White Paper

Intel® Xeon® Scalable Processors

Building for Zettascale with Intel and the Cambridge Open Zettascale Lab

Using the latest Intel technology, Cambridge University is building the next generation of exascale and zettascale class computers to democratise HPC

Contributors

Dr. Paul Calleja
Director of Research Computing Services and the Cambridge Open Zettascale Lab, University of Cambridge

Dr. Oz Parchment
Associate Director Research Computing Services and the Cambridge Open Zettascale Lab, University of Cambridge

Robert Maskell
Director of High-Performance Computing, Modelling and Simulation, Intel

Table of Contents

Executive summary

The modern world is increasingly powered and influenced by vast volumes of data. In order to manage, corral and interpolate that data, today's petascale supercomputers will be superseded by new exascale and zettascale systems. Predictions suggest that 463 exabytes of data will be created each day by 2025, and 26 new exascale computing systems will come online over the next few years.

At least one exascale-class machine already exists - the Frontier system at Oak Ridge National Laboratory (ORNL) in the US. It has a peak performance HPL score of 1.102 Exaflop/s, in excess of a quintillion calculations per second, making it the first true "exascale" supercomputer, and the most powerful computer in the world at the time of writing. There are numerous reports of other exascale machines operational in China, with zettascale systems the next key milestone.

Solving world's biggest problems faster than ever

While the numbers may be almost stupefyingly vast, the reach of these new generation computers is correspondingly broad. Future exascale and zettascale systems will enable scientists, engineers and clinicians in almost all walks of life to tackle global-scale challenges. These will include understanding the secrets of the universe, designing new materials, climate modelling, and developing personalised medicine and future healthcare.

Robert Maskell, Director of High-Performance Computing, Modelling and Simulation at Intel said: "Pretty much everything man-made today has, at some point, been modelled inside a computer. Whether you're a Small to Medium Enterprise doing that on a laptop with a few cores, or a large industrial developing new products, processes or services for competitive advantage. Or you're an academic university undertaking research into the early universe... Everyone is turning theories into computational modelling and iterating those systems."

According to Maskell, the advanced modelling applications that are possible with these systems include 'digital twins'. As he explains: "A major UK newspaper publisher, for example, has a major concern in keeping their printing presses running. They could build a replica, a digital twin of the printing press, and could measure and monitor the performance of that press and its constituent components and compare that to the real-world press. This allows the publisher

The Cambridge Open Zettascale Lab is a joint project between Intel, Dell and the University of Cambridge.

to model and predict failures and ensure that they have the respective replacement parts on the shelf for repairs, thus maintaining uptime."

Here's another example, says Maskell. "The H. J. Heinz Company, famous for its baked beans, uses computation to model the thickness of the actual can. How much of the tin material do they need to create the walls of the can so that when the cans are in boxes they do not deform when those boxes are stacked on pallets, and those pallets are stacked upon each other." Similarly, if you've got a bottle of bleach that falls off the top of a shelf in a supermarket, how thick do the walls of the plastic bottle and cap need to be before it bursts, based on the viscosity of the liquid inside? As you can see, HPC technical computing, modelling and simulation is absolutely pervasive."

While there are innumerable opportunities for advanced modelling, simulation and AI to make a difference in the world of enterprise and consumer product life cycles, there are even more astonishing applications for the power of exascale and zettascale computing.

Dr. Paul Calleja is Director of Research Computing Services and the Cambridge Open Zettascale Lab at the University of Cambridge. As he explains: "Clean energy and nuclear fusion are really good examples of the application of exascale and zettascale computing. A lot of countries are engaged in nuclear fusion research for clean energy." Next-generation compute power is vital in this area to help to study the properties and dynamics of fusion plasma. According to the Argonne Leadership Computing Facility "engineers working with the potential energy source have estimated a window of only 30 milliseconds to control instabilities that can disrupt the energy production process."[1]

The University of Cambridge is also working on a mammoth radio astronomy project called The Square Kilometre Array. This is an international effort to build the world's largest radio telescope, featuring hundreds of dish antennas. Exascale compute power isn't just needed to handle the huge data volumes produced by the telescope, but also to produce images of a fidelity that will satisfy researchers.

The Zettascale Lab - future proofing HPC

Although exascale systems are already possible, and will become increasingly practical, there are considerable barriers to entry. The Cambridge Open Zettascale Lab (COZL) aims to tackle them head on.

"The whole idea behind the Cambridge Open Zettascale Lab," explains Intel's Robert Maskell, "is that we [the UK] can go and spend hundreds of millions of pounds putting vast machines on the ground, but unless we enable the wider ecosystem, build the software and train people with the skill sets and the tools, they're not going to be able to use this machine. It's just going to be sitting there. So, we brought together Intel, Dell and Cambridge - and created COZL. It's a two year co-funded programme, and the idea is to prepare the exa- and zettascale ecosystem as a whole - not just compute power, but also storage, Fabric Interconnects, and software tools."

The Zettascale Lab is based within Research Computing Services at the University of Cambridge, one of the world's oldest universities. The Research Computing Services

[1] https://www.alcf.anl.gov/news/preparing-exascale-eliminating-disruptions-path-sustainable-fusion-energy

unit itself already operates one of the most powerful supercomputers in Europe, powered by Intel and managed by an innovative OpenStack Software Environment.

“Exascale and zettascale computing hold the promise to solve a lot of ‘grand challenge’ problems,” says Dr. Calleja. “The computation provided by such systems is a big step forward compared to most systems today - but those systems have a number of different pain points in their usage. COZL was set up to define those pain points, and try to address them within a design constraint of using - as much as possible - commoditised open standard solutions and open source software. We’re really trying to democratise exascale and zettascale technologies, and then enable those technologies to trickle down to smaller systems – that’s the underlying driver.”

The full sweep of challenges that an aspiring exascale/ zettascale computer operator faces are truly formidable, from staffing and electricity capacity to funding and hardware maintenance. However, the Cambridge Open Zettascale Lab has refined this broad slew into five core areas that need to be addressed.

1. The software environment (Intel oneAPI)

One of the biggest challenges on the road to zettascale will be porting existing application code to machines that incorporate heterogeneous hardware. To maximise the computing power of these machines, software needs to become universal. This not only means scaling across thousands, and ultimately tens of thousands of processors, but across different co mpute architectures.

Intel® oneAPI is a solution to this problem. It’s an open, unified and cross-architecture programming model that can be deployed across CPUs and accelerators (GPUs, FPGAs, and others). The programming model simplifies software development and delivers uncompromising performance for accelerated compute without proprietary lock-in, while enabling the integration of legacy code.

The Intel® oneAPI Rendering Toolkit is a set of open source libraries that enables the creation of high-performance, high-fidelity and cost-effective visualisation applications and solutions. The toolkit supports Intel CPUs and future Xe architectures (GPUs). It includes the award-winning Intel® Embree, Intel® Open Image Denoise, Intel® Open Volume Kernel Library, Intel® OSPRay, Intel® OpenSWR and other components and utilities.

Aurora – HPC and AI at exascale

At the US Department of Energy (DOE) Argonne National Laboratory, engineers are building a supercomputer with a performance of two exaFLOPs – two “quintillion” floating point computations per second. Known as Aurora, this next-generation system will be capable of dramatically advancing scientific research and discovery across projects such as cancer research, climate modelling and cosmological simulation. Aurora uses new Intel technologies designed specifically for the exascale era, including Intel® Xeon® Scalable processors, Intel’s Xe compute architecture, as well as the Intel oneAPI software.

Find out more here.

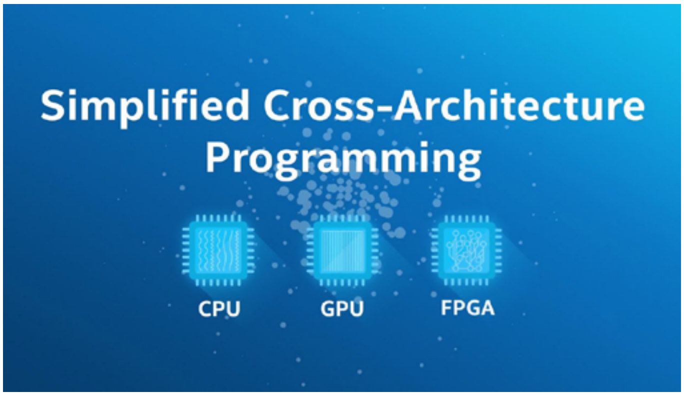

The Cambridge Open Zettascale Lab is a 'centre of excellence' for courses and workshops related to oneAPI, offering courses to train programmers in oneAPI. In addition, the team at Cambridge is porting significant candidate codes to oneAPI, including CASTEP, FEniCS and AREPO.

2. Storage for a zettascale world

When managing huge amounts of data and processing large datasets, how to store that data efficiently is a key concern. Not only does an exascale/zettascale storage solution need to be as fast as possible, but it must also be secure and flexible enough to deal with varying structures and volumes of data. For this problem, the Cambridge Open Zettascale Lab is taking a new approach, incorporating state-of-the-art storage devices such as NVMe disks.

The Lab also plans to focus on solid-state technologies, with initial work emphasising Intel's open-source Distributed Asynchronous Object Storage (DAOS) object store. DAOS has been designed from the ground up to exploit the performance of high-performance hardware platforms to deliver high bandwidth, low latency, and high I/O operations per second (IOPS) storage for HPC applications.

One example is the Cambridge Data Accelerator. It's currently the UK's fastest high-performance computing storage solution, claiming first place in the IO-500 worldwide high-performance computing storage performance rankings in June 2019. It's arguably the shape of storage to come.

3. Connecting the dots with Openstack Middleware

Middleware is another pain point on the road to zettascale computing. "Most traditional supercomputing technologies are based around proprietary solutions," says Dr. Paul Calleja. "If you look at large systems in the US, for example, they are from Cray, the proprietary supercomputing company. They might be using commodity processors and GPUs, but they have a proprietary interconnect. This traditional system architecture is not based on open standards, and the software that they run on is not open - there's always been that tension."

The solution, at least as far as the Zettascale Lab is concerned, is to use Scientific OpenStack, which is capable of managing large-scale high-performance computing systems. A key benefit of Scientific OpenStack is that it supports on-site, cloud-based and hybrid projects, allowing scientists and researchers to deploy clusters on-demand to support specific workflows and use cases. This flexibility will be crucial to securing the future broad-base use cases for exascale and zettascale computers. It will enable enterprise users to easily access the processing power using familiar cloud technologies, as well as tap into existing data standards, and align transparently with established compliance and legal requirements.

As Dr. Paul Calleja explains: "There's a whole range of disciplines within clinical medicine and public health in terms of personalised medicine and genomics that can benefit from exascale and zettascale systems. Of course, these [types of application] need to be delivered within a secure, dynamic, and agile environment, which is where OpenStack becomes very important."

4. Global ambition – Networking at scale

There's a level deeper than the middleware challenge. The next problem that The Cambridge Open Zettascale Lab intends to tackle is ensuring that future networks are capable of delivering data at the scale and speed required by future HPC. To this end, Dr. Calleja and his team plan to investigate multi-vendor, low-latency Ethernet, as well as conducting research into which performance characteristics need to be developed for different types of applications.

The crux is to explore software-defined Ethernet technologies, such as P4. This is a data-plane programming language for network devices that could theoretically allow networks to be tuned to the precise needs of the specific applications using them.

5. Visualisation and qualitative understanding

"These huge supercomputers are just data generation machines," says Dr. Oz Parchment, Associate Director Research Computing Services and the Cambridge Open Zettascale Lab, University of Cambridge. "And the larger the machine the more data they generate. It is incredibly difficult to visualise the results of a simulation. If all you get at the end is 500 terabytes of ones and zeros, it is impossible to analyse your data.

"That's where the power of data visualisation becomes important. The more data you generate, the more necessary

Exascale and zettascale systems will help to model real-world events in as close to real-time as possible.

it becomes to visualise the results in a meaningful way. It's not just about creating a pretty picture - it's giving you an insight that you cannot get from the raw data that has been generated in your simulations."

Exascale-class systems will generate huge data sets that are impossible to understand without further processing, a fact that is already clear from existing 'big data' environments. Learnings from these will evolve new and powerful hardware and software solutions, analysis tools, techniques and algorithms to deliver high-quality visualisations of extreme-scale data sets. The COZL team is already designing solid-state storage test beds that will be combined with large-scale rendering compute engines and HPC graphical software to render exascale and zettascale data sets in real time.

Eventually, The Cambridge Open Zettascale Lab aims to host a remote-access implementation of the resulting suite of visualisation tools. This will enable scientists, engineers and clinicians to easily create interactive, three-dimensional visualisations of remote data.

Urgent computing and energy efficiency

One of the most exciting opportunities in the exascale/zettascale era - and arguably one of the most challenging – is the concept of urgent computing. This is where vast computational power is brought to bear on modelling a real-world event, in as close to real-time as possible.

Examples range from small-scale scenarios, such as modelling the smoke flow in the event of a London Underground fire to allow more efficient evacuation, through to computing the spread of a global pandemic. Other examples include modelling extreme weather events, from hurricanes to floods, droughts to volcanic activity.

"Urgent computing is an area where you need real time modelling to help you understand what's going to happen, or help you to deal with a current scenario," says Dr. Paul Calleja. "The problem is that you need a computer system that can link into multiple governmental emergency service datasets quickly and securely. A traditional old fashioned supercomputer, running 30 year-old middleware designed for physicists and astronomers, is not going to be able to do that for you.

"You need power, but you also need the ability to carve off private spaces that feel comfortable [to stakeholders], and therefore can have sensitive data imported to them in a flexible and agile way. This is why exascale /zettascale and cloud computing are a good mix, you can very quickly stand up ad hoc platforms on demand for any particular activity."

Energy efficiency is also an ongoing challenge, and one that affects all areas of high-performance computing. For example, the former world number one supercomputer, the Fugaku petascale system in Japan, requires a shade under 30,000 kW to operate.

As Dr. Calleja explains: "Every 18 months or so the next generation [of chips] comes out, and that generation is normally twice the performance of the previous one for the same power. It's probably the fastest changing area of computing and it's because of this that you have a problem with everything else.

" That's why we're seeing an ever increasing delta between networking performance, storage performance and raw compute performance, because companies like Intel are just so good at evolving that part of the world. The other areas can't keep up."

Top of the FLOPS

1 FLOP

The speed at which a human brain can solve a basic addition problem.

1 teraFLOP

The teraFLOP barrier was broken by the Intel ASCI Red computer in 1997. ASCI Red's peak performance was 1,340,000,000,000 FLOPS, or 1.34 teraFLOPS[2].

1 petaFLOP

The supercomputer known as Roadrunner topped the TOP500 list and exceeded a petaFLOP of performance in 2008

1 exaFLOP

The Folding@home network was the first to hit the exaFLOP milestone and can draw upon 2.4 exaFLOPS (2,400,000,000,000,000,000 FLOPS) of computing power[3].

1 zettaFLOP

A massive 1,000,000,000,000,000,000,000 (one septillion) FLOPS – a thousand times the performance of an exascale supercomputer. Intel has set an aggressive goal to reach zettascale performance in the data center over the next four to five years.

[2] https://www.energy.gov/science/doe-explainsexascale-computing
[3] https://www.cio.com/article/191667/crowdsourced-supercomputing-zeroes-in-on-covid-19.html

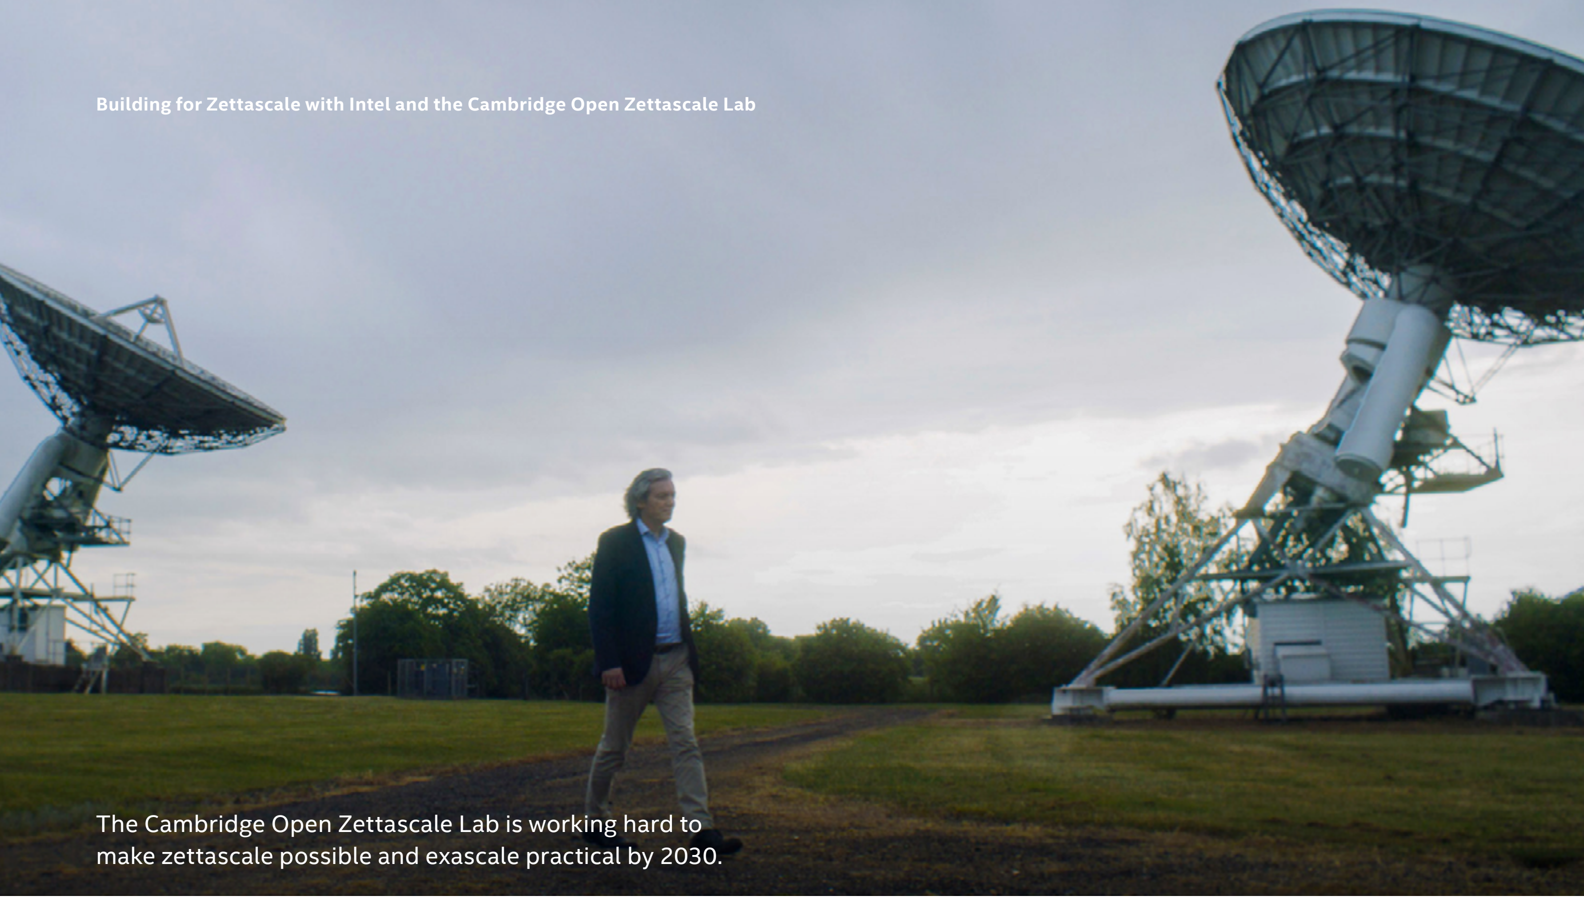
The Cambridge Open Zettascale Lab is working hard to make zettascale possible and exascale practical by 2030.

The democratisation of supercomputing

The overarching aim of The Cambridge Open Zettascale Lab is to achieve the democratisation of supercomputing, trickling new technologies like oneAPI, DAOS and low-latency Ethernet down to petaflop machines. However, even given a complete open source template in the five key areas we've mentioned, any new entrant will need deep pockets to build a true exascale or zettascale machine.

The real key to democratisation in the world of exascale, and eventually zettascale computing, is the use of cloud-style technologies. "The software environment that runs traditional HPC systems or 'supercomputers' is actually awfully old fashioned," says Dr. Calleja. "It's the same environment that's been around for 30 years of traditional large scale parallel supercomputing, and we've moved on.

"Now it's all about the cloud model, with software defined platforms delivered by cloud native APIs. Most traditional supercomputers are just not aligned to that. Whereas one purpose we have using OpenStack is to deploy exascale and zettascale systems within a cloud native environment, which allows a lot more flexibility. It also enables this convergence of traditional simulation workflows, modern AI workloads and data analytic workloads within a single platform."

In the future, we will see exponential increases in the volume of data we need to manage, and the amount of data we create on a daily basis. Scaling this up, the need for more powerful exascale and zettascale computers is inescapable. The benefits of making those systems open source ensures that they will be available to the widest possible selection of scientists, academics, researchers, and businesses.

The mission of the Cambridge Open Zettascale Lab is to "explore, test and advance the next generation of high-performance computers." Working with Intel, and incorporating the latest technologies, exascale platforms will enable us to get a little closer to solving some of the world's biggest problems.

Learn More

You may find the following resources useful:

- Cambridge Open Zettascale Lab
- Intel oneAPI Academic Center of Excellence in the UK
- Intel oneAPI toolkits
- Intel DAOS architecture

Solution provided by

intel®