



NSF Engineering Research
Visioning Alliance

Wave-Based Computing: Key Engineering Bottlenecks

Rapid Report

Wave-Based Computing: Key Engineering Bottlenecks

A rapid report of the Engineering Research Visioning Alliance
Report Finalized March 18, 2026

Based on proceedings from an ERVA event.

This material is based upon work supported by the U.S. National Science Foundation (NSF) under award #2048419. Recommendations expressed in these materials are those of the authors and do not represent the views of the NSF.



Wave-Based Computing: Key Engineering Bottlenecks	1
What is Wave-Based Computing	1
Setting the Stage: Illustrative Use Cases.....	1
AI, Quantum Science and Engineering, and WBC.....	1
Six Key Bottlenecks and Engineering Research Priorities to Address Them	2
01 Nonlinearities	2
02 Storage and Memory (Including Compute-in-Memory)	3
03 Reconfigurability and Programmability	3
04 Incoherent Waves and Modal Diversity.....	4
05 Error Correction, Noise, and Precision	4
06 Practical and Fundamental Limits.....	5

The Engineering Research Visioning Alliance (ERVA) is a neutral convener that helps identify and develop bold and transformative new engineering research directions, directly supporting the nation’s ability to compete in a rapidly changing global economy. Funded by the National Science Foundation (NSF) Directorate for Engineering, ERVA is an engaged partnership that enables an array of voices to impact national engineering research priorities. The initiative convenes, catalyzes, and empowers the engineering community to identify nascent opportunities and priorities for engineering-led innovative, high-impact, cross-domain research that addresses national, global, and societal needs. Learn more at ervacommunity.org.

Copyright 2026 by the Engineering Research Visioning Alliance (ERVA).
All rights reserved.

Wave-Based Computing: Key Engineering Bottlenecks

What is Wave-Based Computing

Wave-Based Computing (WBC) harnesses wave dynamics to process information directly in the physical domain. By operating on continuous analog wavefields, primarily optical and related electromagnetic modalities (microwave, RF, terahertz), and in some cases acoustic or phononic platforms, WBC can exploit native massive parallelism and low-latency propagation. As device concepts, materials, transducers, and chip-scale integration continue to mature, compact WBC units can be embedded into systems to complement or offload parts of conventional electronic pipelines where throughput, energy, or latency are most constrained.

To identify research challenges and current bottlenecks in wave-based computing, 36 experts from academia, industry, and government met on February 24-25, 2026 for a virtual workshop convened by the Engineering Research Visioning Alliance (ERVA), an organization funded by the U.S. National Science Foundation (NSF) Engineering Directorate. The workshop participants considered illustrative use cases and identified engineering research directions that must be addressed to realize the potential of WBC in real-world contexts.

Setting the Stage: Illustrative Use Cases

Prior to the workshop, the event's Thematic Task Force developed several use cases to illustrate how WBC might be highly impactful in the coming years. These use cases, which are summarized below, provided a gateway to imagining potential real-world deployments of WBC.

Imaging and LiDAR pre-processing illustrate WBC's potential at the sensing front end. Direct optical manipulation of image-like fields can realize ultra-low-latency convolutions for machine vision and efficient feature extraction before data ever reaches a digital processor. Because the input may lie anywhere from microwave and terahertz through infrared and visible bands, the same physical approach can, in principle, be tailored across the electromagnetic spectrum to match the application's modality, constraints, and requirements.


Optical data links provide another concrete opportunity. Today's interconnects emphasize detection, digitization, storage, and electronic signal conditioning. Placing computation at the optical receiver, such as a photonic neural network that performs nonlinear equalization and error correction prior to digitization, can reduce conversions and memory traffic, thereby lowering end-to-end latency and energy consumption while improving overall link fidelity.

Finally, optical "smart sensors" integrate photonic computation with phase-sensitive or resonant devices to extract features at the point of sensing. Local analog processing reduces data movement, enables quicker decisions, and can improve energy efficiency in environments where bandwidth is scarce or power budgets are tight.

Consideration of these exemplar use cases by the Thematic Task Force opened the door for discussion of a variety of possible use cases and associated benefits.

AI, Quantum Science and Engineering, and WBC

WBC has great potential to complement and leverage AI and quantum science and engineering. WBC is attractive for signal and image processing. Hence, potential AI applications of WBC include: (i) augmentation of existing electronic AI systems with WBC modules; (ii) use of WBC in perception systems for intelligent machines and robots, and (iii) novel AI models that leverage the physics of WBC systems. Photonic quantum computing provides a WBC pathway for building quantum computers. Conversely, quantum science tools can be used to analyze WBC to understand fundamental limitations, and quantum phenomena can be leveraged to build more powerful WBC systems.



Six Key Bottlenecks and Engineering Research Priorities to Address Them

01 Nonlinearities

Nonlinear transformations are the cornerstone of expressive and trainable wave-based systems, but practical nonlinearities in optical materials are often weak, size- or power-limited, or difficult to cascade. Across the workshop discussions, participants called for nonlinear elements that are giant in magnitude, dynamically tunable, scalable to large systems, programmable in operation, energy-efficient, and ultrafast. A deeper understanding of fundamental bounds and trade-offs among latency, energy, footprint, bandwidth, and noise is required to determine where step-change gains are physically achievable and where hybrid strategies are necessary.

Multiple complementary routes to achieve efficient nonlinearity were identified:

- Using device geometry, nonlinear readout mechanisms, or feedback structures to induce effective nonlinear responses;
- Exploiting engineered nonlocality to concentrate or redistribute fields;
- Leveraging polaritonic platforms to enhance light–matter interaction; and
- Hybrid electronic–photonic schemes that preserve optical speed while improving integrability and cascadability.

Attention also focused on the trade-offs inherent in adding gain. Improvements in cascadability can be offset by noise, power, size, and yield considerations. There was also interest in exploring monolithic opto-electronic-opto (O/E/O) conversion implementations that reduce packaging complexity.

Targeted engineering research efforts may include materials discovery for enhanced light–matter coupling, experimentally validated demonstrations of cascaded nonlinear stages under realistic noise constraints, and co-designed architectures that align nonlinear device physics with trainable, noise-aware models.

Expertise in multiple disciplines is needed to advance this research, including nonlinear materials specialists and device physicists, photonic circuit designers, applied physicists, mixed-signal and control engineers, and teams that couple algorithm design with hardware behavior to enable noise-aware, trainable optical models.

02 Storage and Memory (Including Compute-in-Memory)

A recurring bottleneck in WBC is the limited availability of storage and memory technologies in the same physical modality as the computation. WBC systems could benefit significantly from in-modality storage that reduces data motion between analog wave domains and digital electronics. Key questions include achievable density, read–write complexity, access speed, retention, and energy per operation, as well as how these compare fairly with electronic memory technologies. Photonic compute-in-memory (CIM) is particularly attractive: if practical memory devices and interfaces can be realized on foundry platforms, CIM tiles may deliver system-level throughput and energy advantages by keeping data within the optical path for a greater share of the compute cycle.

Targeted engineering research efforts will require device-level innovation and integration into shared foundry platforms with robust process-design kits, enabling reproducible fabrication and benchmarking. Clear demonstration of computational advantage measured against representative electronic memory hierarchies and workloads will be essential to justify system-level adoption.

Expertise in multiple disciplines is needed to advance this research, including photonic materials scientists and device physicists, computer architects and compute-in-memory system designers, circuit and systems integration engineers, benchmarking and measurement specialists, and foundry engineers supporting process-design-kit development and fabrication infrastructure.

03 Reconfigurability and Programmability

Scaling wave-based systems places a premium on fast, low-loss, and stable reconfiguration of meshes, metasurfaces, and related structures. Without scalable and programmable reconfiguration mechanisms, WBC systems risk remaining task-specific prototypes rather than extensible computing platforms. Even for task-specific systems, some level of reconfigurability is required to compensate for errors, fabrication variations, and other non-idealities. The workshop discussions surfaced goals such as gigahertz-rate phase shifters with wavelength-scale footprints and large programmable circuits that remain tolerant to device variability and optical loss.

Ecosystem gaps are also evident. Mixed-domain electronic design automation (EDA) tools that seamlessly span wave devices and electronic control are immature. No widely adopted compute model or minimal building-block library optimized for photonic hardware currently exists. Calibration at scale can erode energy advantages, and reproducibility is hindered by the absence of modular, pluggable units and robust firmware and compiler layers.

Targeted engineering research efforts may combine advances in reconfigurable materials and actuators with hardware–software co-design, calibration-lite control strategies, standardized modular interfaces, and rigorous test and benchmark practices.

Collaborative expertise across multiple disciplines is needed to advance this research, bringing together photonics and RF designers, control and calibration specialists, CAD/EDA tool developers, algorithms and compiler researchers, system architects, and packaging, reliability, and systems-test engineers.

04 Incoherent Waves and Modal Diversity

Processing by incoherent fields may widen WBC's applicability, simplify control, and lower cost, but design rules and device options remain underspecified. Discussions during the workshop emphasized both the potential of incoherent processing and the opportunity for hybrid coherent-incoherent architectures that exploit distinct advantages in precision, parallelism, and robustness. Key technical questions include:

- Quantifying decoherence tolerance and precision trade-offs;
- Developing brighter, integrated incoherent sources;
- Enabling efficient coupling and routing strategies; and
- Clarifying when integrated versus bulk optical implementations are most appropriate.

Targeted engineering research efforts may address cascading multiple computational layers under incoherent or partially coherent conditions, managing noise accumulation, and defining application niches, such as cost-sensitive sensing or imaging systems, where incoherent processing reduces control overhead without sacrificing required accuracy.

Expertise in multiple disciplines is needed to advance this research, including photonics and statistical optics specialists, optical source and materials scientists, device–system co-design researchers, and engineers focused on integration and packaging.

05 Error Correction, Noise, and Precision

As analog systems scale, noise accumulation, drift, and precision limits can degrade fidelity, particularly when multiple computational layers cascade. Participants emphasized complementary mitigation approaches beyond digitizing every stage:

- Passive filtering and dispersion-engineered error-shaping;
- Physical training strategies for perfect WBC with imperfect components;
- Noise-aware training and architectural choices that exploit shallow, optics-compatible models; and
- Hybrid analog–digital strategies that apply control and limited error checking while retaining optical efficiency.

Targeted engineering research efforts may include developing physically grounded training and calibration frameworks that account for device variability, amplifier noise figures, phase sensitivity, and bit-resolution limits, thereby reducing dependence on idealized digital backpropagation.

Careful accounting of noise, gain, and stability trade-offs over time and temperature will be necessary to rigorously assess end-to-end performance at scale.

Expertise in multiple disciplines is needed to advance this research, including photonics and mixed-signal designers, signal processing and information theory researchers, control and calibration specialists, robust machine learning researchers, and systems testing and validation engineers.

06 Practical and Fundamental Limits

Finally, the field needs fundamental and practical performance limits and bounds, as well as scaling laws for representative tasks. Results along this line include information capacity per unit volume, achievable index tunability, speed and latency limits, cascadability with and without amplification, and the calibration overheads associated with dense wavelength- or spatial-division multiplexing. Fundamental questions of optimal information flow in hybrid wave-digital systems remain insufficiently characterized. It is also important to explore various approaches to physical size reduction in WBC and associated fundamental limits and scaling laws. Because claims of advantage depend on system context, these limits should be assessed alongside hybrid architectures and realistic workloads to identify where wave-based approaches may deliver step-change gains and where they are best positioned as complementary accelerators rather than wholesale replacements.

Targeted engineering research efforts should address developing credible bounds and comparative benchmarks, important for evaluating system-level trade-offs and guiding architecture selection. Increasing efficiency of WBC components through wavelength and mode multiplexing to enable parallel signal transmission and processing should also be explored, as well as using inverse design methodology (or metamaterials) to produce more compact WBC devices.

Expertise in multiple disciplines is needed to advance this research, including applied mathematicians and information theorists, wave physicists and device modelers, systems architects, and researchers developing fair benchmarking methods and performance metrics.

Finally, it should be noted that fruitful and compelling multidisciplinary research directions and efforts are likely to simultaneously impact multiple research challenges across all six bottlenecks. It will be necessary to link materials and device physics to algorithms and architectures; connect systems engineering to mathematical modeling, analysis and optimization; and link fabrication, foundry access, and packaging to benchmarking, modeling, and test.



NSF Engineering Research
Visioning Alliance

Our mission is to identify and develop bold and transformative new engineering research directions and to catalyze the engineering community's pursuit of innovative, high-impact research that benefits society.



**ERVA IS FUNDED BY THE NATIONAL SCIENCE FOUNDATION THROUGH
AWARD NUMBER 2048419**

©2026 Engineering Research Visioning Alliance. All rights reserved.

This material is based upon work supported by the National Science Foundation under Grant # 2048419. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

Ervacommunity.org