


Allen initiative key to LHCb trigger upgrade



Last week, the LHC Experiments Committee formally accepted a proposal for a new first stage of the high-level trigger (HLT) for LHCb. LHCb is one of the four main experiments on the [Large Hadron Collider \(LHC\)](#). It is exploring what happened after the Big Bang that allowed matter to survive and build the Universe we see today.

Like the other experiments on the LHC, LHCb uses a ‘trigger’ system to filter the huge amount of data produced by particle-collision events within its detectors. About 1 in 500 collision events are selected for further analysis. This trigger system is split into two levels: HLT 1, which reduces the data rate from around 40Tbit/s to 1-2 Tbit/s, and HLT 2, which reduces this further to 80 Gbit/s. This is then sent to storage and analysed using the [Worldwide LHC Computing Grid \(WLCG\)](#).

Until now, both HLT 1 and HLT 2 have been carried out using a farm of traditional computer chips called CPUs, which stands for ‘central processing unit’. The new system – set to go into production in 2021 – will see HLT 1 run instead on graphical processing units (GPUs). The highly parallelised structure of

GPUs can make them more efficient than general-purpose CPUs for running algorithms that process large blocks of data in parallel.

Researchers at LHCb have been exploring the potential of GPUs for their trigger systems since around 2013. Building on that foundational work, this new system is the specific result of intense investigations carried out over the last two years, through an initiative called Allen, which is named after the pioneering computer scientist [Frances Elizabeth Allen](#). The three lead developers for the Allen team are, Dorothea vom Bruch, a postdoctoral researcher from the French Laboratory of Nuclear and High-Energy Physics (LPNHE); Daniel Cámpora, a postdoctoral researcher from the University of Maastricht and the Dutch National Institute for Subatomic Physics (Nikhef), who was a PhD student during most of Allen's development, co-supervised between CERN and the University of Sevilla in Spain; and Roel Aaij, a software engineer at Nikhef, who also played a major role in the development and commissioning of LHCb's Run 1 and 2 HLT systems.



The lead developers of the Allen initiative

The Allen team's new system can process 40 Tbit/s, using around 500 NVIDIA Tensor Core GPUs. It matches – from a physics point of view – the reconstruction performance for charged particles achieved on traditional CPUs. It has also been shown that the Allen system will not be limited in terms of memory capacity or bandwidth. Plus, not only can it be used to perform reconstruction, but it can also take decisions about whether to keep or reject collision events.

A diverse range of algorithms has been implemented efficiently on Allen. This demonstrates the potential for GPUs not only to be used as computational accelerators in high-energy physics, but also as

complete and standalone data-processing solutions. Other LHC experiments are also investigating the potential of GPUs; the ALICE experiment already used them in production for their HLT in Run 2.

“We knew that this was an interesting avenue to explore, but we were surprised it worked out so quickly,” says Vladimir Gligorov of LPNHE, who leads LHCb’s Real Time Analysis project. “Over the last two years, the LHCb HLT team made the CPU HLT almost ten times faster, so it could work as planned, which is itself a huge achievement, and then this blue-skies project paid off as well. Now we can have the best of both worlds.”

The Allen initiative has received support through a [CERN openlab](#) project with the Italian company [E4 Computer Engineering](#), which deploys hardware from NVIDIA. [This project](#) provides a testbed for GPU-accelerated applications, with several use cases spread across various LHC experiments.

“Through the CERN openlab project, the team was able to capitalise on E4 Computer Engineering’s expertise and strong links with NVIDIA,” explains Maria Girone, CERN openlab CTO. “This helped ensure the team was supplied with GPUs on which to run tests, and meant there was a good link with the NVIDIA engineers, who provided advice for helping to make the code run as efficiently as possible on the GPUs. This kind of interaction with industry plays an important role in accelerating innovation and helps us to solve the computing challenges posed by the LHC’s ambitious upgrade programme.”

“CERN openlab has played an important role in bringing together various teams across the laboratory and the experiments who are exploring the potential of GPUs,” explains Gligorov. “Seeing that others were exploring this technology too helped give us the confidence to push forward with these investigations. We’re certainly glad we did, as they’ve really paid off.”

This article originally appeared on the [CERN openlab website](#)

/Public Release. View in full [here](#).

:RN, CERN

ig, Engineering, Frances, industry, innovation, LHCb, production, software, technology, trigger, university