

FOR IMMEDIATE RELEASE

Signaloid Joins CERN openlab's Heterogeneous Architectures Testbed to Explore the Future of Scientific Computing

- *British deep-tech company Signaloid joins CERN openlab as an Industrial Member, contributing its probabilistic computing technology to the CERN openlab Heterogeneous Architectures Testbed.*
- *CERN and Signaloid will jointly evaluate whether Signaloid's UxHw® technology can accelerate Monte Carlo event generation for the High-Luminosity Large Hadron Collider.*
- *The collaboration reflects a broader shift towards heterogeneous computing architectures designed to meet the growing demands of next-generation scientific computing.*

Cambridge, UK / Geneva, Switzerland, 30th September 2026 — British computing company Signaloid has joined CERN openlab, the public-private partnership through which CERN evaluates emerging information technologies for scientific computing. As part of the collaboration, CERN and Signaloid will evaluate Signaloid's distribution-extended compute hardware (UxHw®) technology within the CERN openlab Heterogeneous Architectures Testbed, which explores new processor technologies for future scientific computing infrastructures.

The collaboration comes as research organisations worldwide are increasingly adopting heterogeneous computing architectures that combine CPUs, GPUs and specialised accelerators to tackle demanding workloads more efficiently. As the High-Luminosity Large Hadron Collider (HiLumi LHC) approaches, CERN is evaluating a range of emerging computing technologies that could help address its rapidly growing computational requirements.

“Heterogeneous architectures are becoming essential for the HiLumi LHC and CERN openlab is pioneering a model for evaluating next-generation computing technologies such as Signaloid's distribution-extended compute hardware technology. The upcoming deployment of Signaloid's hardware and software stack at CERN openlab illustrates the kind of architectural innovation openlab was created to evaluate.” says Maria Girone, CTO CERN openlab.

Preparing for the High-Luminosity LHC

The Large Hadron Collider (LHC) is the world's largest particle accelerator. By colliding protons at nearly the speed of light, CERN's scientists investigate the fundamental building blocks of matter and seek answers to some of the biggest questions in physics.

Much of this research depends on Monte Carlo simulation, which allows physicists to compare experimental measurements with millions of simulated particle collisions.

Generating these simulations requires repeatedly calculating the same physical processes using different random inputs, making Monte Carlo event generation one of the most computationally-demanding workloads in particle physics.

When the HiLumi LHC begins operation later this decade, the number of recorded collisions will increase dramatically. While this promises unprecedented scientific opportunities, it will also place enormous pressure on CERN's computing infrastructure, with projected computing demand expected to outpace available resources.

UxHw: Beyond the Capabilities of Classical Computers, Available in Production Today

Rather than replacing conventional processors, Signaloid's UxHw® technology is designed to extend heterogeneous computing systems with native support computation directly on digital representations of continuous probability distributions. With this capability, instead of repeatedly executing the same kernel millions of times with different random inputs (so-called sampling), software running on UxHw can perform calculations directly on probability distributions, in a single execution pass, while requiring minimal changes to existing software.

In benchmarking against today's high-performance server platforms, UxHw has demonstrated speed-ups of up to 2,000× for representative workloads ranging from high-energy physics, to regulatory risk modeling for banks, to simulations used in chip design, to robotics. Additional efficiency gains are expected from Signaloid's recently custom ASIC implementations, the first of which taped out in May 2026 using a low-power TSMC fabrication process.

Evaluating a New Computing Architecture

As part of the CERN openlab Heterogeneous Architectures Testbed, CERN and Signaloid will evaluate a representative Monte Carlo event generation workflow based on the Pepper framework for proton-proton collisions producing multiple gluons. The project will assess computational performance, numerical accuracy, and integration effort, helping determine where distribution-extended computing technologies like UxHw can complement existing CPU- and GPU-based scientific computing infrastructure. The joint project is designed to identify where distribution-extended computing adds value within the Monte Carlo event generation pipeline and what practical considerations are involved in integrating the technology into existing high-energy physics software.

"The future of high-performance computing will not be defined by a single processor architecture, but by heterogeneous systems that combine specialised hardware for different classes of computation," said Prof. Phillip Stanley-Marbell, Founder and CEO of Signaloid. *"We're excited that CERN openlab is evaluating UxHw alongside other emerging computing technologies. Experimental particle physics represents one of the most demanding and exciting environments in which to demonstrate its potential."*

Dr. Stefan Roiser, Senior Computing Engineer at CERN, says *“The largest share of LHC computing resources is spent simulating particle collisions. We will explore Signaloid's technology in Monte Carlo event generation, the first step in the simulation chain expected to see substantial cost increases during CERN's upcoming High Luminosity data-taking period. Because event generation relies heavily on multi-dimensional distributions, Signaloid's UxHw technology has strong potential to accelerate this software, helping to meet the forecasted computing budgets during HiLumi LHC, starting in 2030.”*

The collaboration reflects a broader trend across high-performance computing. Around the world, governments and research organisations are investing in heterogeneous computing systems capable of combining multiple processor architectures for increasingly complex scientific and AI workloads, including the UK's planned £750 million AI Research Resource (AIRR) heterogeneous supercomputer.

About Signaloid

Signaloid develops a computing platform that dramatically reduces the runtime and compute infrastructure requirements of workloads ranging from engineering simulations, to robotics and physical AI. Founded by Prof. Phillip Stanley-Marbell, former Professor of Physical Computation at the University of Cambridge, the Signaloid's compute platforms enable software to operate directly on probability distributions through its distribution-extended compute hardware (UxHw®) architecture. Signaloid's technology is available through cloud, on-premises and edge computing platforms and is used by more than 3,000 developers worldwide.

Press contact: press@signaloid.com