results will also include mathematical frameworks explaining noise as a computational resource with complexity-theory extensions accounting for physical constraints, and probabilistic design methodologies for managing variability in large-scale analog systems.

3.5 Key Takeaways

Mathematical foundations are essential for transforming analog computation from promising demonstrations into a predictable, scalable computational paradigm.

- **Canonical Models and Complexity Theory:** Completeness proofs identifying minimal primitive sets will guide systematic hardware development, while formal complexity frameworks will characterize where analog computation provides provable advantages—establishing analog computation as a rigorous discipline comparable to digital theory.
- **Analog Numerical Analysis:** Non-stationary noise models, precision-convergence relationships, and frameworks for noise exploitation will enable reliable algorithm design and predictive optimization, forming a new subdiscipline extending classical numerical analysis to analog-specific error modes.
- **Theory-Practice Co-Development:** Parallel, iterative development where mathematical theory and experimental validation across PRDs 2–6 inform each other will accelerate both theoretical advances and hardware innovation, avoiding sequential dependency bottlenecks.
- **Hybrid Computing:** Understanding continuous-discrete interplay will unlock novel algorithms exploiting strengths of both paradigms, with mathematical frameworks guiding optimal partitioning between analog accelerators and digital precision.
- **Cross-PRD Enablement:** Mathematical foundations underpin all other PRDs: guiding high-performance architectures (PRD 2), providing validated device models (PRD 3), establishing theoretical limits for edge computing (PRD 4), formalizing natural computing paradigms (PRD 5), and informing co-design methodologies (PRD 6).

4 PRD 2: BUILDING HIGH-PERFORMANCE ANALOG SUPERCOMPUTING SYSTEMS

4.1 Introduction and Vision

While analog computing has demonstrated remarkable energy efficiency in edge applications (PRD 4), large-scale analog supercomputing remains largely unexplored. The first practical supercomputers were analog computers designed to solve differential equations[14;15]. Today's analog architectures can leverage these advantages for DOE science applications requiring the continuous dynamics and computational complexity of PDEs, fluid dynamics, geophysical and subsurface simulation, and combustion—precisely where analog's 100–1000× energy efficiency gains are most impactful.

Key Research Questions: *How can we develop large-scale analog HPC systems (100,000+ computational elements) that achieve 100–1000× energy efficiency improvements for DOE science applications while maintaining sufficient numerical accuracy (8–12 bit effective precision)? How do we architect analog systems beyond simple mesh/crossbar structures to support the diverse computational patterns required for PDEs, fluid dynamics, and complex simulations? What design automation tools and manufacturing approaches enable cost-effective deployment of reconfigurable analog HPC platforms at DOE facilities?*

Vision: Establish analog computing as a legitimate HPC paradigm through: (1) demonstrating 100,000-element systems achieving 100–1000× energy efficiency improvements for DOE applications (fluid dynamics, PDE solving, combustion and other multi-physics simulations), (2) developing architectures beyond VMM/crossbar approaches that minimize peripheral ADC/DAC overhead and support diverse computation patterns including hierarchical tree structures and distributed processing, (3) creating field-programmable analog platforms (FPAAs with 10,000+ elements) enabling multiple applications on shared infrastructure with design automation tools reducing configuration time from weeks to hours, (4) establishing validated manufacturing processes for FG-based analog computing in advanced CMOS nodes (16nm and smaller) with monolithic 3D integration achieving 10–100× density improvements, and (5) validating continuous-time analog computation for extended simulations (50–100-year equivalent timescales) with sustained 50–80 dB SNR performance at system level.

4.2 Goals and Desired Outcomes

Realizing analog computing's potential for high-performance scientific computing requires scalable architectures, programmable platforms, and mature manufacturing ecosystems. Main research needs follow.

- **Scalable Analog Architectures:** Develop 100–1000 node analog testbeds for scientific computing (PDE solving, continuous dynamics) demonstrating energy efficiency improvements over digital baselines with validated numerical accuracy and 50–100× accelerated timescale capabilities. Progress toward reconfigurable platforms with 1,000–10,000 programmable elements as foundations for larger-scale deployment.
- **Programmable Platform Development:** Establish FPAA-based reconfigurable HPC platforms (1,000+ programmable elements) achieving 14-bit effective precision through FG programming and mismatch compensation, with design automation tools reducing configuration time from weeks to hours.
- **Manufacturing, Characterization, and Metrology:** Coordinate with the cross-cutting roadmap in Section 2.5 to secure CMOS-compatible processes and statistical PDKs; emphasize HPC-specific characterization needs (runtime calibration for large-scale systems, system-level SNR measurement, accelerated aging protocols) distinct from device-level characterization in PRD 3. Focus on advanced-node integration (28/16nm), monolithic 3D for density, and calibration/runtime metrology tied to effective precision and system SNR.
- **Application Validation:** Demonstrate 1000× acceleration factors for target DOE applications through comprehensive benchmarking quantifying energy-performance tradeoffs at 100kW system scale.
- **Large-Scale Deployment:** Deploy 100,000-element analog HPC systems at DOE facilities integrated with conventional digital infrastructure, achieving sustained 8–12 bit effective precision (50–80 dB SNR) maintained across extended multi-hour to multi-day computations. This long-term goal builds on intermediate demonstrations at progressively larger scales.
- **Advanced Manufacturing:** Mature to production-ready ecosystems with HVM-capable variability ($< 5\%$) and widely accessible PDKs at 16nm and smaller, with validated monolithic 3D integration for analog signal paths to meet HPC density/performance targets (see Section 2.5).
- **Distributed Computing Frameworks:** Develop asynchronous distributed analog computation frameworks

supporting 10,000+ node systems with validated problem partitioning strategies and systematic approaches to hybrid analog-digital integration.

- **Domain Characterization:** Clearly identify application domains achieving 1000×+ analog advantage through comprehensive benchmarking against digital HPC, establishing validated performance models and deployment guidelines.

4.3 State of the Art

Analog computing foundations for HPC-scale systems build on existing demonstrations and infrastructure. Smaller-scale systems are commercially deployed (Mythic AI, Analog Inference) and demonstrated at research scale. Critically, *no dedicated analog fabrication facility is required*—FG devices integrate with standard CMOS processes, with industry foundry partnerships planned to enable university and DOE lab access to existing semiconductor infrastructure. However, significant challenges remain: advanced-node PDK access is expensive and limited, monolithic 3D integration for density improvements requires validation of analog signal integrity, and wafer-scale economics for large builds (1000+ wafers) are uncertain. Cost mitigation strategies through reconfigurable platforms, mask-cost amortization, and 3D integration approaches are addressed in dedicated research directions within this PRD. Additional scaling, integration, and validation challenges prevent deployment of large-scale HPC analog systems (100,000+ elements):

Architectural Constraints Beyond Mesh: Current analog computing heavily emphasizes compute in memory (CiM) and crossbar/mesh architectures optimized for VMM operations. While these architectures achieve high efficiency for specific operations, many scientific computing applications require richer computational patterns (hierarchical tree structures, distributed graph algorithms, dendrite-like computations) that break the classical mesh approach. Peripheral ADC/DAC conversion costs dramatically limit overall speed and efficiency when crossbars are used as digital coprocessors, negating analog advantages. See Section 2.3.2 for interface challenges. Recent demonstrations[7] of an AOC eliminating digital conversions through fixed-point searches achieved 500 TOPS/W efficiency (over 100× better than GPUs) for systems with up to 4,096 weights, but scaling to 100,000+ elements required for HPC while maintaining precision remains an open challenge. This underscores that HPC-scale analog computing requires fundamentally new approaches that minimize interface overhead and co-design algorithms with hardware constraints.

Scalability and Manufacturing Challenges: No validated path exists to 100,000+ element analog systems. Manufacturing analog computing systems in advanced CMOS nodes (28nm, 16nm, smaller) remains unproven at the scales required for HPC. Wafer-scale economics for 1,000+ wafer builds are uncertain. Device variability (20–50% for emerging devices) creates system-level reliability challenges. Process integration for FG devices in advanced nodes requires validation.

Reconfigurability and Programmability Gaps: Current FPAA approaches demonstrate reconfigurability but lack the scale (10,000+ elements) and design automation required for HPC deployment. Design tools for analog computing remain primitive compared to digital electronic design automation (EDA) tools—configuration requires weeks of expert effort rather than hours of automated synthesis. Standard cell libraries for analog/mixed-signal HPC building blocks are nascent. Programmable analog elements (FG) achieve 14-bit precision in older processes (350nm) but validation in advanced nodes is incomplete. While demonstrations[6] of ReRAM-based matrix solving achieving 24-bit precision show promise, scaling such approaches to HPC problem sizes remains to be validated.

Noise Management for Extended Computation: While analog circuits routinely achieve target precision levels, maintaining this performance across extended simulations (50–100-year equivalent timescales) at system level remains unvalidated. Noise modeling for continuous-time large-signal operation is limited. Theoretical frameworks from Section 2.3.1 require experimental validation at HPC scales. Stochastic analysis techniques for general nonlinear dynamical systems are underdeveloped.

Data Movement and System Integration Bottlenecks: Partitioning large problems (PDEs exceeding available compute nodes) across distributed analog systems lacks validated strategies. High-speed data logging (1–10 GB/s) for intermediate results presents significant challenges. Asynchronous distributed computation frameworks for 10,000+ node analog systems do not exist. Interface protocols between analog computation blocks and digital control infrastructure are ad hoc rather than standardized.

Forward Simulation Focus and Inverse Problem Challenges: The research directions outlined in this PRD

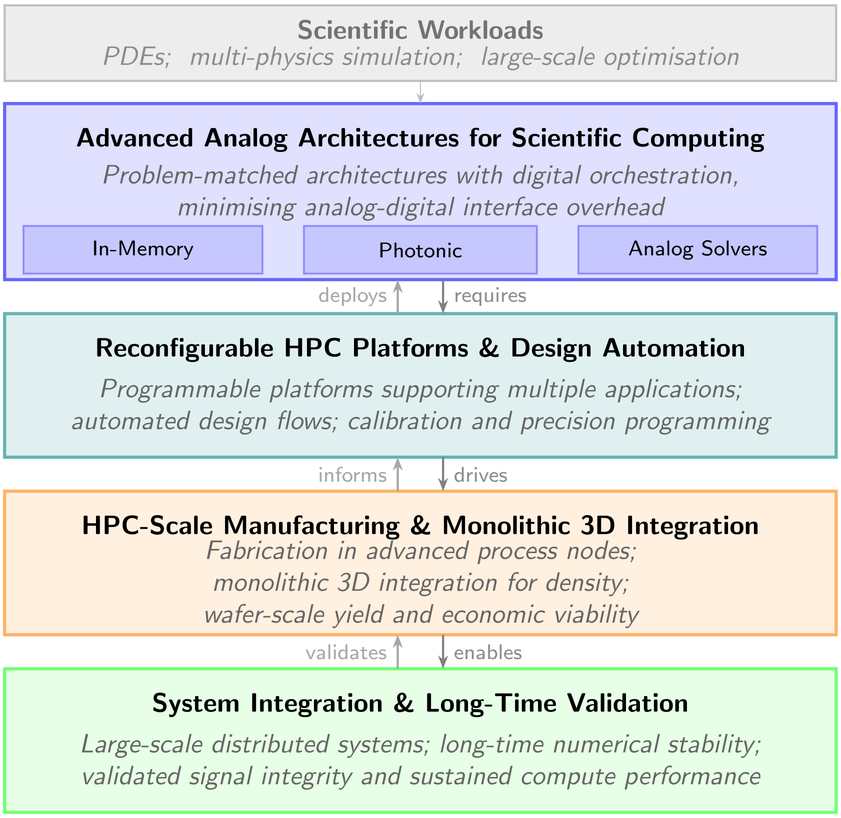

Figure 5: Overview of the four research directions in PRD 2. The stack reflects increasing infrastructure depth from application-facing architectures (top) to system-level validation (bottom); arrows indicate the primary coupling between directions. The three sub-tiles within Advanced Analog Architectures represent representative compute platforms.

emphasize forward simulation capabilities—solving PDEs, fluid dynamics, and multi-physics models where initial conditions and parameters are given. Inverse problems (parameter estimation, data assimilation) that require inferring unknown parameters from observational data present distinct computational challenges. Gradient-based inverse methods require sensitivity information typically obtained through finite differences (numerically unstable at 8–12 bit precision due to subtractive cancellation), forward sensitivities (viable only for problems with small parameter counts), or adjoint methods (most efficient but requiring reverse-mode differentiation likely infeasible in continuous-time analog computation). While forward simulation represents a transformative capability for DOE applications, extending analog HPC to gradient-limited inverse problems with large parameter spaces remains an open research challenge requiring investigation of alternative inversion strategies beyond traditional gradient-based approaches.

4.4 Research Directions

Four research directions structure this PRD, as illustrated in Figure 5. Advanced analog architectures define the compute primitives and problem-matching strategies, working in close co-design with reconfigurable platform research that provides programmability and design automation. HPC-scale manufacturing and 3D integration delivers the physical density needed for large-scale deployment, and system-level integration and long-time validation demonstrates sustained performance for DOE applications.

4.4.1 Advanced Analog Architectures for Scientific Computing

Mesh and crossbar architectures optimized for VMM operations[72–76] represent only one architectural pattern. DOE scientific applications demand diverse computational structures: hierarchical processing for tree-structured algorithms, distributed graph computations, dendrite-like dynamics for neural emulation[77], and spatial locality for PDE solving. Current approaches using crossbars as digital coprocessors incur prohibitive peripheral ADC/DAC costs. HPC analog systems require native architectures matching problem structure rather than forcing all computations through mesh operations.

Research objectives center on developing analog computing architectures beyond VMM/crossbar: (1) hierarchical structures for tree computations and multi-resolution processing, (2) distributed architectures with adaptive routing minimizing interface overhead, (3) continuous-time dynamics for PDE solving with natural spatial locality[78], (4) architectural frameworks for cortex-inspired processing and LLM inference, (5) hybrid analog-digital partitioning strategies exploiting analog for continuous dynamics while using digital for discrete control, (6) analog memory technologies and analog network-on-chip (NoC) architectures enabling analog systems to scale without analog-to-digital conversions, (7) current-based analog computing approaches offering integrated solutions for computation, data conversion, and input/output operations. This work should leverage mathematical foundations from PRD 1 for ODE/PDE implementations, and interfaces minimizing ADC/DAC usage must be designed as discussed in Section 2.3.2. Research should target operations beyond inner products: correlations[79–81], Gaussian mixtures[82], learning rules[79;80;83–85], and continuous-time integration.

Research in this area may lead to validated architectural patterns for multiple DOE application classes, with demonstrated 10× reduction in peripheral interface overhead through hierarchical and distributed designs. These results will also include proof-of-concept testbeds (100–1000 nodes) revealing how interface overhead accumulates at system scale—challenges invisible in single-node demonstrations—along with performance models quantifying analog advantages and identifying problem classes unsuitable for analog, and architectural guidelines informing co-design frameworks (PRD 6).

4.4.2 Programmable and Reconfigurable HPC Platforms

HPC infrastructure investments demand reconfigurability to serve multiple applications and accommodate algorithm evolution. Device mismatch fundamentally challenges analog circuit design (Section 2.3.1), requiring programmable compensation. FG devices in standard CMOS eliminate mismatch through compensatory charge programming, achieving high precision with minimal temperature/power-supply sensitivity[86–90]. Current FPAA demonstrations[16;91] prove reconfigurability but lack scale (10,000+ elements) where calibration overhead dominates design complexity—a challenge absent in hundred-element systems—and design automation. DOE needs multi-application platforms with rapid reconfiguration enabled by automated design flows.

Research objectives address scaling FPAA architectures to 10,000+ programmable elements with system-on-chip (SoC) integration supporting heterogeneous computation modes (continuous/discrete time/amplitude). Automated design flows from high-level specifications (PDE descriptions, neural network architectures) to analog circuit configurations must be developed, reducing expert configuration time from weeks to hours[54;92]. Analog/mixed-signal standard cell libraries[93] for HPC building blocks enabling design-space exploration[55] must be created. Design automation for analog HPC[92;94] must be advanced to be comparable to digital EDA tool maturity. Research must validate FG programming stability (14-bit precision, $<1\%$ drift) across temperature and aging in advanced nodes, leveraging device primitives from PRD 3 and natural computing approaches from PRD 5.

Research in this area may lead to a reconfigurable testbed (1,000–10,000 elements) with design automation tools achieving hour-scale configuration and validated programming precision (12–14 bits, $<1\%$ drift over 5+ years). These results will also include a standard cell library with analog HPC primitives demonstrated across multiple DOE applications proving infrastructure reuse, and tool ecosystems enabling non-expert deployment.

4.4.3 HPC-Scale Manufacturing and Monolithic 3D Integration

Analog HPC systems leverage existing semiconductor foundry infrastructure—no dedicated analog fabrication facility is required. FG devices integrate into standard CMOS processes, with the primary challenge being process integration validation in advanced nodes (28nm, 16nm, smaller) rather than building new manufacturing capabilities; research must therefore establish industry foundry partnerships to enable university and DOE lab access to existing fabrication capacity.

Building on the cross-cutting manufacturing foundations in Section 2.5, HPC-scale analog systems introduce unique manufacturing challenges distinct from smaller-scale implementations. Deploying 100,000-element systems requires 1000+ wafers, fundamentally changing manufacturing economics: at advanced process nodes (28nm, 16nm, smaller), mask set costs become comparable to total wafer costs, making reconfigurability essential to amortize non-recurring engineering expenses across multiple platform variants and applications. Monolithic 3D stacking—critical for achieving 10–100× density improvements—must preserve analog signal integrity through vertical interconnects while maintaining compatibility with foundry back-end processes. HPC deployment requires yield prediction at unprecedented scales where even small per-device failure rates (e.g., 0.1%)—negligible in thousand-element demonstrations—compound into system-level reliability challenges requiring new statistical models and redundancy strategies.

Research objectives include addressing HPC-specific manufacturing challenges beyond general analog device fabrication, with particular emphasis on advanced packaging and 3D integration.

Monolithic 3D integration for analog signal: Validate monolithic 3D integration specifically for FG-based analog HPC: demonstrate vertical analog signal paths through 3-5 stacked planes with copper interconnects, characterize inter-plane capacitance and parasitic effects in scaled processes[88], and establish thermal management strategies for high-density 3D stacks operating at HPC power levels. This approach is critical for achieving the 10–100× density improvements needed for HPC-scale systems.

Wafer-scale economics and reconfigurability: Develop wafer-scale economic models quantifying mask-set cost amortization strategies through FPAA-based reconfigurability, enabling cost-effective deployment across diverse DOE applications.

Create system-scale yield models translating device-level statistics (from Section 2.5) into 100,000-element system reliability predictions, enabling design-for-manufacturing tradeoffs between redundancy, calibration overhead, and raw device yield. Develop manufacturing test strategies for wafer-scale systems including probe card designs, statistical sampling methodologies, and accelerated burn-in protocols suited to analog HPC rather than digital memory arrays. Cross-reference PRD 3 for FG device physics and Section 2.4.1 for co-design methodologies integrating manufacturing constraints into algorithm mapping.

Research in this area may lead to a monolithic 3D demonstration with 3–5 stacked analog planes achieving 10–50× density improvement with characterized signal integrity and thermal profiles. These results will also include wafer-scale economic analysis quantifying cost-per-TOPS across mask-set amortization scenarios, system-scale yield prediction models enabling confident design of 100,000-element systems given device-level statistics, manufacturing test methodologies validated on multi-wafer builds, and PDK extensions capturing 3D-specific design rules and parasitic models integrated with foundry partnerships (PRD 6).

4.4.4 System-Level Integration and Long-Time Validation

Continuous-time analog computation's fundamental advantage lies in long-time numerical stability without degradation. Analog systems can run 50–100-year equivalent simulations with $<$ 1 s temporal resolution[95], enabling previously intractable long-term multi-physics models and cortical development predictions. System-level challenges—noise management at 50–80 dB SNR[16], problem partitioning for PDEs exceeding available nodes, asynchronous distributed communication, high-speed data logging (1–10 GB/s)—must be validated at scale to realize DOE HPC deployment. Computing cost scales proportionally with speedup: a 100 W real-time system requires 100 kW for 1000× acceleration, necessitating HPC-class cooling infrastructure.

Research objectives address validating long-time numerical stability for multi-year equivalent simulations without numerical degradation or reset cycles. Problem partitioning strategies for PDEs with grid spacing exceeding available compute nodes must be developed, including subsampling approaches exploiting analog's ideal intermediate numerics followed by residual computation for fine-grid refinement. High-speed data logging using FG memory and hybrid analog-digital storage must be designed, carefully selecting intermediate values to store without impacting computation timing. Asynchronous distributed computation frameworks for large-scale systems (10,000+ nodes) supporting both continuous-amplitude and discrete-event (neuromorphic) approaches[96] with validated handshaking protocols must be created. Research must demonstrate sustained target SNR at system level managing noise effects per Section 2.3.1, and must benchmark analog HPC against digital baselines for target applications (see Section 8.4.5 for benchmark frameworks and PRD 4 for edge deployment considerations).

Research in this area may lead to long-time simulation demonstrations (10+ year equivalent timescales) with validated numerical stability and sustained 50–80 dB SNR. These results will also include a distributed system testbed (100+ nodes) demonstrating asynchronous coordination protocols that emerge as critical requirements only in multi-node systems, with frameworks designed to scale to 10,000+ nodes; high-speed data logging (1–10 GB/s) using FG memory with selective storage strategies; and a benchmarking suite quantifying analog advantages for DOE applications with a deployment roadmap for facility integration (PRD 4, PRD 6).

4.5 Key Takeaways

High-performance analog computing represents a transformative opportunity for DOE computational science, requiring coordinated advances across architecture, manufacturing, and system integration:

- **Four Integrated Research Directions:** HPC-scale analog computing requires (1) advanced architectures beyond mesh/crossbar supporting diverse computational patterns, (2) reconfigurable platforms with design automation enabling rapid multi-application deployment, (3) scalable manufacturing in advanced nodes with monolithic 3D integration, and (4) system-level validation for extended timescales with distributed computation frameworks.
- **Ideal Application Targets:** DOE science applications combining continuous dynamics, appropriate noise tolerance, and massive computational demand—including multi-decade multi-physics models and century-scale fluid dynamics—present ideal targets where analog delivers 100–1000× energy efficiency gains.
- **Economic Viability:** Large-scale deployment (100,000-element) is achievable through reconfigurable designs amortizing mask costs across applications, advanced process nodes, and monolithic 3D integration providing essential density improvements.
- **Hybrid Architectures:** Establishing analog computing as a legitimate HPC paradigm requires hybrid analog-digital architectures that exploit analog for continuous dynamics while using digital for discrete control and data management.
- **Cross-PRD Integration:** Mathematical foundations (PRD 1) provide theoretical grounding; device primitives (PRD 3), natural computing (PRD 5), and co-design methodologies (PRD 6) enable practical realization; while edge computing (PRD 4) benefits from shared device advances and manufacturing processes.

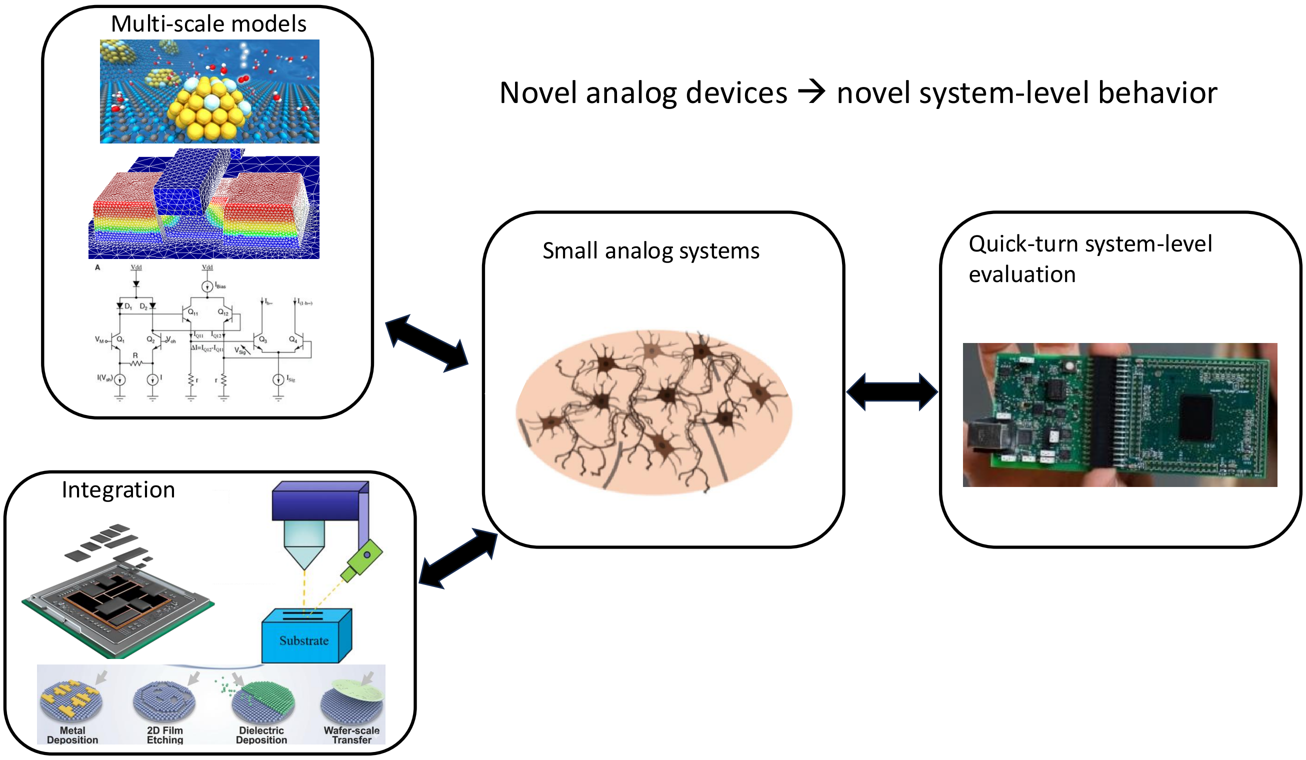

Figure 6: Device-to-system integration flow for analog computing. Novel device primitives require predictive models and abstraction layers to enable system-level design. The flow emphasizes co-design between device physics understanding, fabrication processes, circuit modeling, and system validation—essential for translating device-level advantages into deployable computing systems.

5 PRD 3: IDENTIFYING NOVEL DEVICE PRIMITIVES FOR NEXT-GENERATION ANALOG SYSTEMS

5.1 Introduction and Vision

Many novel physical computing platforms have demonstrated significant advantages over digital CMOS electronics at the scale of individual devices or small systems[97–102], yet these promising research results have not translated to improved high-performance computers. As illustrated in Figure 6, a fundamental challenge exists: device-level performance does not necessarily translate to system-level advantages. Understanding whether a novel computing technology primitive will provide value requires knowing how it scales[39]—not merely achieving functionality at smaller dimensions or building more devices, but understanding how devices fit into larger systems, how fabrication and performance depend on system context, and how sub-components interact during operation.

Realizing analog computing's potential requires addressing fundamental physical limits while recognizing that any practical computing system will require substantial digital CMOS infrastructure for control, data management, and general-purpose computation[103]. While analog hardware may accelerate specific computational kernels such as matrix multiplication or dynamical system simulation, broader functionality requires efficient interfacing with digital circuitry (see Section 2.3.2).

For next-generation analog computing devices to translate into transformative computing systems, we must enable the design of devices and small-scale systems with predictable scalability and efficient integration with existing digital electronics, without sacrificing their inherent advantages[104]. This requires devices to be abstractable, with models enabling the design of realistic systems performing practically relevant algorithms. Robust fabrication processes compatible with mainstream chip manufacturing are essential[37], as is comprehensive co-design across the entire stack (see Section 2.4.1 and PRD 6).

Key Research Questions: *How can we develop predictive multi-scale models that bridge device physics understanding with system-level design requirements? What fabrication and characterization standards are needed to ensure analog devices scale reliably from laboratory prototypes to deployable systems? Can we establish validation frameworks that quantify analog device advantages within realistic hybrid analog-digital architectures, accounting for interface overheads and system-level constraints?*

Vision: This PRD envisions establishing device primitives as the foundational building blocks for scalable analog computing systems through: (1) developing reproducible fabrication processes enabling predictable device performance across scales, (2) creating multi-fidelity modeling frameworks that capture essential device physics in computationally efficient forms suitable for system design, (3) establishing standardized characterization protocols providing statistical device models for circuit simulation, (4) building validation testbeds demonstrating device integration within hybrid architectures at increasing scales, and (5) enabling co-design between device physics, circuit architectures, and algorithms to realize measurable system-level advantages in DOE applications.

5.2 Goals and Desired Outcomes

Translating novel analog devices into practical computing systems requires validated models, reproducible fabrication processes, and systematic integration methodologies. Main research needs follow.

- **Device Fabrication and Characterization:** Establish reproducible fabrication protocols for analog device classes (memristive, phase-change, electrochemical) achieving $<$10% device-to-device variability with multi-laboratory validation, and create PDKs including statistical models for yield-aware circuit design, in alignment with the cross-cutting manufacturing and metrology effort in Section 2.5.
- **Predictive Device Modeling:** Develop validated compact device models integrated into standard circuit simulators achieving $<$20% error across operating ranges, and establish multi-fidelity modeling frameworks with 100× computational speedup.
- **System-Level Validation:** Build testbeds integrating 100–1000 analog devices with digital infrastructure to quantify performance and interface overheads, while establishing standardized benchmarking protocols measuring energy efficiency, latency, accuracy, and scalability across device technologies.
- **Manufacturing Maturity:** Achieve industry-grade fabrication maturity for analog devices comparable to digital CMOS, with high-volume manufacturing capabilities and $<$5% variability, demonstrating scalability from laboratory prototypes to production systems (10^6+ devices) while maintaining energy-efficiency advantages (see Section 2.5).
- **Comprehensive Device Ecosystems:** Develop comprehensive device libraries spanning diverse physical mechanisms (electronic, magnetic, optical, electrochemical) with validated models and automated co-design tools integrating device physics, circuit architectures, and algorithms for minimal-intervention tradeoff exploration.
- **Deployed System Advantages:** Deploy analog computing accelerators in DOE HPC facilities demonstrating order-of-magnitude energy efficiency or throughput improvements for specific application classes, with predictive system-level performance models accurately forecasting advantages while accounting for fabrication variability and interface overheads.

5.3 State of the Art

Emerging analog computing devices exploit diverse physical mechanisms—memristive switching, phase transitions, electrochemical reactions, spin dynamics—offering potential advantages in energy efficiency and computational density. However, translating device-level demonstrations into scalable computing systems faces fundamental challenges that have limited broader deployment despite decades of research. Recent advances now position the field to overcome these barriers: modern characterization tools (automated metrology, AI-driven materials discovery) dramatically accelerate device optimization, while coordinated co-design methodologies (PRD 6) replace the isolated device-centric efforts that characterized previous decades, enabling systematic translation from device to system performance.

Scalability and Manufacturing Limitations: Laboratory demonstrations of analog devices rarely scale beyond small prototypes due to fabrication variability, reproducibility challenges, and lack of CMOS-compatible processes. Device-to-device variability often exceeds 20–50%, making it difficult to predict system-level performance from individual device measurements. Unlike digital CMOS with mature process design kits and high-volume manufacturing, analog devices lack standardized fabrication protocols that can be transferred across laboratories or scaled to production volumes. This fabrication uncertainty prevents confident assessment of whether device-level advantages will persist at scales where practical computing benefits might be realized[37;39].

This challenge is particularly acute given that custom analog integrated circuits (ICs) have already demonstrated remarkable success in DOE scientific instrumentation, where they meet stringent specifications with ultra-low-noise ($\leq$ 1000 e^- root mean square (RMS)) and low-power ($\leq$ 1mW/channel) operation. Specialized readout ASICs have been developed for sensors in ATLAS at CERN, DUNE, and neutrino physics experiments such as nEXO (operating below 100 mK with radiation tolerance). Ultra-high-speed photon science applications at LCLS and LCLS-II at SLAC have produced ASICs with sub-20 ps timing resolution and high-resolution (14–16b), low-jitter (sub-100 fs) signal acquisition. These successes demonstrate that when analog design expertise is properly applied to domain-specific requirements, transformative capabilities emerge. However, such systems remain largely custom designs by analog specialists. Extending this success to broader analog computing applications requires addressing fundamental challenges in programmability, design automation, and workforce development.

Physical Understanding and Modeling Deficiencies: The fundamental physical processes limiting analog device operation remain poorly characterized. Critical questions persist: *What are the thermodynamic limits on operating energy for analog non-volatile memory? How do principal noise mechanisms scale with device dimensions and operating conditions? What kinetic processes govern switching dynamics and device lifetime?* Understanding these complex, frequently coupled mechanisms is challenging, resulting in few reliable physics-based models suitable for circuit and chip designers. Current device models, which guide digital CMOS manufacturing, are overly simplistic for analog electronics and provide limited guidance for exploiting analog features or mitigating issues. A communications-centric approach relating signal-to-noise ratio to functional correctness metrics could provide valuable design frameworks (see Section 2.3.1 for rigorous mathematical treatment of noise in analog systems).

Device Variability and Noise Management: Emerging analog devices exhibit heightened sensitivity to variability, noise, and poorly controlled kinetic processes compared to digital circuits[105]. Analog operation over quasi-continuous ranges makes devices vulnerable to perturbations from thermal fluctuations, defects, fabrication variations, and environmental factors. While some analog systems like RF blocks have achieved commercial success (notably in the thriving analog IC industry with 40–50% margins for sensor interfaces, data converters, and voltage regulators), general-purpose or large-scale analog computing faces more stringent accuracy, scalability, and programmability requirements. Statistical device models capturing variability effects are essential for yield-aware system design but remain underdeveloped for most emerging analog device technologies.

A critical bottleneck is the limited pool of engineering talent across the analog ecosystem—system architects, IC designers, verification engineers, test engineers, and system developers. Analog design methodology has changed little over decades, remaining largely the domain of analog "artists" with expertise primarily in component-level design (*e.g.*, ADCs, DACs). Even fewer designers possess the system-level analog architecture expertise needed for computing applications. While digital computing achieved its "very large scale integration (VLSI) revolution" through standardization and automation (microprocessors, FPGAs, GPUs with associated synthesis tools), analog electronics needs an equivalent transformation in computing scope, applicability, and usability. Key barriers include: programmable and mismatch-insensitive design techniques need generalization (floating-gate devices offer promise but require wider adoption); configurability analogous to FPGAs remains largely in research stages (*e.g.*, large-scale FPAAs); and synthesis and automation tools comparable to digital EDA remain rare exceptions. Addressing these ecosystem challenges through research in programmable analog design, synthesis tools, and educational programs is essential for analog computing to achieve broad impact.

Analog-Digital Interface Bottlenecks: Specialized analog devices may outperform digital systems for specific mathematical operations but rarely execute complete algorithms independently, necessitating digital interfaces. ADC and DAC conversion introduces substantial energy, footprint, and latency overheads that can negate analog advantages[36]. Effectively coupling analog and digital signal processing with minimal energy and latency loss remains a key scientific challenge. This requires innovations in data representation and manipulation, including hybrid approaches where digital transistors emulate analog behavior, analog devices operate in digital modes, or complex analog systems incorporate broader computational abstractions (see Section 2.3.2 for comprehensive interface discussion).

5.4 Research Directions

5.4.1 Device Physics and Modality-Specific Fabrication

While Section 2.5 establishes cross-cutting fabrication and characterization protocols, each analog device class presents unique physics-driven fabrication challenges requiring specialized understanding. Unlike digital CMOS where variability arises primarily from statistical dopant fluctuations and line-edge roughness, analog device variability stems from complex multiphysics—filament formation stochasticity in memristors, phase boundary roughness in transition-metal oxides, ion distribution inhomogeneity in electrochemical devices, domain wall pinning in magnetic systems. Achieving the common manufacturing targets ($<$10% variability, multi-lab reproducibility) demands device physics understanding linking atomic-scale mechanisms to fabrication process parameters. This research direction addresses the fundamental question: *what physical mechanisms differentiate device classes, how do these mechanisms determine fabrication sensitivities, and which devices offer computational primitives genuinely unavailable through conventional CMOS?*

Research objectives center on developing physics-informed fabrication strategies tailored to distinct device modalities, providing the device-specific foundation for the standardized protocols in Section 2.5. For **filamentary memristors** (ReRAM), research must characterize how forming conditions—voltage ramp rates, current compliance, ambient atmosphere—control Joule heating profiles, interface redox reactions, and defect migration to create conductive pathways with reproducible concentration gradients and phase distributions[106;107]; establish switching voltage distributions and their correlation with retention/endurance. For **phase-change devices** (vanadium oxide, niobium oxide), research must characterize how film deposition parameters—temperature, pressure, stoichiometry—determine crystallographic phase boundaries, grain structures, and Mott transition sharpness[98]; quantify how phase boundary roughness translates to device-to-device switching threshold variation. For **electrochemical devices** (electrochemical random access memory (ECRAM)), research must map ion mobility and redox kinetics dependencies on electrolyte composition, electrode work functions, and interfacial barrier heights; establish design rules preventing electrochemical crosstalk in dense arrays. For **spin-based and magnetic devices**, research must characterize domain wall pinning site distributions, thermal activation barriers for switching[99], and write-error-rate dependencies on material defects and grain boundaries[108].

Research must critically assess which devices offer computational primitives unavailable through conventional CMOS: true quantum or thermal noise sources for hardware random number generation, intrinsic nonlinear dynamics from many-body electron interactions (Mott insulators, superconducting junctions), or native analog operations (e.g., temporal integration in ECRAM) providing genuine efficiency advantages rather than alternative implementations. Design-of-experiments methodologies mapping fabrication parameters to statistical device performance across modalities must be developed, informing process optimization. Modality-specific noise mechanisms—thermal fluctuations, telegraphic noise from defect switching, flicker noise from charge trapping—must be characterized to establish physical origins and enable mitigation strategies beyond post-processing "denoising" that may negate energy advantages[109]. Modality-adapted templates must be contributed to the standardized characterization repositories defined in Section 2.5.

Research in this area may lead to physics-based fabrication process maps for memristive, phase-change, electrochemical, and spin-based device classes achieving reproducibility targets ($<$10% variability) with documented sensitivity to process parameters. These results will also include validated scaling from laboratory prototypes (10–100 devices) to arrays (1000+ devices) demonstrating maintained performance distributions; noise characterization linked to physical mechanisms enabling physics-informed mitigation; statistical device models capturing modality-specific behaviors integrated into multi-lab validation protocols; and identification of computational primitives with quantified advantages over CMOS baselines informing technology selection for PRD 2 and PRD 6.

5.4.2 Predictive Multi-Scale Device Models

Circuit and system designers require predictive device models that are computationally efficient yet capture essential physics governing analog operation. Emerging analog devices exhibit nonlinear, history-dependent, and stochastic behaviors arising from complex multiphysics—ion transport, phase transitions, defect dynamics—that cannot be described by conventional compact models developed for digital CMOS. The computing challenge is not to perform exhaustive atomic-scale simulations, but to develop reduced-order representations that abstract essential device physics into forms tractable for circuit simulation and system-level design while maintaining predictive accuracy. As illustrated in Figure 7, multi-fidelity frameworks must bridge high-fidelity physics understanding with practical design tools.

Research objectives center on developing modeling frameworks that use high-resolution physics simulations to parameterize computationally efficient models for system design, taking a computing-focused approach to materials science where understanding complex device physics (ion/vacancy kinetics in ECRAM and ReRAM devices,

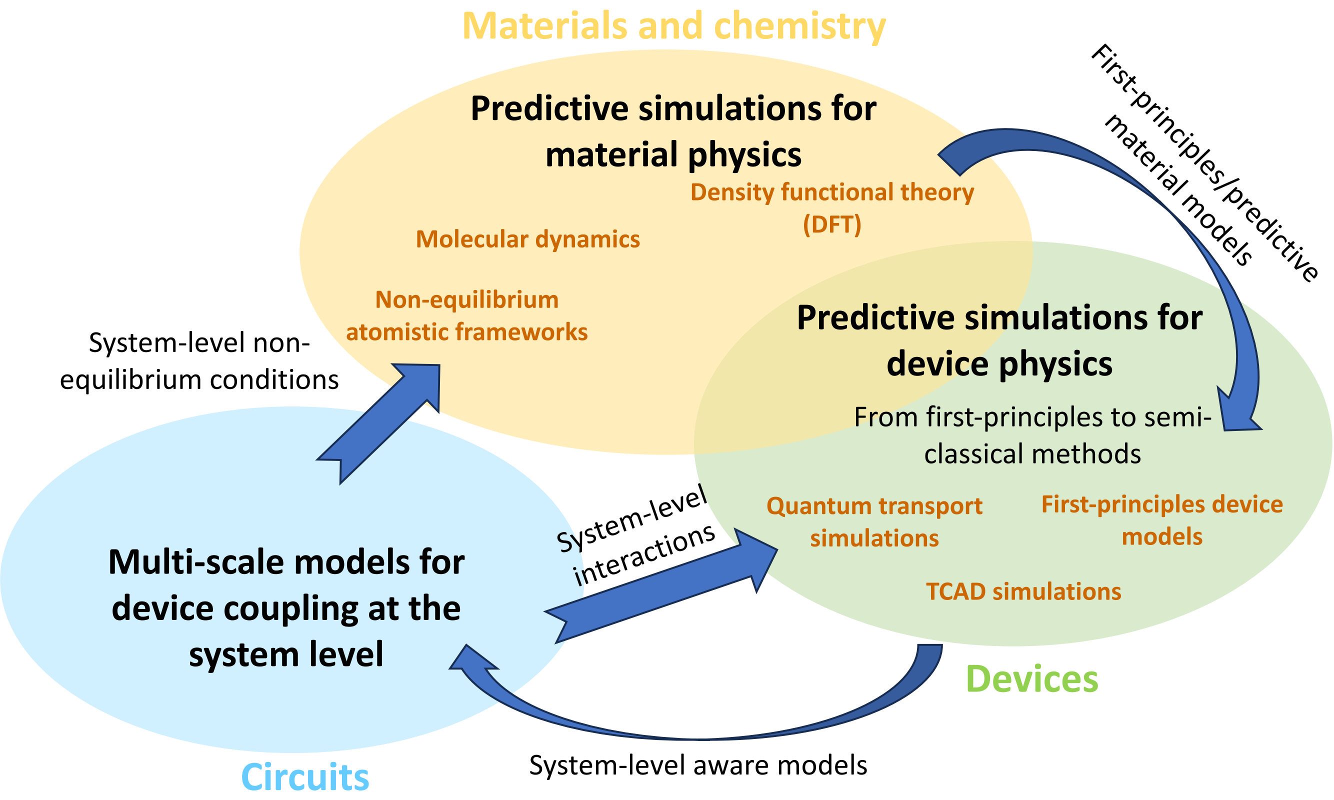

Figure 7: Multi-scale modeling framework integrating device physics, circuit-level interactions, and system-level performance for analog computing. High-fidelity simulations inform reduced-order compact models suitable for circuit simulation, enabling efficient system design while capturing essential device physics including variability, thermal effects, and non-equilibrium dynamics. The framework enables computational co-design from materials to systems.

phase transformations in valence-change memory (valence change memory (VCM)), defect-mediated conduction) is essential not for fundamental materials discovery but to inform compact behavioral models capturing energy-delay tradeoffs, noise characteristics, and history-dependent dynamics for circuit simulation. For instance, understanding how voltage-driven ion/vacancy movement creates or dissolves conductive filaments with specific oxygen vacancy distributions enables reduced-order models predicting switching characteristics without simulating every atomic-scale process. Advanced computational methods (accelerated molecular dynamics, machine-learning potentials, multiscale techniques) enable simulation of non-equilibrium thermodynamics and long-timescale kinetics, providing physics insight to parameterize reduced-order models achieving 100×+ computational speedup while maintaining predictive accuracy for system-level design.

System-level models must also account for device coupling within computing systems, where analog devices operating over quasi-continuous ranges with behavior governed by non-equilibrium thermodynamics are vulnerable to perturbations from thermal gradients, electrical crosstalk, and magnetic coupling. Predictive models must capture these system-level interactions—thermal, electric, and magnetic gradients, quantum-mechanical coupling—not just isolated device characteristics, enabling co-design of analog circuits accounting for thermal management, signal integrity, and noise propagation (see Section 2.4.1 and PRD 6).

Models must also address defect impacts from a computing perspective: *how do defects affect device-to-device variability, reliability metrics, and yield?* While even single defects significantly influence advanced gate-all-around field-effect transistors (GAAFETs) and beyond-CMOS devices, in analog devices where energy differences between quasi-continuous states are small, defects critically impact variability and reliability (see Section 2.3.1 for mathematical frameworks). Models must predict statistical distributions of device performance accounting for defect populations, enabling yield-aware circuit design and reliability assessment.

Validation methodologies must compare model predictions against experimental data across device sizes, operating conditions, fabrication variations, and aging effects, ensuring models remain predictive beyond calibration datasets. Validated models must be integrated into standard circuit simulators (SPICE, Verilog-A) and system-level design tools, enabling analog circuit designers to leverage device models as effectively as digital designers use CMOS models.

Research in this area may lead to validated compact models for multiple analog device classes integrated into commercial circuit simulators, with multi-fidelity simulation frameworks achieving 100× speedup versus atomic-scale methods while maintaining <20% prediction error. These results will also include physics-informed machine learning models and system-level models capturing thermal coupling, crosstalk, and variability propagation, as well as model libraries enabling automated analog circuit synthesis and co-optimization of device physics with circuit architectures (PRD 6).

5.4.3 System-Level Integration and Validation

Projecting whether analog devices provide system-level advantages requires assessing how they integrate within complete computing systems including digital infrastructure, transduction interfaces, and realistic application workloads. Applications spend significant time across diverse computational primitives—mathematical operations, memory access, control flow—making system performance depend not only on analog device characteristics but critically on integration costs: interface energy and latency, data movement overheads, and architectural constraints. Fine-grained integration of simple analog devices incurs heavy transduction costs that may negate device-level advantages, while coarse-grained accelerators sharing interface circuitry across many devices introduce dataflow restrictions limiting applicability. Quantifying these integration costs through experiments and modeling spanning small analog device arrays to complete hybrid analog-digital systems is essential for determining whether and where analog approaches provide value (see Section 2.3.2 for comprehensive interface discussion and PRD 4 for edge computing applications).

Research objectives address building testbeds at increasing scales (100, 1000, 10,000+ devices) integrating analog devices with digital infrastructure to quantify energy, latency, accuracy, and throughput for specific computational kernels relevant to DOE applications. Testbeds must measure not just analog device performance in isolation but complete system metrics accounting for ADC/DAC overhead, data movement costs, control logic, and algorithmic accuracy degradation. Architectural frameworks identifying optimal integration strategies must be developed: when does fine-grained integration justify transduction costs versus coarse-grained acceleration minimizing interface overhead? Alternative integration approaches must be compared, including in-memory computing (CiM), hybrid analog-digital devices enabling nearly transductionless operation, and complex analog systems combining dissimilar devices to handle multiple functions while minimizing digital interfaces.

Standardized benchmarking protocols must enable fair comparison across analog device technologies and against digital baselines (see Section 8.4.5 for comprehensive benchmarking discussion). Benchmarks must capture system-level metrics—end-to-end energy efficiency, wall-clock time, solution accuracy—not just device-level specifications. Application-specific test cases organized by computational problem type must be created: differential equation solving for fusion, multi-physics modeling, and reactor dynamics; matrix operations for scientific machine learning, quantum chemistry, and linear system solving; continuous optimization and control for PDE-constrained inverse problems and real-time feedback systems; and discrete combinatorial optimization for graph-based problems, resource allocation, and configuration space exploration. This organization enables assessment of analog advantages across diverse computational paradigms in realistic DOE contexts (see PRD 5 for natural computing device applications).

Co-design methodologies should integrate device physics, circuit architectures, and algorithms to maximize system-level performance. This includes exploring innovations in signal handling—adaptive control systems compensating for device drift and variability, error-correction schemes maintaining computational accuracy despite analog imprecision, calibration frameworks enabling reliable long-term operation. Architectural innovations exploiting analog device properties must be investigated: analog devices operating in digital modes when beneficial, reconfigurable hybrid systems adapting to workload characteristics, and hierarchical architectures balancing analog acceleration with digital control and general-purpose computation.

Research in this area may lead to validation testbeds at multiple scales (100, 1000, 10,000+ devices) with a standardized benchmark suite demonstrating quantified performance for DOE-relevant kernels. These results will also include architectural frameworks identifying optimal integration strategies with demonstrated 10×+ system-level energy efficiency improvements for multiple application classes, and co-design tools enabling automated exploration of device-circuit-algorithm tradeoffs (PRD 6), with clear identification of application domains where analog provides advantages versus limitations where digital remains superior.

5.5 Key Takeaways

Realizing transformative analog computing systems requires bridging the gap from device-level demonstrations to scalable, deployable systems through three interconnected research directions:

- **Scalable Fabrication and Characterization:** Establishing reproducible manufacturing processes with statistical device models enables confident yield prediction and circuit design. Understanding device physics—phase transitions, defect dynamics, electrochemical processes—from a computing perspective informs process optimization for predictable performance at scale.
- **Predictive Multi-Scale Modeling:** Computationally efficient device models capturing essential physics are critical for system design. High-fidelity simulations (e.g., ab initio and molecular dynamics) inform reduced-order representations suitable for circuit simulation, enabling co-design without atomic-scale computational cost.
- **System-Level Integration and Validation:** Testbeds quantifying performance within realistic hybrid architectures must account for analog-digital interface overheads—particularly ADC/DAC conversion energy, footprint, and latency costs that can negate device-level advantages. Standardized benchmarks enable fair comparisons across technologies and against digital baselines.
- **Interface Cost Minimization:** Critical to realizing analog advantages is developing integration strategies that reduce conversion overheads through architectural innovations such as coarse-grained acceleration, hybrid analog-digital devices, and in-domain analog processing.
- **Cross-PRD Enablement:** Validated device primitives with fabrication processes and models enable ultra-efficient edge computing (PRD 4), natural computing substrates (PRD 5), high-performance architectures (PRD 2), and co-designed systems (PRD 6).

Success demands integrated progress across all three research directions: reproducible fabrication provides devices with predictable properties, validated models enable confident system design, and testbeds demonstrate whether device-level advantages translate to measurable system-level benefits for DOE applications.

6 PRD 4: REDEFINING THE EDGE: INTEGRATING ANALOG COMPUTE WITH SENSING AND TRANSDUCTION

6.1 Introduction and Vision

Analog computing's exceptional energy efficiency makes it particularly compelling for edge and sensor-integrated applications, where computation occurs close to data sources to reduce latency and energy dissipation[110;111]. Next-generation sensors present two extremes: large-scale scientific experiments producing massive data volumes (e.g., SLAC's Linac Coherent Light Source reaching 1 TB/s by 2029[112]), and ubiquitous sensing platforms for Internet of Things (IoT)[113], smart infrastructure[114], and personalized healthcare[115] requiring years of continuous operation under stringent energy budgets. Both scenarios demand new paradigms: high-throughput sensors need data bandwidth reduction, while energy-constrained sensors require orders-of-magnitude efficiency improvements.

Key Research Questions: *How can analog computing address the mounting challenge of data proliferation from next-generation, high-density, and high-throughput sensors? What role can end-to-end co-design play in understanding the fundamental limits and scaling challenges of analog computing at the edge? Can integrating analog computing at the edge or within the sensing platform address these mounting challenges of data proliferation and stringent operating energy budgets?*

Vision: This PRD envisions redefining edge computing through new analog computing paradigms that integrate computation directly with sensing and transduction, achieving: (1) elimination of traditional digital conversion bottlenecks by exploiting physics at the sensing-transduction interface, (2) transformative order-of-magnitude improvements in energy efficiency, latency, and data bandwidth, (3) machine learning and artificial intelligence approaches that reduce data bandwidth requirements while maintaining scientific fidelity, (4) end-to-end co-design methodologies across sensing modalities (optical, acoustic, mechanical), and (5) in-domain analog computing systems that minimize unnecessary domain conversions.

6.2 Goals and Desired Outcomes

Analog edge computing offers transformative opportunities for DOE scientific applications where real-time processing, extreme energy constraints, or bandwidth limitations make conventional digital approaches impractical. Main research needs follow.

- **Fundamental Limits and Co-Design:** Develop and validate physics-based models capturing fundamental limits of information transduction and signal acquisition in analog edge systems, and establish end-to-end co-design methodologies integrating sensing, transduction, and analog computation across sensing modalities (optical, acoustic, mechanical), building on the general mathematical foundations in PRD 1 and cross-cutting co-design methodologies in PRD 6.
- **Application Demonstrations:** Demonstrate proof-of-concept analog edge computing systems for DOE applications (high-energy physics detectors, fusion diagnostics, quantum system control) achieving 10× improvements in latency and/or energy efficiency, with initial in-domain analog computing prototypes minimizing domain conversions.
- **Benchmarking Frameworks:** Define analog edge computing benchmark datasets and standardized evaluation metrics for three sensing platform categories—active, passive, and self-powered—in coordination with the cross-cutting benchmarking infrastructure described in PRD 6.
- **Large-Scale Deployment:** Deploy analog edge computing solutions at scale in DOE scientific facilities, achieving order-of-magnitude reductions in data requirements while preserving scientific information content, with mature ecosystems of standardized components, interfaces, and design tools enabling non-expert deployment.
- **Theoretical Foundations at the Edge:** Establish a comprehensive theoretical framework for analog information processing at the edge that unifies concepts from information theory, statistical physics, and signal processing, complementing the general mathematical foundations in PRD 1 with edge-specific notions of information transduction, bandwidth reduction, and sensing-constrained computation.
- **Adaptive Edge Intelligence:** Achieve real-time analog learning and adaptation at the edge with energy budgets approaching fundamental thermodynamic limits, enabling rapid deployment for new sensing applications.

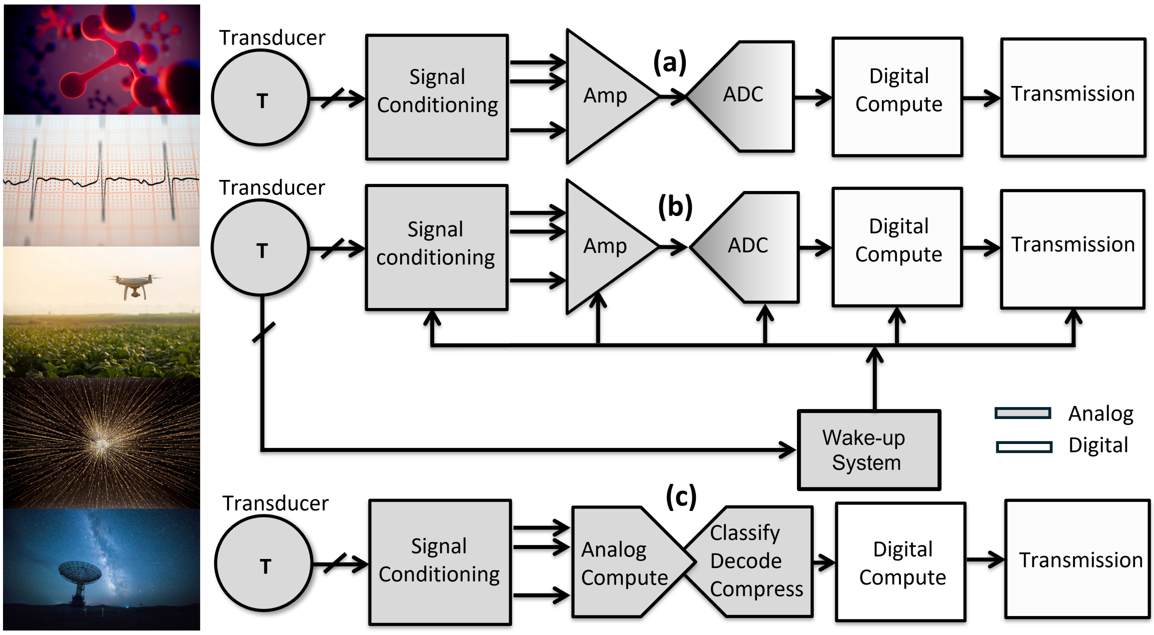

Figure 8: In-sensor and edge computing applications and architectures:(a) conventional approach where AI and low-latency data reduction algorithms are implemented in digital; (b) wake-up based approach where analog computation is used to power-gate the conventional architecture; and (c) direct approach where analog computation is inserted within the sensor signal processing architecture.

6.3 State of the Art

6.3.1 Conventional Digital Edge Computing Paradigm

The widely used paradigm for designing in-sensor and edge computing platforms is based on a standard signal processing architecture, as depicted in Fig. 8(a). This architecture comprises an analog signal conditioning unit that interfaces with a transducer or a transducer array, followed by low-noise analog amplifiers, analog-to-digital converters (ADCs), digital processors, and subsequently a data logging or data transmission unit. Any edge-AI algorithm for data compression, signal classification, or detection is implemented on the digital processor or a co-processor which could be a graphical processing unit (GPU), a tensor processing unit (tensor processing unit (TPU)) or a neural processing unit (neural processing unit (NPU)).

While this approach provides the flexibility to deploy various low-latency algorithms tailored to specific application requirements, it is fundamentally constrained by:

- **Bandwidth bottlenecks:** ADC conversion rates and digital processing throughput limit data acquisition rates.
- **Energy inefficiency:** Digitization and digital processing consume significant power, limiting battery-operated deployments.
- **Latency:** Sequential analog-to-digital conversion and digital processing introduce delays incompatible with real-time control applications.
- **Resolution tradeoffs:** Higher ADC resolution comes at the cost of speed and power consumption.
- **size, weight, and power (SWaP) constraints:** Size, weight, and power limitations restrict deployment in harsh or resource-constrained environments.
- **Data deluge:** High-rate sensors generate data volumes that overwhelm storage, transmission, and processing capabilities.

6.3.2 Current Analog Edge Computing Approaches

Current analog computing approaches for integrating intelligence at the edge and addressing SWaP challenges are highlighted in Fig. 8(b)-(c).

Wake-Up and Power-Gating Architectures: One approach involves using an analog wake-up for triggering and power-gating as shown in Fig. 8(b). This method leverages the low-latency and energy-efficient properties of analog computing to detect a triggering event (e.g., energy detection or change detection), which then activates the more power-hungry conventional architecture to sample the event at a higher resolution. This approach has been central to current state-of-the-art near-zero and always-on systems[116–118]. However, it is only advantageous for processing infrequent or rare events (e.g., speech activated units, seismic activity or intrusion detection), where energy savings can be achieved by keeping the conventional architecture in deep-sleep mode. Wake-up latency and settling-time requirements of amplifiers and converters also limit precise event recording, and the accuracy and response time of the analog wake-up system determine false-alarm and true-positive rates.

In-Line Analog Processing: Several approaches integrate analog computing and processing directly within the conventional signal acquisition flow, as shown in Fig. 8(c). In these approaches, analog computing is directly applied prior to any domain conversion (such as optical-to-electronic conversion) or data conversion (such as analog-to-digital conversion) to implement tasks like event sensing[16;119], decoding[90], correlation[120], feature extraction[121], or classification[122–124]. These architectures can be regarded as a *smart* analog-to-digital converter, which significantly relaxes the bandwidth requirements on the digital processor. Analog artifacts introduced by the frontend (mismatch and non-linearity) are typically corrected by the back-end digital processor using digitally-assisted analog approaches[125] or auxiliary calibration techniques[126;127]. While promising, current implementations are limited in scalability, resolution, and generalizability across different sensing modalities. No rigorous theoretical framework yet unifies sensing physics, transduction, and computation to characterize fundamental performance limits or define optimal operating regimes. Consequently, end-to-end co-design methodologies remain undeveloped, and system design proceeds empirically rather than from principled joint optimization of the sensing-computation interface.

6.3.3 Sensing and Edge Platform Taxonomy

Current in-sensor and edge computing approaches are becoming increasingly impractical for addressing the challenges of signal representation, domain conversion (e.g., optical to electrical), and input-output data transfer. Specific operating constraints, such as harsh environments, further limit the applicability of conventional electronic approaches. One effective way to identify opportunities for analog computing at the edge is to classify sensing platforms based on their energy and powering requirements, as shown in Fig. 9(a)-(c).

DOE scientific facilities present unique requirements that differentiate them from commercial applications: operation in harsh environments with high radiation flux, extreme temperatures, and strong electromagnetic interference; event-based sensing paradigms for detecting rare phenomena in particle physics and astrophysics experiments; requirements for graceful degradation in autonomous systems deployed in fusion reactors or deep subsurface monitoring; and the need for asynchronous event processing where sensors activate only upon detecting significant changes. These constraints require targeted research investments in radiation-hard analog devices, noise-tolerant processing architectures, and adaptive algorithms designed for non-stationary, asynchronous data streams across all platform types.

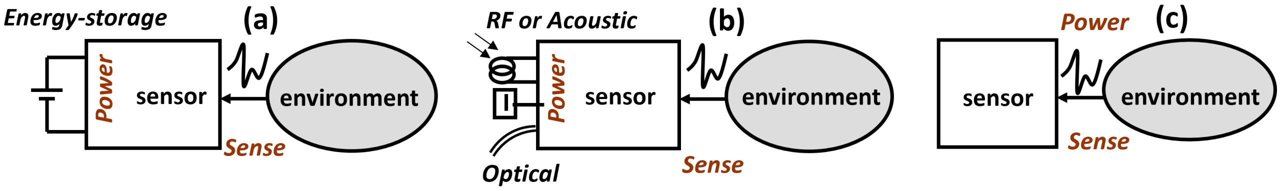

Figure 9: Classification of sensing and edge compute platforms based on powering requirements: (a) Active sensors with on-board battery or energy-source; (b) Passive sensors where an extrinsic interrogation signal is used to measure events of interest in the environment; and (c) Self-powered sensors where the operating energy is harvested from the signal being sensed.

Active Sensing Platforms: The first category, shown in Fig. 9(a), includes platforms where the energy source or battery is integrated on the platform. On-board power enables continuous operation, supports complex edge-AI algorithms, and allows both trigger-less streaming and triggered readout systems, enabling implementation of state-of-the-art computational models[9;128] without stringent energy constraints. However, operational lifetime is constrained by battery capacity or energy source availability, and high data rates from dense sensor arrays can overwhelm processing and storage capabilities. Beyond maximizing energy efficiency, key opportunities for analog computing include:

- **Data compression with fidelity preservation:** Balancing the trade-off between processing or compressing high volumes of sensor data while retaining information essential for decision-making or control.
- **Real-time, low-latency operation:** Supporting high-speed control and decision-making where analog-to-digital conversion latency is impractical (e.g., control systems for quantum computing[129], precision sensing[130;131], or advanced imaging[132]).
- **High-dimensional analog signal processing:** Exploiting analog computation to eliminate interference, prevent data loss, or identify novel signatures that may be buried in noise or destroyed by ADC quantization[133;134].

The central challenge is developing lightweight analog computing architectures that can process or compress high-volume sensor data in real-time while preserving scientific information content and supporting diverse sensing modalities across DOE mission areas.

Passive Sensing Platforms: The second category (Fig. 9(b)) consists of passive sensors without onboard energy sources, relying on backscattered interrogation signals (radio-frequency, optical, or acoustic) to detect environmental changes. Examples include radio-frequency identification (RFID) sensors[135], Fiber-Bragg grating sensors[136;137], and surface-acoustic-wave sensors. Their passive nature enables cost-effective, miniaturized deployment at scale, ideal for harsh environments with battery-free operation in inaccessible locations; however, they cannot perform onboard processing—all computation occurs at the interrogation source, with signal quality dependent on interrogation strength and environmental coupling. Key opportunities for analog computing with passive sensors include:

- **In-domain processing:** Exploiting within-domain analog computing (optical, acoustic, electromagnetic) to perform signal modulation, encoding, or feature extraction without domain conversion.
- **Passive analog encoding:** Using passive analog structures (diffraction gratings, metamaterials, acoustic resonators) to encode sensor information directly into backscattered signals.
- **Backend analog processing:** Implementing analog signal processing at the interrogation receiver to extract sensor signatures from complex backscattered signals with high throughput and energy efficiency.

The central challenge is avoiding domain conversion (e.g., optical-to-electrical) and instead relying on within-domain passive architectures, like diffraction gratings[138] or passive reflectors[139], to maximize energy coupling from the environment. These architectures necessitate sophisticated backend processing to extract signatures of interest, making end-to-end co-design of analog computing with digital signal reconstruction essential (see PRD 6).

Self-Powered Sensing Platforms: The third category (Fig. 9(c)) harvests operating energy directly from the sensed signal. Examples include self-powered mechanical and temperature sensors[140] and optical sensors[141]. The asynchronous nature of self-powering makes these platforms ideal for detecting rare events in energy-constrained or inaccessible environments, with event-driven operation naturally filtering data; however, harvested energy is often insufficient to power amplification, conversion, and wireless transmission, limiting applications to data forensics and post-processing rather than real-time operation. Key opportunities for analog computing in self-powered sensors include:

- **Ultra-low-power analog processing:** Implementing analog feature extraction, thresholding, or pattern matching that operates within the harvested energy budget.
- **Analog non-volatile storage:** Using analog memory technologies (floating-gate, phase-change) to record signal characteristics without power-hungry ADCs or digital logic.
- **Energy-aware computation:** Developing analog computing primitives that exploit the physics of energy harvesting transducers to perform computation during the transduction process itself.
- **Asynchronous analog processing:** Creating analog circuits that operate asynchronously, triggered by signal events rather than clocked synchronously, matching the sporadic nature of harvested energy.

The central challenge is matching the physics of transduction with computing primitives and non-volatile storage to ensure that events of interest can be recorded and retrieved during post-processing, while balancing the energy budget between sensing, processing, storage, and data retrieval without compromising signal fidelity.

6.4 Research Directions

To address these challenges, we identify three major research directions organized around: (1) understanding fundamental limits through end-to-end co-design, (2) developing analog learning and adaptation capabilities at the edge, and (3) creating novel devices and architectures. These directions specialize the cross-cutting theoretical and co-design themes of PRD 1 and PRD 6 to in-sensor and edge deployments and build on the sensing-platform taxonomy summarized in Fig. 9, which highlights distinct opportunities and constraints for active, passive, and self-powered platforms.

6.4.1 Understanding Fundamental Limits through End-to-End Co-Design

To overcome current limitations and realize the full potential of analog edge computing, we must develop a fundamental understanding of the physical limits governing information transduction, signal acquisition, and analog computation at the edge. This requires moving beyond empirical system design to establish rigorous theoretical frameworks grounded in physics and information theory.

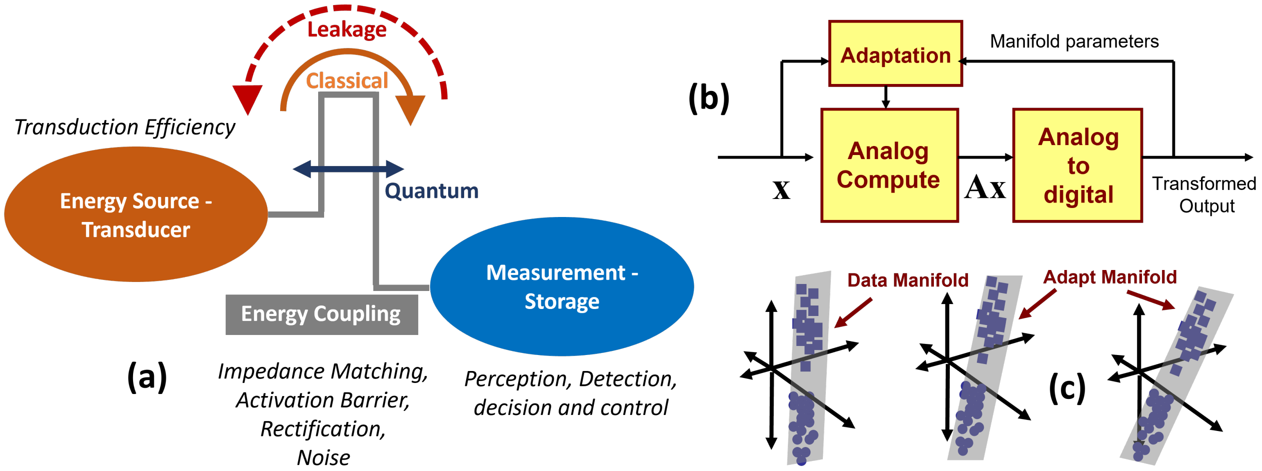

Figure 10: Research opportunities for analog computing at the edge: (a) Understanding the fundamental limits of analog computing paradigms at the interface of transduction, coupling and measurement; and (b)-(c) Integrating learning at the edge where analog compute is used to identify and track the high-dimensional sensor data manifold.

Understanding practical and theoretical challenges for analog in-sensor compute requires investigating the interface and coupling phenomena that connect the process of transduction to the process of measurement. This is illustrated in Fig. 10(a) where the objective is to match the physics of the transducer to the physics of measurement[140;142]. Fig. 10(a) shows key performance metrics that can be optimized using an end-to-end co-design approach (see Section 2.4.1 and PRD 6 for comprehensive co-design methodologies). Research objectives address multiple interconnected challenges:

Noise and nonlinearity exploitation: Research must address how noise, device and system non-linearity, and other analog artifacts impact transduction and how these physical phenomena can be exploited for computation (see Section 2.3.1 for mathematical frameworks for noise, variability, and their mitigation), including developing physics-based models that accurately predict noise propagation and its impact on computational fidelity.

Information-theoretic limits: Research must explore analog computing paradigms that can maximize information flow and operate at the fundamental limits of information transduction, signal acquisition, and measurement, establishing theoretical bounds on achievable information rates for different sensing modalities and analog processing architectures.

Energy–information tradeoffs: Research must also explore tradeoffs governing the rate of information flow and energy dissipation, developing analytical frameworks that relate sensing resolution, processing complexity, and energy consumption to guide optimal system design, identifying operating regimes where analog approaches achieve orders-of-magnitude advantages over digital counterparts.

Materials and device innovation: New materials and hardware integrated at the edge must maximize the energy coupling shown in Fig. 10(a), focusing on device primitives that naturally implement useful computational operations (e.g., correlation, convolution, thresholding) in the analog domain.

DOE testbeds and validation: DOE-specific platforms and experiments[143] must be established to validate and demonstrate new approaches, creating edge-focused benchmark problems derived from real DOE applications that feed into the shared benchmarking infrastructure in PRD 6 and enable fair comparisons between competing approaches.

In-domain analog computing: Within-domain analog computing (e.g., all-optical processing, acoustic processing) must be investigated to achieve high computational throughput and energy efficiency while avoiding costly domain conversions, developing theory and design methodologies for in-domain processing across optical, acoustic, mechanical, and electromagnetic sensing modalities.

Generalizable paradigms: General-purpose and scalable analog computing paradigms that can be applied to different signal and transduction domains must be explored, identifying common computational primitives and architectural patterns that enable reuse across applications.

Research in this area may lead to validated theoretical models of fundamental limits with co-design methodologies demonstrating 10× improvements in energy efficiency or latency, as well as in-domain processing prototypes for multiple sensing modalities and a benchmark suite for analog edge computing evaluation.

6.4.2 Analog Learning and Adaptation at the Edge

As transducer arrays (MEMS microphone arrays, RF antenna arrays, optical sensor arrays) grow in size and density, opportunities arise to exploit spatial statistics and redundancies across the array. However, extracting relevant features while suppressing noise and irrelevant information in the analog domain—before digitization—remains a fundamental challenge. Biological sensory systems provide compelling existence proofs: they achieve remarkable efficiency through tight integration of sensing and computation, local adaptation, and manifold learning—with some systems, such as chemoreceptors, operating close to the thermodynamic and information-theoretic limit[144–146] (see PRD 5 for natural and bio-inspired computing approaches).

In biological systems, sensory neurons interfacing with transducers extract salient spatio-temporal features before encoding information as spikes for transmission to higher-level processing centers[147]. Unlike traditional analog interfaces, these neuro-sensory interfaces can *adapt* and *learn* to reduce data bandwidth by extracting only relevant information[148;149]. This architecture can be viewed as a high-dimensional analog-to-information converter. An analogous approach can be applied to edge computing, where analog computation transforms the input signal before quantization using ADCs[85]. As illustrated in Fig. 10(b), the information in high-dimensional sensor data often lies on a low-dimensional manifold. The analog computing architecture must identify and track this manifold in space-time and frequency, extracting discriminatory features for compression and classification.

The key challenge is that the analog computing frontend must adapt to track the data manifold as it drifts due to changes in signal distribution or due to analog device artifacts such as environmental drifts and temperature fluctuations. Due to compression and subsequent quantization, some information might be irreversibly lost. The challenge is to identify analog computing elements (algorithms, devices, and architecture) amenable to continual learning and adaptation at a precision where analog-to-information conversion can be successful.

Research objectives address multiple interconnected challenges:

Adaptive analog feature extraction: Analog circuit architectures must be developed that can learn and adapt feature extractors in situ without requiring retraining in the digital domain, investigating mechanisms for local, unsupervised learning rules implementable in analog hardware (e.g., Hebbian-like plasticity, competitive learning).

Manifold tracking algorithms: Algorithms for tracking low-dimensional manifolds in high-dimensional sensor data using analog computation must be developed, including online adaptation strategies that account for non-stationarity in signal statistics and hardware drift.

Noise exploitation strategies: Research must explore strategies for noise exploitation and resolution enhancement including noise-shaping[147], stochastic resonance[150], and non-linear phenomena[151], determining when and how injected or intrinsic noise can improve analog information processing.

Spatio-temporal processing: Analog architectures must be developed to exploit spatial redundancy across transducer arrays and temporal correlations in sensor signals, designing efficient analog correlation and convolution primitives for array processing.

Continual learning in analog: Continual learning paradigms must be investigated for analog systems that can adapt to new tasks and environments without catastrophic forgetting, developing metrics to quantify learning performance and generalization in analog systems.

Hardware-algorithm co-design: Learning algorithms must be co-designed with analog hardware primitives to maximize efficiency, identifying computational operations that are naturally efficient in analog (e.g., multiply-accumulate, nonlinear activation functions, winner-take-all) and building learning algorithms around these primitives.

Research in this area may lead to demonstrated analog learning systems with 100× data bandwidth reduction while preserving task-relevant information, including continual learning prototypes adapting to non-stationary environments, as well as validated noise exploitation strategies with co-designed learning architectures achieving 10× energy efficiency gains over digital baselines.

6.4.3 Novel Devices and Architectures for Analog Edge Computing

To support new analog computing paradigms at the edge, new devices, compute architectures that can efficiently incorporate these devices, and energy-efficient interfaces to the architecture will be required. This necessitates exploration of unconventional computing approaches that go beyond traditional digital computing methods while addressing the unique constraints of edge deployments: extreme energy efficiency, compact form factors, operation in harsh environments, and real-time processing.

Research objectives span multiple interconnected areas:

Emerging device primitives: Emerging computing and storage devices such as memristors, tunneling devices, phase-change materials, and spintronic devices must be explored to achieve scale and robustness while operating at high energy efficiencies, investigating device characteristics that enable native analog operations (e.g., multiply-accumulate, nonlinear activation, correlation) with minimal energy overhead.

Compute-in-memory for edge: The synergy between state-of-the-art practices like CiM, analog, and digital computing and how these practices can be adapted and optimized for edge computing constraints must be explored (see PRD 2 for comprehensive discussion of CiM and mesh architectures for HPC), developing edge-specific CiM architectures that balance programmability, precision, and energy efficiency.

Task-driven architecture optimization: Problem complexity and figures of merit relevant to edge tasks (information bandwidth, recognition rates, false alarm rates, latency, energy per inference) must be defined and used to optimize analog computing architectures, developing analytical frameworks that map application requirements to optimal architectural choices.

Few-shot and zero-shot learning architectures: Analog computing architectures and operations must be designed to minimize the data required for calibration and training using one-shot or zero-shot learning paradigms, investigating meta-learning approaches implementable in analog hardware that enable rapid task adaptation with minimal examples.

Noise-aware design: Device and system noise must be exploited to design analog processors at the edge, developing metrics that incorporate noise into physical modeling and identifying computational tasks where noise can be beneficial (e.g., stochastic sampling, exploration in optimization, regularization), creating design methodologies that co-optimize signal and noise characteristics.

Analog–digital trade-off analysis: Analog computing dynamics and representations must be rigorously analyzed to compare trade-offs versus digital design, developing quantitative models that predict when analog approaches will outperform digital alternatives for specific tasks and constraints (energy budget, latency requirements, precision needs).

Benchmarking and validation: End-to-end benchmarking of analog computing at the edge and its benefits compared to traditional approaches must be conducted, establishing standardized benchmarks derived from DOE applications (e.g., particle detection, fusion diagnostics, quantum control) that enable fair comparisons across different analog architectures and against digital baselines.

In-domain computing modalities: In-domain computing modalities beyond electronics, such as optical computing, acoustic computing, and mechanical computing, must be explored, developing theory and design principles for in-domain processing that maximizes energy coupling and minimizes conversions, and investigating hybrid approaches that combine multiple physical domains.

Programmability and software integration: The programmability of analog computing systems and their compatibility with existing and proven algorithms must be ensured by leveraging the common analog intermediate representation (intermediate representation (IR)) and compiler toolchain defined in PRD 6; developing edge-specific high-level programming abstractions and runtime systems that enable non-experts to deploy analog edge computing solutions; and creating standardized interfaces between analog and digital components.

Research in this area may lead to novel analog edge computing device prototypes achieving 100× energy efficiency improvements with validated architectures for DOE application domains, as well as a benchmark suite spanning diverse sensing modalities with demonstrated few-shot learning systems and programmable platforms featuring open-source toolchains.

6.5 Key Takeaways

This PRD addresses the critical challenge of data proliferation from next-generation sensors through analog computing integrated directly with sensing and transduction. Three interconnected research directions provide a comprehensive path forward:

- **Fundamental Limits through Co-Design:** Establishing theoretical frameworks for understanding the physical limits of analog information processing at the edge, including energy-information tradeoffs and optimal coupling between sensing and computation.
- **Analog Learning and Adaptation:** Developing algorithms and architectures for in situ learning that can adapt to non-stationary environments and extract task-relevant features in the analog domain, achieving bandwidth reduction while preserving scientific fidelity.
- **Novel Devices and Architectures:** Creating specialized analog computing primitives and system architectures optimized for edge constraints of energy, latency, form factor, and harsh environments—including radiation-hard devices for DOE facilities.
- **DOE-Specific Applications:** Real-time feature extraction from particle detectors, sub-microsecond feedback for quantum systems, plasma diagnostics for fusion stabilization, and massive deployments for climate and environmental monitoring represent high-impact targets.
- **Cross-PRD Integration:** Edge computing leverages device primitives (PRD 3), connects to mathematical foundations for information-theoretic limits (PRD 1), complements high-performance systems (PRD 2), and benefits from natural computing approaches (PRD 5) and co-design methodologies (PRD 6).

Success requires addressing key scientific challenges: understanding noise and nonlinearity as both limiting factors and computational resources; developing rigorous theoretical frameworks linking information theory, physics, and computation; and creating end-to-end co-design methodologies that simultaneously optimize sensing, analog processing, and digital interfaces.

7 PRD 5: HARNESSING, ENHANCING, AND INSPIRING EFFICIENT COMPUTATION THROUGH NATURAL SYSTEMS

7.1 Introduction and Vision

Natural systems—from biological cells to physical spin networks—demonstrate remarkable computational capabilities surpassing conventional computing in energy efficiency, adaptability, and robustness. Modern electronic components are individually more efficient than neurons, yet paradoxically consume orders of magnitude more power when assembled into AI systems. Biological and physical systems perform complex computations at energy scales orders of magnitude lower by harnessing natural dynamics: chemical reaction networks, coupled oscillators, organoid assemblies, magnetic spin interactions. For select tasks, these approaches avoid inefficient mappings such as binarization, providing massive concurrency at low energy costs. Chemical and biological substrates also enable computation in environments fundamentally incompatible with electronics—inside cells, within tissues, and in molecular systems—opening entirely new application domains.

Key Research Questions: *How can we harness computational principles from natural systems to overcome the energy and scaling limitations of conventional computing? What programming abstractions, control mechanisms, and interfaces are needed to make natural computing practical for DOE applications? Can we establish theoretical foundations that predict capabilities and guide co-design of natural computing substrates across physical, chemical, and biological modalities?*

Vision: This PRD envisions establishing natural computing as a transformative paradigm complementing conventional computation through: (1) achieving biological-level energy efficiency for broad application classes, (2) developing intuitive programming tools enabling non-specialists to deploy natural computing solutions, (3) creating self-organizing adaptive systems that leverage inherent physical dynamics, (4) seamlessly integrating natural computing into mainstream computational ecosystems, and (5) establishing natural computing as an enabling technology for DOE's scientific mission.

7.2 Goals and Desired Outcomes

The development of natural computing requires progressively ambitious, coordinated efforts to establish it as a transformative paradigm complementing conventional computation. Main research needs follow.

- **Theoretical Foundations:** Develop theoretical frameworks predicting performance bounds for diverse natural computing approaches.
- **Programming and Control:** Develop comprehensive programming tools (high-level abstractions) for selected natural computing modalities (physics-based, chemical, biological) that target the common analog IR and compiler toolchain defined in PRD 6; contribute substrate-specific backends, runtime control, and robust bidirectional interfaces enabling hybrid system demonstrations with conventional computing infrastructure.
- **Scaled Demonstrations:** Demonstrate scaled prototype implementations with practical relevance beyond laboratory proofs of concept, establishing initial fabrication standards and reproducibility protocols for key natural computing elements, while leveraging the cross-cutting manufacturing and metrology framework in Section 2.5 with modality-appropriate assays and biosafety.
- **Benchmarking Frameworks:** Establish standardized characterization and benchmarking protocols enabling fair comparisons across natural computing modalities and against conventional computing baselines, building on Sections 2.5 and 2.4.3.
- **Mainstream Integration:** Enable non-specialists to readily deploy natural computing solutions through seamless integration into mainstream computational ecosystems, with comprehensive design automation integrated with the standardized IR and compiler toolchain from PRD 6, spanning theory, programming, and hardware implementation.
- **Advanced Energy Efficiency:** Achieve biological-level energy efficiency approaching thermodynamic limits for broad application classes, with deployed natural computing accelerators at DOE HPC facilities achieving order-of-magnitude improvements in energy-per-operation for target workloads.
- **Adaptive Systems:** Develop self-organizing, adaptive computing systems that leverage emergent dynamics and learn from environmental feedback, establishing natural computing as an indispensable mainstream paradigm complementing digital approaches for high-impact scientific and engineering challenges.

- **Novel Application Domains:** Deploy chemical and biological computing in environments inaccessible to electronics (inside cells, within tissues, harsh chemical conditions) for applications in smart therapeutics, environmental monitoring, and synthetic biology.

7.3 State of the Art

Natural computing encompasses diverse approaches—physics-based systems (Ising machines, coupled oscillators), biological/chemical substrates (DNA circuits, organoids, reaction networks), brain-inspired analog circuits, and engineered dynamical systems. While recent progress demonstrates compelling advantages for specific tasks, fundamental limitations prevent broader adoption and scaling to DOE-relevant problem sizes.

Programmability and Abstraction Gaps: Current natural computing systems require expert-level manipulation of low-level physical, chemical, or biological parameters. High-level programming abstractions analogous to conventional languages are largely absent. Problem mapping requires manual translation from computational goals to substrate-specific configurations (Hamiltonian parameters, reaction concentrations, network topologies), severely limiting accessibility and development speed.

Scalability and Fabrication Challenges: Laboratory demonstrations rarely scale beyond small prototypes. Physics-based systems face device imperfections, parameter drift, and connectivity constraints. Biological/chemical systems struggle with long-term stability, reproducibility, and contamination. Brain-inspired analog circuits cannot replicate the brain's massive fan-out and adaptive plasticity within current technology constraints. Scaling from dozens to thousands or millions of computational elements remains an open challenge across all modalities.

Control, Robustness, and Noise Management: Natural systems exhibit inherent stochasticity, variability, and sensitivity to environmental factors (temperature, electromagnetic interference, biological evolution). Effective control with limited observability and actuation points is poorly understood. Verification and validation techniques for analog, stochastic computation are immature. While noise can sometimes be harnessed constructively (stochastic resonance, annealing), systematic methods to manage or exploit variability are lacking.

Interface and Integration Deficiencies: Efficient bidirectional transduction between natural computing elements and conventional digital electronics remains a critical bottleneck. Signal conversion across physical domains (chemical concentrations to voltages, optical to electrical, quantum to classical) introduces noise, latency, and energy overhead that can negate natural computing advantages. Lack of standardized interfaces prevents seamless integration with existing HPC infrastructure.

Learning Paradigms and Biological Efficiency: Artificial neural networks typically rely on top-down global optimization requiring centralized weight updates and full network information propagation. In contrast, biological neurons adapt locally in a bottom-up manner, using only information available at individual synapses without requiring global coordination. This local learning paradigm is dramatically more energy-efficient and inherently tolerant to defects and damage—addressing critical vulnerabilities of traditional analog computation. Additionally, biological systems achieve computational sophistication through combinatorial interactions of simple building blocks, as exemplified by gene regulatory networks and molecular signaling pathways, where complex behaviors emerge from relatively simple components operating under local rules. Understanding and replicating these principles could unlock new pathways to energy-efficient, robust analog computing.

7.4 Research Directions

Four research directions structure this PRD, as illustrated in Figure 11. Theoretical foundations provide unified frameworks, performance bounds, and local learning formalisms across all substrate modalities. Programming infrastructure translates these frameworks into domain-specific languages (DSLs) and control mechanisms that target a shared analog IR. Devices and architectures realise the three substrate families at scale and interface them with conventional computing. Benchmarking and scalability assessment closes the loop, driving hardware priorities and validating deployment readiness.

7.4.1 Theoretical Foundations and Modeling

Natural computing systems lack the comprehensive theoretical frameworks that guide conventional computing design. Without rigorous mathematical theories, we cannot predict performance bounds, understand fundamental limits, or systematically optimize designs. Analogous to how quantum information theory revolutionized quantum

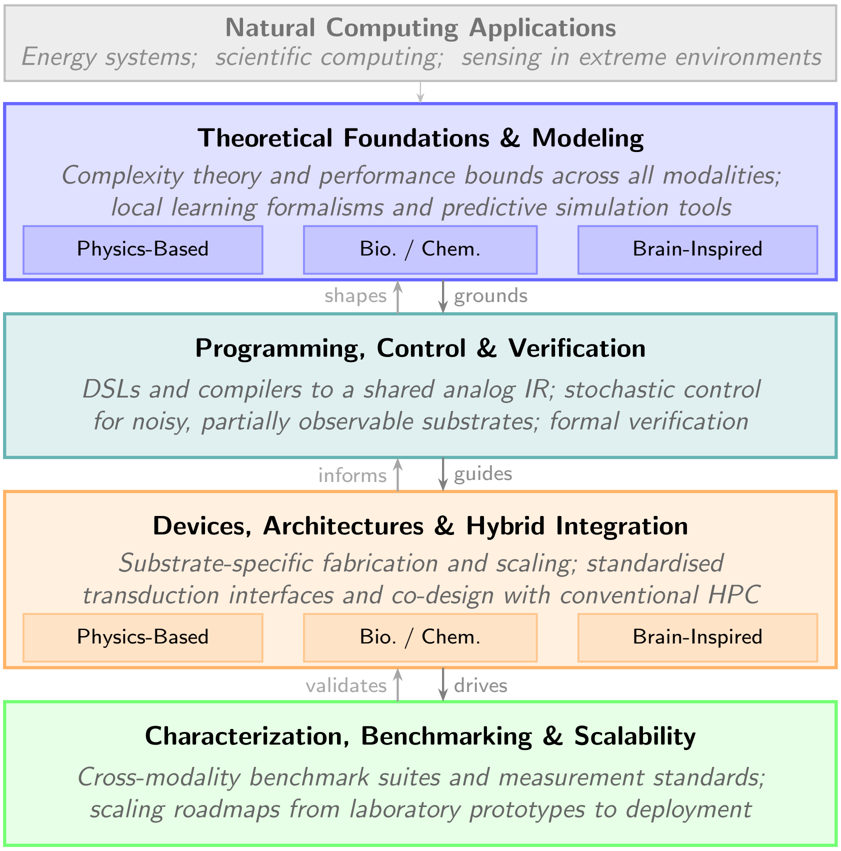

Figure 11: Overview of the four research directions in PRD 5. Theoretical foundations ground the programming layer, which in turn guides device and architecture development across three substrate modalities (physics-based, biological/chemical, and brain-inspired). Benchmarking drives hardware priorities and validates scalability toward deployment.

computing, natural computing requires unified theories encompassing physics-based, biological, and chemical computation.

Research objectives include developing comprehensive mathematical frameworks unifying diverse natural computing approaches, capable of describing memory models (fixed-point attractors, sequential memories), information-theoretic properties of analog systems, and complexity measures for physics-based and biological computation; characterizing how biological systems achieve computational sophistication through combinatorial interactions of simple building blocks—such as gene regulatory networks and molecular signaling pathways—and formalizing principles of local, bottom-up learning that enable energy efficiency without requiring global optimization; establishing clear performance bounds and fundamental limits, including theories linking energy dissipation, information processing capacity, and system robustness; and creating predictive simulation and emulation tools that accurately model complex dynamics before costly fabrication, enabling design space exploration and verification.

Research in this area may lead to a unified theoretical framework providing design principles across modalities with performance bounds and complexity measures enabling rigorous comparison across natural and conventional computing, as well as validated simulation tools achieving 10× faster design cycles with energy-information-robustness trade-off theories to guide system design.

7.4.2 Programming, Control, and Verification

Current systems demand manual manipulation of substrate-specific parameters, creating barriers for DOE application scientists. Controlling stochastic, partially observable systems with delayed feedback requires fundamentally new control theories beyond conventional approaches.

Research objectives include developing intuitive, expressive domain-specific languages (DSLs) for natural computing that directly express concepts such as reaction networks, cellular interactions, physical dynamics, and emergent collective behaviors, and that compile to the common analog IR from PRD 6; creating substrate-specific backends and runtime adapters translating programs to diverse platforms (physics-based, chemical, biological) while optimizing for each platform's unique characteristics (low energy, massive parallelism, inherent adaptivity); advancing control theory and reinforcement learning methods for systems with largely unobservable, noisy, and delayed feedback—particularly critical for biological and chemical substrates; establishing rigorous verification and validation techniques for stochastic analog computation, including formal methods adapted to continuous dynamics; and developing techniques for precisely guiding self-organizing dynamics toward desired computational states through targeted perturbations and engineered feedback loops.

Research in this area may lead to high-level programming languages leveraging the unified toolchain to achieve 100× programming time reduction, enabling application scientists to deploy natural computing without substrate expertise; control methods operating with limited observability; verification techniques for analog stochastic computation; and demonstrations of feedback-guided self-organization for multiple modalities.

7.4.3 Devices, Architectures, and Hybrid Integration

Natural computing's promise depends on practical implementations spanning diverse physical, chemical, and biological substrates. Each modality offers unique advantages: physics-based systems leverage intrinsic dynamics for solving hard problems (optimization, differential equations); chemical and biological substrates approach thermodynamic efficiency limits; brain-inspired circuits promise massive parallelism. However, isolated substrate development is insufficient—seamless integration with conventional computing through hybrid architectures is essential for DOE adoption.

Research objectives include advancing physics-based computing systems—coherent Ising machines[152], coupled oscillator networks, optical computing, and analog differential equation solvers—by addressing device imperfections and parameter drift, developing precise fabrication and calibration methods, creating efficient problem-to-Hamiltonian mapping techniques, designing architectures that scale to large problem sizes while managing connectivity and signal integrity, and mitigating environmental sensitivity (thermal fluctuations, electromagnetic interference); developing chemical and biological substrates, including DNA strand-displacement circuits[153;154] implementing arbitrary reaction networks with analog signal processing, protein-based circuits that achieve robust adaptation in living cells[155;156], organoid-based computation, in vitro neural networks, and non-DNA substrates (protein conformational changes, lipid membrane dynamics, metabolic fluxes), while advancing molecular analog computing for smart therapeutics[157;158] that detect specific disease states and autonomously generate therapeutic responses by differentially weighting multiple molecular signatures and biosensors[159] with sophisticated decision-making capabilities, and addressing scalability[160], long-term stability, reproducibility, efficient readout mechanisms, and cytomorphic compilation backends[161] translating chemical models to analog circuits for real-time simulation via the shared compiler infrastructure from PRD 6; bridging the gap between theoretical models and practical hardware by developing sophisticated input/output systems for faithful signal transduction, control mechanisms that enable external manipulation of system parameters despite limited observability and inherent stochasticity, adaptive control strategies that learn system dynamics from noisy measurements, feedback protocols that maintain desired computational states, and interface standards enabling reproducible experimental validation; advancing brain-inspired analog circuits through alternative materials such as memristors[40;162] and disordered memristive networks[163], 3D integration and heterogeneous computing for brain-like connectivity, addressing analog component reliability, power consumption, and scalability challenges[164], and driving truly analog implementations—beyond digital neuromorphic chips such as TrueNorth, Loihi2, and SpiNNaker—that achieve the brain's fan-out and adaptive plasticity; and creating efficient interfaces between natural and conventional computing (see Section 2.3.2), including chemo/bio-electronic interfaces with high-fidelity signal conversion (chemical concentrations ↔ voltages, optical ↔ electrical), digital-physical hybrid frameworks integrating physics-based accelerators (Ising machines, oscillator networks) with CPU/GPU infrastructure, clear interface standards ensuring interoperability, and co-design methodologies (Section 2.4.1 and PRD 6) for hybrid systems.

Chemical and biological substrates offer a distinctive advantage: they operate in environments inaccessible to conventional electronics—inside cells, within biological tissues, and in molecular self-assembly or DNA data storage systems. Biological chemical computation achieves costs per logical operation six orders of magnitude lower than those of electronic computers, a proof of principle that evolution has optimized these processes for energy conservation. Research in this direction may also benefit from connections to DOE's Biological and Environmental Research (BER) program, particularly the Biological Systems Science Division (BSSD), which addresses fundamental design principles of biological systems relevant to understanding and engineering biological computation.

Research in this area may lead to testbed demonstrations spanning physics-based, chemical/biological, and hybrid modalities for DOE applications, achieving 1000× energy efficiency improvements for optimization and simulation tasks; standardized interface protocols with fabrication standards achieving 10× improved reproducibility; and scalable architectures (10,000+ elements) with validated hybrid systems achieving synergistic performance.

7.4.4 Characterization, Benchmarking, and Scalability

Without standardized benchmarks and measurement techniques, natural computing progress cannot be rigorously assessed or compared across modalities and against conventional approaches. The current lack of metrics, reproducibility challenges, and unclear scaling paths hinder both research prioritization and DOE adoption decisions. Community-driven standards are essential to transition from ad-hoc demonstrations to systematic engineering.

Research objectives include developing community-driven benchmark suites analogous to NeuroBench[165] and MLPerf[166], tailored to natural computing's unique capabilities and operational regimes (see Section 8.4.5 for comprehensive benchmarking discussion); creating problem sets capturing aspects beyond traditional metrics: energy efficiency (operations/watt, solutions/joule), solution quality for optimization tasks, robustness to noise and variability, adaptability, and convergence time; establishing measurement techniques accurately quantifying performance in inherently analog, stochastic systems; creating comprehensive comparative frameworks enabling meaningful evaluation across physics-based, chemical/biological, and brain-inspired approaches, and against conventional digital methods; addressing scalability systematically by developing architectural principles, fabrication reliability standards, and communication protocols (see Section 2.3.3) to transition from laboratory prototypes to DOE-relevant scales while maintaining computational coherence and efficiency; and establishing standards for integration with conventional digital systems, including physical/logical interfaces and software frameworks.

Research in this area may lead to benchmark suites covering optimization, differential equations, pattern recognition, and DOE-specific applications with standardized measurement protocols achieving 10× improved reproducibility; comparative frameworks enabling rigorous cross-modality evaluations; scaling roadmaps identifying bottlenecks; and demonstrated scalability from laboratory (10–100 elements) to practical systems (10,000+ elements).

7.5 Key Takeaways

Natural computing addresses fundamental limitations of conventional computing through four interconnected research directions, with key points summarized below.

- **Theoretical Foundations:** Predictive models, performance bounds, and unified frameworks across physics-based, chemical, and biological modalities provide the scientific grounding for natural computing—including understanding how biological systems achieve computational sophistication through combinatorial interactions of simple building blocks.
- **Programming Infrastructure:** High-level languages, automated compilation, and robust control methods for stochastic systems enable non-experts to deploy natural computing solutions, dramatically lowering barriers to adoption.
- **Practical Implementations:** Physics-based accelerators (Ising machines, oscillators), chemical/biological substrates (DNA circuits, organoids), brain-inspired analog circuits, and hybrid architectures seamlessly integrating with conventional computing provide diverse pathways to achieving natural computing's potential.
- **Evaluation Methodologies:** Benchmarks, measurement standards, and scaling roadmaps enable rigorous progress assessment and guide DOE investment toward highest-impact opportunities.
- **DOE Mission Applications:** Energy systems (real-time power control, grid optimization), scientific computing (plasma dynamics, multi-physics modeling, molecular dynamics), and sensing/control in extreme environments represent transformative targets for natural computing.
- **Cross-PRD Integration:** Natural computing leverages mathematical foundations (PRD 1), device primitives (PRD 3), complements high-performance (PRD 2) and edge computing (PRD 4), and benefits from co-design methodologies (PRD 6).

Success requires coordinated efforts across all four directions—isolated advances are insufficient. Co-design spanning theory, programming, devices, and hybrid integration is essential to realize transformative impact.

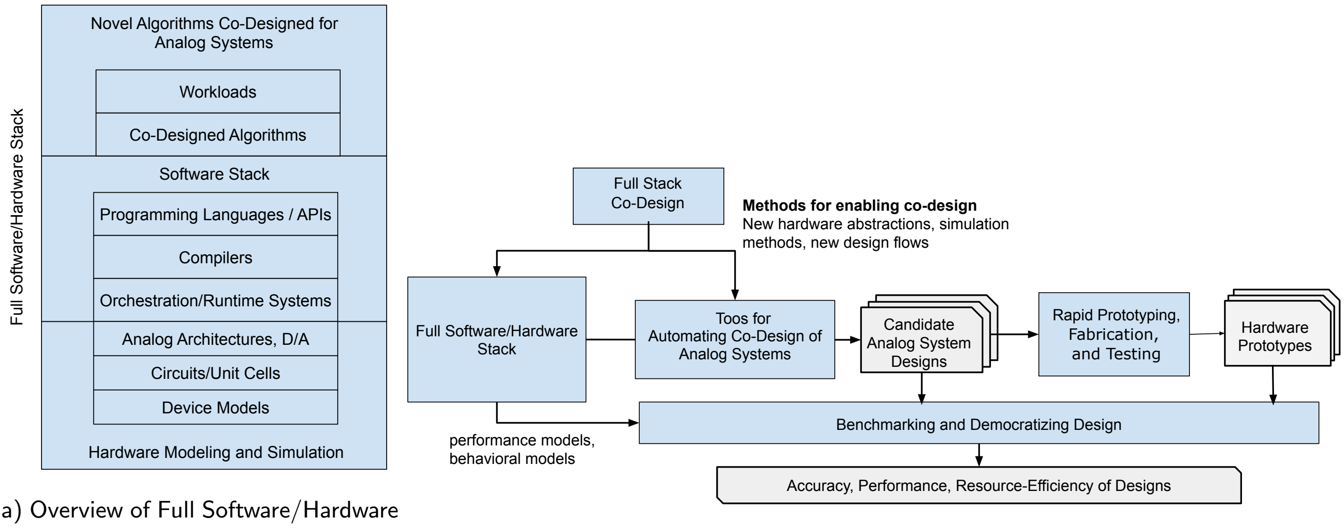

(a) Overview of Full Software/Hardware Stack

(b) Co-design Flow

Figure 12: Key Technologies for Cross-Cutting Multi-Granularity Design

8 PRD 6: CROSS-CUTTING MULTI-GRANULARITY CO-DESIGN METHODS

8.1 Introduction and Vision

As detailed in Section 2.4, co-design methodologies are essential for translating analog computing innovations into practical systems. Hardware-software co-design principles, hardware-aware algorithm development, and standardized benchmarking frameworks provide the foundation for integrated system development. However, realizing these principles for analog computing requires addressing specific research gaps and developing new tools, methodologies, and infrastructure.

This PRD identifies the research directions needed to build a comprehensive co-design ecosystem for analog computing. While digital computing benefits from mature design automation, programming models, and evaluation frameworks developed over decades, analog computing lacks equivalent infrastructure. The challenges are fundamental: continuous computation models defy traditional abstraction layers, device non-idealities propagate through the stack, and application-specific trade-offs require simultaneous optimization across materials, circuits, architectures, and algorithms.

Key Research Questions: *How do we build multi-granularity modeling frameworks that capture analog behavior from device physics to system architectures? What programming abstractions and compiler toolchains enable productive development while preserving analog efficiency? Can we develop automated design space exploration tools that navigate the complex trade-offs inherent in analog systems? How do we establish benchmarking infrastructure that enables fair comparison across diverse analog computing approaches?*

Vision: Figure 12a illustrates the full software/hardware stack that must be addressed, from algorithms and programming interfaces through architectures, circuits, and device models. Figure 12b depicts the integrated co-design flow encompassing modeling, software development, automated tools, and benchmarking. Success requires bridging fundamental research gaps in each area while ensuring the components work together as a cohesive ecosystem.

8.2 Goals and Desired Outcomes

The overarching goal is to develop a comprehensive co-design ecosystem that enables efficient development, optimization, and deployment of analog computing systems across the full spectrum from HPC platforms to edge devices, supporting cross-layer design space exploration, providing reusable tools and methodologies, and accelerating the path from innovative devices to practical systems. Main research needs follow.

- **Unified Modeling Frameworks:** Establish multi-resolution simulation frameworks accurately modeling analog hardware from device physics through system architectures, with validated models for major analog computing paradigms achieving 10× speedup over SPICE-level simulation while maintaining fidelity for algorithm development.

- **Software Stack Foundations:** Develop domain-specific languages, compiler toolchains, and runtime systems for analog computing platforms, demonstrating automated compilation and deployment of scientific workloads with correctness guarantees through formal verification methods adapted to continuous computation models.
- **Automated Co-Design Tools:** Create AI-guided design space exploration tools for joint hardware-software optimization across multiple stack layers, demonstrating 100× reduction in designer time for co-optimizing hybrid analog-digital systems.
- **Rapid Prototyping Platforms:** Establish rapid prototyping platforms enabling quick cycles from concept to physical validation with seamless analog-digital integration.
- **Benchmarking Infrastructure:** Establish community-driven benchmark suites with standardized metrics, reference implementations, open-source hardware models, and cloud-accessible testbeds for major analog computing technologies, drawing on domain-specific benchmark problems defined in the application-focused PRDs (e.g., high-performance analog computing in PRD 2 and edge and in-sensor computing in PRD 4).
- **End-to-End Automation:** Achieve push-button design flows automatically synthesizing optimized analog computing systems from high-level application specifications, enabling design space exploration encompassing 1000+ hardware/software co-design points within practical timeframes.
- **Cross-Platform Portability:** Develop hardware-agnostic intermediate representations and compilation frameworks enabling algorithm portability across diverse analog computing platforms with automated platform-specific optimization and mature verification methodologies.
- **Community Ecosystem:** Build a vibrant research community with shared tools, datasets, benchmarks, and fabrication access, reducing barriers to entry and enabling researchers to focus on innovation, with demonstrated 10× time-to-prototype improvements and commercial adoption by industry partners.

8.3 State of the Art

Current approaches to analog computing system design are predominantly serialized and fragmented across disciplines. Materials scientists and device physicists develop novel devices often in isolation for years before circuit designers and architects begin developing peripherals and system architectures. This protracted serial process means device technologies may take decades to reach commercially viable systems—a timeline incompatible with the rapid pace of scientific computing needs.

Modeling and Simulation Limitations: Existing simulation tools operate at specific abstraction levels with limited cross-layer integration. Device-level SPICE simulations provide high fidelity but are computationally prohibitive for system-scale exploration (billions of cycles required for software development). Higher-level behavioral models often oversimplify critical analog phenomena—variability, parasitics, nonlinearity, noise—leading to poor predictive accuracy at scale. Frameworks such as CrossSim[167] address some challenges for specific analog paradigms (e.g., crossbar arrays) but lack generalizability across diverse analog computing approaches. The absence of standardized interfaces between modeling layers forces researchers to develop bespoke solutions, resulting in substantial duplicated effort and hindering model reuse across the community.

Software Stack Gaps: Unlike digital systems with mature abstractions (instruction set architectures (ISAs), compilers, runtime systems), analog computing lacks foundational software infrastructure. Most analog systems require manual low-level programming with limited portability. Compiler technologies for analog are nascent, with few attempts at hardware-aware optimization or automated mapping of algorithms to analog hardware. Formal verification methods designed for discrete Boolean logic do not transfer to continuous analog computation models, leaving correctness verification largely unaddressed. The absence of domain-specific languages and standardized programming models forces each research effort to reinvent basic software infrastructure.

Co-Design Tool Deficiencies: Existing electronic design automation tools (EDA) focus on single abstraction layers with limited cross-stack optimization capabilities. Manual co-design efforts, while successful in limited demonstrations, require deep expertise spanning materials through algorithms—skills rarely possessed by individual researchers and not systematically taught. AI-guided design tools for analog are emerging, including evolutionary algorithms for device-level synthesis[168], reinforcement learning for circuit design and optimization[169–171], generative AI for circuit characterization[172;173], physics-informed neural networks as device/circuit surrogates[174], and hybrid ML methods for material optimization[175;176]. However, these remain specialized to narrow problem domains (e.g., specific circuit topologies or device classes) rather than providing general-purpose design space exploration across the full analog computing stack. The lack of integrated toolchains spanning the full stack forces protracted design cycles and limits exploration of novel analog computing architectures.

Rapid Prototyping Limitations: Analog system development faces long iteration cycles due to limited rapid prototyping infrastructure. While digital FPGAs revolutionized digital design by enabling quick iterations, analog equivalents remain nascent. Reconfigurable FPAAs[16;54;92] exist but have limited scale and routing resources. Chiplet-based approaches[177] promise modular composition but lack mature analog-specific interconnect standards. Analog standard cell libraries[178;179] and associated synthesis tools[55;92;94] are emerging but not yet production-ready. Hybrid analog-digital integration within commercial FPGAs[180–183] provides limited analog capabilities focused on narrow use cases. These prototyping platforms remain fragmented without common abstractions or interfaces, requiring significant expertise to use effectively and limiting accessibility for algorithm developers and domain scientists.

Benchmarking and Evaluation Challenges: Standardized evaluation frameworks are addressed in Section 8.4.5.

8.4 Research Directions

This section outlines five interconnected research directions that together comprise a comprehensive co-design framework for analog computing systems. Each direction addresses critical gaps identified in the state of the art while building toward an integrated ecosystem.

8.4.1 Hardware Modeling and Simulation Frameworks

Hardware modeling and simulation frameworks are foundational to co-design, enabling prediction of system performance, evaluation of design trade-offs, and optimization before costly physical fabrication. Models must capture complex continuous signal dynamics—including nonlinearity, noise, variability, and temperature dependencies—across multiple abstraction levels from device physics to system architectures. Current tools either provide high-fidelity device simulation (prohibitively expensive for system-scale exploration) or oversimplified behavioral models (insufficient for accurate prediction). Multi-resolution frameworks that balance accuracy and computational efficiency are essential for practical co-design.

Research objectives include developing parameterized, multi-resolution modeling frameworks that span device, circuit, architecture, and system levels with standardized interfaces between layers; enabling co-simulation of analog and digital elements within unified frameworks; creating scalable models supporting tens to hundreds of core analog compute units with various interconnect technologies (digital-CMOS, analog, photonics); establishing generic interface standards for analog-digital integration compatible with emerging protocols (e.g., Compute Express Link (CXL) for chiplet communication); and extending frameworks such as Crosssim[167] to represent emerging device/material characteristics, hardware-in-the-loop simulation, and seamless digital co-simulation integration. Models should be amenable to optimization algorithms and design space exploration tools while maintaining sufficient fidelity for algorithm development.

Research in this area may lead to multi-resolution simulation frameworks achieving 10× to 100× speedup versus SPICE while maintaining $<10\%$ error for key metrics, with open-source model libraries demonstrating scalability to system-level designs, and reduced designer time by 5× through model reuse and standardized interfaces.

8.4.2 Software Stack Development

Software infrastructure—programming languages, compilers, runtime systems—is critical for making analog hardware accessible and usable for scientific applications. Consistent with its cross-cutting scope, this PRD is the canonical home for intermediate representations and compiler toolchains that are referenced by other PRDs (e.g., PRD 1, PRD 4, PRD 5). Without robust software stacks, each analog system requires manual low-level programming, limiting adoption and preventing algorithm portability across platforms.

Research objectives include developing domain-specific languages (DSLs) and programming application programming interfaces (APIs)—both standalone and embedded in Python/Julia—enabling high-productivity programming of analog systems, applying algorithm design principles from PRD 1 to practical programming; creating analog intermediate representations (IRs) analogous to Low Level Virtual Machine (LLVM) IR that implement semantic foundations for representation developed in PRD 1, supporting hardware-agnostic mid-end optimizations while targeting diverse analog backends; designing hardware-aware compilers that map algorithms to analog hardware while applying compensating optimizations for non-idealities (e.g., gain compensation, noise mitigation based on noise-aware optimization frameworks from PRD 1); building runtime systems for hardware characterization, calibration leveraging precision-convergence theory from PRD 1, configuration, and orchestration of hybrid analog-digital computations; establishing formal verification methods adapted to continuous computational models building on PRD 1's formal models and canonical computation frameworks, including computer algebra systems for mathematical rewriting and dynamical systems analysis for equivalence checking; developing multi-chip mapping algorithms for heterogeneous analog-digital systems; creating correctness-preserving program transformation techniques ensuring semantic preservation across abstraction layers; and extending initial efforts in analog high-level synthesis (HLS) for design space exploration[55] to production-quality toolchains.

Research in this area may lead to complete software toolchains demonstrating automated compilation from high-level DSLs to hardware with hardware-aware optimizations achieving 2× to 10× improvements versus naive mappings; formal verification frameworks with reusable compilation passes demonstrating portability across analog hardware types; and accessible programming tools enabling domain scientists to utilize analog computing without deep hardware expertise.

8.4.3 Automated Co-Design Tools and Methods

Analog microelectronics design is labor-intensive, requiring skilled expertise, extensive simulations, fabrication, and testing cycles. Design space exploration to identify optimal hardware/software configurations is fundamentally a constrained multi-objective optimization problem in high-dimensional continuous and discrete parameter spaces. The fundamental challenge is navigating vast, non-convex design spaces with continuous and discrete parameters spanning multiple abstraction levels—finding Pareto-optimal designs that balance competing metrics (fidelity, energy efficiency, area, programmability) while satisfying physical and manufacturing constraints. The problem structure includes: (1) high-dimensional search spaces where exhaustive exploration is intractable; (2) cross-layer coupling where changes at device level propagate through circuit, architecture, and algorithm levels, creating complex dependencies; (3) non-convexity with local minima requiring global optimization strategies; (4) knowledge representation challenges for encoding design expertise and constraints in tractable mathematical frameworks; and (5) parameter variability requiring uncertainty quantification for robust designs. This connects to established applied mathematics research areas: optimization theory for multi-objective non-convex problems, uncertainty quantification for parameter variability and manufacturing tolerances, model order reduction for creating tractable surrogates of complex SPICE simulations, and formal verification for ensuring correctness under varied operating conditions. Automated design space exploration tools that can simultaneously optimize across multiple stack layers—algorithms, architectures, circuits, devices—are essential to accelerate development and discover Pareto-optimal designs. Manual co-design is intractable for complex systems; building on PRD 1's computational models and completeness proofs that identify minimal primitive sets, co-design tools must find optimal combinations and configurations of these primitives. Currently, most co-design constraints are not fully automated, representing a significant opportunity for transformative progress[9;184].

Research objectives include developing modular co-design software tools with version control, collaboration features, and training datasets enabling multi-layer optimization; building tools capable of white-box optimization using gradient-based methods with interpretable models (leveraging differentiable simulation frameworks) and black-box optimization for systems too complex to differentiate, using simulator-only methods; applying AI-guided design space exploration methods for discovering Pareto-optimal designs across the multi-objective landscape[185–187]; and integrating automated testing and validation methods ensuring performance specifications under varied conditions, building on formal verification frameworks adapted to continuous computation models (developed in PRD 1).

Research in this area may lead to automated design space exploration tools demonstrating 100× reduction in designer time with discovery of Pareto-optimal designs achieving 10× to 100× energy efficiency improvements for specific workloads, as well as open-source co-design toolchains supporting diverse analog paradigms with community datasets enabling AI-guided tool training and end-to-end integration.

8.4.4 Hardware-Algorithm Co-Optimization and Rapid Prototyping

The fundamental challenge is abstraction design: digital programming models (discrete, sequential, deterministic) fundamentally mismatch analog phenomena (continuous, parallel, stochastic). This is a programming language and formal methods research problem. Analog computing's potential can only be fully realized through algorithms specifically designed to exploit continuous signals, natural parallelism, and inherent nonlinearity—rather than simply porting digital algorithms. Building on algorithm design principles from PRD 1, hardware-aware algorithms must account for and leverage analog-specific phenomena: noise, device variability, drift, and nonidealities discussed in Section 2.3.1. This applies across diverse analog computing paradigms including physics-based, chemical, and biological substrates (see PRD 5 for natural computing approaches). Conversely, hardware should be tuned to support algorithmic requirements. This bidirectional co-optimization is crucial where physical properties fundamentally impact computation.

Recent work[7] demonstrates the transformative power of well-designed abstractions: a differentiable digital twin of an AOC achieving over 99% correspondence with physical hardware enabled training models specifically adapted to analog characteristics, while a co-designed fixed-point abstraction provided noise robustness. This illustrates that abstraction design—defining the right mathematical interfaces and programming models—is as critical as hardware design itself, enabling fundamentally new computational paradigms. The challenge is creating abstractions that preserve analog advantages while being tractable for algorithm developers, implementing semantic foundations for representation developed in PRD 1.

Rapid prototyping capabilities—through reconfigurable platforms (FPAAs), chiplet ecosystems, and rapid IC fabrication—are essential to validate and iterate co-optimized designs, reducing development cycles from years to weeks. Currently, analog design requires highly trained engineers, long engineering times, and significant financial resources, with fabrication requiring extended lead times. Reconfigurable hardware and synthesis tools can dramatically accelerate prototyping, similar to how FPGAs revolutionized digital design.

Research objectives include developing abstraction layers defining interfaces between algorithm, architecture, circuit, and device that preserve analog advantages (continuous-time operation, natural parallelism, physical dynamics) while providing tractable programming models; creating mathematical frameworks for hardware-aware algorithms that exploit continuous computation, building on algorithm design principles from PRD 1; establishing variability-tolerant abstractions accommodating analog non-idealities (noise, drift, device mismatch discussed in Section 2.3.1) while maintaining composability; developing clear, formal abstractions encoding best practices and enabling cross-disciplinary collaboration between materials scientists, circuit designers, and algorithm developers; implementing hardware-aware algorithms applying these abstractions through white/gray-box optimization methods using gradient information from interpretable models, black-box methods for complex systems, and on-chip training and dynamic fine-tuning mechanisms adapting to hardware imperfections in real time; creating rapid prototyping platforms including large-scale reconfigurable FPAAs[16] with synthesis tools enabling rapid design-test-iterate cycles, chiplet ecosystems with standardized interfaces preserving analog signal integrity while enabling modular composition, hybrid analog-digital integration platforms with clear partitioning abstractions, and testbeds validating minimal primitive sets identified through completeness analysis in PRD 1; establishing cross-stack abstractions and interfaces enabling reusable layer-to-layer interactions, including common ISAs, IRs, and runtime interfaces building on semantic foundations developed in PRD 1; developing calibration methods and error compensation techniques counteracting device mismatch, drift, and environmental variations; building analog-friendly application datasets driving co-development and providing consistent comparison bases; and fostering workforce development creating generalists with cross-stack expertise spanning materials through algorithms.

Research in this area may lead to hardware-aware algorithms demonstrating 10× to 100× efficiency improvements for specific workloads with reconfigurable platforms enabling rapid prototyping cycles (weeks versus years); functional analog HLS tools, chiplet ecosystems with commercially available components, and open-source toolchains enabling broad community access; and a workforce trained in cross-disciplinary co-design methodologies.

8.4.5 Benchmarking and Evaluation Methodologies

Effective engineering and scientific progress require the ability to compare competing designs on well-defined benchmark problems, as demonstrated by SPEC, HPL, and MLPerf for digital systems. These benchmarks focus research efforts, enable fair comparisons, and simplify research by providing standard problems. Analog computing faces unique evaluation challenges: performance depends on device-specific behaviors often insufficiently documented; researchers select idiosyncratic benchmarks complicating cross-work comparisons; simulation approaches vary widely affecting reproducibility; and model/hardware access limitations prevent validation. As highlighted in Section 2.4.3, developing community-driven benchmarks and standardized evaluation methodologies is essential for assessing progress and guiding efforts. Two prerequisites are critical: (1) shared, fair evaluation frameworks for co-designed hardware; and (2) reuse and dissemination of models among industry and academia stakeholders.

Research objectives include developing robust benchmark suites with standardized problems sensitive to important analog dynamics (nonlinearity, noise, variability) that provide meaningful insights rather than "too easy" tests; establishing standardized metrics appropriate to continuous computation (energy efficiency, latency, throughput, accuracy, area); creating reference implementations and simulation infrastructure usable across the analog computing community with consistent configuration approaches; building open-source model libraries with common interfaces for circuits and devices, enabling model reuse and reducing duplicated effort; developing validation procedures assessing model fidelity across design scenarios; establishing cloud-based access to hardware demonstrators for DOE and trusted partners, making hardware availability a funding requirement (similar to artifact evaluation for AI publications); creating labeled datasets of good/bad designs enabling training of automated tools and knowledge transfer; developing frameworks enabling reuse and dissemination of compact models for system/application-level optimization with standardized interfaces; and establishing DOE-supported shared benchmarking infrastructure serving the analog computing research community, providing multi-level benchmark suites spanning devices, circuits, and full systems tailored to scientific computing applications (multi-physics modeling, materials simulation, particle physics, fusion energy), standardized measurement protocols ensuring reproducibility across implementations, community-accessible testbeds enabling researchers to validate approaches without requiring custom hardware access, and reference datasets and problem formulations representing DOE mission applications. This shared infrastructure accelerates scientific discovery by enabling fair comparisons, reducing barriers to entry for new researchers, and focusing community efforts on problems of greatest relevance to DOE's scientific mission rather than fragmented, application-agnostic benchmarks.

Research in this area may lead to a community-adopted benchmark suite comparable to SPEC/MLPerf with representative workloads spanning HPC, edge computing, and scientific applications; open-source simulation infrastructure reducing researcher model development time by 5× to 10×, with cloud-accessible testbeds enabling reproducible research; and policy frameworks for open-sourcing benchmarks and models, accelerating field entry and improving progress assessment through consistent evaluation metrics.

8.5 Key Takeaways

Realizing the transformative potential of analog computing for scientific applications requires a comprehensive co-design ecosystem spanning the full stack from materials and devices through algorithms and applications. This priority research direction addresses five interconnected challenges:

- **Hardware Modeling Frameworks:** Multi-abstraction-level frameworks that balance accuracy and computational efficiency enable system-scale exploration without prohibitive simulation costs, achieving 10–100× speedup versus SPICE while maintaining fidelity.
- **Software Infrastructure:** Programming languages, compilers, and runtime systems adapted to continuous computation models make analog computing accessible to domain scientists without requiring deep hardware expertise.
- **Automated Design Space Exploration:** AI-guided optimization tools accelerate discovery of Pareto-optimal designs across multiple stack layers, reducing designer time by 100× for complex co-optimization problems.
- **Rapid Prototyping Platforms:** Hardware-aware algorithms co-optimized with analog systems and reconfigurable platforms (FPAAs, chiplets) reduce development cycles from years to weeks, analogous to how FPGAs transformed digital design.
- **Standardized Benchmarking:** Community-driven benchmarks and evaluation methodologies enable fair comparison, knowledge sharing, and rigorous progress assessment across diverse analog computing approaches.
- **Cross-PRD Integration:** Co-design methodologies provide the integration fabric enabling all other PRDs to translate advances into deployable systems—accelerating mathematical foundations (PRD 1), high-performance architectures (PRD 2), device primitives (PRD 3), edge computing (PRD 4), and natural computing (PRD 5).

Success requires not just technical advances but cultural shifts: open-source tool and model sharing, cross-disciplinary workforce development, and community-driven standards for evaluation and validation. The methodologies developed here will transfer to other emerging paradigms—quantum, neuromorphic, photonic—facing similar cross-layer optimization challenges.

9 CONCLUSIONS

The six PRDs outlined in this report collectively define a comprehensive research agenda for analog computing that addresses fundamental scientific questions while establishing practical pathways toward deployment for DOE mission applications. This concluding section synthesizes the anticipated scientific and technological impacts across all PRDs, highlighting how coordinated investment in these interconnected research directions can transform computational capabilities for scientific discovery.

9.1 Scientific Impact

Advancing analog computing will yield transformative scientific contributions spanning theoretical foundations, computational capabilities, and cross-disciplinary understanding in important fields such as applied mathematics, computer science, and AI.

At the theoretical level, research in mathematical foundations (PRD 1) will establish analog computation as a rigorous discipline comparable to digital computing theory. Completeness proofs, complexity characterizations, and formal computational models will provide the theoretical grounding needed to guide hardware development and algorithm design across all application domains. The development of analog numerical analysis—a new subdiscipline extending classical numerical analysis to account for analog-specific error modes such as noise, drift, and device variability—will enable reliable algorithm design with predictive performance models. Beyond classical computation theory, research into noise as a computational resource will provide mathematical understanding of when and how stochastic processes can enhance optimization and sampling, transforming what is typically viewed as a limitation into a design advantage.

These theoretical advances will enable previously intractable computational capabilities. High-performance analog systems (PRD 2) will enable long-time simulations for multi-physics models such as fluid dynamics—computations where digital systems struggle with numerical degradation. Analog accelerators for differential equations, optimization, and matrix operations (PRD 3) will achieve order-of-magnitude improvements in energy efficiency for computational kernels central to SC mission space and beyond. Edge computing systems (PRD 4) will enable real-time feature extraction from particle detectors operating at terahertz data rates, and low-latency analog control achieving sub-microsecond feedback for quantum error correction at DOE experimental facilities. Natural computing approaches (PRD 5) will transform power systems through real-time analog control and planning optimization, accelerate dynamics simulations, and enable computational science over longer timescales and larger systems than currently feasible.

This research program will also establish new paradigms for understanding computation itself. The investigation of natural computing substrates—from coupled oscillator networks and chemical reaction systems to biological neural circuits—will bridge physics, information theory, and biology, revealing fundamental principles of how physical systems process information. Research into the thermodynamic limits of sensing and computation will advance understanding of energy-information relationships applicable beyond analog computing. The continuous formulations of discrete problems developed through mathematical foundations research may reveal hidden combinatorial structure in optimization problems, analogous to how the Riemann zeta function illuminates prime number distributions.

9.2 Technological Impact

The technological advances enabled by this research program will directly benefit DOE scientific facilities while establishing foundations for broader adoption of analog computing.

The most immediate technological impact will be deployable analog computing systems for DOE applications. High-performance analog systems (PRD 2) targeting 100,000-element scale will integrate with conventional digital HPC infrastructure at national laboratories, providing validated acceleration for specific computational workloads including AI for science. Device primitives research (PRD 3) will deliver standardized fabrication processes, process design kits, and circuit simulation models that lower barriers to analog system development and enable broader research community participation. Edge computing advances (PRD 4) will enable order-of-magnitude improvements in data handling for scientific facilities including light sources, colliders, and observatories, with robust sensing and computation capabilities for operation in extreme environments where conventional electronics fail.

Critical to realizing these capabilities is the development of comprehensive design infrastructure. Co-design methodologies (PRD 6) will provide automated tools that dramatically reduce barriers to analog system development, enabling smaller research groups to innovate without requiring deep hardware expertise. Reconfigurable platforms analogous to FPGAs for digital design will provide accessible prototyping capabilities that accelerate development cycles from years to weeks. Open-source toolchains, model libraries, and benchmark suites will create a vibrant ecosystem where researchers build upon rather than duplicate each other's work, while cloud-accessible hardware testbeds will ensure broad participation including institutions lacking fabrication resources.

The research program will also establish analog computing as an enabling technology for applications beyond its immediate scope. Natural computing approaches (PRD 5) will provide novel co-processors and alternative computational modalities for heterogeneous computing systems, dramatically reducing energy consumption of HPC facilities to address critical sustainability challenges. Edge computing capabilities will enable wearable and implantable medical devices with onboard intelligence, massive-scale structural health monitoring with self-powered sensors, and distributed sensor networks with years of autonomous operation. The co-design methodologies developed may transfer to other emerging paradigms—quantum, neuromorphic, and photonic computing—that face similar cross-layer optimization challenges.

Finally, sustained investment in this research program will develop the workforce essential for long-term progress. Graduate programs and textbooks in analog computational theory will train a next-generation workforce, while cross-disciplinary training programs will create engineers with end-to-end expertise spanning materials through algorithms. This human capital development addresses critical skills gaps that currently hinder analog computing adoption across industry and national laboratories, ensuring sustained innovation beyond the initial research program.

9.3 Interdependencies and Integration

The six PRDs are deeply interdependent, and realizing their full impact requires coordinated efforts across the portfolio. Mathematical foundations (PRD 1) provide the theoretical grounding that guides hardware development and validates complexity predictions across all application domains. Device primitives (PRD 3) supply the physical building blocks that high-performance systems (PRD 2) and edge computing platforms (PRD 4) require for practical implementation. Natural computing (PRD 5) explores alternative substrates and computational principles that may yield transformative capabilities beyond conventional approaches. Co-design methodologies (PRD 6) provide the integration fabric (the tools, abstractions, and benchmarks) that enables all other PRDs to translate research advances into deployable systems.

Success ultimately depends on establishing a virtuous cycle where theoretical advances guide experimental validation, hardware demonstrations inform theoretical refinements, and co-design tools accelerate this bidirectional feedback. The research directions outlined in this report define the coordinated program needed to sustain this cycle and realize analog computing's potential for DOE scientific missions.

NOTIONAL QUESTIONS FROM THE CALL FOR POSITION PAPERS

Prior to the workshop, the organizing committee identified a set of notional questions to frame the discussions and guide the submission of position papers. These questions were designed to define the vision for the future of analog computing for science and explore the breadth of challenges and opportunities in the field. The accepted position papers addressed various subsets of these questions and are archived on OSTI.gov (`https://doi.org/10.2172/2506701`); the workshop discussions were structured around synthesizing answers and identifying research directions.

The notional questions posed to the community were:

1. **Application domains and metrics:** For which applications does analog computation demonstrate superiority, and for what metrics?
2. **Energy efficiency:** How does analog computation's energy efficiency compare to digital computation in various applications?
3. **Enabling technologies:** What new materials, devices, systems, design software, *etc.*, are needed to enable the future analog computation applications?
4. **Learning from biology:** What can be learned from biology's reliance on a mixture of analog and digital computation and applied to science and engineering problems? Examples of biological computing hardware include regulatory reaction networks and neural tissue.
5. **Mathematical frameworks:** Is there a cross-cutting mathematical framework for analog computation?
6. **Scalability:** How do we scale analog computing systems to solve large-scale problems?
7. **Fundamental limits:** What are the limits of analog computing and how do we approach this limit using practical devices, circuits, and systems?
8. **Programming and software:** How do we program analog computing systems? Do we need new programming models and compilers? What does the software stack look like?
9. **Benchmarking and standards:** What benchmarks and standards are necessary to evaluate and compare the performance of analog computing systems? How can we establish a common framework for assessing the capabilities and limitations of different analog computing approaches?
10. **Hybrid systems and sensing:** How can hybrid systems that combine analog and digital computing be designed to exploit the strengths of both approaches? What are the challenges in developing efficient interfaces between analog and digital components? What opportunities or challenges does analog computing offer for integration with sensing devices?
11. **Heterogeneous systems:** How do we design extremely heterogeneous systems for large-scale and edge systems?
12. **Noise and robustness:** How do we address noise, variability, and robustness issues in analog computation? How do we leverage these non-idealities?
13. **Interdisciplinary collaboration:** What role can interdisciplinary collaboration play in advancing analog computation? How can fields as diverse as physics, materials science, biology, neuroscience, and computer science contribute to the development of analog computing?
14. **Workforce development:** What programs and curricula must be developed to train the scientific workforce in analog computation?

The six PRDs presented in this report represent the community's response to these questions, synthesizing insights from position papers, workshop discussions, and collaborative synthesis sessions. While not every question is addressed explicitly in isolation, the research directions collectively provide a comprehensive roadmap for advancing analog computing to meet the needs of DOE scientific computing applications.

WORKSHOP ORGANIZERS

Name	**Affiliation**	**Role**
Sara Achour	Stanford Univ	Organizer
Natalia Berloff	Cambridge	Organizer
Suma George Cardwell	Sandia	Organizer
Shantanu Chakrabartty	Washington Univ St Louis	Organizer
Hal Finkel	DOE/ASCR	Sponsor
Marco Fornari	DOE/ASCR	Sponsor
Jennifer Hasler	Georgia Tech	Organizer
Siddharth Joshi	Notre Dame	Organizer
Margaret Lentz	DOE/ASCR	Sponsor
Kalyan Perumalla	DOE/ASCR	Sponsor
Robinson Pino	DOE/ASCR	Sponsor
David Rabson	DOE/ASCR	Sponsor
Jaijeet Roychowdhury	UC Berkeley	Organizer
David Soloveichik	UT Austin	Co-Chair
Bill Spotz	DOE/ASCR	Sponsor
Antonino Tumeo	PNNL	Co-Chair

REPORT CONTRIBUTORS

Sara Achour*, Luca Buonanno, Suma Cardwell*, Francesco Caravelli, Shantanu Chakrabartty*, Austin Ellis-Mohr, Saugata Ghose, Jennifer Hasler*, Michael Huang, Connah Johnson, Siddharth Joshi*, Vasileios Kalantzis, Ang Li, Ruipeng Li, Juan Pedro Mendez Granado, Narasinga Rao Miniskar, Shashank Misra, Benjamin Parpillon, David Soloveichik*, Alec Talin, Corinne Teeter, Zoltan Toroczkai, Antonino Tumeo*, Hakan Tureci, Logan Wright, Tianyao Xiao, Zhi Jackie Yao

**Organizing Committee member*

IN-PERSON ATTENDEES

Full Name	Institution
Pamela Abshire	University of Maryland
Sara Achour	Stanford University
Sapan Agarwal	Sandia National Laboratories
Luca Buonanno	Hewlett Packard Enterprise (HPE)
William Cannon	Pacific Northwest National Laboratory (PNNL)
Francesco Caravelli	Los Alamos National Laboratory (LANL)
Suma Cardwell	Sandia National Laboratories
Shantanu Chakrabartty	Washington University in St. Louis
Douglas Durian	University of Pennsylvania
Austin Ellis-Mohr	University of Illinois
Ben Feinberg	Sandia National Laboratories
Hal Finkel	Department of Energy (DOE)
Marco Fornari	Department of Energy (DOE)
Tony (Tong) Geng	University of Rochester
Saugata Ghose	University of Illinois
Luisa Gonzalez	Lawrence Berkeley National Laboratory (LBL)
Peter Graf	National Renewable Energy Laboratory (NREL)
Jennifer Hasler	Georgia Institute of Technology
Aaron Holder	Department of Energy (DOE)
Michael Huang	University of Rochester
James Ignowski	Hewlett Packard Labs
Jean Anne Incorvia	University of Texas at Austin
Sumit Jha	Florida International University (FIU)
Connah Johnson	Pacific Northwest National Laboratory (PNNL)
Siddharth Joshi	University of Notre Dame
Vasileios Kalantzis	IBM
Margaret Lentz	Department of Energy (DOE)
Ang Li	Pacific Northwest National Laboratory (PNNL)
Rui Peng Li	Lawrence Livermore National Laboratory (LLNL)
Yang Liu	Lawrence Berkeley National Laboratory (LBL)
Soumyajit Mandal	Brookhaven National Laboratory (BNL)
Juan Mendez	Sandia National Laboratories
Narasinga Rao Miniskar	Oak Ridge National Laboratory (ORNL)
Marc Miskin	University of Pennsylvania
Shashank Misra	Sandia National Laboratories
Todd Munson	Department of Energy (DOE)
Wilkie Olin-Ammentorp	Argonne National Laboratory (ANL)
Benjamin Parpillon	Fermi National Accelerator Laboratory (Fermilab)
Kalyan Perumalla	Department of Energy (DOE)
Robinson Pino	Department of Energy (DOE)
David Rabson	Department of Energy (DOE)
Yihui Ren	Brookhaven National Laboratory (BNL)
Hava Siegelmann	University of Massachusetts
Andrew Sloss	Vaire
David Soloveichik	University of Texas at Austin
Tathagata Srimani	Carnegie Mellon University
Alec Talin	Sandia National Laboratories
Corinne Teeter	Sandia National Laboratories
Chris Thachuk	University of Washington

Full Name	**Institution**
Luke Theogarajan	University of California, Santa Barbara (UCSB)
Zoltán Toroczkai	University of Notre Dame
Antonino Tumeo	Pacific Northwest National Laboratory (PNNL)
Hakan Tureci	Princeton University
Elisabetta Valiante	1QBit
Marat Valiev	Department of Energy (DOE)
Chris Winstead	Utah State University
Logan Wright	Yale University
Tianyao Xiao	Sandia National Laboratories
Zhi (Jackie) Yao	Lawrence Berkeley National Laboratory (LBL)
Qiaoyan Yu	National Science Foundation (NSF)

GLOSSARY

ADC	analog-to-digital converter.
AI	artificial intelligence.
AOC	analog optical computer.
API	application programming interface.
ASCR	Advanced Scientific Computing Research.
CiM	compute in memory.
CMOS	complementary metal-oxide-semiconductor.
CPU	central processing unit.
CXL	Compute Express Link.
DAC	digital-to-analog converter.
DAE	differential-algebraic equation.
DFT	density functional theory.
DNA	Deoxyribonucleic Acid.
DOE	Department of Energy.
DRC	design rule check.
DSL	domain-specific language.
ECP	Exascale Computing Project.
ECRAM	electrochemical random access memory.
EDA	electronic design automation.
FG	floating-gate.
FPAA	field-programmable analog array.
FPGA	field-programmable gate array.
GAAFET	gate-all-around field-effect transistor.
GPU	graphics processing unit.
HLS	high-level synthesis.
HPC	high-performance computing.
HPL	High-Performance Linpack.
HVM	high-volume manufacturing.
IC	integrated circuit.
IoT	Internet of Things.
IR	intermediate representation.
ISA	instruction set architecture.
LLM	large language model.
LLVM	Low Level Virtual Machine.
LVS	layout versus schematic.
MEMS	micro-electro-mechanical systems.
MHD	magnetohydrodynamics.
ML	machine learning.
MW	megawatt.
NOR	NOT-OR.
NP	nondeterministic polynomial time.

NPU	neural processing unit.
NVM	non-volatile memory.
ODE	ordinary differential equation.
PDE	partial differential equation.
PDK	process design kit.
PRD	priority research direction.
PTAT	proportional to absolute temperature.
R&D	research and development.
RAM	random access memory.
ReRAM	resistive random access memory.
RF	radio frequency.
RFID	radio-frequency identification.
RMS	root mean square.
SAT	Boolean satisfiability.
SC	Office of Science.
SNR	signal-to-noise ratio.
SoC	system-on-chip.
SPEC	Standard Performance Evaluation Corporation.
SPICE	Simulation Program with Integrated Circuit Emphasis.
SWaP	size, weight, and power.
TOPS	tera operations per second.
TPU	tensor processing unit.
VCM	valence change memory.
VLSI	very large scale integration.
VMM	vector-matrix multiplication.

REFERENCES

[1] L. Belostotski, A. Uddin, A. Madanayake, and S. Mandal. A survey of analog computing for domain-specific accelerators. *Electronics*, 14(16):3159, 2025. https://doi.org/10.3390/electronics14163159.

[2] Electronic Associates, Inc. *EAI PACE TR-10 Analog Computer Operators Handbook*. Electronic Associates, Inc., 1960. https://archive.org/details/bitsavers_eaitr10AC6ogComputerOperatorsHandbook_11241073. Publication No. AC6020-1.

[3] T. Fischer and B. Ulmann. *THE ANALOG THING – FIRST STEPS*. anabrid GmbH, 2.01 edition, 2023. https://the-analog-thing.org/THAT_First_Steps.pdf.

[4] S. Köppel, B. Ulmann, L. Heimann, and D. Killat. Using analog computers in today's largest computational challenges. *Advances in Radio Science*, 19:105–116, 2021. https://doi.org/10.5194/ars-19-105-2021.

[5] J. Liang, N. Udayanga, A. Madanayake, S. I. Hariharan, and S. Mandal. An offset-cancelling discrete-time analog computer for solving 1-D wave equations. *IEEE Journal of Solid-State Circuits*, 56 (9):2881–2894, 2021. https://doi.org/10.1109/JSSC.2021.3074003.

[6] P. Zuo, Q. Wang, Y. Luo, R. Xie, S. Wang, Z. Cheng, L. Bao, Z. Wang, Y. Cai, R. Huang, and Z. Sun. Precise and scalable analogue matrix equation solving using resistive random-access memory chips. *Nature Electronics*, 2025. https://doi.org/10.1038/s41928-025-01477-0.

[7] K. P. Kalinin, J. Gladrow, J. Chu, J. H. Clegg, D. Cletheroe, D. J. Kelly, B. Rahmani, G. Brennan, B. Canakci, F. Falck, M. Hansen, J. Kleewein, H. Kremer, G. O'Shea, L. Pickup, S. Rajmohan, A. Rowstron, V. Ruhle, L. Braine, S. Khedekar, N. G. Berloff, C. Gkantsidis, F. Parmigiani, and H. Ballani. Analog optical computer for AI inference and combinatorial optimization. *Nature*, 645: 354–361, 2025. https://doi.org/10.1038/s41586-025-09430-z.

[8] J. Büchel, A. Vasilopoulos, W. A. Simon, I. Boybat, H. Tsai, G. W. Burr, H. Castro, B. Filipiak, M. Le Gallo, A. Rahimi, V. Narayanan, and A. Sebastian. Efficient scaling of large language models with mixture of experts and 3D analog in-memory computing. *Nature Computational Science*, 5(1):13–26, 2025. https://doi.org/10.1038/s43588-024-00753-x.

[9] J. Ang, A. A. Chien, S. D. Hammond, A. Hoisie, I. Karlin, S. Pakin, J. Shalf, and J. S. Vetter. Reimagining codesign for advanced scientific computing: Report for the ASCR workshop on reimagining codesign. *USDOE Office of Science Tech. Rep.*, 2022. https://doi.org/10.2172/1822199.

[10] U.S. Department of Energy Office of Science. New Frontiers: Post-Exascale Computing Research and Development. Oak Ridge National Laboratory, 2024. https://www.olcf.ornl.gov/2024/09/06/newfrontiers/. Project Director: Christopher Zimmer.

[11] K. Yelick. Beyond Exascale Computing. Keynote Address, Salishan Conference on High Speed Computing, 2024.

[12] S. Kumar, X. Wang, J. P. Strachan, Y. Yang, and W. D. Lu. Dynamical memristors for higher-complexity neuromorphic computing. *Nature Reviews Materials*, 7(7):575–591, 2022. https://doi.org/10.1038/s41578-022-00434-z.

[13] K. Berggren, Q. Xia, K. K. Likharev, D. B. Strukov, H. Jiang, T. Mikolajick, D. Querlioz, M. Salinga, J. R. Erickson, S. Pi, F. Xiong, P. Lin, C. Li, Y. Chen, S. Xiong, B. D. Hoskins, M. W. Daniels, A. Madhavan, J. A. Liddle, J. J. McClelland, Y. Yang, J. Rupp, S. S. Nonnenmann, K.-T. Cheng, N. Gong, M. A. Lastras-Montaño, A. A. Talin, A. Salleo, B. J. Shastri, T. F. de Lima, P. Prucnal, A. N. Tait, Y. Shen, H. Meng, C. Roques-Carmes, Z. Cheng, H. Bhaskaran, D. Jariwala, H. Wang, J. M. Shainline, K. Segall, J. J. Yang, K. Roy, S. Datta, and A. Raychowdhury. Roadmap on emerging hardware and technology for machine learning. *Nanotechnology*, 32(1), 2021. https://doi.org/10.1088/1361-6528/aba70f.

[14] C. Mead. Neuromorphic electronic systems. *Proceedings of the IEEE*, 78(10):1629–1636, 1990. https://doi.org/10.1109/5.58356.

[15] R. Chawla, A. Bandyopadhyay, V. Srinivasan, and P. Hasler. A 531 nw/mhz, 128/spl times/32 current-mode programmable analog vector-matrix multiplier with over two decades of linearity. In

Proceedings of the IEEE 2004 Custom Integrated Circuits Conference (IEEE Cat. No. 04CH37571), pages 651–654. IEEE, 2004. https://doi.org/10.1109/CICC.2004.1358910.

[16] J. Hasler. Large-scale field programmable analog arrays. *IEEE Proceedings*, 108(8):1283–1302, 2020. https://doi.org/10.1109/JPROC.2019.2950173.

[17] S. Kim, J. Hasler, and S. George. Integrated floating-gate programming environment for system-level ICs. *IEEE Transactions VLSI*, 24(6):2244–2252, 2016. https://doi.org/10.1109/TVLSI.2015.2504118.

[18] S. Director and R. Rohrer. The Generalized Adjoint Network and Network Sensitivities. *IEEE Trans. Ckt. Theory*, 16:318–323, 1969. https://doi.org/10.1109/TCT.1969.1082965.

[19] C. Robert and G. Casella. *Monte-Carlo Statistical Methods*. Springer, 2000.

[20] C. W. Gardiner. *Stochastic Methods: A Handbook for the Natural and Social Sciences*. Springer Series in Synergetics. Springer, fourth edition, 2009.

[21] D. K. C. MacDonald. *Noise and Fluctuations: An Introduction*. Dover Books on Physics. Dover, 2006.

[22] D. Gillespie. A General Method for Numerically Simulating the Stochastic Time Evolution of Coupled Chemical Reactions. *Comput. Phys.*, 22:403–434, 1976. https://doi.org/10.1016/0021-9991(76)90041-3.

[23] J. Roychowdhury, D. Long, and P. Feldmann. Cyclostationary noise analysis of large RF circuits with multitone excitations. *IEEE Journal of Solid-State Circuits*, 33(3):324–336, 1998. https://doi.org/10.1109/4.661198.

[24] A. Demir, A. Mehrotra, and J. Roychowdhury. Phase Noise in Oscillators: a Unifying Theory and Numerical Methods for Characterization. *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications*, 47(5):655–674, 2000. https://doi.org/10.1109/81.847872.

[25] K. Y. Camsari, R. Faria, B. M. Sutton, and S. Datta. Stochastic p-bits for invertible logic. *Physical Review X*, 7(3):031014, 2017. https://doi.org/10.1103/PhysRevX.7.031014.

[26] K. Y. Camsari, B. M. Sutton, and S. Datta. P-bits for probabilistic spin logic. *Applied Physics Reviews*, 6 (1), 2019. https://doi.org/10.1063/1.5055860.

[27] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. Optimization by simulated annealing. *Science*, 220 (4598):671–680, 1983. https://doi.org/10.1126/science.220.4598.671.

[28] J. D. Smith, A. J. Hill, L. E. Reeder, B. C. Franke, R. B. Lehoucq, O. Parekh, W. Severa, and J. B. Aimone. Neuromorphic scaling advantages for energy-efficient random walk computations. *Nature Electronics*, 5(2):102–112, 2022. https://doi.org/10.1038/s41928-021-00705-7.

[29] J. B. Aimone, W. Severa, and J. D. Smith. Synaptic sampling of neural networks. In *2023 IEEE International Conference on Rebooting Computing (ICRC)*, pages 1–9. IEEE, 2023. https://doi.org/10.1109/ICRC60800.2023.10386770.

[30] L. Buesing, J. Bill, B. Nessler, and W. Maass. Neural dynamics as sampling: a model for stochastic computation in recurrent networks of spiking neurons. *PLoS Computational Biology*, 7(11):e1002211, 2011. https://doi.org/10.1371/journal.pcbi.1002211.

[31] L. Rehm, C. C. M. Capriata, S. Misra, J. D. Smith, M. Pinarbasi, B. G. Malm, and A. D. Kent. Stochastic magnetic actuated random transducer devices based on perpendicular magnetic tunnel junctions. *Physical Review Applied*, 19(2):024035, 2023. https://doi.org/10.1103/PhysRevApplied.19.024035.

[32] S. Misra, L. C. Bland, S. G. Cardwell, J. A. C. Incorvia, C. D. James, A. D. Kent, C. D. Schuman, J. D. Smith, and J. B. Aimone. Probabilistic neural computing with stochastic devices. *Advanced Materials*, 35(37):2204569, 2023. https://doi.org/10.1002/adma.202204569.

[33] A. Maicke, J. Arzate, S. Liu, J. Kwon, J. D. Smith, J. B. Aimone, S. Misra, C. Schuman, S. G. Cardwell, and J. A. C. Incorvia. Magnetic tunnel junction random number generators applied to dynamically tuned probability trees driven by spin orbit torque. *Nanotechnology*, 35(27):275204, 2024. https://doi.org/10.1088/1361-6528/ad3b01.

[34] Z. Chen, Z. Xiao, M. Akl, J. Leugring, O. Olajide, A. Malik, N. Dennler, C. Harper, S. Bose, H. A. Gonzalez, M. Samaali, G. Liu, J. Eshraghian, R. Pignari, G. Urgese, A. G. Andreou, S. Shankar, C. Mayr, G. Cauwenberghs, and S. Chakrabartty. ON-OFF neuromorphic Ising machines using Fowler-Nordheim annealers. *Nature Communications*, 16(1):3086, 2025. https://doi.org/10.1038/s41467-025-58231-5.

[35] T. Wang, L. Wu, P. Nobel, and J. Roychowdhury. Solving combinatorial optimisation problems using oscillator based Ising machines. *Natural Computing*, pages 1–20, 2021. https://doi.org/10.1007/s11047-021-09845-3.

[36] R. Sehgal, T. Thareja, S. S. Xie, C. Ni, and J. P. Kulkarni. A bit-serial, compute-in-SRAM design featuring hybrid-integrating ADCs and input dependent binary scaled precharge eliminating DACs for energy-efficient DNN inference. *IEEE J. Solid-State Circuits*, 58(7):2109–2124, 2023. https://doi.org/10.1109/jssc.2023.3235210.

[37] M. Lanza, G. Molas, and I. Naveh. The gap between academia and industry in resistive switching research. *Nature Electronics*, 2023. https://doi.org/10.1038/s41928-023-00954-8.

[38] J. Teich. Hardware/software codesign: The past, the present, and predicting the future. *Proceedings of the IEEE*, 100 (Special Centennial Issue):1411–1430, 2012. https://doi.org/10.1109/JPROC.2011.2182009.

[39] Semiconductor Research Corporation and Semiconductor Industry Association. Decadal plan for semiconductors. Technical report, Semiconductor Research Corporation, 2021.

[40] W. Wan, R. Kubendran, C. Schaefer, S. B. Eryilmaz, W. Zhang, D. Wu, S. Deiss, P. Raina, H. Qian, B. Gao, S. Joshi, H. Wu, G. Cauwenberghs, and H.-S. P. Wong. A compute-in-memory chip based on resistive random-access memory. *Nature*, 608(7923):504–512, 2022. https://doi.org/10.1038/s41586-022-04992-8.

[41] S. George, S. Kim, S. Shah, J. Hasler, M. Collins, F. Adil, R. Wunderlich, S. Nease, and S. Ramakrishnan. A programmable and configurable mixed-mode FPAA SoC. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 24(6):2253–2261, 2016. https://doi.org/10.1109/TVLSI.2015.2504119.

[42] J. Hasler. The rise of SoC FPAA devices. In *2022 IEEE Custom Integrated Circuits Conference (CICC)*, pages 1–8. IEEE, 2022. https://doi.org/10.1109/CICC53496.2022.9772820.

[43] J. Feldmann, N. Youngblood, M. Karpov, H. Gehring, X. Li, M. Stappers, M. Le Gallo, X. Fu, A. Lukashchuk, A. S. Raja, J. Liu, C. D. Wright, A. Sebastian, T. J. Kippenberg, W. H. P. Pernice, and H. Bhaskaran. Parallel convolutional processing using an integrated photonic tensor core. *Nature*, 589 (7840):52–58, 2021. https://doi.org/10.1038/s41586-020-03070-1.

[44] T. Inagaki, Y. Haribara, K. Igarashi, T. Sonobe, S. Tamate, T. Honjo, A. Marandi, P. McMahon, T. Umeki, K. Enbutsu, O. Tadanaga, H. Takenouchi, K. Aihara, K.-i. Kawarabayashi, K. Inoue, S. Utsunomiya, and H. Takesue. A coherent Ising machine for 2000-node optimization problems. *Science*, 354(6312):603–606, 2016. https://doi.org/10.1126/science.aah4243.

[45] N. Stroev and N. G. Berloff. Analog photonics computing for information processing, inference, and optimization. *Adv. Quantum Technol.*, 6(9):2300055, 2023. https://doi.org/https://doi.org/10.1002/qute.202300055.

[46] J. M. Shainline, S. M. Buckley, R. P. Mirin, and S. W. Nam. Superconducting optoelectronic circuits for neuromorphic computing. *Physical Review Applied*, 7(3):034013, 2017. https://doi.org/10.1103/PhysRevApplied.7.034013.

[47] N. G. Berloff, M. Silva, K. Kalinin, A. Askitopoulos, J. D. Töpfer, P. Cilibrizzi, W. Langbein, and P. G. Lagoudakis. Realizing the classical XY Hamiltonian in polariton simulators. *Nature materials*, 16(11): 1120, 2017. https://doi.org/10.1038/nmat4971.

[48] T. Leleu, F. Khoyratee, T. Levi, R. Hamerly, T. Kohno, and K. Aihara. Scaling advantage of chaotic amplitude control for high-performance combinatorial optimization. *Communications Physics*, 4(1):1–10, 2021. https://doi.org/10.1038/s42005-021-00768-0.

[49] M. Davies, N. Srinivasa, T.-H. Lin, G. Chinya, Y. Cao, S. H. Choday, G. Dimou, P. Joshi, N. Imam, S. Jain, Y. Liao, C.-K. Lin, A. Lines, R. Liu, D. Mathaikutty, S. McCoy, A. Paul, J. Tse, G. Venkataramanan, Y.-H. Weng, A. Wild, Y. Yang, and H. Wang. Loihi: A neuromorphic manycore processor with on-chip learning. *Ieee Micro*, 38(1):82–99, 2018. https://doi.org/10.1109/MM.2018.112130359.

[50] P. A. Merolla, J. V. Arthur, R. Alvarez-Icaza, A. S. Cassidy, J. Sawada, F. Akopyan, B. L. Jackson, N. Imam, C. Guo, Y. Nakamura, B. Brezzo, I. Vo, S. K. Esser, R. Appuswamy, B. Taba, A. Amir, M. D. Flickner, W. P. Risk, R. Manohar, and D. S. Modha. A million spiking-neuron integrated circuit with a

scalable communication network and interface. *Science*, 345(6197):668–673, 2014. https://doi.org/10.1126/science.1254642.

[51] K. M. Cherry and L. Qian. Scaling up molecular pattern recognition with DNA-based winner-take-all neural networks. *Nature*, 559(7714):370–376, 2018. https://doi.org/10.1038/s41586-018-0289-6.

[52] N. Srinivas, J. Parkin, G. Seelig, E. Winfree, and D. Soloveichik. Enzyme-free nucleic acid dynamical systems. *Science*, 358(6369):eaal2052, 2017. https://doi.org/10.1126/science.aal2052.

[53] A. A. Green, J. Kim, D. Ma, P. A. Silver, J. J. Collins, and P. Yin. Complex cellular logic computation using ribocomputing devices. *Nature*, 548(7665):117–121, 2017. https://doi.org/10.1038/nature23271.

[54] M. Collins, J. . Hasler, and S. George. An open-source toolset enabling analog–digital software codesign. *Journal of Low Power Electronics Applications*, 6(1):1–15, 2016. https://doi.org/10.3390/jlpea6010003.

[55] J. Hasler and C. Hao. Programmable analog system benchmarks leading to efficient analog computation synthesis. *ACM Trans. Reconfigurable Technol. Syst.*, 17(1), 2024. https://doi.org/10.1145/3625298.

[56] C. E. Shannon. Mathematical theory of the differential analyzer. *Journal of Mathematics and Physics*, 20 (1-4):337–354, 1941. https://doi.org/10.1002/sapm1941201337.

[57] O. Bournez, D. S. Graça, and A. Pouly. Polynomial time corresponds to solutions of polynomial ordinary differential equations of polynomial length. *Journal of the ACM (JACM)*, 64(6):1–76, 2017. https://doi.org/10.1145/3127496.

[58] J. Hasler and E. Black. Physical computing: Unifying real number computation to enable energy efficient computing. *Journal of Low Power Electronics and Applications*, 11(2), 2021. https://doi.org/10.3390/jlpea11020014.

[59] E. L. Post. Introduction to a general theory of elementary propositions. *American Journal of Mathematics*, 43(3):163–185, 1921. https://doi.org/10.2307/2370324.

[60] J. Hasler. Analog abstraction, computation, and numerical analysis. In *2018 IEEE International Symposium on Circuits and Systems (ISCAS)*, pages 1–5. IEEE, 2018. https://doi.org/10.1109/ISCAS.2018.8351646.

[61] K. Shirriff. A computer built from NOR gates: inside the Apollo guidance computer. https://www.righto.com/2019/09/a-computer-built-from-nor-gates-inside.html, 2019.

[62] D. Soloveichik, G. Seelig, and E. Winfree. DNA as a universal substrate for chemical kinetics. *Proceedings of the National Academy of Sciences*, 107(12):5393–5398, 2010. https://doi.org/10.1073/pnas.0909380107.

[63] J. Hopfield. Neurons with graded responses have collective computational properties like those of two-state neurons. *Proceedings of the National Academy of Sciences*, 81:3088–3092, 1984. https://doi.org/10.1073/pnas.81.10.3088.

[64] M. Ercsey-Ravasz and Z. Toroczkai. Optimization hardness as transient chaos in an analog approach to constraint satisfaction. *Nature Physics*, 7(12):966–970, 2011. https://doi.org/10.1038/nphys2105.

[65] A. Lucas. Ising formulations of many NP problems. *Frontiers in Physics*, 2:5, 2014. https://doi.org/10.3389/fphy.2014.00005.

[66] S. Conte and C. de Boor. *Elementary Numerical Analysis: An Algorithmic Approach*. McGraw Hill, 1980.

[67] J. Hasler. Starting framework for analog numerical analysis for energy-efficient computing. *Journal of Low Power Electronics and Applications*, 7(3), 2017.

[68] T. Murphy. GradIEEEnt half decent. http://tom7.org/grad/murphy2023grad.pdf, 2023.

[69] V. Kalantzis, M. S. Squillante, C. W. Wu, A. Gupta, S. Ubaru, T. Gokmen, and L. Horesh. Solving sparse linear systems via flexible GMRES with in-memory analog preconditioning. In *2023 IEEE High Performance Extreme Computing Conference (HPEC)*, pages 1–7. IEEE, 2023. https://doi.org/10.1109/HPEC58863.2023.10363502.

[70] Y. Saad. *Iterative methods for sparse linear systems*. Society for Industrial and Applied Mathematics, 2003. https://doi.org/10.1137/1.9780898718003.

[71] J. Liesen and Z. Strakos. *Krylov subspace methods: principles and analysis*. Numerical Mathematics and Scie, 2013.

[72] J. Hasler and L. Akers. Implementation of analog neural networks. In *Annual International Conference on Computers and Communications*, pages 32–38, 1991. https://doi.org/10.1109/PCCC.1991.113788.

[73] J. Hasler, C. Diorio, B. A. Minch, and C. A. Mead. Singlet transistor learning synapses. In *Advances in Neural Information Processing Systems 7*, pages 817–824, 1994.

[74] S. Y. Kung. *VLSI Array Processors*. Prentice-Hall, 1987.

[75] J. A. JaJa. *An Introduction to Parallel Algorithms*. Addison-Wesley, 1992.

[76] G. A. Selim. *The Design and Analysis of Parallel Algorithms*. Prentice Hall, 1989.

[77] S. Ramakrishnan, R. Wunderlich, J. Hasler, and S. George. Neuron array with plastic synapses and programmable dendrites. *IEEE Transactions on Biomedical Circuits and Systems*, 7(5):631–642, 2013. https://doi.org/10.1109/TBCAS.2013.2282616.

[78] J. Hasler and P. R. Ayyappan. An analog architecture and algorithm for efficient convolutional neural network image computation. *Journal of Low Power Electronics and Applications*, 15(3), 2025. https://doi.org/10.3390/jlpea15030037.

[79] M. Cohen and A. Andreou. Current-mode subthreshold MOS implementation of the Herault-Jutten autoadaptive network. *IEEE Journal of Solid-State Circuits*, 27(5):714–727, 1992.

[80] E. Vittoz and X. Arreguit. CMOS integration of Herault-Jutten cells for separation of sources. *Analog VLSI Implementation of Neural Systems*, 1989.

[81] K. Rashed, A. Undavalli, S. Chakrabartty, A. Nagulu, and A. Natarajan. 12.3 a scalable and instantaneously wideband 5gs/s RF correlator based on charge thresholding achieving 8-bit ENOB and 152 TOPS/w compute efficiency. In *2024 IEEE International Solid-State Circuits Conference (ISSCC)*, volume 67, pages 222–224, 2024. https://doi.org/10.1109/ISSCC49657.2024.10454315.

[82] S.-Y. Peng, J. Hasler, and D. V. Anderson. An analog programmable multidimensional radial basis function based classifier. *IEEE Transactions on Circuits and Systems I*, 54(10):2148–2158, 2007. https://doi.org/10.1109/TCSI.2007.905642.

[83] M. Cohen and A. Andreou. Analog CMOS integration and experimentation with an autoadaptive independent component analyzer. *IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing*, 42(2):65–77, 1995.

[84] J. Hasler and L. Akers. Circuit implementation of a trainable neural network using the generalized Hebbian algorithm with supervised techniques. In *International Joint Conference on Neural Networks*, pages 160 – 165 vol.1, 1992.

[85] A. Gore and S. Chakrabartty. A min–max optimization framework for designing $\sigma\delta$ learners: Theory and hardware. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 57(3):604–617, 2010. https://doi.org/10.1109/TCSI.2009.2025002.

[86] V. Srinivasan, G. J. Serrano, J. Gray, and P. Hasler. A precision CMOS amplifier using floating-gate transistors for offset cancellation. *IEEE Journal of Solid-State Circuits*, 42(2):280–291, 2007.

[87] V. Srinivasan, G. Serrano, C. Twigg, and Hasler. Floating-gate-based programmable CMOS reference. *IEEE Transactions CAS I*, 55(11):3448 – 3456, 2008. https://doi.org/10.1109/TCSI.2008.925351.

[88] C. Schlottmann and Hasler. A highly dense, low power, programmable analog vector-matrix multiplier: the FPAA implementation. *IEEE Journal on Emerging CAS*, 1(3):403–411, 2011. https://doi.org/10.1109/JETCAS.2011.2165755.

[89] S. Shah and J. Hasler. Tuning of multiple parameters with a BIST system. *Journal of Low Power Electronics Applications*, 64(7):1772–1780, 2017. https://doi.org/10.1109/TCSI.2017.2652123.

[90] M. Gu and S. Chakrabartty. A 100 pJ/bit, (32,8) CMOS Analog Low-Density Parity-Check Decoder Based on Margin Propagation. *IEEE Journal of Solid-State Circuits*, 46(6):1433–1442, 2011. https://doi.org/10.1109/JSSC.2011.2134550.

[91] P. Kumar, A. Nandi, A. Saha, K. S. P. Teja, R. Das, S. Chakrabartty, and C. S. Thakur. ARYABHAT: A digital-like field programmable analog computing array for edge AI. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 71(5):2252–2265, 2024. https://doi.org/10.1109/TCSI.2024.3349776.

[92] A. Ige, L. Yang, H. Yang, J. Hasler, and C. Hao. Analog system high-level synthesis for energy-efficient reconfigurable computing. *Journal of Low Power Electronics and Applications*, 13(4), 2023. https://doi.org/10.3390/jlpea13040058.

[93] J. Hasler. Scalable analog standard cells for mixed-signal processing and computing. In *GOMAC*, 2024.

[94] A. Ige and J. Hasler. ASHES 1.5: Analog computing synthesis for FPAAs and ASICs. In *2025 Design, Automation & Test in Europe Conference (DATE)*, pages 1–6, 2025. https://doi.org/10.23919/DATE64628.2025.10993219.

[95] Kirchhoff Institute for Physics, Heidelberg University. The FACETS Project. https://facets.kip.uni-heidelberg.de/, 2017. Project ended 31 August 2010; page last updated 17 July 2017; accessed 2026-03-18.

[96] J. Hasler and H. B. Marr. Finding a roadmap to achieve large neuromorphic hardware systems. *Frontiers in Neuroscience*, 7, 2013. https://doi.org/10.3389/fnins.2013.00118.

[97] F. Zangeneh-Nejad, D. L. Sounas, A. Alu, and R. Fleury. Analogue computing with metamaterials. *Nature Reviews Materials*, 6(3):207–225, 2021. https://doi.org/10.1038/s41578-020-00243-2.

[98] M. Boniardi, M. Baldo, M. Allegra, and A. Redaelli. Phase change memory: A review on electrical behavior and use in analog in-memory-computing (A-IMC) applications. *Advanced Electronic Materials*, 10(12), 2024. https://doi.org/10.1002/aelm.202400599.

[99] M. R. Zhao, B. Gao, J. S. Tang, H. Qian, and H. Q. Wu. Reliability of analog resistive switching memory for neuromorphic computing. *Applied Physics Reviews*, 7(1), 2020. https://doi.org/10.1063/1.5124915.

[100] A. A. Talin, J. Meyer, J. Li, M. Huang, M. S. Schwacke, H. W. Chung, L. Xu, E. J. Fuller, Y. Li, and B. Yildiz. Electrochemical random-access memory: Progress, perspectives, and opportunities. *Chem. Rev.*, 125(4):1962–2008, 2025. https://doi.org/10.1021/acs.chemrev.4c00512.

[101] S. Kumar, J. P. Strachan, and R. S. T. Williams. Chaotic dynamics in nanoscale NbO_2 Mott memristors for analogue computing. *Nature*, 548(7667):318–321, 2017. https://doi.org/10.1038/nature23307.

[102] A. Ankit, I. E. Hajj, S. R. Chalamalasetti, G. Ndu, M. Foltin, R. S. Williams, P. Faraboschi, W.-m. W. Hwu, J. P. Strachan, K. Roy, and D. S. Milojicic. PUMA: A programmable ultra-efficient memristor-based accelerator for machine learning inference. In *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*, 2019. https://doi.org/10.1145/3297858.3304049.

[103] K. Asifuzzaman, N. R. Miniskar, A. R. Young, F. Liu, and J. S. Vetter. A survey on processing-in-memory techniques: Advances and challenges. *Memories - Materials, Devices, Circuits and Systems*, 4:100022, 2023. https://doi.org/10.1016/j.memori.2022.100022.

[104] B. Murmann. Mixed-signal computing for deep neural network inference. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 29(1):3–13, 2021. https://doi.org/10.1109/tvlsi.2020.3020286.

[105] N. Semenova, X. Porte, L. Andreoli, M. Jacquot, L. Larger, and D. Brunner. Fundamental aspects of noise in analog-hardware neural networks. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 29 (10), 2019. https://doi.org/10.1063/1.5120824.

[106] J. Li, A. Appachar, S. L. Peczonczyk, E. T. Harrison, A. V. Ievlev, R. Hood, D. Shin, S. Yoo, B. Roest, K. Sun, K. Beckmann, O. Popova, T. Chiang, W. S. Wahby, R. B. Jacobs-Godrim, M. J. Marinella, P. Maksymovych, J. T. Heron, N. Cady, W. D. Lu, S. Kumar, A. A. Talin, W. Sun, and Y. Li. Thermodynamic origin of nonvolatility in resistive memory. *Matter*, 7(11):3970–3993, 2024. https://doi.org/10.1016/j.matt.2024.07.018.

[107] A. Chen. Area and thickness scaling of forming voltage of resistive switching memories. *IEEE Electron Device Letters*, 35(1):57–59, 2014. https://doi.org/10.1109/led.2013.2288262.

[108] K. Liu, S. Lee, S. Yang, O. Delaire, and J. Wu. Recent progresses on physics and applications of vanadium dioxide. *Materials Today*, 21(8):875–896, 2018. https://doi.org/10.1016/j.mattod.2018.03.029.

[109] M. Rao, H. Tang, J. Wu, W. Song, M. Zhang, W. Yin, Y. Zhuo, F. Kiani, B. Chen, X. Jiang, H. Liu, H.-Y. Chen, R. Midya, F. Ye, H. Jiang, Z. Wang, M. Wu, M. Hu, H. Wang, Q. Xia, N. Ge, J. Li, and J. J. Yang. Thousands of conductance levels in memristors integrated on CMOS. *Nature*, 615(7954): 823–829, 2023. https://doi.org/10.1038/s41586-023-05759-5.

[110] F. Zhou and Y. Chai. Near-sensor and in-sensor computing. *Nature Electronics*, 3(11):664–671, 2020. https://doi.org/10.1038/s41928-020-00501-9.

[111] A. Passian and N. Imam. Nanosystems, edge computing, and the next generation computing systems. *Sensors*, 19(18):4048, 2019. https://doi.org/10.3390/s19184048.

[112] J. Thayer, Z. Chen, R. Claus, D. Damiani, C. Ford, M. Dubrovin, V. Elmir, W. Kroeger, X. Li, S. Marchesini, V. Mariani, R. Melcchiori, S. Nelson, A. Peck, A. Perazzo, F. Poitevin, C. P. O'Grady, J. Otero, O. Quijano, M. Shankar, M. Uervirojnangkoorn, R. Veraldi, M. Weaver, C. Weninger, S. Yamajala, C. Wang, and C. H. Yoon. Massive scale data analytics at LCLS-II. In *EPJ Web of Conferences*, volume 295, page 13002. EDP Sciences, 2024. https://doi.org/10.1051/epjconf/202429513002.

[113] O. B. Akan, E. Dinc, M. Kuscu, O. Cetinkaya, and B. A. Bilgin. Internet of everything (IoE) - from molecules to the universe. *IEEE Communications Magazine*, 61(10):122–128, 2023. https://doi.org/10.1109/MCOM.001.2200594.

[114] H. Salehi, R. Burgueño, S. Chakrabartty, N. Lajnef, and A. H. Alavi. A comprehensive review of self-powered sensors in civil infrastructure: State-of-the-art and future research trends. *Engineering Structures*, 234:111963, 2021. https://doi.org/https://doi.org/10.1016/j.engstruct.2021.111963.

[115] H. C. Ates, P. Q. Nguyen, L. Gonzalez-Macia, E. Morales-Narváez, F. Güder, J. J. Collins, and C. Dincer. End-to-end design of wearable sensors. *Nat Rev Mater*, 7:887–907, 2022. https://doi.org/10.1038/s41578-022-00460-x.

[116] R. H. Olsson, C. Gordon, and R. Bogoslovov. Zero and near zero power intelligent microsystems. *Journal of Physics: Conference Series*, 1407(1):012042, 2019. https://doi.org/10.1088/1742-6596/1407/1/012042.

[117] V. Mangal and P. R. Kinget. Clockless, Continuous-Time Analog Correlator Using Time-Encoded Signal Processing Demonstrating Asynchronous CDMA for Wake-Up Receivers. *IEEE Journal of Solid-State Circuits*, 55(8):2069–2081, 2020. https://doi.org/10.1109/JSSC.2020.2980526.

[118] S. Jeong, Y. Chen, T. Jang, J. M.-L. Tsai, D. Blaauw, H.-S. Kim, and D. Sylvester. Always-on 12-nw acoustic sensing and object recognition microsystem for unattended ground sensor nodes. *IEEE Journal of Solid-State Circuits*, 53(1):261–274, 2017. https://doi.org/10.1109/JSSC.2017.2728787.

[119] G. Gallego, T. Delbrück, G. Orchard, C. Bartolozzi, B. Taba, A. Censi, S. Leutenegger, A. J. Davison, J. Conradt, K. Daniilidis, and D. Scaramuzza. Event-based vision: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(1):154–180, 2020. https://doi.org/10.1109/TPAMI.2020.3008413.

[120] K. Rashed, A. Undavalli, S. Chakrabartty, A. Nagulu, and A. Natarajan. A Scalable and Instantaneously Wideband RF Correlator Based on Margin Computing. *IEEE Journal of Solid-State Circuits*, 59(11): 3612–3626, 2024. https://doi.org/10.1109/JSSC.2024.3465244.

[121] S. Joshi, C. Kim, S. Ha, Y. M. Chi, and G. Cauwenberghs. 21.7 2pj/MAC 14b 8x8 linear transform mixed-signal spatial filter in 65nm CMOS with 84db interference suppression. In *2017 IEEE International Solid-State Circuits Conference (ISSCC)*, pages 364–365. IEEE, 2017. https://doi.org/10.1109/ISSCC.2017.7870412.

[122] S. Dutta, W. Chakraborty, J. Gomez, K. Ni, S. Joshi, and S. Datta. Energy-efficient edge inference on multi-channel streaming data in 28nm HKMG FeFET technology. In *2019 symposium on VLSI technology*, pages T38–T39. IEEE, 2019. https://doi.org/10.23919/VLSIT.2019.8776525.

[123] S. Chakrabartty and G. Cauwenberghs. Sub-Microwatt Analog VLSI Trainable Pattern Classifier. *IEEE Journal of Solid-State Circuits*, 42(5):1169–1179, 2007. https://doi.org/10.1109/JSSC.2007.894803.

[124] J. Hasler and S. Shah. Security implications for ultra-low power configurable SoC FPAA embedded systems. *Journal of Low Power Electronics and Applications*, 8(2), 2018. https://doi.org/10.3390/jlpea8020017.

[125] B. Murmann. Digitally assisted analog circuits. *IEEE Micro*, 26(2):38–47, 2006. https://doi.org/10.1109/MM.2006.33.

[126] S. Shapero and P. Hasler. Mismatch characterization and calibration for accurate and automated analog design. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 60(3):548–556, 2013. https://doi.org/10.1109/TCSI.2012.2215741.

[127] S. Chakrabartty, R. K. Shaga, and K. Aono. Noise-shaping gradient descent-based online adaptation algorithms for digital calibration of analog circuits. *IEEE Transactions on Neural Networks and Learning Systems*, 24(4):554–565, 2013. https://doi.org/10.1109/TNNLS.2012.2236572.

[128] B. Fleming, I. Shipsey, M. Demarteau, J. Fast, S. Golwala, Y.-K. Kim, A. Seiden, J. Hirschauer, G. Sciolla, O. Palamara, K. Scholberg, J. Cooley, D. McKinsey, C. Chang, B. Flaugher, S. Demers, M. Pepe-Altarelli, F. Lanni, R. Rusack, R. Guenette, J. Monroe, L. Winslow, P. Krizan, A. Geraci, K. Irwin, G. Carini, M. Newcomer, M. Artuso, C. Haber, D. Acosta, T. Bose, M. Begel, M. Narain, D. Dwyer, A. Connolly, A. Sonnenschein, R. Henning, K. Dawson, L. Newburgh, M. Reece, N. Serra, N. Akchurin, S. Eno, P. Rumerio, R. Zhu, J. Raaf, A. Pocar, J. Asaadi, H. Lippincott, G. Giovanetti, A. Lita, F. Sefkow, G. Campbell, A. Sushkov, R. Walsworth, A. Grassellino, A. Dragone, M. Garcia-Sciveres, T. Shaw, J. Thom-Levy, A. Tricoli, P. Merkel, W. Ketchum, J. Zhang, P. O'Connor, G. Karagiorgi, C. Ashton, L. Chaterjee, P. Lee, D. Hornback, H. Marsiske, M. Shinn, K. Byrum, G. Crawford, E. Lessner, D. Nevels, G. Rai, K. Turner, and R. Ruchti. Basic research needs for high energy physics detector R&D. Technical report, USDOE Office of Science (SC), High Energy Physics (HEP), 2019.

[129] A. Potočnik. How to scale the electronic control systems of a quantum computer. *Nature Electronics*, 8: 3–4, 2025. https://doi.org/10.1038/s41928-024-01331-9.

[130] B. Stray, A. Lamb, A. Kaushik, J. Vovrosh, A. Sherlock, J. Winch, F. Hayati, D. Boddice, A. Stabrawa, A. Niggebaum, M. Langlois, Y.-H. Lien, S. Lellouch, S. Roshanmanesh, K. Ridley, G. de Villiers, G. Brown, T. Cross, G. Tuckwell, A. Faramarzi, N. Metje, K. Bongs, and M. Holynski. Quantum sensing for gravity cartography. *Nature*, 602:590–594, 2022. https://doi.org/10.1038/s41586-021-04315-3.

[131] C. D. Panda, M. J. Tao, M. Ceja, J. Khoury, G. M. Tino, and H. Müller. Measuring gravitational attraction with a lattice atom interferometer. *Nature*, 631:515–520, 2024. https://doi.org/10.1038/s41586-024-07561-3.

[132] MAGIS Collaboration. The MAGIS-100 experiment: Exploring the quantum frontier with atom interferometry, 2023. https://magis.fnal.gov/about/physics/. Accessed: 2023-10-31.

[133] H. Lee, W.-T. Kim, J. Kim, M. Chu, and B.-G. Lee. A compressive sensing CMOS image sensor with partition sampling technique. *IEEE Transactions on Industrial Electronics*, 68(9):8874–8884, 2020. https://doi.org/10.1109/TIE.2020.3018053.

[134] A. Fazel, A. Gore, and S. Chakrabartty. Resolution enhancement in $\sigma\delta$ learners for superresolution source separation. *IEEE Transactions on Signal Processing*, 58(3):1193–1204, 2009. https://doi.org/10.1109/TSP.2009.2034909.

[135] H. Anam, S. M. Abbas, I. B. Collings, and S. Mukhopadhyay. A PDMS/MWCNTs RFID flexible tag with advanced resonator design for read range enhancement in IoT monitoring systems. *Scientific Reports*, 15: 9686, 2025. https://doi.org/10.1038/s41598-025-86773-7.

[136] J. Li and M. Zhang. Physics and applications of raman distributed optical fiber sensing. *Light: Science & Applications*, 11:128, 2022. https://doi.org/10.1038/s41377-022-00811-x.

[137] J. Bonefacino, H.-Y. Tam, T. S. Glen, X. Cheng, C.-F. J. Pun, J. Wang, P.-H. Lee, M.-L. V. Tse, and S. T. Boles. Ultra-fast polymer optical fibre bragg grating inscription for medical devices. *Light: Science & Applications*, 7:17161, 2018. https://doi.org/10.1038/lsa.2017.161.

[138] J. He, H. Liu, D. Zhao, J. S. Mehta, C.-W. Qiu, F. Sun, J. Teng, and K. Huang. High-order diffraction for optical superfocusing. *Nature Communications*, 15:7819, 2024. https://doi.org/10.1038/s41467-024-52256-y.

[139] J. Lee. Quantum imaging with ultra-thin metasurfaces. *Light: Science & Applications*, 14:150, 2025. https://doi.org/10.1038/s41377-025-01830-0.

[140] D. Mehta, K. Aono, and S. Chakrabartty. A self-powered analog sensor-data-logging device based on fowler-nordheim dynamical systems. *Nature Communications*, 11(1):5446, 2020. https://doi.org/10.1038/s41467-020-19292-w.

[141] N. Horiuchi. Self-powered sensors. *Nature Photonics*, 12:644, 2018. https://doi.org/10.1038/s41566-018-0290-5.

[142] C. L. Degen, F. Reinhard, and P. Cappellaro. Quantum sensing. *Rev. Mod. Phys.*, 89:035002, 2017. https://doi.org/10.1103/RevModPhys.89.035002.

[143] Fermi National Accelerator Laboratory. Quantum sensing and applications, 2023. https://quantum.fnal.gov/research/quantum-sensing-and-applications/. Accessed: 2023-10-31.

[144] H. C. Berg and E. M. Purcell. Physics of chemoreception. *Biophysical Journal*, 20(2):193–219, 1977. https://doi.org/10.1016/S0006-3495(77)85544-6.

[145] W. Bialek and S. Setayeshgar. Physical limits to biochemical signaling. *Proceedings of the National Academy of Sciences*, 102(29):10040–10045, 2005. https://doi.org/10.1073/pnas.0504321102.

[146] M. Bauer and W. Bialek. Information bottleneck in molecular sensing. *PRX Life*, 1(2):023005, 2023. https://doi.org/10.1103/PRXLife.1.023005.

[147] S. Chakrabartty, B. Raman, and C. S. Thakur. Sensing-to-learn and learning-to-sense: Principles for designing neuromorphic sensors. In N. V. Thakor, editor, *Handbook of Neuroengineering*, pages 1563–1607. Springer Nature Singapore, 2023. https://doi.org/10.1007/978-981-16-5540-1_122.

[148] W. L. Shew, W. P. Clawson, J. Pobst, Y. Karimipanah, N. C. Wright, and R. Wessel. Adaptation to sensory input tunes visual cortex to criticality. *Nature Physics*, 11:659–663, 2015. https://doi.org/10.1038/nphys3370.

[149] J. Carriot, G. McAllister, H. Hooshangnejad, I. Mackrous, K. E. Cullen, and M. J. Chacron. Sensory adaptation mediates efficient and unambiguous encoding of natural stimuli by vestibular thalamocortical pathways. *Nature Communications*, 13:2612, 2022. https://doi.org/10.1038/s41467-022-30348-x.

[150] A. Dodda, A. Oberoi, A. Sebastian, T. H. Choudhury, J. M. Redwing, and S. Das. Stochastic resonance in MoS2 photodetector. *Nature Communications*, 11:4406, 2020. https://doi.org/10.1038/s41467-020-18195-0.

[151] R. Kononchuk, J. Cai, F. Ellis, R. Thevamaran, and T. Kottos. Exceptional-point-based accelerometers with enhanced signal-to-noise ratio. *Nature*, 607:697–702, 2022. https://doi.org/10.1038/s41586-022-04904-w.

[152] T. Honjo, T. Sonobe, K. Inaba, T. Inagaki, T. Ikuta, Y. Yamada, T. Kazama, K. Enbutsu, T. Umeki, R. Kasahara, K.-i. Kawarabayashi, and H. Takesue. 100,000-spin coherent Ising machine. *Science Advances*, 7(40):eabh0952, 2021. https://doi.org/10.1126/sciadv.abh0952.

[153] D. Y. Zhang and G. Seelig. Dynamic DNA nanotechnology using strand-displacement reactions. *Nature Chemistry*, 3(2):103, 2011. https://doi.org/10.1038/nchem.957.

[154] L. Qian, E. Winfree, and J. Bruck. Neural network computation with DNA strand displacement cascades. *Nature*, 475(7356):368–372, 2011. https://doi.org/10.1038/nature10262.

[155] S. K. Aoki, G. Lillacci, A. Gupta, A. Baumschlager, D. Schweingruber, and M. Khammash. A universal biomolecular integral feedback controller for robust perfect adaptation. *Nature*, 570(7762):533–537, 2019. https://doi.org/10.1038/s41586-019-1321-1.

[156] T. S. Gardner, C. R. Cantor, and J. J. Collins. Construction of a genetic toggle switch in Escherichia coli. *Nature*, 403(6767):339–342, 2000. https://doi.org/10.1038/35002131.

[157] S. Li, Q. Jiang, S. Liu, Y. Zhang, Y. Tian, C. Song, J. Wang, Y. Zou, G. J. Anderson, J.-Y. Han, Y. Chang, Y. Liu, C. Zhang, L. Chen, G. Zhou, G. Nie, H. Yan, B. Ding, and Y. Zhao. A DNA nanorobot functions as a cancer therapeutic in response to a molecular trigger in vivo. *Nature Biotechnology*, 36(3): 258–264, 2018. https://doi.org/10.1038/nbt.4071.

[158] Y. Wang, I. Baars, I. Berzina, I. Rocamonde-Lago, B. Shen, Y. Yang, M. Lolaico, J. Waldvogel, I. Smyrlaki, K. Zhu, R. A. Harris, and B. Högberg. A DNA robotic switch with regulated autonomous display of cytotoxic ligand nanopatterns. *Nature Nanotechnology*, 19(9):1366–1374, 2024. https://doi.org/10.1038/s41565-024-01676-4.

[159] D. Y. Zhang, A. J. Turberfield, B. Yurke, and E. Winfree. Engineering entropy-driven reactions and networks catalyzed by DNA. *Science*, 318(5853):1121–1125, 2007. https://doi.org/10.1126/science.1148532.

[160] M. P. Stumpf, T. Thorne, E. De Silva, R. Stewart, H. J. An, M. Lappe, and C. Wiuf. Estimating the size of the human interactome. *Proceedings of the National Academy of Sciences*, 105(19):6959–6964, 2008. https://doi.org/10.1073/pnas.0708078105.

[161] J. K. Medley, J. Teo, S. S. Woo, J. Hellerstein, R. Sarpeshkar, and H. M. Sauro. A compiler for biological networks on silicon chips. *PLoS Computational Biology*, 16(9):e1008063, 2020. https://doi.org/10.1371/journal.pcbi.1008063.

[162] F. Caravelli, F. C. Sheldon, and F. L. Traversa. Global minimization via classical tunneling assisted by

collective force field formation. *Science Advances*, 7(52), 2021. https://doi.org/10.1126/sciadv.abh1542.

[163] F. Caravelli, G. Milano, C. Ricciardi, and Z. Kuncic. Mean field theory of self-organizing memristive connectomes. *Annalen der Physik*, 535(8), 2023. https://doi.org/10.1002/andp.202300090.

[164] V. Baccetti, R. Zhu, Z. Kuncic, and F. Caravelli. Ergodicity, lack thereof, and the performance of reservoir computing with memristive networks. *Nano Express*, 5(1):015021, 2024. https://doi.org/10.1088/2632-959x/ad2999.

[165] J. Yik, K. Van den Berghe, D. den Blanken, Y. Bouhadjar, M. Fabre, P. Hueber, W. Ke, M. A. Khoei, D. Kleyko, N. Pacik-Nelson, A. Pierro, P. Stratmann, P.-S. V. Sun, G. Tang, S. Wang, B. Zhou, S. H. Ahmed, G. V. Joseph, B. Leto, A. Micheli, A. K. Mishra, G. Lenz, T. Sun, Z. Ahmed, M. Akl, B. Anderson, A. G. Andreou, C. Bartolozzi, A. Basu, P. Bogdan, S. Bohte, S. Buckley, G. Cauwenberghs, E. Chicca, F. Corradi, G. de Croon, A. Danielescu, A. Daram, M. Davies, Y. Demirag, J. Eshraghian, T. Fischer, J. Forest, V. Fra, S. Furber, P. M. Furlong, W. Gilpin, A. Gilra, H. A. Gonzalez, G. Indiveri, S. Joshi, V. Karia, L. Khacef, J. C. Knight, L. Kriener, R. Kubendran, D. Kudithipudi, S.-C. Liu, Y.-H. Liu, H. Ma, R. Manohar, J. M. Margarit-Taulé, C. Mayr, K. Michmizos, D. R. Muir, E. Neftci, T. Nowotny, F. Ottati, A. Ozcelikkale, P. Panda, J. Park, M. Payvand, C. Pehle, M. A. Petrovici, C. Posch, A. Renner, Y. Sandamirskaya, C. J. S. Schaefer, A. van Schaik, J. Schemmel, S. Schmidgall, C. Schuman, J.-s. Seo, S. Sheik, S. B. Shrestha, M. Sifalakis, A. Sironi, K. Stewart, M. Stewart, T. C. Stewart, J. Timcheck, N. Tömen, G. Urgese, M. Verhelst, C. M. Vineyard, B. Vogginger, A. Yousefzadeh, F. T. Zohora, C. Frenkel, and V. J. Reddi. The NeuroBench framework for benchmarking neuromorphic computing algorithms and systems. *Nature Communications*, 16(1):1545, 2025. https://doi.org/10.1038/s41467-025-56739-4.

[166] V. J. Reddi, C. Cheng, D. Kanter, P. Mattson, G. Schmuelling, C. Wu, B. Anderson, M. Breughe, M. Charlebois, W. Chou, R. Chukka, C. Coleman, S. Davis, P. Deng, G. Diamos, J. Duke, D. Fick, J. S. Gardner, I. Hubara, S. Idgunji, T. B. Jablin, J. Jiao, T. S. John, P. Kanwar, D. Lee, J. Liao, A. Lokhmotov, F. Massa, P. Meng, P. Micikevicius, C. Osborne, G. Pekhimenko, A. T. R. Rajan, D. Sequeira, A. Sirasao, F. Sun, H. Tang, M. Thomson, F. Wei, E. Wu, L. Xu, K. Yamada, B. Yu, G. Yuan, A. Zhong, P. Zhang, and Y. Zhou. MLPerf inference benchmark. *CoRR*, abs/1911.02549, 2019. http://arxiv.org/abs/1911.02549.

[167] Sandia National Laboratories. CrossSim: Crossbar Simulator. https://cross-sim.sandia.gov/, 2019.

[168] C. Mattiussi and D. Floreano. Analog genetic encoding for the evolution of circuits and networks. *IEEE Transactions on Evolutionary Computation*, 11(5):596–607, 2007. https://doi.org/10.1109/TEVC.2006.886801.

[169] K. Settaluri, A. Haj-Ali, Q. Huang, K. Hakhamaneshi, and B. Nikolic. AutoCkt: Deep reinforcement learning of analog circuit designs. In *2020 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, pages 490–495. IEEE, 2020. https://doi.org/10.23919/DATE48585.2020.9116200.

[170] A. F. Budak, P. Bhansali, B. Liu, N. Sun, D. Z. Pan, and C. V. Kashyap. DNN-Opt: An RL Inspired Optimization for Analog Circuit Sizing Using Deep Neural Networks. In *Proceedings of the 58th ACM/IEEE Design Automation Conference (DAC)*, 2021. https://doi.org/10.1109/DAC18074.2021.9586139.

[171] A. Mirhoseini, A. Goldie, M. Yazgan, J. W. Jiang, E. Songhori, S. Wang, Y.-J. Lee, E. Johnson, O. Pathak, A. Nova, J. Pak, A. Tong, K. Srinivasa, W. Hang, E. Tuncer, Q. V. Le, J. Laudon, R. Ho, R. Carpenter, and J. Dean. A graph placement methodology for fast chip design. *Nature*, 594(7862): 207–212, 2021. https://doi.org/10.1038/s41586-021-03544-w.

[172] R. Liang, A. Agnesina, G. Pradipta, V. A. Chhabria, and H. Ren. CircuitOps: An ML infrastructure enabling generative AI for VLSI circuit optimization. In *2023 IEEE/ACM International Conference on Computer Aided Design (ICCAD)*, pages 1–6. IEEE, 2023. https://doi.org/10.1109/ICCAD57390.2023.10323611.

[173] L. Chen, Y. Chen, Z. Chu, W. Fang, T.-Y. Ho, R. Huang, Y. Huang, S. Khan, M. Li, X. Li, Y. Li, Y. Liang, J. Liu, Y. Liu, Y. Lin, G. Luo, Z. Shi, G. Sun, D. Tsaras, R. Wang, Z. Wang, X. Wei, Z. Xie, Q. Xu, C. Xue, J. Yan, J. Yang, B. Yu, M. Yuan, E. F. Y. Young, X. Zeng, H. Zhang, Z. Zhang, Y. Zhao, H.-L. Zhen, Z. Zheng, B. Zhu, K. Zhu, and S. Zou. The dawn of AI-native EDA: Promises and challenges of large circuit models. *arXiv preprint arXiv:2403.07257*, 2024.

[174] M. Raissi, P. Perdikaris, and G. E. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019. https://doi.org/10.1016/j.jcp.2018.10.045.

[175] K. P. Patel, A. Maicke, J. Arzate, J. Kwon, J. D. Smith, J. B. Aimone, J. A. C. Incorvia, S. G. Cardwell, and C. D. Schuman. AI-guided framework for the design of materials and devices for magnetic-tunnel-junction-based true random number generators. *Communications Engineering*, 4(1):43, 2025. https://doi.org/10.1038/s44172-025-00376-8.

[176] S. G. Cardwell, J. D. Smith, K. Patel, A. Maicke, J. Arzate, S. Liu, J. Kwon, C. R. Allemang, D. C. Crowder, S. Misra, F. S. Chance, C. D. Schuman, J. A. C. Incorvia, and J. B. Aimone. AI-guided codesign for novel computing paradigms. In *Proceedings of the 30th Asia and South Pacific Design Automation Conference*, pages 849–856, 2025. https://doi.org/10.1145/3658617.3703314.

[177] National Institute of Standards and Technology. CHIPS NAPMP NOFO 2: Frequently asked questions. https://www.nist.gov/chips/chips-rd-funding-opportunities/chips-napmp-nofo-2/frequently-asked-questions-chips-napmp-nofo, 2025. Accessed 2025-02-20.

[178] P. O. Mathews, P. R. Ayyappan, A. Ige, S. Bhattacharyya, L. .Yang, and J. Hasler. A 65nm CMOS analog programmable standard cell library for mixed-signal computing. *IEEE Transactions on VLSI*, 2024. https://doi.org/10.1109/TVLSI.2024.3432916.

[179] J. Hasler, P. R. Ayyappan, A. Ige, and P. Mathews. A 130nm CMOS programmable analog standard cell library. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 71(6):2497–2510, 2024. https://doi.org/10.1109/TCSI.2024.3355070.

[180] A. Arora, A. Bhamburkar, A. Borda, T. Anand, R. Sehgal, B. Hanindhito, P.-E. Gaillardon, J. Kulkarni, and L. K. John. CoMeFa: Deploying compute-in-memory on FPGAs for deep learning acceleration. *ACM Transactions on Reconfigurable Technology and Systems*, 16(3):1–34, 2023. https://doi.org/10.1145/3603504.

[181] E. Nurvitadhi, R. D'Souza, and M. Won. Real performance of FPGAs tops GPUs in the race to accelerate AI. *Intel-White Paper, intel.com*, 2021.

[182] Xilinx, Inc. AI Engines and Their Applications. White Paper WP506 (v1.0.2), Xilinx, Inc., 2018. Accessed 2025-3-10.

[183] A. Gajjar, L. Zhao, O. Eldash, A. Natarajan, R. Jean, X. Sheng, G. Pedretti, P. Faraboschi, J. Ignowski, and L. Buonanno. Enhancing FPGAs with analog in-memory computing macros. In *Proceedings of the 2025 ACM/SIGDA International Symposium on Field Programmable Gate Arrays*, pages 179–179, 2025. https://doi.org/10.1145/3706628.3708832.

[184] M. Descour, D. Stracuzzi, J. Tsao, J. Weeks, A. Wakeland, D. Schultz, and W. Smith. AI-enhanced co-design for next-generation microelectronics: Innovating innovation (workshop report). Technical report, Sandia National Lab.(SNL-NM), Albuquerque, NM (United States), 2021.

[185] K. S. NS and P. Li. Pareto optimization of analog circuits using reinforcement learning. *ACM Transactions on Design Automation of Electronic Systems*, 29(2), 2024. https://doi.org/10.1145/3640463.

[186] K. Touloupas and P. P. Sotiriadis. LoCoMOBO: A local constrained multiobjective Bayesian optimization for analog circuit sizing. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 41(9):2780–2793, 2022. https://doi.org/10.1109/TCAD.2021.3121263.

[187] S. Yin, R. Wang, J. Zhang, X. Liu, and Y. Wang. Fast surrogate-assisted constrained multiobjective optimization for analog circuit sizing via self-adaptive incremental learning. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 42(7):2080–2093, 2023. https://doi.org/10.1109/TCAD.2022.3221694.

U.S. DEPARTMENT
of ENERGY

Office of
Science