

# FURTHER DOWN THE RABBIT HOLE



The European Organisation for Nuclear Research—CERN—is synonymous with the world’s brightest minds, cutting-edge research and groundbreaking discoveries. **Lars Lorenz** interviews **Dr Kevin Vella** (Faculty of ICT) about the University of Malta’s involvement at CERN and its game-changing tech contribution to the ALICE experiment.

The largest particle physics laboratory in the world is in Geneva. CERN operates the Large Hadron Collider (LHC), a huge particle collider located in a circular, 27km long tunnel deep beneath the Franco-Swiss border. Over the years, the University of Malta (UoM) has contributed to the development of software technology for experiments at CERN, most recently ALICE (A Large Ion Collider Experiment).

Humanity has been questioning its origins for centuries, but we have never been as close to answers as we are now. To understand what happened about 13.8 billion years ago, ALICE uses the LHC to accelerate lead ions to extremely high speeds. These

particles then collide at a specific location in the tunnel, momentarily generating a quark-gluon plasma, which is thought to be the state of matter immediately following the big bang. ALICE’s carefully located detectors pick up data about these events, which needs to be stored for eventual analysis by physicists worldwide. Such large high energy physics experiments tend to produce data at immense rates approaching a terabyte per second. It is a tall order to compress and store data at such speed, which is why CERN draws on some of the most powerful computers and storage facilities available.

Until the turn of the 21<sup>st</sup> century, CERN, like most other supercomputer users, relied on specialised high

performance components, which were hard to develop and expensive to maintain. Rapid advances in personal computing technology rendered this practice ineffective, and new supercomputers were designed around standard hardware components instead. Back in 1999 the UoM, through Dr Vella’s work, contributed to the case for the adoption of industry-standard ethernet to connect hundreds of rack-mounted PCs in the data acquisition system for CERN’s nascent ATLAS (A Toroidal LHC Apparatus) experiment. Today, many of the components that power the world’s supercomputers can be bought at a high street computer store. The loss in performance is tolerable when compared to the



savings, and the use of readily-available commodity components not only facilitates maintenance, but also simplifies upgrade planning. The parts are combined into a supercomputing cluster, which uses software to harness thousands of PCs working together as a single, extremely powerful computer.

With terabyte after terabyte pouring out of ALICE's fourteen detectors, there is no way all of it could be stored permanently. Hence the supercomputing facility is responsible for reconstructing events from the raw data streams, filtering out uninteresting

*CERN and its member states have created a giant distributed computing and storage infrastructure, or grid, that spans the globe*

information and compressing the remaining events for long-term storage and analysis, all while keeping up with the flood of data that continues to emanate from ALICE's detectors.


This system needs to sustain a data reduction factor of ten on a total of fourteen data sources which deliver a combined input that will exceed one terabyte (the size of an average hard disk) per second by 2020. Its continued development involves the combined effort of hundreds of software developers, engineers and physicists.

The reduced data stream, arriving at a leisurely pace of tens or hundreds of gigabytes per second, is stored in CERN's data centre for eventual event reconstruction and analysis. Luckily, the advent of the Internet has provided an innovative solution for the long-term storage and analysis of experimental data. CERN and its member states have created a giant distributed computing and storage infrastructure, or grid, that spans the globe. The grid enables geographically dispersed researchers to access and analyse experimental data remotely and at their convenience, while located at their home institution. This facility extends

to the Mediterranean region owing to EUMEDGRID, an initiative funded by the EU under the 6<sup>th</sup> and 7<sup>th</sup> Framework Programmes which saw the UoM lead technical requirements analysis across thirteen Mediterranean countries.

While Malta is not a CERN member state, the UoM was recently granted its first formal associate membership of one of CERN's principal experiments: ALICE. As part of this collaboration, Dr Kevin Vella, Dr Keith Bugeja, and Kevin Napoli (Department of Computer Science, Faculty of ICT) are developing scheduling software for the cluster that will be employed in ALICE's third run, slated for 2020.

A fresh computing infrastructure, dubbed O<sup>2</sup> for Online-Offline computing, is being designed and developed at CERN to handle Run 3 input data rates of 1.1 terabytes per second. It aims for improved operating efficiency by dynamically scheduling 'online' data reduction jobs and 'offline' long-term data analysis jobs simultaneously on the same cluster. This is an ambitious target for what is ultimately a data acquisition system with tight deadlines to meet, but if successful it could influence ➤



Left: Dr Keith Bugeja. Photo by Jean Claude Vancell

Right: One of the supermodules of the Time of Flight under assembly at CERN. Photo by Saba, A. for CERN

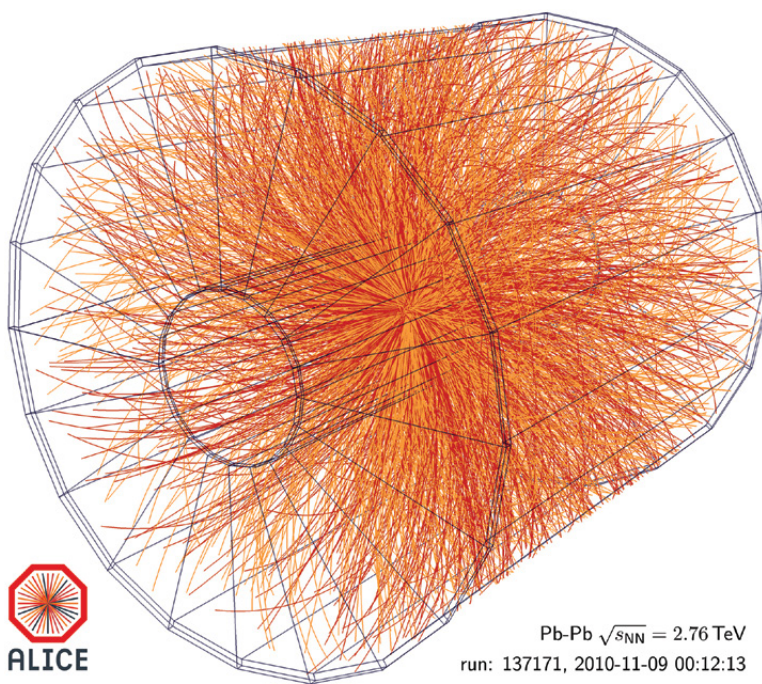


The ALICE Control Centre. Photo by Arora, H. for CERN

the design of future data acquisition systems at a massive scale. The Apache open source software foundation's Mesos cluster scheduler meets several O<sup>2</sup> requirements, and is being evaluated for adoption as the 'operating system' for the new ALICE cluster. During a recent visit to CERN supported by OpenLab, Kevin Napoli, who is completing an M.Sc. in Computer Science at the UoM, delivered the new job submission and control system for the Mesos-powered O<sup>2</sup> candidate.

Latency, the amount of time needed to travel from one location on a network to another, is the bane of our lives even while computing. The transfer of data is necessary when two or more computers in a cluster are working on the same problem simultaneously. Latency increases with network distance and traffic load, and at a supercomputer's scale it can become a hindrance to efficient operation. To mitigate this problem, the UoM team is developing algorithms and tools that automatically and unobtrusively characterise the topology, or shape, of a cluster, and estimate the traffic load across the network. This information is fed to an enhanced Mesos, which uses it to determine where to locate new jobs on the cluster.

Helping to unravel the Universe's



ALICE event display of a Pb-Pb collision at 2.76A TeV. Photo by Weber, Steffen Georg; Andronic, Anton for CERN

best-kept secrets is a fascinating task, but when it comes to computing, time is of the essence. Indeed, technology advances at an incredible pace. Moore's law, an observation which has successfully predicted the rate of technological development for the last fifty years, sees the density at minimum cost per transistor in an integrated circuit double every two years. If

Gordon Moore's foresight were to keep true for another decade, the power of CERN's present-day supercomputers could be had for the price of an average smartphone then, just as today's smartphone outperforms NASA's 1969 Moon landing computer by orders of magnitude. In this broad sense, supercomputers can provide us with a glimpse of the future, today. **T**




## MEET THE TEAM



**KEVIN NAPOLI**  
COMPUTING CERN OPENLAB STUDENT

“My summer experience at CERN was remarkable. The sharing of knowledge among students and researchers was the highlight of the trip. During the openlab programme we attended lectures about security, machine learning, computer hardware, software optimisation and lots more, many of which are topics not covered at our home university. We also played a role in the ALICE experiment and I worked on my project alongside top notch computer scientists. Another positive aspect of the programme were the various trips to companies and universities in Switzerland. Being able to say that I have worked at CERN is something I will value throughout my career.”

Stacks of CDs containing ALICE data (~20 km or 20 million CDs)



A Concorde flies at 15km



# How big is Big Data?

Mont Blanc peaks at 4.8 km

