

# SCIENTIFIC COMPUTING WORLD

Computing solutions for  
scientists and engineers

Summer 2022  
Issue #180

**High performance computing**  
Facilitating access to HPC


**Laboratory informatics**  
AI drug development

**Modelling and simulation**  
Cern satellite success



**DYNAMIC  
DESIGNS**  
MULTIBODY  
SIMULATION  
TAKES OFF

Register by October 14 for Big Discounts!



# You. Dallas. SC22.

**NOVEMBER 13-18**

One of the world's largest gatherings of HPC researchers and professionals, SC22 is an incredible week of learning new skills, forging new friendships, solving complex problems, and viewing next-gen technology.


A new, custom-built Digital Experience is available if you prefer to attend virtually.



The International Conference for High Performance  
Computing, Networking, Storage, and Analysis

## Register Today!

[sc22.supercomputing.org](https://sc22.supercomputing.org)



# SC22

Dallas, TX | hpc  
accelerates.

Sponsored by:



IEEE  
COMPUTER  
SOCIETY



sighpc

# Contents

Summer 2022 • Issue 180

## High performance computing

- A quantum leap** **4**  
As quantum technology matures, several different hardware technologies are showing promise, finds **Robert Roe**
- Facilitating access to HPC** **7**  
HPC integrators and cloud providers play an integral role in facilitating access to computing architectures for scientific research and discovery
- Tech Focus: Cloud** **10**  
A roundup of cloud technologies available to researchers
- Tech Focus: Software tools** **12**  
A round-up of software tools available for HPC users
- Climate in the clouds** **14**  
Climate research is becoming increasingly important and driving new investment in HPC systems and support from cloud computing



## Laboratory informatics

- Drug discovery** **16**  
AI is driving new methods in the search for new drugs and medicines, reports **Robert Roe**
- Genomics** **20**  
Graph databases and open source software are enabling researchers to derive greater insights from genomics data, finds **Sophias Ktori**
- Accelerating drug pipelines** **24**  
The rise of AI and quantum technology are driving new methods of analysis for new drugs

## Modelling and simulation

- Multibody simulation** **26**  
Multibody Dynamics simulations are now a vital tool for today's engineers to understand and improve the performance of a vast range of dynamic systems, reports **Gemma Church**

## Resources

- News** **30**  
A roundup of the latest news stories in scientific computing
- Suppliers' directory** **34**

### Editorial and administrative team

**Editor:** Robert Roe  
robert.roe@europascience.com  
**Managing editor:** Annabel Oia  
editor.scw@europascience.com  
**Head of content:** Mark Elliott  
mark.elliott@europascience.com  
**Specialist reporters:** Sophia Ktori, Gemma Church  
**Design:** David Houghton, Justin Zwierzanski

### Advertising team

**Advertising manager:** Lexi Taylor  
lexi.taylor@europascience.com +44 (0) 1223 221039  
**Production manager:** David Houghton  
david.houghton@europascience.com +44 (0) 1223 221034

### Corporate team

**Managing director:** Warren Clark

Scientific Computing World is published by Europa Science Ltd,  
4 Signet Court, Cambridge, CB5 8LA • ISSN 1744-8026  
Tel: +44 (0) 1223 211170 • Fax: +44 (0) 1223 213385  
Web: www.researchinformation.info

**Subscriptions:** Free registration is available to qualifying individuals. Register online at [www.scientific-computing.com](http://www.scientific-computing.com). Subscriptions £180 a year for six issues to readers outside registration requirements. Single issue £20. Orders to ESL, SCW Circulation, 4 Signet Court, Swann Road, Cambridge, CB5 8LA, UK. Tel: +44 (0) 1223 221 030. Fax: +44 (0) 1223 213 385. ©2022 Europa Science Ltd.

While every care has been taken in the compilation of this magazine, errors or omissions are not the responsibility of the publishers or of the editorial staff. Opinions expressed are not necessarily those of the publishers or editorial staff. All rights reserved. Unless specifically stated, goods or services mentioned are not formally endorsed by Europa Science Ltd, which does not guarantee or endorse or accept any liability for any goods and/or services featured in this publication.

US copies: Scientific Computing World (ISSN 1356-7853/USPS No 018-753) is published bi-monthly for £180 per year by Europa Science Ltd, and distributed in the USA by DSW, 75 Aberdeen Rd, Emigsville PA 17318-0437. Periodicals postage paid at Emigsville PA. Postmaster: Send address corrections to Scientific Computing World PO Box 437, Emigsville, PA 17318-0437.

**Cover:** Timothy Dry/Shutterstock.com

**Subscribe for free at**  
[www.scientific-computing.com/subscribe](http://www.scientific-computing.com/subscribe)



**Robert  
Roe**  
Editor

This issue of *Scientific Computing World* has a strong focus on the emerging technologies and applications that are driving scientific discovery and the tools that help facilitate access to computing frameworks. The rise in emerging computing paradigms, such as AI and quantum computing, will shape the future of how scientists carry out research. But accessing these technologies and putting them to best use is a challenge. Cloud computing is one way researchers can gain access, however, it is just as important for the development of new technologies to have multiple viable options available to users, so they can make use of these resources in a way that suits them or their organisations.

Emerging technologies and applications will continue to rise, but adoption relies on access and ease of use, not just delivering the highest performance.

On page 4, the HPC coverage begins with a feature looking at the development of quantum computing and its integration with classical computing systems – including HPC. Starting on page 8, there is a feature exploring the role of integrators and cloud computing providers in facilitating access to HPC AI.

On pages 10 and 12, respectively, we have this issue's two tech-focus articles, which detail cloud and software tools. Rounding out the HPC section of this issue we have an article that looks at weather and climate simulation and the role that cloud computing plays in helping researchers to gain access to advanced computing architectures.

Coverage of the laboratory informatics market begins on page 16, with a feature looking at the use of AI in modern drug discovery workflows. On page 20, Sophia Ktori takes a look at the development of genomics software and how this is accelerating scientists' ability to analyse massive data sets.

The modelling and simulation feature on page 24 explores how multibody simulation is used across many engineering disciplines and helps engineers to address challenges in the design and development process.

 [@scwmagazine](https://twitter.com/scwmagazine)



# LABWARE®

Automate Your Laboratory  
with the Global Leader for  
**LIMS** and **ELN**

[www.labware.com](http://www.labware.com)



**LABWARE®**  
*Results Count*

# Subscribe for free\*

**SCIENTIFIC  
COMPUTING  
WORLD**

## Do you compute?


The only global publication for scientists and engineers using computing and software in their daily work



Do you subscribe? Register for free now!  
[scientific-computing.com/subscribe](https://scientific-computing.com/subscribe)

\*Registration required

Also published  
by Europa  
Science



# A quantum leap



Production Perig/Shutterstock

**QUANTUM TECHNOLOGY IS GOING THROUGH A PERIOD OF RAPID DEVELOPMENT, WITH SEVERAL TECHNOLOGIES DRIVING THE ADOPTION OF THIS EMERGING COMPUTING FRAMEWORK, FINDS ROBERT ROE**

One of the biggest stumbling blocks in the development of quantum hardware is with the qubits themselves. This varies depending on the underlying technology used to create the qubits, but they are often error-prone and difficult to control, making quantum computers unstable and highly complex systems.

Advancing the technology requires larger quantum computers that can be scaled up and integrated with the cloud or existing classical computing systems. Scale is, therefore, of paramount importance in delivering real-world scientific insight.

There are many ways to build these systems, depending on the type of technology used. Universal Quantum, for example, is trying to develop the world's first million qubit quantum computer using a technology called 'trapped ion'.

Dr Luuk Earl, quantum engineer at Universal Quantum, highlighted the path the company has taken since its inception in 2018. 'It's a company that spun out from a research group at the University of Sussex. Two senior scientists, Professor Winfried Hensinger and Dr Sebastian Weidt, decided that the research they were doing was promising for quantum computing. They formed this company with a bit of venture capital funding.

'The real aim is to make quantum computers that can solve real-world interesting problems,' Earl continued. 'That is the main difference between other quantum computing companies and us at the moment. We're just focusing on that big-scale stuff. We're not interested in making toy models that can do some interesting science but aren't going to impact humanity. We're focused on that point where quantum computing is useful to everyone, which is a big challenge.'

And that is why a million qubit system is such an important goal, as this is the point at which some scientists and researchers believe that quantum computing systems will start to impact science and engineering. Earl noted that scientists at Universal Quantum had done some modelling of the resources required to solve particular problems.

'One of them is synthesising a particular chemical for fertiliser. They've done some simulations of how many resources you



need to do the quantum chemistry for that to make an efficient process. And that kind of comes out around a million qubits.

'Of course, there's always more to the story,' stressed Earl. 'We trade-off three things in a real-world application. One of them is the number of qubits, which is, of course, always important. There are also the error rates and coherence time. So if you have really good error rates, you can probably do fewer qubits, if you have long coherence times you can do with a few less qubits. So it's a bit of a rough number. But that's the kind of order of magnitude where quantum computers become interesting.'

### Trapped ion quantum systems

Universal Quantum is developing its quantum computers based on trapped ion technology. A simple explanation would be that ions are trapped and precisely controlled using electromagnetic fields. Each ion levitates above the surface of a silicon microchip. The idea behind trapped ion systems is the ions are relatively easy to control, as they are all precisely the same shape and size.

'A trapped ion system is what you'd initially think of if you think of a quantum computer,' said Earl. 'The qubits you use are naturally quantum systems; an atom is the closest thing we can come to a single quantum bit. Whereas a superconducting qubit is kind of an analogy to a qubit.'

'The big benefit is that every qubit is identical because every atom is identical

to every other atom, the energy levels are very well defined, and especially with trapped ions, we can control the position and environment of those atoms very precisely,' Earl continued. 'By tuning electrode voltages, you can move the qubits around the surface of a chip very precisely; this means we have ultimate control over what these qubits are doing and how they're interacting.'

Another important consideration that goes into developing these systems is the availability and price of the components used to control the qubits. For example, in the Universal Quantum system, the company uses lasers and microwaves to control the qubits.

Earl noted it is important to focus on developing systems with technology that is available today. 'From the work – we've done at Sussex and the preliminary work – we've done at Universal Quantum already, we think we can build million qubit machines using technology that already

**"We're all very passionate about making real-world impactful systems as soon as possible. The quickest way to do that is using technology that already exists"**

exists. We're all very passionate about making real-world impactful systems as soon as possible. The quickest way to do that is using technology that already exists.'

### Quantum accelerators

Founded in 2019, Quantum Brilliance has a very different take on the development of quantum computing. The venture-backed company develops quantum computers using a diamond substrate to help boost the reliability of the qubits and increase coherence time. The goal of Quantum Brilliance is to enable mass deployment of quantum technology to propel industries to harness edge and supercomputing applications.

The first generation of the company's technology has already been installed in Passey Supercomputing Centre, which is exploring how this technology might be used alongside high performance computing (HPC) systems in the future.

Mark Mattingley-Scott, managing director, EMEA for Quantum Brilliance, explains why the company opted for this radically different quantum technology. 'What diamond does is give you coherence for free. So you get qubits if you make qubits in diamond the right way. They maintain quantum coherence, even at room temperature. What it means is all the stuff you have to do with other quantum computing technologies, like keep it cold, or ensure it's under a really high vacuum, or use precise lasers to get photons aligned, all those things fall away.'

This is an important distinction from other quantum systems, as it means the Quantum Brilliance prototype systems can be smaller and more easily integrated with existing computing systems. They operate at room temperature and do not complex systems or advanced cooling.

'I was actually at the Pawsey Supercomputing Centre this afternoon and I saw our quantum computer,' Mattingley-Scott said. 'It's a 6u device, so it's a little bit higher than the standard 19-inch rack unit. That's the first generation of machines. Inside this device is a small piece of diamond with the qubits on it. And some simple optoelectronics to interact with that – stuff you would find in a 5g antenna mast. We're working on miniaturising that, and we're pretty sure we can get down to a graphics-card-size accelerator within the next few years.'

'Once you've got something like a graphics accelerator, then you're in the same world as a normal Graphics Processing Unit (GPU) or Tensor Processing Unit (TPU),' Mattingley-Scott added. 'You can start to put these things



mapman/shutterstock

→ in large quantities in a standard compute environment.' Another benefit to diamond-based quantum systems is they can be used for edge computing systems, such as in robotics and autonomous vehicles.

This is relevant to high performance computing HPC and research centres in general because it allows quantum computing to be more easily integrated with classical computing architectures. This could help drive the adoption of the technology and allow more scientists and researchers to get access to quantum technology.

But scaling these systems to the point of mass adoption is still some way off. There are significant challenges facing today's quantum computing developers. One significant challenge is getting to the point where you've got enough qubits to provide a measurable performance improvement in some way. The second challenge is how to integrate a quantum computer with classical computing? While Quantum Brilliance wants to connect Quantum Processing Units (QPUs) to classical systems directly, some other organisations want to connect these systems via the cloud.

However, Mattingley-Scott thinks this is a mistake. 'Most companies are using the

**"If you talk to almost all the hardware vendors, there will not be isolated quantum computers, at least for the foreseeable future – the next few decades"**

cloud, so they're looking at a cloud hybrid execution model. We believe – and I think history bears this out – the QPU needs to be physically as close to the other classical compute devices, like CPUs and GPUs, as possible.

'We envisage a future in which you'll go into your computing centre and pull a blade out – maybe it'll be a CPU blade, maybe it'll be a GPU blade, or maybe it'll be a hybrid, and it'll hopefully be a Quantum Brilliance QPU sat next to AMD or Nvidia or Intel CPUs and GPUs.

'Quantum computing must operate in a quantum-classical hybrid – it has to be the case,' Mattingley-Scott continued. 'If you talk to almost all the hardware vendors, there will not be isolated quantum computers churning away doing stuff, and then delivering their results, at least for the foreseeable future – the next few decades. It is all going to be hybrid.'

'In which case, bite the bullet and put your QPU actually in an accelerator card. Next to the GPU, next to the CPU. And then you're not worried about data throughput, latency times and interaction times,' Mattingley-Scott concluded.

## Quantum in the cloud

Quantinuum, on the other hand, is a company that has embraced the use of the cloud to help facilitate access to its prototype quantum systems. Quantinuum's H1 generation of quantum computers is already commercially available. The Quantinuum H1 generation, currently consisting of two computers, the H1-1 and the H1-2, are fully accessible over the cloud and compatible with a variety of software frameworks.

Tony Uttley, president and chief operating officer at Quantinuum, highlights the company's growth from both hardware and software provider. 'Quantinuum is the combination of Cambridge Quantum with Honeywell Quantum Solutions. Honeywell Quantum Solutions did a lot of work directly with the products Cambridge Quantum Computing was making.

'What we found as we were working together as separate companies, was that most people who are developing hardware will extrapolate away from the metal layer,' Uttley explained. 'They will make a separation to protect IP, and you can't get the full integrated benefit if you have that separation layer. We realised we could make fully integrated solutions based upon both the application layer on top of the middleware on top of our hardware.'

However, although the platform is based on the integration of two distinct companies, they also choose to make the software platform inclusive. 'The applications, the operating system that we develop, is designed to work on everybody's hardware,' Uttley said. 'And as a real practical example, we are one of the biggest users of IBM's quantum computers in the world. IBM is also an investor in Quantinuum.'

## Making use of quantum

While the hardware stack continues to mature, scientists and researchers are now getting access to software development tools to create quantum algorithms and quantum simulators, or emulators, that allow them to simulate how a quantum computer might work in a classical system. This allows researchers to start to develop expertise and test out how applications might benefit them in the future.

'A lot of the algorithmic work is in imagining this future where you don't have to worry about qubits and how they

interact. Because all of that has been "taken care of" by universal fault tolerance,' said Uttley. 'That's a decade away. So the key is, what do you do in the intervening time? How do you make progress? Can you do things with some of these earlier systems? And the answer is, yes you can.'

However, this requires users to begin to think about their problems differently,' stresses Uttley. 'What I mean by that is, don't think about what the problem is, and how you abstract that into the system. You need to think about what these systems are good at. And how do I use that for these kinds of problems?'

## Not all qubits are created equally

One critical aspect of the development of these quantum systems, particularly in the early days where coherence and error rates make these systems relatively unstable, is that different hardware architectures are more suited to different problems. A simple example of this would be a large number of qubits with a low coherence versus a much smaller number of qubits with a high coherence time. Another factor is the amount of communication required between qubits.

'If you're trying to simulate a molecule, then ostensibly it depends on how that molecule is shaped, believe it or not,' Uttley said. 'This is because what happens in a superconducting quantum computer and semiconducting ones are similar, where they have an architectural property that's called "nearest neighbour". This means the qubits are physically manufactured on a piece of silicon.'

Uttley gave an example of a grid or several rows of qubits where the qubits can easily communicate with their nearest neighbour, but where communications from one side of the grid to the other take much longer and adds additional error to the system.

'There are molecules where the shape of the molecule itself is kind of a nearest neighbour interaction,' Uttley continued. 'A nearest neighbour molecule running on a nearest neighbour quantum computer actually can work pretty effectively. But if you have complex molecules, where you need these qubits to talk to every other qubit arbitrarily, then our trapped ion hardware works well. This is because we can physically transport our qubits so that any one can talk to any other one without introducing any additional error.'

'It's that kind of deep knowledge about both the problem and the way these systems work that allows us to know which hardware or platform will be most suitable for a given problem,' Uttley concluded. ■



# Maximising the value of HPC

**INTEGRATORS AND CLOUD PROVIDERS HELP FACILITATE ACCESS TO HPC AND DELIVER ADDITIONAL EXPERTISE AND SUPPORT, WHICH HELPS SCIENTISTS TO EFFECTIVELY USE COMPUTING RESOURCES, FINDS ROBERT ROE**

HPC integrators can help scientists and researchers to access and manage their high performance computing (HPC) resources to get the most out of their computing infrastructure. Increasingly, this includes hybrid cloud environments and managed services, which also help to support the upkeep and usability of the systems.

But for the service provider to determine the right level of support for each organisation requires knowledge of the existing stack, the use application portfolio and the level of expertise of the user community

Andy Dean, sales director at OCF, commented: 'It depends on how mature their usage of HPC is within that specific customer. So if somebody has been using HPC for a number of years, this could be the second, third, fourth iteration of the HPC environment; they tend to have a very good understanding of the workload they've got.'

Dean described how mature users can fall into two further categories. Users with a specific small set of applications want to maximise the performance of those key applications. 'In that scenario, it becomes a conversation about how do we achieve their performance requirements – if they have requirements for how quickly they want to run a model, for example,' said Dean.

Alternatively, there could be wide and varied workloads where the customer wants to try and provide a balanced

infrastructure that meets a wide number of requirements, but potentially is less specialised than a system focused on just a handful of applications. 'Or maybe they've got a very varied workload,' Dean added. 'Then we might look at how we will build an environment to try and keep all of the users happy. In that scenario, you want to deliver a balanced environment that works for the largest number of users.'

When the customer or its user community are new to HPC, then engagement focuses on understanding the application portfolio and how the users intend to make use of the HPC resources. 'Other potential customers are very new to HPC, or really looking to expand their environment and aren't necessarily in a position to understand precisely what all of their users are after and what they require,' noted Dean. 'For

**"Other potential customers are very new to HPC, or really looking to expand their environment and aren't necessarily in a position to understand precisely what all of their users are after"**

those kinds of customers, it can be more like an interview-type process that's more of a consultancy engagement, rather than a technical pre-sales engagement.'

Dean states that it is important to understand exactly what the users currently have access to and how



Gorodenkoff/shutterstock



→ the resources are being consumed. 'We need to build a picture of the organisation that we can build into a report that can be used to help shape what their next system might look like. It depends on how well the customer feels they understand the requirements and then we kind of go down one of the two approaches.'

Integrators no longer just provide hardware and support services but increasingly deliver fully managed services. In the future, this may also extend to application support and optimisation. Cloud providers deliver similar hardware and managed service agreements, but it remains to be seen whether they can provide the same level of expertise in HPC-specific applications, hardware and software frameworks.

'Our business is moving further up the stack,' said Dean. 'We're initially involved in projects around deploying HPC. Now when we are deploying a system, we typically use our own OCF steel stack that's based on a number of open source technologies. But in addition to the integration side of things, and support of that, we're also getting a lot more involved in managed services, and helping to manage those environments.'

'I can see, as time goes on, that we are being asked more specific questions about applications, end-user management and things like that. I can imagine that's a direction where things are going, that we're being asked to do more.'

Dean also noted that while this is possible in some centres with a more monolithic application portfolio, many HPC centres have a large set of applications

that prohibits optimisation of each one individually by their integrator partner.

'Some users have hundreds of applications, so there's not much point in really spending a lot of time optimising for each application; you're trying to build something that works for everyone. It's more on the commercial side of things; we find that users – maybe they're running

**“When it comes to data centres and HPC, the Nordics are relatively cold, and so we take that benefit and make use of it for HPC and AI”**

two or three engineering applications, let's say – and we work closely with them to make sure we're picking the right hardware initially to make sure we're getting the best out of that application. There can be other work to help with the workflow side of things as well,' Dean concluded.

### The role of cloud computing

At a time when energy prices are soaring and sustainability is becoming an increasingly complex problem for many organisations, data centre provider atNorth is supporting scientists and researchers with its HPC, GPU and AI data centres, which are based on energy-efficient hardware and renewable energy sources with additional heat recovery, improving the cost of delivering highly complex computing services.

The Nordic data centre company says it

offers environmentally responsible, power-efficient, cost-optimised data centre hosting facilities and high-performance computing services. It describes its HPC resources as sustainable, highly scalable and fully delivered as a service, enabling scientists and researchers to focus on their applications without having to worry about the underlying HPC infrastructure. atNorth also recently announced the availability of its new GPU-as-a-service (GPUaaS) solution. This new service is aimed at scientists who want to accelerate deep learning, machine learning and HPC workloads that are suitable for large-scale use of GPUs.

The company has data centres based in Iceland and Sweden that are specifically designed and optimised for HPC and AI computing. According to atNorth, these resources are delivered through managed services that can be tailored to meet customer requirements, based on their level of experience and the type of service level agreements (SLA) they require. These managed services can be scaled up or down as necessary to provide general capacity for everyday operations and the ability to cloudburst or quickly scale up operations as needed.

Guy D'Hauwers, sales director — HPC and AI, atNorth, commented: 'The speed at which technology innovation is moving is often incalculable, and much of this is due to digitalisation and the rise of extreme data-hungry applications to fuel the transformation. Today's data-driven businesses are reinventing the way in which they work and recognise they need a new type of partner that can help them achieve next-generation computing power,





with great connectivity and infrastructure built on high precision and sustainability.

'Our GPUaaS solution not only multiplies the HPC and AI capacity, delivering energy and cost-efficient service, but it also operates as a full tech stack solution that does the legwork for our clients, so their data scientists, engineers, developers and researchers can focus on their increasingly important day job – from building solutions and services to gathering insights.

D'Hauwers added: 'Many companies are adjusting their business models to secure stakeholders' trust and safeguard long-term profitability. Digitalisation has had a massive impact on our climate, yielding an ever-increasing demand for electricity and rising carbon emissions because of its acceleration. Many businesses rely on technology and data to drive value to their customers, whilst also recognising this can come at a cost to the environment. Therefore, many businesses are migrating their data centre footprint to atNorth's site in Iceland. Businesses must be exploring the best possible ways to walk the talk by adopting best practices when it comes to reducing the digital footprint of their IT and operations.'

GPUs are designed for high-density workloads, such as advanced calculations for AI, natural language processing, scientific simulations and risk analysis. The nature of these applications, in addition to rising costs associated with using public cloud services and increased pressure on sustainability, has tasked many organisations with the challenge of finding new alternatives to ensure continuity with high-performance

applications in a cost-effective and energy-efficient way. atNorth's new GPUaaS will offer a much larger capacity, according to D'Hauwers. Its sites already deliver a total capacity equivalent to 125,000 A100 GPUs, with plans afoot to double this in the next 12 – 18 months.

But to sustainably grow the capacity for its HPC and AI systems, atNorth has had to be very careful about how it designs its data centres, taking advantage of the climate in the Nordics and available renewable energy sources to deliver highly efficient, HPC and AI infrastructure.

'We rely on the fantastic weather of the Nordics,' stated D'Hauwers. 'When it comes to data centres and HPC, the Nordics are relatively cold, and so we

**"We constantly take whatever actions are possible to make it as sustainable as possible"**

take that benefit and make use of it for HPC and AI. Our users benefit from this infrastructure and deliver faster science and simulation projects. And as a result, they are capable of improving the time to market and doing this in a very sustainable way, driving towards carbon neutrality where possible.

'And, because we do everything from the ground up from the data centres, not just from the HPC side, but the entire integration of it as a fully managed service to the user, this allows the user to focus on their own business, their core business,' added D'Hauwers.

'We use 100 per cent renewable energy,' stated D'Hauwers. 'In Iceland, this comes primarily from geothermal energy. But, importantly, the goal is to use much less of it. Because "the energy you don't use is the most sustainable".'

Rising energy prices across Europe are making cloud an even more attractive proposition than running an in-house data centre, as users can effectively outsource their IT requirements to countries with lower energy costs. D'Hauwers added: 'So they bring [their project] to us, and they see the delta of so much less energy used, and then, on top of that, the energy is 100 per cent renewable.

'With the current evolution of energy prices in mainland Europe, the cost of energy has been driven up substantially, with gas prices increasing and so on. The pricing [is going] through the roof'.

But D'Hauwers was keen to stress it was not just renewable energy sources but atNorth's focus on efficient systems that can reduce overall energy consumption, which is key to the company's vision. 'The whole stack we deliver, the data centres are efficiently built for HPC and AI, [as well as] the way the cooling is done and the recovery of heat.

'For example, what we do in Stockholm is recover the heat and sell it to the municipality to heat tens of thousands of houses. We constantly take whatever actions are possible to make it as sustainable as possible.

'Of course, these AI systems and big supercomputers use a lot of energy, but if you can reduce that to the minimum, there is a huge benefit for research organisations,' said D'Hauwers. ■



# Tech Focus: Cloud

## A ROUNDUP OF CLOUD TECHNOLOGY PROVIDERS THAT SUPPORT RESEARCHERS USING HPC

Cloud computing provides huge potential to scientists and researchers, who can use the technology to access computing resources or new and emerging technology.

Cloud also facilitates collaboration, helps organisations to scale quickly and can provide security and ease of use for domain experts accessing complex computing architectures – enabling researchers to get the most out of their investment in computing resources.

Cloud technology has now reached a level of maturity that is making it appealing to HPC users. Whether using public or hybrid cloud, the technology offers unprecedented flexibility for users who can create or 'spin-up' nodes with specific architectural requirements, use cloud bursting to increase the capacity of their in-house infrastructure – or it can increase the agility of a company that shares data over multiple sites.

Cloud enables organisations to access emerging technologies such as quantum computing hardware without the investment in prototype technologies. Users can adopt a strategy to learn and understand how a computing



aoripiza/shutterstock

system can impact their business using a pay-per-use model, which enables them to evaluate new technology and then scale as necessary.

One aspect of designing and procuring HPC systems in the past was the need to create a balanced architecture. This means looking at the kind of applications that will be run on a particular cluster to try and match the requirements of applications with the technologies that are needed. For example,

some workloads require large memory nodes, high-speed storage or interconnects, or high-performance storage.

Cloud HPC allows people setting up this infrastructure to make more efficient decisions, particularly if they are cloud bursting or developing a hybrid cloud strategy – as they can build their in-house resources to cater for 80 per cent of the user requirements while using the cloud to provide GPUs or specific node architectures that

suit a small number of users.

This allows all applications to benefit from this balanced architectural approach, while still being able to cater to the specialised applications that have more niche requirements.

### Other products

Based on **Alibaba** Cloud infrastructure, Alibaba Cloud Elastic High Performance Computing (E-HPC) is an end-to-end public cloud service. E-HPC provides individual users, education and research

institutions, and public institutions with a fast, elastic and secure cloud compute platform that interconnects with Alibaba Cloud products.

**atNorth** is a Nordic data centre services company offering environmentally responsible, power-efficient, cost-optimised data centre hosting facilities and high performance computing (HPC) services. atNorth offers sustainable and extremely scalable HPC resources fully delivered as a service, enabling our customers to focus on their simulation applications and calculations without having to worry about the underlying HPC infrastructure.

**AWS** provides the most elastic and scalable cloud infrastructure to run your HPC applications. With virtually unlimited capacity, engineers, researchers and HPC system owners can innovate beyond the limitations of on-premises HPC infrastructure. AWS delivers an integrated suite of services that provides everything needed to quickly and easily build and manage HPC clusters in the cloud to run the most compute-intensive workloads across various industry verticals.

**Cirrascale Cloud Services** is a premier cloud services provider of deep learning infrastructure solutions for autonomous vehicles, medical imaging, natural language processing and other deep learning workflows. The company was designed to focus on helping clients choose the right platform and performance criteria for their cloud service needs.

LMX Cloud from **Define Tech** is a comprehensive Cloud HPC cluster management stack that supports a broad range of workloads and software environments, enabling organisations with an agile and scalable IT infrastructure. One of its many key features that speaks to HPC users in particular is the ability to 'compose' or dynamically

configure their HPC resource when demand dictates. With LMX Cloud, HPC users and IT admins can auto-provision resources from pools of compute, GPU, FPGA, NVMe and storage-class memory in seconds and scale up or out as needed – all from a single, easy-to-use management interface, and this composable HPC feature is also compatible with job schedulers, so can be automated. HPC is all about scale and speed.

**Google Cloud's** (HPC) solutions are easy to use, built on the latest technology and cost-optimised to provide a flexible and powerful HPC foundation that clears the way for innovation. Google Cloud enables users to scale their team and use pre-configured HPC virtual machines (VMs) to get jobs started quickly and with predictable performance. Get deeper insights and explore your results using Google's AI and machine learning (ML) capabilities.

The **Grey Matter** Connected Cloud is a comprehensive pathway to ensure you connect to the cloud with confidence. Their specialist Cloud Solutions Team can help you build a cloud strategy and work with you to transform your business with the right licensing and cloud configuration, mobile devices for business, end-to-end cloud migration services, and post-deployment training and support.

**Gompute** provides a flexible HPC platform for CAE workflows and simulations. Gompute's compute node capacity delivered in the service is bare metal, equipped with a high-speed, low-latency interconnect and large memory options. Private Cloud Access dedicated hardware to outsource 'steady' capacity or burst individual projects. With a private cloud from Gompute you can get a tailored solution for short or long terms.

**H66cloud** from **Hydro66** provides a mature enterprise grade cloud environment,

instant launch, high performance with GPU options and zero maintenance. The company says there is no single point-of-failure, 100 per cent guaranteed uptime, no upfront costs and the opportunity to cancel at any time. The customer controls whether to pay in five-minute increments and only for what they run, or longer commitments for known workloads. Real-time technical support is available around the clock when you need a helping hand.

The **Azure** HPC OnDemand Platform, or azhop, delivers an end-to-end deployment mechanism for a complete HPC cluster solution in Azure. Industry standard tools, such as Terraform, Ansible and Packer, are used to provision and configure this environment. Each environment contains an Open OnDemand Portal for unified user access, remote shell access, remote visualisation access, job submission, file access and more, an active directory for user authentication and domain control, a PBS job scheduler and Azure Cycle Cloud to handle autoscaling of PBS nodes through PBS integration.

**Nimbix** offers cloud and on-premises HPC, giving engineers and scientists access to infrastructure and the software needed to build, compute, analyse, scale and deploy simulation and AI/ML/DL applications for faster, more powerful, less expensive cloud computing.

The Nimbix Supercomputing Suite is a set of flexible and secure as-a-service HPC solutions. This as-a-service model for HPC, AI and Quantum in the cloud provides customers with access to one of the broadest HPC and supercomputing portfolios, from hardware to bare metal as a service, to the democratisation of advanced computing in the cloud across public and private data centres.

**Penguin Computing's** Cloud Technology practice is focused on delivering software-defined architectures that enable you to run your workloads regardless of where your compute or data resources reside. The company suggests these platforms deliver the advances of a Cloud 2.0 world, where workloads are delivered on simultaneously addressable resources. Its goal is to enable you to run workloads everywhere as a seamless user experience by removing the complexities of workload portability, inclusive workflows, data locality and remote visualisation.

Open hybrid cloud is a recommended strategy for architecting, developing and operating a hybrid mix of applications, delivering a truly flexible cloud experience with the speed, stability and scale required for digital business transformation. The flexibility to run your applications across environments – from bare metal to VMs, edge computing, private cloud and public clouds – without having to rebuild applications, retrain people, or maintain disparate environments is the outcome of implementing an open hybrid cloud strategy.

**ScaleCloud Enterprise** from **Scalematrix** is designed to address the common trade offs in cloud environments for compute-intensive workloads. The product features top-of-the-line Intel processors and HPE servers housed in cabinet technology, the company says.

**UberCloud** – a cloud simulation platform for engineers – helps engineers run their simulation tools with high performance and reliability in the cloud. The company says its self-service software platform lets you create scalable cloud clusters, all while using the native GUI of Ansys, COMSOL, CST, NUMECA and more. There is no loss of features with simplistic web portals that only support batch use cases. ■

# Tech Focus: Software tools

## A ROUNDUP OF SOFTWARE TOOLS AVAILABLE TO SCIENTISTS USING HPC AND AI SOFTWARE

There are a wide range of software tools available to HPC users. This article focuses on freely available or open-source tools that scientists can use to improve their software performance or increase software portability. While there are a huge number of different categories of tools available to the HPC community, the exascale projects in the US and Europe are focused on the development of open source software, or software that can facilitate the use of a wide range of resources.

The trend towards open source – or at least collaboratively produced, freely available HPC tools – helps to harness the expertise of a fragmented software ecosystem to accelerate exascale development. This is particularly relevant for the development of the first applications for exascale, as these science codes will be the first applications to run at this kind of scale.

Hardware diversity is the driving trend for portability and the development of tools can make it easier for scientists to make use of the wide variety of different potential exascale HPC architectures. LLVM, Raja, Kokkos and SYCL



Gorodenko/shutterstock

are all examples of software tools currently being used by the US Department of Energy National Labs in the development of the Exascale Computing Project (ECP). While these tools support different aspects of the HPC software stack, they share a common goal in promoting access to a wide range of resources and help scientists increase the portability of their applications.

LEGaTO is an example of a European-funded software framework for exascale computing. The software toolset has recently been released to the public and was designed to accelerate the use of heterogeneous resources, as well as specific application areas, such as machine learning, healthcare and IoT applications for smart cities.

The current iteration of the EU-funded DEEP projects, 'DEEP-SEA', started on 1 April 2021 and will help to underpin the European Processor Initiative (EPI), which is developing hardware

for exascale systems.

DEEP-SEA will deliver the programming environment for future European exascale systems, adapting all levels of the software stack to support highly heterogeneous compute and memory configurations. While this project is only just starting development, the goals are to allow code optimisation across existing and future architectures and systems. The software stack includes low-level drivers, computation and communication libraries, resource management and programming abstractions with associated run-time systems and tools.

### Software tools

**Arm Developer**<sup>1</sup> provides a suite of software tools to help port and optimise applications, including porting and optimising HPC applications for Arm and the Arm SVE. These tools are split into groups based on their application area: biosciences, chemistry and

materials, computational fluid dynamics (CFD), high-energy physics, weather and climate, benchmarks and mini-apps and visualisation.

**Intel**<sup>2</sup> provides several software tools aimed at helping developers optimise HPC applications and software, including frameworks for AI and data analytics running on Intel architecture. This includes open-source HPC platform software through OpenHPC, Intel Parallel Studio XE, Intel Distribution for Python and Intel oneAPI Toolkits.

**LEGaTO**<sup>3</sup> (Low Energy Toolset for Heterogeneous Computing) is a programming framework designed to support heterogeneous systems. The toolset enables scientists to make use of CPU, GPU and FPGA resources that can offload specific tasks to different acceleration technologies through its own run-time system.

After three years of research, the various elements of the European-funded



project software toolset have been integrated together to facilitate the porting of future use cases to the energy-efficient LEGaTO hardware/ software platform.

The **LLVM** Project<sup>4</sup> is a collection of modular and reusable compiler and toolchain technologies written in C++. The LLVM Core libraries provide a modern source- and target-independent optimiser, along with code generation support for many popular CPUs. These libraries are built around a well-specified code representation known as the LLVM intermediate representation (LLVM IR).

The LLVM code representation provides in-memory compiler IR, and a human-readable assembly language representation. This allows LLVM to provide a powerful intermediate representation for efficient compiler transformations and analysis, while providing a natural means to debug and visualise the transformations.

The **Kokkos** C++ Performance Portability EcoSystem<sup>5</sup> is a production-level solution for writing modern C++ applications in a hardware-agnostic way. It is part of the US Department of Energies ECP. The EcoSystem consists of multiple libraries addressing the primary concerns for developing and maintaining applications in a portable way. The three main

components are the Kokkos Core Programming Model, the Kokkos Kernels Math Libraries and the Kokkos Profiling and Debugging Tools. The Nvidia HPC Software Development Kit (SDK)<sup>6</sup> includes compilers, libraries and software tools essential to maximising developer productivity and the performance and portability of HPC applications.

The **Nvidia** HPC SDK<sup>6</sup> C, C++ and Fortran compilers support GPU acceleration of HPC modelling and simulation applications with standard C++ and Fortran, OpenACC directives and CUDA. GPU-accelerated math libraries maximise performance on common HPC algorithms, and optimised communications libraries enable standards-based multi-GPU and scalable systems programming.

**OpenCL** (Open Computing Language)<sup>7</sup> is an open, royalty-free standard for cross-platform, parallel programming of diverse accelerators found in supercomputers, cloud servers, personal computers, mobile devices and embedded platforms. OpenCL improves the speed and responsiveness of a wide spectrum of applications in numerous market categories, including professional creative tools, scientific and medical software, vision processing and neural network training and inferencing.

Together with the OpenCL 3.0 specification, the working

group has released an early initial Khronos OpenCL SDK that developers can use to easily begin OpenCL coding. The SDK is open sourced on the Khronos GitHub under the Apache 2.0 license and will be continuously updated and expanded. This initial SDK release includes a new OpenCL guide, headers including vendor extensions, some small sample programs to illustrate how to use the SDK build system (with Cl), and an ICD Loader that will soon support installable development layers.

**OpenMP**<sup>8</sup> is a specification for a set of compiler directives, library routines and environment variables that can be used to specify high-level parallelism in Fortran and C/C++ programs. OpenMP allows users to create, manage, debug and analyse parallel programs while helping to support portability. The directives extend the C, C++ and Fortran base languages with single program multiple data (SPMD) constructs, tasking constructs, device constructs, worksharing constructs and synchronisation constructs, and they provide support for sharing, mapping and privatising data.

**OpenHPC**<sup>9</sup> is a Linux Foundation Collaborative Project whose mission is to provide a reference collection of open-source HPC software components and best practices, lowering barriers to deployment, and advancement and use of modern HPC methods and tools.

OpenHPC v2.0 was the most recent significant update that targets support for two new major OS distro versions: CentOS8 and OpenSUSE Leap 15. As the OpenHPC 2.x series targets major new distro versions, please note it is not intended to be backwards compatible with the previous OpenHPC 1.3.x series. OpenHPC, v2.3 is the current update intended primarily to enable resource manager support with the newer

hwloc included in RHEL 8.4.

**RAJA**<sup>10</sup> is a software library of C++ abstractions, developed at Lawrence Livermore National Laboratory (LLNL), which enables architecture and programming model portability for HPC applications. RAJA has two main goals: to enable application portability with manageable disruption to existing algorithms and programming styles; and to achieve performance comparable to using common programming models (for example, OpenMP and CUDA.)

RAJA is part of a portability suite that includes other tools: CHAI<sup>11</sup>, UMPIRE<sup>12</sup> and CAMP<sup>13</sup>. All these tools are developed by LLNL and are freely available on GitHub.

**SYCL**<sup>14</sup> (pronounced 'sickle') is a royalty-free, cross-platform abstraction layer that enables code for heterogeneous processors to be written using standard ISO C++, with the host and kernel code for an application contained in the same source file.

First introduced in 2014, SYCL is a C++ based heterogeneous parallel programming framework for accelerating HPC, machine learning, embedded computing, and compute-intensive desktop applications on a wide range of processor architectures, including CPUs, GPUs, FPGAs and tensor accelerators. ■

#### References

1. [developer.arm.com/solutions/hpc](https://developer.arm.com/solutions/hpc)
2. [intel.ly/3b9clr5](https://intel.ly/3b9clr5)
3. [legato-project.eu/softwarecomponents](https://legato-project.eu/softwarecomponents)
4. [llvm.org/](https://llvm.org/)
5. [github.com/kokkos/kokkos](https://github.com/kokkos/kokkos)
6. [developer.nvidia.com/hpc-sdk](https://developer.nvidia.com/hpc-sdk)
7. [www.khronos.org/opencl/](https://www.khronos.org/opencl/)
8. [www.openmp.org/specifications/](https://www.openmp.org/specifications/)
9. [openhpc.community/](https://openhpc.community/)
10. [computing.llnl.gov/projects/raja](https://computing.llnl.gov/projects/raja) managing-application-portability-next-generation-platforms
11. [computing.llnl.gov/projects/chai](https://computing.llnl.gov/projects/chai) copy-hiding-application-interface
12. [computing.llnl.gov/projects/umpire](https://computing.llnl.gov/projects/umpire)
13. [github.com/LLNL/camp](https://github.com/LLNL/camp)
14. [www.khronos.org/sycl/](https://www.khronos.org/sycl/)



Nasir Rizniak/shutterstock

# Weathering the storm

Weather and climate simulation can now be accentuated through the use of cloud computing, writes **Robert Roe**



With heatwaves scorching much of Europe over the past weeks, the importance of accurate weather and climate predictions has been brought sharply into focus. Weather and climate simulation are traditional HPC applications that require vast amounts of computational power, and this makes weather and climate simulation an expensive business as the barrier to entry is a massively powerful supercomputer, in addition to the skills and expertise in software to ensure the users can get the best performance out of their applications.

Just last month the US National Oceanic and Atmospheric Administration (NOAA) inaugurated its newest weather and climate supercomputers with an operational run of the National Blend of Models. The new supercomputers, first announced in February 2020 with a contract award to General Dynamics Information Technology (GDIT), provide a significant upgrade to the computing capacity, storage space and interconnect speed of the nation's Weather and Climate Operational Supercomputing System.

US Secretary of Commerce Gina Raimondo commented: 'Accurate weather and climate predictions are critical to informing public safety, supporting local economies and addressing the threat of climate change. Through strategic and sustained investments, the US is reclaiming a global top spot in high performance computing to provide more accurate and timely climate forecasts to the public.'

NOAA Administrator Dr Rick Spinrad added: 'More computing power will enable NOAA to provide the public with more detailed weather forecasts further in advance. Today's supercomputer implementation is the culmination of years of hard work by incredible teams across



Limitech/Shutterstock.com

NOAA – everyone should be proud of this.'

Enhanced computing and storage capacity will allow NOAA to deploy higher-resolution models to capture small-scale features like severe thunderstorms, more realistic model physics to capture the formation of clouds and precipitation, and a larger number of individual model simulations to better quantify model certainty. The result is even better forecasts and warnings to support public safety and the national economy.

Ken Graham, director of NOAA's National Weather Service, stated: 'Researchers are developing new ensemble-based forecast models at record speed, and now we have the computing power needed to implement many of these substantial advancements to improve weather and climate prediction.'

However, there is now a way to harness the power of weather and climate simulation without the huge investment in state-of-the-art supercomputers. Cloud providers such as atNorth can deliver large-scale HPC platforms with the latest in processor and GPU technology for any scientists and researchers to use. These resources can be deployed quickly and easily to ensure

they get access to the resources they need when they need them and only pay for the resources that are used. This means a cluster can be scaled up for brief periods of high usage and then scaled back when that period of peak usage subsides.

## Developing the weather and climate ecosystem

It was recently announced that the European Centre for Medium-Range Weather Forecasts (ECMWF) is starting to make some of its Integrated Forecasting System (IFS) available on an open-source basis to facilitate collaboration on the code.

This decision aims to bring the weather and climate forecasting community together so they can share their expertise and improve the quality of the models.

Some of the code is already available on the internet: a GitHub space has been created to host open-source IFS components. The merits of moving the full IFS forecast model to open source in the future will be reviewed this year, in consultation with ECMWF's Member States.

The main aim of the move is enhanced collaboration. Removing restrictions on redistribution could make working with ECMWF more attractive to collaborators. Contributing to open-source codes could also be more attractive to academic partners.

Another aim is greater efficiency. There currently is a mix of bespoke licences and open-source material that is time-consuming to handle. Some journals also require open access to the codes used.

A third reason is positioning ECMWF and its Member States at the centre of

**“Having a lower price and better performance from a sustainable source is a win-win situation. We love the cost savings as well as the sustainability, and so do our clients”**

international efforts on emerging high performance computing architectures. Making IFS code open source should encourage work on the IFS by computational science experts in academia and vendors.

The IFS GitHub space contains code released primarily to support pre-existing collaborations. While available to anyone, code in the IFS space is not generally supported. It currently contains the following elements: ecRad, ecTrans, FIAT, CLOUDSC, CLOUDSC2 and ecBundle. These tools help to predict spectral storms and cloud microphysics, among other features.

### Weather in the clouds

Today, the world's leading weather-forecasting organisations usually process weather data four times a day, which, depending on the speed of execution, can take around two hours at a time. This can be very expensive when using the cloud, resulting in companies only using cloud capacity for eight out of 24 hours, leaving a lot of downtime still to be paid for.

Tomorrow.io began its business providing weather intelligence. Their weather and climate security platform help organisations – from Ford, Porsche and Uber to Google and Fox Sports – to predict the business impacts of weather and solve their climate

security challenges. By producing business insights and action plans calibrated to a minute-by-minute forecast, they can automate decision-making at scale, helping countries, businesses and individuals better manage their weather-related challenges.

To achieve forecasting and analysis at speed and scale, Tomorrow.io needed to harness the power of high performance computing (HPC) to run its numerical weather prediction models.

As a start-up they used the public cloud, which gave them instant access together with a degree of flexibility. But as the business expanded and grew its operational and research offering, their HPC costs became increasingly expensive and more complicated to maintain; they found they were paying a premium for a service that did not always deliver what they needed, and though they considered other public cloud options, the potential for making savings proved very limited. Getting the power they needed to analyse data for the industry standard four times a day was proving an issue, with Tomorrow.io needing to increase data processing time so that it became a near-constant requirement.

The team at Tomorrow.io approached atNorth to see if they could power their next phase of research and development

(R&D) more cost-effectively. They wanted to change the rules of weather forecasting, analysing data constantly, in real time, to provide on-the-spot calculations that would allow their high-end clients to make business-critical decisions with confidence.

As a meteorological organisation with a strong focus on environmental, social and governance (ESG) issues, Tomorrow.io was interested in atNorth's cloud capabilities, but also in its sustainable, 100 per cent renewable energy-powered data centres, and their commitment to circular economy principles. Keenly aware of its carbon footprint and ready to minimise it, Tomorrow.io was looking for a solution that would not only save money for clients, but would also be good for the planet.

Luke Peffers, chief weather officer at Tomorrow.io, commented: 'We are a growing company with massive scale as our goal. One of the consequences of this is that we run large jobs for our customers that can eat away at our margins. atNorth offers us an excellent solution at a discount so we can scale as an organisation and pass the cost savings to our customers. Having a lower price and better performance from a sustainable source is a win-win situation. We love the cost savings as well as the sustainability, and so do our clients.'

## New White Paper now online

**VIEW  
FOR  
FREE\***



### Flexible, scalable and sustainable: The new power behind tomorrow's weather predictions

The world's leading weather-forecasting companies usually process weather data four times a day, which, depending on the speed of execution, can take around two hours at a time. This can be very expensive when using the cloud – resulting in companies only using cloud capacity for eight out of 24 hours, leaving a lot of downtime still to be paid for.

**SCIENTIFIC  
COMPUTING  
WORLD**

[www.scientific-computing.com/white-papers](http://www.scientific-computing.com/white-papers)

\*Registration required





# Unlocking new treatments with AI

**SCIENTISTS AND RESEARCHERS ARE USING AI TO HELP ACCELERATE THE DISCOVERY OF NEW DRUGS FOR A WIDE VARIETY OF DIFFERENT MEDICAL APPLICATIONS**

Drug discovery is undergoing a radical evolution of its capabilities due to the growing use of computational methods, including artificial intelligence (AI) and machine learning (ML) methods. These increasingly ubiquitous methods are driving companies to find new ways to develop candidate drug compounds and are opening the path toward more personalised medicine initiatives in the future.

One method that has gained a lot of popularity is harnessing AI and ML to find potential drug molecules faster. In a paper that will be presented at the International Conference on Machine Learning (ICML), MIT researchers developed a geometric deep-learning model called EquiBind that is 1,200 times faster than one of the fastest existing computational molecular docking models, QuickVina2-W, in successfully binding drug-like molecules to proteins.

The researchers found that a geometric deep-learning model is faster and more accurate than state-of-the-art computational models, reducing the chances and costs of drug trial failures.

EquiBind is based on its predecessor, EquiDock, which specialises in binding two proteins using a technique developed by the late Octavian-Eugen Ganea, a recent MIT Computer Science

and Artificial Intelligence Laboratory researcher and Abdul Latif Jameel Clinic for Machine Learning in Health (Jameel Clinic) postdoc, who also co-authored the EquiBind paper.

Before drug development can even take place, drug researchers must find promising drug-like molecules that can bind or 'dock' properly onto certain protein targets in a process known as drug discovery. After successfully docking to the protein, the binding drug, also known as the ligand, can stop a protein from functioning. If this happens to an essential protein of a bacterium, it can kill the bacterium, conferring protection to the human body.

However, the process of drug discovery can be costly both financially and computationally, with billions of dollars poured into the process and over a decade of development and testing



Darren Baker/Shutterstock.com

before final approval from the US Food and Drug Administration (FDA).

Moreover, 90 per cent of all drugs fail once they are tested in humans due to having no or too many side effects. One of the ways drug companies recoup the costs of these failures is by raising the prices of the drugs that are successful.

The current computational process for finding promising drug candidate molecules goes like this: most state-of-the-art computational models rely upon heavy candidate sampling coupled with methods like scoring, ranking, and fine-tuning to get the best 'fit' between the ligand and the protein. Hannes Stärk, lead author of the paper and a first-year graduate student advised by Regina Barzilay and Tommi Jaakkola in the MIT Department of Electrical Engineering and Computer Science, likens typical ligand-to-protein binding methodologies

to 'trying to fit a key into a lock with a lot of keyholes.' Typical models time-consumingly score each 'fit' before choosing the best one. In contrast, EquiBind directly predicts the precise key location in a single step without prior knowledge of the protein's target pocket, which is known as 'blind docking'.

Unlike most models that require several attempts to find a favourable position for the ligand in the protein, EquiBind already has built-in geometric reasoning that helps the model learn the underlying physics of molecules and successfully generalise to make better predictions when encountering new, unseen data.

The release of these findings quickly attracted the attention of industry professionals, including Pat Walters, the chief data officer for Relay Therapeutics. Walters suggested the team try their model on an already existing drug and protein for lung cancer, leukaemia, and gastrointestinal tumours. Whereas most of the traditional docking methods failed to bind the ligands that worked on those proteins successfully, EquiBind succeeded. 'EquiBind provides a unique solution to the docking problem that incorporates both pose prediction and binding site identification,' Walters says. 'This approach, which leverages information from thousands of publicly available crystal structures, has the potential to impact the field in new ways.'

'We were amazed that while all other methods got it completely wrong or only got one correct, EquiBind was able to put it into the correct pocket, so we were very happy to see the results for this,' added Stärk.

While EquiBind has received a great deal of feedback from industry professionals that has helped the team consider practical uses for the computational model, Stärk hopes to find different perspectives at the upcoming ICML in July.

'The feedback I'm most looking forward to is suggestions on how to improve the model further,' he says. 'I want to discuss with those researchers ... to tell them what I think can be the next steps and encourage them to go ahead and use the model for their own papers and for their own methods ... we've had many researchers already reaching out and asking if we think the model could be useful for their problem.'

This work was funded, in part, by the Pharmaceutical Discovery and Synthesis consortium; the Jameel Clinic; the DTRA Discovery of Medical Countermeasures Against New and Emerging threats program; the DARPA

**"We strongly believe, after several decades of stagnated investment and innovation, cardiovascular disease is re-emerging with a newfound interest"**

Accelerated Molecular Discovery program; the MIT-Takeda Fellowship; and the NSF Expeditions grant Collaborative Research: Understanding the World Through Code.

### Cardiovascular disease

CardiaTec Biosciences recently announced that it had secured a £1.4million pre-seed investment led by Laidlaw Scholars Ventures and APEX Ventures with participation from Crista Galli Ventures, o2h ventures and Cambridge Enterprise.

The AI drug target discovery company, which specialises in cardiovascular disease, was co-founded in 2021 by a trio of AI academics and ambitious alumni from the University of Cambridge.

Of the three – Raphael Peralta (CEO), Thelma Zablocki (COO), and Namshik Han (CTO) – Raphael and Thelma are graduates of the University of Cambridge MPhil in bioscience enterprise.

Dr Han is an academic in AI applications for target and drug discovery, and he holds positions at the University of Cambridge as head of AI, at the Milner Therapeutics Institute and associate faculty of the Cambridge Centre for AI in Medicine.

The company is developing a target discovery platform leveraging AI to make sense of large-scale multi-omic cardiovascular data. As opposed to conventional singular omic analysis, CardiaTec's proprietary platform unravels relationships that span across every level of biology, from gene variation, methylation and expression, to their connection to proteomic and metabolomic functions to understand disease development best.

Dr Han said: 'Recent advances in artificial intelligence are generating novel ways to interpret multi-omic data. I am excited to lead CardiaTec's technology strategy to establish a new paradigm for understanding the pathophysiology of cardiovascular diseases.'

Raphael Peralta, CEO of CardiaTec, said: 'We strongly believe, after several decades of stagnated investment and innovation, cardiovascular disease is →

**"This approach, which leverages information from thousands of publicly available crystal structures, has the potential to impact the field in new ways"**



Gorodenkoff/Shutterstock.com

→ re-emerging with a newfound interest, driven not only by the increasing requirements to fulfil the unmet need as it persists as the world-leading cause of death, but in the application of AI in being able to drive new and meaningful insights to help meet patients' needs. Therefore, CardiaTec finds itself incredibly well placed to help drive innovation forwards within this space, now supported by a great syndicate of investors.'

### Expanding AI

It has recently been announced that AI-powered drug discovery company Healx has announced that is opening new labs at Chesterford Research Park. Healx, specialises in using AI models to find candidate drugs to help treat rare diseases. This is particularly important in this field as rare diseases can be much harder to treat as there is insufficient investment in these conditions due to the small number of people suffering from them. In addition, rare diseases are often not well studied and there is a limited understanding of many of the aspects necessary to support a drug discovery programme. Healx aims to change this through the application of AI. The company uses AI models helps to solve these challenges by analysing millions of drug and disease data points to find novel connections that could be turned into new treatment opportunities.

Dr Neil Thompson, Chief Scientific Officer, Healx said: 'Historically, we have worked through partnerships to access

**"As our team and operations have scaled, we started looking to secure our own labs in order to support our growing portfolio of disease projects and to expand the proprietary data types we use in our AI platform"**

the experimental systems we require for our preclinical and clinical programmes, but, as our team and operations have scaled, we started looking to secure our own labs in order to support our growing portfolio of disease projects and to expand the proprietary data types we use in our AI platform.

'The Chesterford Research Park facility is ideal for us and will play a pivotal role in our vision for the next generation of drug discovery for rare diseases. The modern, purpose-built lab has the flexibility to support our scientific and technical ambitions and the location - close to our Cambridge headquarters and enable us to expand our team with experienced local talent. Importantly, the Park Management staff have been very supportive in enabling a speedy acquisition, and we are looking forward to getting up and running in the space,' Thompson continued.

The fully fitted, high-quality suites are

served by high-speed data connectivity, a dedicated external plant area and Cat 6 cabling. These ready-made labs include demountable benching, an open plan write-up area, plus a self-contained single office which means that Healx can 'plug and play' without needing to initiate an extensive fit-out.

Focused on finding novel treatments for patients with rare diseases, Healx is deeply passionate about the application of AI to improve access to treatment and the new lab will allow the Healx team to accelerate the discovery and validation of potential new therapies for a range of rarer conditions.

Julian Cobourne, head of regional offices, Aviva Investors, joint owners of Chesterford Research Park with Uttlesford District Council commented: 'We are thrilled to provide space for a new breed of healthcare technology companies, like Healx, to grow amongst our existing community of cutting-edge, global life science companies. With its new lab space in the Park's innovative community, Healx will be able to continue its crucial research, rapid expansion and recruitment.'

Chesterford Research Park provides a superb environment for innovative research and development for both established and early-stage biotechnology and pharmaceutical companies and is home to life science and technology innovators including Arecor, AstraZeneca, Cambridge Epigenetix, Charles River Laboratories, Domainex, Illumina and Lonza. ■



# Case study: Secure data sharing – making collaborative drug discovery a reality

CDD vault helps lab users overcome data challenges in modern drug discovery workflows

'We had spreadsheets all over the place, and data from different projects that were just separated in different folders. It got to the point where we didn't know where to put our data, or where to later find it.' This is one of the common issues that has led global biotechs of all sizes to Collaborative Drug Discovery (CDD), a software provider for research and development data management.

Drug discovery is data driven, and that data underpins every scientific and commercial decision, which could ultimately spell the difference between the success and failure of a research and developmental program for new drugs. Yet in today's labs the handling and management of data doesn't necessarily maximise its value or usability.

Scientists commonly store and manage their data in unsecure, often difficult to find disparate documents and spreadsheets. While this method might be okay for a lone scientist working in a vacuum, it is not likely to represent a smart approach for collaborative scientists working in drug discovery or in other chemical or biological fields that rely on the ability to store, recall, process and share large amounts of data quickly.

CDD Vault acts as a central smart warehouse for all drug discovery data, explains Kellan Gregory, the informatics firm's head of product excellence. 'Our platform offers a comprehensive set of core utilities to allow lab members to access all of their results

data, in its contextual format.' This means the ability to handle any type of data. 'As well as being able to capture numbers and text, we can also capture the native file in line with the data.' A formula builder, tools for activity and physical chemistry property calculations, the ELN and dynamic visualisation tools then provide extra layers of intuitive analyses.

Julio Martin is director and head of the Kinetoplastid Discovery Performance Unit (DPU) at GlaxoSmithKline (GSK) R&D's Tres Cantos Open Lab Foundation, a ground-breaking PPP initiative set up at GSK's dedicated diseases of the developing world (DDW) research facility at Tres Cantos in Madrid. The Open Lab Foundation supports collaboration by giving external partners access to GSK compounds, infrastructure and drug discovery expertise, with a view to accelerating research in multiple areas from target discovery and validation, to compound screening lead identification, and optimisation. The organisation chose CDD Vault as its data management platform for all Tres Cantos anti-kinetoplastic screening data generated internally and through external partnerships.

'This kind of hosted solution means there is no need to have to navigate firewalls, and it also scales up in parallel with project growth. Using CDD Vault means we can put more of our internal resources into scientific research, rather than have increased costs associated with setting up, maintaining and upgrading

complex platforms,' Julio says. 'Researchers can interrogate the database for their own and externally submitted data, and be confident of confidentiality and security for every user.'

Importantly, CDD Vault can be set up quickly, with no specialist IT input, and is very agile. 'The days of heavy, custom solutions are numbered, Kellan states. They are expensive, difficult to roll out, and can require extensive ongoing IT expertise. In contrast, we can get a new customer up and running with CDD Vault in minutes. The CDD team manages all of the infrastructure and provides ongoing support for any additional configurations required at any point. The ability to become productive in a short amount of time is really big plus point for our customers.'

And with UK government figures released in April 2018 indicating that more than four

in 10 of all UK businesses – and 72 per cent of large businesses – suffered a cyber breach or attack in the last 12 months, CDD Vault is highly secure. Built into industry standard SSAE 16 Type II certified cloud storage, and with two-factor authentication and IP tracking on the ground, CDD Vault is designed to minimise the chance of a hack either from the outside, or from inside the client's organisation. James Moe is president, CEO and co-founder of Oligomerix, a small biotech company interested in understanding the role of tau protein in neurodegenerative diseases, and the discovery and development of treatments for Alzheimer's disease and related tauopathies. Moe comments: 'We're using CDD Vault as a way of storing our chemical structures, as a way of searching them, as a way of storing all of our assay data. We're also using it for doing calculations, for analysing our data, for creating reports and communicating the data to others, and then also, very importantly, for working securely with collaborators. So it's been instrumental for all of those purposes. Another primary concern is having a database where we have better security over our molecules.'

**Find out more at:**  
[www.collaborativedrug.com](http://www.collaborativedrug.com)

#### Reference:

[https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment\\_data/file/702074/Cyber\\_Security\\_Breaches\\_Survey\\_2018\\_-\\_Main\\_Report.pdf](https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/702074/Cyber_Security_Breaches_Survey_2018_-_Main_Report.pdf)

**"Using CDD Vault means we can put more of our internal resources into scientific research, rather than have increased costs associated with setting up, maintaining and upgrading complex platforms"**

# Genomic potential

GENOMICS SOFTWARE IS HELPING TO TRANSFORM LARGE VOLUMES OF UNSTRUCTURED DATA INTO ACTIONABLE KNOWLEDGE WITH OPEN-SOURCE TOOLS AND DATA ARCHITECTURE THAT IS APPLICABLE FOR CLINICAL GENOMICS AT SCALE, WRITES SOPHIA KTORI

Tartila/shutterstock



Today's DNA sequencing technologies now make it possible to sequence whole human genomes cost effectively and with speed. Sequencing initiatives are generating vast volumes of data that – theoretically – give scientists a starting point to drill down into individual patient genomes in the hunt for disease-related variants, and also to mine collectives of huge public datasets to aid our understanding of the genetic basis of disease, unpick disease mechanisms, identify drug and diagnostic targets, and stratify patients for clinical trials and personalised medicine.

In practice, analysing this wealth of genomic data, in the context of associated biological and clinical data, is challenging. Gene variants identified through genotyping studies are stored in variant call format (VCF) files, but deriving patterns and insight from these files and connecting disparate data types isn't necessarily intuitive. And with relational datasets generated through large public and private initiatives (containing potentially millions of variants from many thousands of individuals) there are immediate issues associated with scale, as well as with how one can formulate the

**“If you've ever worked with a relational database, you have to typically join data across lots of tables”**

right queries. In contrast with relational databases, graph databases can help to transform large-volume unstructured data into actionable knowledge, explains Alicia Frame, senior director, graph data science at Neo4j. 'In the case of genomic research, the key problem is how to integrate the large volumes of highly heterogeneous data and gain maximum insight,' she says. This is whether for diagnosis, personalised therapies or drug development, she is keen to stress. 'Graph databases are an ideal way to represent biomedical knowledge and offer the necessary flexibility to keep up with scientific progress. Using graph databases, a well-designed data model and query can deliver in seconds what previously took days of manual analysis.'

Graph platforms are effectively a way of representing and storing data as connected concepts, Frame explains. 'You can think of the graph as built on nodes that are concepts and then the relationships that connect them,' she says. 'In 'everyday speak', we might well

consider the nodes as nouns. So, in the genomics or bioinformatics space, these 'nouns' are the genes, chemicals, diseases, variants and phenotypes. And then, of course, the relationships between them are effectively the verbs, which connect the concepts. It's – kind of – a real-world systems biology model.'

Under the Neo4j platform, the data is stored in the same way that the 'nouns' and 'verbs' relate to each other in biology, says Frame, so getting the data you want back out is very intuitive. In a relational database, where everything is stored as rows and columns, you need to join the data – and that means spending a lot of time thinking about how the computer stores that data and trying to map how to connect it. Cypher lets a domain expert query far more naturally for patterns in the data. 'So the user can literally ask the database to find chemicals that bind to receptors for particular genes that are associated with a particular disease,' says Frame. This makes it very easy to effectively express a 'mental model' and phrase the questions naturally and retrieve the relevant information from the underlying database.

'If you've ever worked with a relational database, you have to typically join data across lots of tables,' she says. 'The more complex the query, the more complex it is to join the proverbial dots in the table. The more joins you have, the slower it is and the more difficult it is to write the query,' Frame acknowledges. 'Use a labelled property graph model based on nodes and relationships and there is no need to consider joins between tables, because the data is already joined.' It also becomes intuitive to add new data as it is derived.

#### **Open-source and user-friendly**

Graph databases also make it much easier to build applications for every end user – think again, clinicians and researchers – and, at the back end, it becomes relatively easy for the person building the graph to maintain the resource, update it and deliver it to those end users.

Neo4j has focused on making the open-source platform easily accessible and user-friendly for novices and smaller initiatives. 'For the community edition, we offer the database, plugins for data science and visualisation tools,' explains Frame. 'If you are a researcher or an individual, you can download our database and our software from our website for free. In fact, many groups start there.' The pivot point between the free, open-source version and the commercial enterprise platform will depend on the volume of data and the number of people who will

**“Graph databases are an ideal way to represent biomedical knowledge and offer the necessary flexibility to keep up with scientific progress”**

be using the system, she adds. 'One of the primary differences between the free community version and the enterprise system is parallelisation. The community platform will use up to four cores, whereas users of the enterprise platform can tap unlimited numbers of cores for faster computation when datasets are really huge and speed is important.'

In fact, many public genomic datasets are already encoded as graph databases. 'The NCBI, for example, has downloadable graph representations of many of its public databases,' Frame says. 'We also have a 'graphs for good' programme, through which we offer the commercial, enterprise software for free to nonprofits, charities, researchers and academics in order for them to do their research; we also licence the database and the plugins to drug discovery companies such as Novo Nordisk.'

The most obvious – although not the only – challenge associated with managing and analysing genomics data is its scale, comments Ignacio (Nacho) Medina, CTO of Zetta Genomics and founder of the open-source computational biology (OpenCB) platform. OpenCB is a bioinformatics suite that is designed to allow genotypic data management and analysis on a scale relevant to the massive sets of genome sequencing results that the research and clinical communities are generating. Medina describes the platform as a full stack open-source software solution, enabling large-scale genomic data storage, indexing, analysis and visualisation.

#### **Scalable genomics research**

The need for a dedicated, genomics-focused platform became increasingly evident to Medina more than a decade ago with the emergence of next-generation sequencing technologies and with the application of genotyping – not just for basic disease research, but also in clinical settings for potential applications in disease diagnosis and the development of personalised medicine.

As the first scalable solution enabling genotypes – recorded in variant call







→ format – to be stored in a variant database, OpenCB is a high-performance solution for indexing and analysing many hundreds of thousands of samples, he believes.

Medina, who has been Head of the Computational Biology lab on the HPC team at the University of Cambridge since 2015, conceptualised and founded the OpenCB project while working in Spain during 2012. Within a few years, the platform was gaining the attention of some major genomics research initiatives. 'At first, it was just a prototype – very small – but this was enough to raise the attention of EBI, the University of Cambridge and Genomics England in 2015, which adopted and contributed significantly to its development,' he says. During this period, Medina remained the platform's architect and has led the design and development of OpenCB. 'Today, OpenCB also includes a metadata and clinical database, fine-grained security management and a knowledge database, representing a complete genome data interpretation platform,' Medina notes.

As an open-source platform, OpenCB is accessible and free-of-charge for any organisation looking to manage and analyse genomics data in a non-regulated setting. In 2019, Medina spun Zetta Genomics out of Genomics England

and the University of Cambridge to commercialise the OpenCB technology as XetaBase – a regulated, clinically validated and technically supported data architecture and software solution that is applicable for clinical genomics data management and evaluation at large scale.

'Zetta Genomics is, effectively, the commercial venture established to extend the scope of OpenCB, and XetaBase – OpenCB's commercial name – was created and launched in 2020,' says Medina. 'XetaBase is now becoming a certified platform that meets the regulatory requirements for data in clinical settings, while also addressing the need for customer support and implementation skills 'built-in'. It's offered as a software and through a service model, so we provide updates, fixes and training, along with ongoing support.'


Medina remains the CTO of Zetta Genomics, which is now also the main

**“Regulatory and security issues aside, clinical labs face particular challenges with respect to how you deal with patients”**

contributor to OpenCB. In June 2021, Zetta won £2.5 million in VC seed-funding. This investment is being focused on growth, improving performance, stability and implementing new analysis. Some is also enhancing the company's partnership network while it expands from the UK to open both Spanish and US offices. Resource is also being focused on talent; securing additional team members with software, development and commercialisation expertise. Importantly, the OpenCB and XetaBase data architecture supports regulatory governance for clinical and genomic data management, including NHS digital security and privacy policies.

'Regulatory and security issues aside, clinical labs face particular challenges with respect to how you deal with patients' genotyping test data,' Medina explains. These challenges relate to the sheer numbers of tests that are performed and the volumes of data generated but, also, the almost inevitable shortfall in human resources to analyse all the data for each patient in the hunt for a gene variant that might be the pathogenic cause of a disease.

Another challenge that the OpenCB platform and XetaBase address is one of data sharing between scientists. Typically, if a clinician identifies a new disease-



pinkneyer/Shutterstock.com

related variant that explains pathogenesis and disease symptoms, that finding may stay buried in the clinician's notes.

'In some cases they can submit or publish their findings but, even if that happens, it can take as long as 12 to 18 months for peer review and publication,' says Medina. 'Clinicians really need to be able to share their findings – with all of the patient data-related regulations in place – across hospitals. With the new federation feature, XetaBase will finally address that need to make findings available within minutes, not months.'

XetaBase is cloud-hosted and this simplifies data management and scalability, with a huge emphasis on making data secure and, effectively, available in real time.

'You may have several gigabytes of genotypic and other contextual data and metadata per patient,' explains Medina. 'The server for our platforms runs in the cloud and so this fact allows customers to easily scale to their needs, [supporting] tens or hundreds of thousands of patients in some cases, while we take care of and provide all the services that they need for the platform.'

Importantly, the OpenCB platform is built on a fileless infrastructure. 'Other solutions rely on a file-based system, but

**"We take all of the billions of mutations from large datasets and put them into one index on the system to enable incredibly fast analyses"**

then how can you easily search across, say, 20,000 files to look for a disease-related variant that matches that of your patient?' he asks. 'In OpenCB, in contrast, all of the genomic variant data is aggregated in one indexed database. The largest example we have is Genomics England – for which there are about 140,000 whole genomes in one single installation, accounting for about 300 terabytes of data,' says Medina. 'And this fileless system means that, despite this massive volume and breadth of data, we can scan the whole database within minutes or scan any patient or the entire family in a few seconds.'

In fact, the OpenCB architecture makes it possible to include hundreds of different pieces of information relevant to each genetic variant and still query the whole platform. 'One analogy we can use to help explain this is Google,' says Medina. 'When

you search for something on Google, the system doesn't search through the one trillion pages of content individually. Rather, Google has every page indexed so, when you query Google, you query that index and it takes just milliseconds.

'We have done something similar with OpenCB. We take all of the billions of mutations from large datasets and put them into one index on the system to enable incredibly fast analyses.'

And, of course, this is critically relevant whether the query is for insight into one patient, such as searching for patients with the same mutation, but also for the disease researcher who might be querying different variants across all of the different samples, Medina adds.

The ultimate vision is for a platform such as OpenCB and XetaBase to help reduce drug development times, increase the speed of disease diagnosis and aid decision-making for patients.

'My goal for the next five years is to demonstrate we can have a significant impact on research and healthcare, and realistically help to reduce drug development times by potentially years,' says Medina. 'We also want to enable researchers to communicate their findings in a secure way, so that they can reanalyse data and ensure no patient is forgotten.' ■■

# Emerging software accelerates drug pipelines

# The rise of AI and quantum technology are driving new methods of analysis for new drugs, finds Robert Roe



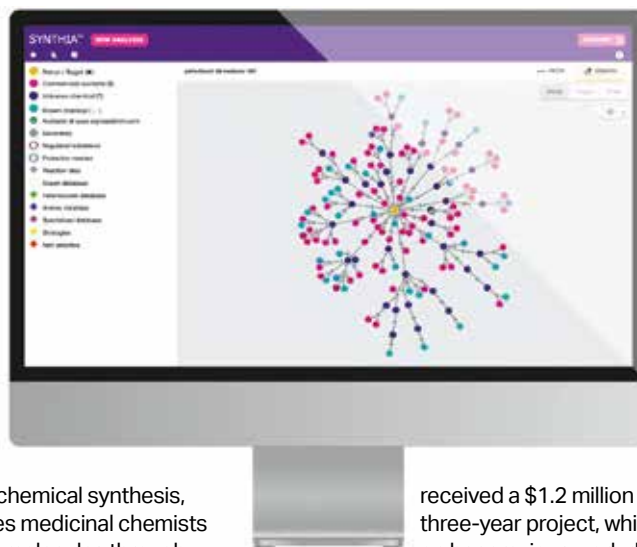
While drug discovery is an arduous and expensive process, computational chemistry software is helping researchers to streamline and accelerate that investigative pipeline. More specifically, AI-powered software algorithms and quantum chemistry are uncovering new pathways to synthesise chemical compounds and help scientists find new targets to evaluate.

Recent advances in AI-based software and quantum chemistry offer unprecedented opportunities to help speed up this stage of drug discovery and get effective medicine to patients much faster than has been possible using conventional methods.

This is a hugely impactful area of research when you consider the cost of drug development. An estimated £1.15bn and up to 12 years is spent designing a new drug and the majority of these (76%) are small molecules. While it appears easy to group this collection of molecules based on their molecular weight; in reality, they are a vastly diverse set of compounds. Most pharmaceuticals are considered small molecules, with the exception of proteins, such as insulin, and other biological medical products. One advantage that small molecule drugs have over large molecule biologics is that many small molecules can be taken orally, whereas biologics generally require injection or another parenteral administration.

But from a chemistry standpoint – particularly when concerning the synthesis of new compounds – small molecules are a hugely disparate category of compounds. There are potentially millions of active compounds and so sifting through the available data and creating new drugs is arduous, however critical, and everyone is searching for ways to make this process take less time and be more efficient.

At the heart of small molecule drug



discovery is chemical synthesis, which involves medicinal chemists creating new molecules through a complex, step-by-step process. Despite decades of research, this remains a lengthy, laborious procedure – and is still a critical bottleneck for advancing new medicines to the clinic.

Applying new software approaches such as artificial intelligence (AI) and quantum computing could open up new possibilities for researchers to accelerate drug discovery pipelines. While quantum computing may still be in its initial stages, there are some early applications on the horizon. AI-fuelled drug discovery software is much more established and is available today for a variety of different applications.

### \$1.2 million grant for quantum chemistry

Yuan Ping, who is Associate Professor of Chemistry and Biochemistry at UC Santa Cruz, leads one of eight research projects funded by the US Department of Energy (DOE) to advance the development of modelling and simulation software for the chemical sciences. This is one example of an early-stage project that hopes to explore the use of quantum software. Ping's group

received a \$1.2 million grant for the three-year project, which builds on her previous work developing

computational tools for predicting essential properties of materials and molecules based on quantum dynamics. Working with co-principal investigator Ravishankar Sundararaman at Rensselaer Polytechnic Institute, Ping's team plans to develop computational techniques and massively-parallel software for the rapidly emerging field of 'spin chemistry'.

'We can apply the same theoretical framework and computational tools in a different context to study the effects of spin dynamics on chemical reactions,' says Ping. 'If you can control the spin properties, you can change the products or increase the efficiency of chemical reactions, but realising the promise of spin chemistry requires a fundamental understanding of the mechanisms involved in the effects of spin on charge transfer and chemical reactivity.'

Ping's group will develop computational chemistry techniques for predictive modelling of spin dynamics and spin-dependent charge transfer, paving the way toward a detailed mechanistic understanding of spin chemistry.

## Open-source toolkit for computational chemists

While this spin chemistry grant is predominantly focused on high-performance computing, there are quantum-based chemistry software tools available to the public today. For example, Quantinuum recently announced the release of its InQuanto software as a standalone platform, bringing together the latest quantum

**“Recent advances in AI-based software and quantum chemistry offer unprecedented opportunities to help speed up this stage of drug discovery”**



computing tools in a single application InQuanto enables users to mix and match the latest quantum algorithms, advanced subroutines and chemistry-specific noise-mitigation techniques to make the best use of today's quantum computers. The platform also helps computational chemists to break down larger, industrially relevant systems into smaller fragments that can run on today's small-scale quantum machines. It uses Quantinuum's open-source toolkit TKET to reduce the computational requirements for electronic structure simulations and maximise performance across the widest range of quantum devices and simulators.

Patrick Moorhead, who is CEO and Chief Analyst of Moor Insights and Strategy, comments: 'Quantum computing offers a path to rapid and cost-effective development of new molecules and materials that could unlock novel answers to some of the biggest challenges we face. The way to ensure progress is to start prototyping now – using real-world use cases, so that methods are tailored to solving actual needs of the industry. InQuanto is built to enable exactly this.'

#### AI-powered retrosynthesis

Retrosynthesis helps to simplify the organic synthesis of a target compound. This has several potential benefits, which include

revealing different synthetic routes to create the target compound and cheaper synthesis procedures based on easily available precursor compounds. This is achieved by breaking a target compound into smaller constituent parts to develop a structural simplification.

Developing drug molecules is a slow, iterative process that involves sifting through thousands of candidates to find the one that's most suitable. Researchers need to carry out large numbers of experiments across several rounds of optimisation.

Merck has launched its retrosynthesis software, Synthia Retrosynthesis Software, which helps to accelerate the development of new drugs by harnessing the power of AI to assist chemists in completing retrosynthesis on the compounds being worked on. Retrosynthesis software provides invaluable information that can recommend the best possible route to execute – minimising cost, the number of steps needed and creating the best chance of making the required molecule with the desired properties. This can dramatically reduce the time it takes for a chemist to think of a viable route in the lab.

Engineered by chemists and computer scientists for more than 15 years, Synthia Retrosynthesis Software is powered by sophisticated algorithms that can help experts access and use the vast amounts of

data on chemical synthesis collated across many decades of research.

The tool works by harnessing the potential of advanced algorithms powered by more than 100,000 expert-coded reaction rules – painstakingly sifting through retrosynthetic possibilities while, at the same time, examining what has been done, what could be done and what starting materials are available.

This process begins with a researcher identifying a new drug target, such as an enzyme involved in a disease-critical pathway. The search then starts for potential drug candidates that can specifically interact with the chosen target and produce the required effect but, to become a successful drug, a molecule will also need a host of other properties, including non-toxicity, solubility and stability.

This is where optimisation and streamlining of the workflow can occur. Merck's Synthia Retrosynthesis Software helps scientists to find different routes to make their molecules. It can be used to explore new and known solutions and eliminate approaches that won't work. Ultimately, this software narrows down the most promising pathways, while also reducing the amount of unnecessary work undertaken by chemists who would traditionally conduct experiments manually to evaluate these different options. ■

## New White Paper now online

VIEW  
FOR  
FREE\*



### Reducing Risk & Improving Synthesis Outcomes with SYNTHIA Retrosynthesis Software

Addressing the many intrinsic and external challenges to chemical synthesis has led to a highly complex system of regulations that involve numerous academic and governmental organisations in modern times



[www.scientific-computing.com/white-papers](http://www.scientific-computing.com/white-papers)

SCIENTIFIC  
COMPUTING  
WORLD

\*Registration required


# Moving with the times

## THE WORLD OF MULTIBODY DYNAMICS SIMULATION IS CHANGING, WRITES GEMMA CHURCH


Additive manufacturing (AM) is finally opening up new possibilities for engineers to create components and subsystems enabling rapid prototyping and novel designs that can help to accelerate innovation. Software plays a key role in the maturation of this market with rapid increases in the complexity of parts and the materials and methods used to create new components with AM.

Multibody Dynamics (MBD) simulations are now a vital tool for today's engineers to help understand and improve the performance of a vast range of dynamic systems, helping them to address many ongoing challenges in the design and development process.

Francois Barral, Director of Simulia Multibody Systems Industry Process Consulting at Dassault Systèmes, explains: 'Product development teams face increased pressure from global competition and government regulations, which require product development organisations to create new, energy-efficient, safe and reliable products



Flexible bodies in a gearbox shown in Adams



The new Adams Modeler UX makes Adams more accessible and reduced prep time by at least 20 per cent for this robot

in ever-shorter time frames. Products are also becoming more complex, with electronic and software controls combined with the mechanical system. To meet these many demanding challenges, organisations are increasing their use of MBD simulation, which helps them replace physical testing with virtual prototypes.'

MBD simulation can be applied to almost any branch of mechanical engineering, but it is regularly used across automotive, railway and wind-energy industrial sectors. And these multibody systems must often work across a broad range of applications and scenarios, all of which can affect their functional requirements.

For example, an automotive braking system may need to work across numerous vehicle types from many manufacturers. In construction, a machine may need to excavate different types of material or undergo rigorous safety tests.


Japan-based crane manufacturer Tadano, for example, uses multibody dynamics to help accelerate and validate its R&D process, while providing a safe environment for operators.

Barral explains: 'A crane boom that is equal to an arm of a crane can cause boom-crane deflection and load wing when it attempts to lift a load. To

understand the physical property of a load swing and how cranes are used at construction sites, it is essential to consider the control method of a load swing. However, designers are not allowed to enter construction sites where cranes are at work due to 'hazardous areas.'

Using Dassault Systèmes' Simpack multibody dynamic simulation software, Tadano could simulate large-scale 3D vehicle models, reduce a company's manufacturing prototype costs and time for rework 'significantly', according to Barral. 'It was crucial for engineers to generate and analyse a physical model for structures based on accurate design information. Simpack is now being used to develop a model for major types of rough terrain cranes that are a mainstay product of Tadano.'

MBD is even helping our washing machines run more efficiently. Recently, Altair MotionSolve was used to simulate the motion dynamics of the washer drum, including the balance ring and the suspension system. An idealised motor model was used to define the various spin speeds of the washing machine during agitation and spin cycles. Combined with Altair's other simulation and modelling tools, consumer appliance company Mabe



▲ A simulation of landing gear

increased the capacity of its washing machines by 35 per cent, improved the energy factor by 24 per cent and water factor by 52 per cent, while simultaneously reducing the cost per cubic foot by 10 per cent and the product development cycle time by 25 per cent.

EnginSoft's RecurDyn solution also helped Toyota's Forklift division, which was having difficulties predicting the behaviour of the vehicle when carrying loads and driving through obstacles. The truck was modelled in RecurDyn, including flexible bodies to model the forks and the built-in tyre module, to accurately predict the behaviour of a fully-loaded truck driving over obstacles.

Danujan Sivanesan, a Senior Project Engineer specialising in multibody dynamics at Enginsoft, says: 'RecurDyn reliably analysed the critical conditions normally understood during the testing phase and computed the dynamic loads between the forks and the loads accurately using RecurDyn Geo-contact technology. This helped the customer reduce their product development time, enabling them to iterate through multiple designs and improve the performance and quality of their product.'

This range of applications and interested industries continues to constantly expand. 'Markets are starting to adopt this type of technology now to design better products and have not traditionally been operating in this space before,' according to Chris Harduwar, VP of Business Development at Maplesoft.

'[Engineers are] exploring rigid and flexible multibody dynamics as a way to assess the interactions on components such as loads, tensions and torques and develop the specifications and limitations for a machine to operate,' he says. 'At Maplesoft, our customers often work to improve performance with components

that fit into a larger machine or address issues that span multiple domains.

All industries have 'some common and representative problems that can be addressed and solved through the use of MBD simulations,' according to Jim Ryan, VP of Model-Based Systems Solutions at Altair. 'MBD can really help improve the positioning to make sure that parts move as intended – one example is with those systems involving precisely controlled movements such as active suspensions and precise pointing systems such as telescopes and radar.'

MBD simulations can also help

**"Using MBD, engineers could understand how fast the robot can go, how much payload it can handle, and how to optimise the design given the operating conditions"**

engineers calculate the sizing forces and torques needed to move the system with the correct actuators. You can also manage accelerations and peak forces, reduce wear and avoid friction- or contact-induced damage or calculate loads to minimise stresses, fatigue and, ultimately, failed operations or broken products.

MapleSim, for example, is used to develop high-speed robotics, including custom components between the joints or using motors that have not been defined yet. Engineers need to understand the torques and forces that are applied to those motors across the full design and including the interaction of the movement of all the components.

MapleSim was recently used to design a six-degree-of-freedom cable robot

and help understand the entire system's behaviour. 'Using MBD, engineers could understand how fast the robot can go, how much payload it can handle, and how to optimise the design given the operating conditions,' says Harduwar.

### Challenging times

Engineers are still faced with challenges when running MBD simulations. Methodically constructing very large models where potentially hundreds of components may move separately from each other is one such issue.

Ryan explains: 'It is quite easy and common to model such connections incorrectly, in which case the simulation might not proceed, or it might proceed but yield incorrect results, consistent with the CAE engineer's much-used axiom: garbage in, garbage out!'

The easy re-use of pre-constructed CAD assemblies (3D geometries) to retain the relative starting positions of all the components in a multibody system, together with all the mass and inertia properties, is another challenge to accurately predicting the motion.

'Even though multibody dynamics modelling has improved across several decades, it can still be challenging for engineers using MBD to properly model 100 per cent of the complex physics involved, due to such phenomena as part-to-part contacts and friction, especially involving flexible bodies and elastic parts of the system like belts, ropes and cables,' adds Ryan.

Mounting system complexity puts engineers in a difficult position when it comes to MBD simulations. Harduwar says: 'The engineers are inevitably left with a trade-off between the level of realism or fidelity of a model, compared to the computational and human engineering effort of defining the elements within the simulation.'

The motion control of the manipulators in a robotic arm, for example, is 'particularly challenging' to simulate, according to Harduwar. 'This is where system-level modelling becomes more effective – where equations of all the systems and components are combined and handled by a simulation tool.'

As product complexity continues to increase, this situation will continue to impact engineers in the near-term where trade-offs are often made. Harduwar explains: 'The challenge is doing something at the right fidelity level that will take the appropriate amount of time to get the right level of results.'

This is where system-level modelling using MBD has advantages over finite





→ element analysis (FEA) or computational fluid dynamics (CFD) analysis. Using system-level modelling and multibody dynamics, different types of equations are used where each component and the interactions between components are described. This is a more efficient way for engineers to understand the behaviour of certain systems in response to the force acting upon it.

Barral explains: 'These integrated capabilities enable designers and engineers to collaborate on replicating real-world scenarios and analysing multiple design options quickly and accurately. MBD simulations can help engineers study sub [systems] and complete systems, optimise critical components and systems earlier in the development cycle, avoid costly last-minute changes during physical testing and achieve overall efficiency and performance goals.'

As such, MBD is not only a more computationally efficient way to simulate a dynamic system, but also to simulate the rate at which engineers can get answers on a whole system when using system-level modelling.

## Visualisation matters

Visualisation tools are vital for MBD simulations; enabling engineers to analyse results quicker, thanks to the animations generated. Ryan says: 'People expect to see full 3D geometries realistically moving on their screens – and they can with MBD. And in the spirit of 'a picture can be worth a thousand words', it is also true that 'an animation can be worth 1e6 pictures.'

Visual animations, for example, can help engineers understand rates of change for the system as it moves. 'Especially when limits or targets are involved, it could be essential to know if actual values for velocities or accelerations are exceeding those limits or targeted values – for example, with a cruise-control system in a car that is set to make sure the car doesn't exceed 100kph,' Ryan explains.

Plot and graphs are helpful, while complementary visualisation aids are 'essential in those cases where motion-related quantities cannot easily be perceived or even accurately measured

**"People expect to see full 3D geometries realistically moving on their screens – and they can with MBD"**



▼ VTD and Adams - truck sensor perception as cab tilts and oscillates

through animations alone,' according to Ryan. 'For instance, an animation typically will not show the time-history of specific forces in the system to enable an engineer to spot the peak value of a force and also the time at which it occurs... which could be vitally important to the product's design.'

## Future motion

Advancing electrification will continue to impact mechanical simulations and MBD is helping ease this transition. Ryan explains: 'For years, as the products that companies make have continually grown – and accelerated – in complexity, we have seen an increased need to perform simulations of smarter electro-mechanical products involving all kinds of logic and feedback control systems.'

This is 'especially applicable' to systems in motion, according to Ryan. 'This is to help ensure that they move as intended – especially when safety considerations are important such as when developing something like antilock braking systems (ABS) and traction control systems (TCS) for different types of ground vehicles.'

MBD is also increasingly being used earlier in the design and development lifecycle. Previously, control system engineers and multibody dynamics engineers would develop their separate systems in isolation and then try to put them together at some point later in the development cycle – maybe as late as the first prototype stage.

But this is still 'much later than desired,' according to Ryan, who adds: 'Fortunately, we are now able to help companies break down the silos that have long existed between engineers and engineering teams who are focused on different disciplines – electrical

versus mechanical versus controls, for instance. These streamlined workflows are now helping companies to minimise bottlenecks in their product development cycle so that they can continue to accelerate time-to-market.'

AI and machine learning are complementary technologies that are really finding their feet in the MBD realm. Ryan says: 'Leveraging artificial intelligence/machine learning (AI/ML), especially using neural networks to create reduced order models (ROMs) of subsystems to accelerate simulation throughput, is another trend.'

'We're also seeing greater interest in digital twins in general and finding VR and AR applications to manage and monitor real-time operations of fleets of machines,' Harduwar adds.

Environmental concerns will continue to push MBD simulation adoption, a trend that will affect not just MBD but the entire simulation landscape.

Barral explains: 'Trends in sustainability, green energy and autonomous – and driver-assisted – vehicles are really important factors in the increased use of multibody simulation.'

'Moreover, industries are subject to irreversible 'trends' such as climate-change mitigation redefining operations. Climate issues are accelerating quickly and, with global regulation pushing for a more circular economy in line with market expectations, aligning to the sustainability agenda won't be easy.'

'However, the first step to reaching this goal is creating products in a fully collaborative, virtual environment. Once you have all of this information available in a single immersive experience, it becomes so much easier to visualise your data in the right context and then apply it for the best results.' ■

# FREE ON-DEMAND WEBCAST\*

## AVAILABLE NOW

### CDS Built for the Lab - Driving Ease of Use



#### SPEAKERS

**Darren Barrington-Light**  
Senior Marketing Manager  
Thermo Fisher Scientific

**Zoe Julian**  
Product Marketing Manager  
Thermo Fisher Scientific

The laboratory is a centralised hub of activity from sample analysis through to result generation. Setting your laboratory up for success requires selecting the right chromatography data system (CDS) and ensuring you realise its full capabilities to deliver maximum benefit to your business. That means you need software that is easy to learn and easy to use, enabling you to do more with less, eliminate errors and reduce out-of-specification results while keeping up with ever-evolving standards and regulations. Discover the latest release of Thermo Scientific Chromeleon CDS, bringing new capabilities to every laboratory.

#### Key Learning Objectives:

- Learn how to super-streamline your data analysis workflows
- Discover best practices for ensuring efficient data processing with smarter reporting
- Uncover how to make compliance easier and eliminate non-compliant spreadsheets from your laboratory
- Understand how to centrally and remotely maintain your CDS for maximised performance

**SCIENTIFIC  
COMPUTING  
WORLD**

**ThermoFisher**  
SCIENTIFIC

[www.scientific-computing.com/webcasts](http://www.scientific-computing.com/webcasts)

\*This webcast is FREE to view, though registration is required

# A roundup of the latest news stories in scientific computing

## LABORATORY INFORMATICS

### Francis Crick Institute receives £1bn to advance biomedical research

The Medical Research Council (MRC), Cancer Research UK (CRUK) and Wellcome Trust will provide The Francis Crick Institute with £1bn in funding over the next seven years to help expand the research institute's role.

This will support the ambition to make the UK into a place where researchers work at the forefront of global innovation. Research at the 'Crick' has studied the evolution of lung cancer, revealing how therapies can outpace the disease, and helped to discover a new vaccine against tuberculosis.

The UK's Prime Minister Boris Johnson, comments: 'The UK's Francis Crick Institute is at the centre of forging ground-breaking advances to beat diseases like cancer and dementia sooner, helping deliver major improvements to diagnosis and treatment, as well as preventing infection in the first place.'

'Thanks to £1bn in new funding, the Crick can go further to propel scientific discovery forward, harnessing British ingenuity, supporting new innovative companies to grow, and cementing the UK's place as a science superpower,' Johnson adds.

The Francis Crick Institute is a national flagship for biomedical research. It was formed in 2015 to help us understand more about how living things work to help transform treatment, diagnosis and aid the prevention of



Alex Teung/shutterstock

human diseases such as cancer, heart disease, infections and neurodegenerative diseases.

The institute also generates economic opportunities for the UK, with 10 companies launching from research projects undertaken across the years.

Professor John Iredale, MRC Executive Chair, says: 'The Crick has been a flagship discovery biomedical science centre since its formation in 2015. This funding from MRC, CRUK and Wellcome will continue to support the institute in advancing its world-class biomedical research and helping to solve scientific challenges.'

'Since its founding, the Crick has already produced many important advances in human

health and disease, spanning cancer, Covid-19, neurodegeneration and embryo development, and we're proud to continue supporting its ground-breaking research.'

Dr Michael Dunn, Director of Discovery Research at Wellcome, comments: 'The Francis Crick Institute is a world-leading research organisation, committed to ground-breaking discovery research, pioneering interdisciplinary research and training early-career researchers from all around the world. We are delighted to continue supporting the Crick's work, alongside MRC and CRUK, and this funding comes at a crucial time for discovery research. The Crick's bold ideas will make a real difference to human health.'

## MODELLING AND SIMULATION

### The first Cern-driven satellite to study radiation effects on electronics

Celesta, the first Cern-driven satellite, successfully entered orbit during the maiden flight of Europe's Vega-C launch vehicle. It was launched by the European Space Agency from the French Guiana Space Centre (CSG) on 13 July 2022. The satellite deployed smoothly and transmitted its first signals in the afternoon.

Weighing one kilogram and measuring 10 centimetres on each of its sides, Celesta (Cern Latchup

and Radmon experiment student satellite) is a 1U CubeSat designed to study the effects of cosmic radiation on electronics.

The satellite carries a Space RadMon – a miniature version of a well-proven radiation monitoring device deployed in Cern's Large Hadron Collider (LHC). Celesta has been sent into an Earth orbit of almost 6,000 kilometres. 'Right in the middle of the inner Van Allen belt, Celesta will survey an

unusual orbit where radiation levels are at their highest,' explains Markus Brugger, Head of the Cern Experimental Areas group and initiator of both the Charm and Celesta projects in the context of the R2E (Radiation to Electronics) initiative.

Ruben Garcia Alia, R2E project leader, says: 'The Space RadMon is a flagship example of how Cern technologies can have applications beyond particle

physics experiments. Based entirely on standardised, ultra-sensitive components selected and calibrated by Cern, and mostly in Cern facilities, the Space RadMon is a lightweight and low-power instrument, ideal for future risk-tolerant space missions.'

'If Celesta is successful, the Space RadMon could even be adapted to satellite constellations as a predictive maintenance tool to anticipate the →



# White Papers

now available online

VIEW  
FOR  
FREE\*



## Considerations for the digital laboratory

Thermo Fisher

Automation, robotics and digital technologies are transforming laboratory processes in life sciences, R&D and manufacturing, and driving a new era in data-focused disease and drug discovery.

## Accelerating the use of emerging technology

RCH Solutions

RCH Solutions supports the adoption of novel computing technologies offering a complete suite of strategic advisory and purpose-built managed services to help R&D teams in the Life Sciences modernise their compute environment

## Reducing Risk & Improving Synthesis Outcomes with SYNTHIA Retrosynthesis Software

Merck KGAA

We've come a long way from the earliest days of synthesis planning. Though ambitions of alchemical transmutation from base elements such as lead, to precious metals such as gold, have long since passed, the modern chemist still labors over a different kind of bubbling cauldron – one of cheminformatics.

## Engineering reaches for the cloud

RESCALE

Computational engineering is set for strong growth – and the use of the cloud is burgeoning

\*Registration required

[www.scientific-computing.com/white-papers](http://www.scientific-computing.com/white-papers)

SCIENTIFIC  
COMPUTING  
WORLD



robert.s.shutterstock

→ necessary renewal of satellites.’

A radiation model of the Celesta satellite was also tested in Charm, a Cern mixed-field facility capable of reproducing the radiation environment of low Earth orbit. The mission will be an important validation of this capability at the facility. ‘Capable of testing satellites all at once, rather than component by component, Charm is a unique installation

worldwide, remarkably different from other irradiation test facilities. It offers a simple, low-cost alternative and the possibility to assess system-level effects,’ says Salvatore Danzeca, Charm facility coordinator.

The success of this satellite is the result of a fruitful partnership between Cern and the University of Montpellier, which involved many students from both institutions

and radiation effect specialists from Cern. Celesta is based on the CSUM radiation-tolerant platform. It will be operated from the CSUM control centre. The European Space Agency provided the launch slot in the framework of its small satellite programme.

Enrico Chesta, Cern’s Aerospace and Environmental Applications Coordinator in the Knowledge Transfer Group, adds:

‘On a mission to make space more accessible, Celesta is an exciting example of how Cern expertise can have a positive impact on the aerospace industry.

‘With this mission, Cern displays its low-cost solutions for measuring radiation and testing satellites against it – thus providing universities, companies and start-ups with the means to realise their space ambitions.’

## MODELLING AND SIMULATION

# Collaboration explores quantum computing for semiconductor research

Quantinuum has announced a global collaboration with materials technology leader JSR Corporation of Japan to explore the application of quantum computing methods in semiconductor research.

The collaboration brings together JSR’s world-leading materials scientists with quantum computing experts at Quantinuum in Japan, Europe and the USA. The joint team will use the state-of-the-art InQuanto software platform to explore methods using quantum computers to model semiconducting materials, such as metal complexes and transition metal oxides.

These materials are essential to microelectronics. It is hoped new modelling methods using quantum computers may achieve accurate predictions of their physical properties that, in the future, could accelerate the identification of new candidate molecules and materials and open the way to future microelectronic device paradigms.

Rei Sakuma, Principal Researcher of the Materials Informatics Initiative of JSR, says: ‘We are delighted to have formed this new

collaboration with Quantinuum, which builds on our previous work together. The Quantinuum team continues to lead the quantum computing hardware and software field, complementing our scientists’ deep expertise in materials innovation. Our aspiration is to develop materials that can enrich society and the environment. Quantinuum’s software platform InQuanto is already helping our team to better understand how quantum computing may help us accelerate our path towards that ambitious goal.’

The collaboration will focus on developing quantum algorithms and methods based on dynamical mean-field theory (DMFT). This approach could provide a more accurate understanding of the electronic properties of complex organic and inorganic materials in the real world, such as optical absorption and conductivity, which could pave the way for future progress in the silicon-based information age.

Quantinuum and JSR will use InQuanto to explore new methods to model these complex molecular systems and defect subsystems. The new methods discovered will be incorporated

into InQuanto, and will become available for the use of other scientists and researchers using the software platform.

Ilyas Khan, CEO of Quantinuum, says: ‘The work we do with JSR is at the absolute cutting edge of materials science using quantum computers and we are thrilled to continue our relationship. This work will further develop InQuanto’s functionality, ensuring that new developments will become available to other users. This is the value of such a collaboration: JSR’s scientists know materials science, we know quantum computing and the scientific community benefits.’

InQuanto was recently launched as a stand-alone platform and brings together the latest algorithms, methods and noise mitigation techniques used by molecular and materials scientists and researchers on quantum computers and emulators. It will give JSR’s scientists and researchers a greater understanding of the capabilities of quantum computers in their path towards quantum advantage in computational chemistry.


## MODELLING AND SIMULATION

## Oxford Quantum Circuits raises £38m Series A funding

Oxford Quantum Circuits (OQC) has announced that it has raised £38m in the first close of an ongoing Series A investment round. The investment is the UK's largest-ever Series A in quantum computing and will accelerate research and development and fuel expansion in the Asia-Pacific region. The round was co-led by Lansdowne Partners, one of Europe's leading investment firms, and The University of Tokyo Edge Capital Partners (UTEC), Japan's largest deep-tech VC fund. British Patient Capital and existing investors, Oxford Science Enterprises (OSE) and Oxford Investment Consultants (OIC), also participated.

Ilana Wisby, CEO of OQC, comments: 'This initial close is the UK's largest-ever Series A in quantum computing – demonstrating our investors' confidence in our ability to lead the global quantum industry. It's testament to the significant technological and commercial progress we have achieved in recent months, thanks to our world-class team. It is also the first step in our international expansion, bringing quantum to our customers' fingertips – wherever they are in the world.'

OQC's patented 3D architecture, the Coaxmon, combines unparalleled scalability with world-class performance. This funding follows the February 2022 launch of OQC's latest system, Lucy, on Amazon Braket – making it the first European quantum



Bartłomiej K. Wroblewski/shutterstock

company on AWS. The company's technology can be used by businesses in financial services, pharma and logistics to increase their competitive edge.

Lenny Chin, principal at UTEC, says: 'Quantum computing promises to be the next frontier of innovation, and OQC, with its state-of-the-art Coaxmon technology, aims to integrate the forefront of modern physics into our everyday lives. UTEC is honoured to be part of OQC's mission of making quantum technology accessible to all and will support OQC's expansion into

Asia-Pacific through collaborations with academia, including the University of Tokyo, and partnerships with Japan's leading financial and tech corporations.'

This Series A investment will be used to scale the company's quantum systems and its private quantum computing-as-a-service offering. It will also consolidate OQC's position in Europe and accelerate international expansion in the Asia-Pacific region, including the Japanese market – a hotspot for financial services eager to realise the advantages of quantum computing.

## MODELLING AND SIMULATION

## German start-up MachineWare announces RISC-V simulator

MachineWare aims to disrupt semiconductor design with the introduction of its high-speed functional simulator, SIM-V.

Headquartered in Aachen and emerging from stealth mode in May, MachineWare's SIM-V combines simulation performance with customisability, for applications ranging from embedded devices to supercomputers. SIM-V enables software developers to test full software stacks – including firmware, operating system kernel and complex user-space applications, such as (Java-) virtual machines or rich graphical environments – in real time.

Lukas Jünger, MachineWare

Managing Director and co-founder, says: 'Our mission is to equip RISC-V software developers with the tools they need to deliver safe and secure software stacks on-schedule and glitch-free.'

Today's hardware-software systems are becoming increasingly complex – with even tiny edge systems executing millions of lines of code. SIM-V allows software developers to interactively debug the most complex designs without the need for physical hardware – even before first prototypes are available. Integrating SIM-V into continuous integration systems minimises test execution times, saves compute resources and allows developers to continue their



Dmitriy Rybin/shutterstock

work sooner. 'Human errors are unavoidable and critical bugs that compromise system safety and security are bound to appear in every project,' adds Lukas.

'Correct system functionality can only be ensured through extensive testing and rigorous

verification. However, automated, cross-architecture continuous-integration systems are still a major resource drain on many software teams. With SIM-V, complex test suites can be set up, executed and scaled up, all before getting near the hardware.'



# Suppliers' Directory

Find the suppliers you need quickly and easily  
[www.scientific-computing.com/suppliers](http://www.scientific-computing.com/suppliers)

## HIGH-PERFORMANCE COMPUTING

**Altair Engineering, Inc**  
 (248) 614-2400  
[info@altair.com](mailto:info@altair.com)  
[www.altair.com](http://www.altair.com)  
 ■ Software

**Boston Limited**  
 +44 (0)1727 876 100  
[sales@boston.co.uk](mailto:sales@boston.co.uk)  
[www.boston.co.uk](http://www.boston.co.uk)  
 ■ Systems Integrator

**Calyos SA**  
 0032 4 75 73 54 65  
[newbusiness@calyos-tm.com](mailto:newbusiness@calyos-tm.com)  
[www.calyos-tm.com](http://www.calyos-tm.com)  
 ■ Cooling

**Cluster Vision**  
 +31 20 407 7550  
[marketing@clustervision.com](mailto:marketing@clustervision.com)  
[www.clustervision.com](http://www.clustervision.com)  
 ■ Software

**Eurotech Spa**  
 +39 0433 485 411  
[welcome@eurotech.com](mailto:welcome@eurotech.com)  
[www.eurotech.com](http://www.eurotech.com)  
 ■ Systems Integrator

**Numerical Algorithms Group (NAG)**  
 +44 (0)1865 511 245  
[katie.ohare@nag.co.uk](mailto:katie.ohare@nag.co.uk)  
[www.nag.co.uk](http://www.nag.co.uk)  
 ■ Software

**OCF**  
 +44 (0) 114 257 2200  
[scw@ocf.co.uk](mailto:scw@ocf.co.uk)  
[www.ocf.co.uk](http://www.ocf.co.uk)  
 ■ Systems Integrator

## LABORATORY INFORMATICS

**Abbott Informatics**  
 +44 (0)161 711 0340  
[AlSalesUK@abbott.com](mailto:AlSalesUK@abbott.com)  
[www.abbottinformatics.com](http://www.abbottinformatics.com)  
 ■ LIMS  
 ■ Scientific Document Management Systems

**Accelerated Technology Laboratories**  
 +1 910 673 8165 (outside of US)  
 800 565 LIMS (5467) US and Canada  
[info@atlab.com](mailto:info@atlab.com)  
[www.atlab.com](http://www.atlab.com)  
 ■ Bioinformatics  
 ■ LIMS

**ACD/Labs**  
 +1 (416) 368-3435  
[info@acdlabs.com](mailto:info@acdlabs.com)  
[www.acdlabs.com](http://www.acdlabs.com)  
 ■ Data management  
 ■ Drug discovery  
 ■ Analysis tools

**Amphora Research Systems**  
 +44 (0)845 230 0160  
[info@amphora-research.com](mailto:info@amphora-research.com)  
[www.amphora-research.com](http://www.amphora-research.com)  
 ■ Electronic Laboratory Notebooks

**Autoscribe**  
 +44 (0) 118 984 0610  
[info@autoscribeinformatics.com](mailto:info@autoscribeinformatics.com)  
[www.autoscribeinformatics.com](http://www.autoscribeinformatics.com)  
 ■ Data Management  
 ■ Electronic Laboratory Notebooks  
 ■ Scientific Document Management Systems

**Bioviva**  
 +1 858 799 5000  
[www.3ds.com/products-services/bioviva](http://www.3ds.com/products-services/bioviva)  
 ■ LIMS  
 ■ Electronic Laboratory Notebooks

**Eusoft**  
 +39 080 5426799  
[info@eusoft.it](mailto:info@eusoft.it)  
[www.eusoft.it](http://www.eusoft.it)  
 ■ LIMS

**IDBS**  
[info@idbs.com](mailto:info@idbs.com)  
[www.idbs.com](http://www.idbs.com)  
 ■ Electronic Laboratory Notebooks

**InfoChem**  
 +49 89 2153 7160  
[info@infochem.de](mailto:info@infochem.de)  
[www.infochem.de](http://www.infochem.de)  
 ■ Cheminformatics

**LabWare**  
[infoEU@labware.com](mailto:infoEU@labware.com)  
[www.labware.com](http://www.labware.com)  
 ■ LIMS

**Osthus GmbH**  
 +49 241 943 140  
[office@osthus.de](mailto:office@osthus.de)  
[www.osthus.de](http://www.osthus.de)  
 ■ Electronic Laboratory Notebooks  
 ■ Cheminformatics  
 ■ LIMS

**Siemens**  
 +1 322 536 2139  
[industrial-it.swe@siemens.com](mailto:industrial-it.swe@siemens.com)  
[www.siemens.com/simaticit-rdsuite](http://www.siemens.com/simaticit-rdsuite)  
[www.siemens.com/industrial-it/lims](http://www.siemens.com/industrial-it/lims)  
 ■ Electronic Laboratory Notebooks  
 ■ LIMS

**Thermo Fisher Scientific**  
 +44 (0)161 942 3000  
[marketing.informatics@thermofisher.com](mailto:marketing.informatics@thermofisher.com)  
[thermofisher.com](http://thermofisher.com)  
[www.thermoscientific.com/informatics](http://www.thermoscientific.com/informatics)  
 ■ LIMS  
 ■ Scientific Document Management Systems  
 ■ Chromatography Data Systems  
 ■ Spectroscopy Software

## MODELLING AND SIMULATION

**Altair Engineering, Inc**  
 (248) 614-2400  
[info@altair.com](mailto:info@altair.com)  
[www.altair.com](http://www.altair.com)  
 ■ Software



**Esteco**  
 +39 040 3755548  
[marketing@esteco.com](mailto:marketing@esteco.com)  
[www.esteco.com](http://www.esteco.com)

ESTECO is an independent software provider, highly specialized in numerical optimization and simulation data management with a sound scientific foundation and a flexible approach to customer needs.



**ACD/Labs**  
 +1 (416) 368-3435  
[info@acdlabs.com](mailto:info@acdlabs.com)  
[www.acdlabs.com](http://www.acdlabs.com)

ACD/Labs is a leading provider of scientific software for R&D. We help our customers assemble digitized analytical, structural, and molecular information for effective decision-making, problem solving, and product lifecycle control.



**Osthus GmbH**  
 +49 241 943 140  
[office@osthus.de](mailto:office@osthus.de)  
[www.osthus.de](http://www.osthus.de)

Our experts in life science and industrial R&D help you to transform information into knowledge by custom data integration projects for LIMS/ELN/Bio&Chem Informatics.



**OCF**  
 +44 (0) 114 257 2200  
[scw@ocf.co.uk](mailto:scw@ocf.co.uk)  
[www.ocf.co.uk](http://www.ocf.co.uk)

At OCF we work with you to meet your high performance data processing, data management and data storage needs

# Subscribe for free\*

**SCIENTIFIC  
COMPUTING  
WORLD**

## Do you compute?


The only global publication for scientists and engineers using computing and software in their daily work



**Do you subscribe? Register for free now!**  
[scientific-computing.com/subscribe](https://scientific-computing.com/subscribe)

\*Registration required

**Also published  
by Europa  
Science**



# LABWARE®

Automate Your Laboratory  
with the Global Leader for  
**LIMS** and **ELN**

[www.labware.com](http://www.labware.com)

