

## Welcome

**Prof. Peter V Coveney**  
Principal Investigator &  
Coordinator



I am delighted to welcome you to the inaugural issue of the Verified Exascale Computing for Multiscale Applications (VECMA) Newsletter. In the seven months that have elapsed since the beginning of this exciting project, much has taken place and so much more is yet to come.

Our kick-off meeting in July attracted over 30 participants with contributing talks pertinent to applications in areas such as refugee migrations, climate change, fusion, medical and advanced functional materials. Indeed, in an era where science is afflicted by a “reproducibility crisis”, with increasing concern that most contemporary published research findings are unreliable, VECMA aims to create a unified European Verification, Validation, and Uncertainty Quantification (VVUQ) toolkit for exascale computing which will facilitate the adoption of numerical simulations as trusted tools of decision-making in areas of application as disparate as the above.

Our presence is already noticeable through journal publications, contributions to international conferences, and participation in other events. Last autumn we spoke at the Materials Research Society Meeting (MRS) in Boston, the leading conference in Materials Science across the world, and also at SC18,

the International Conference for High Performance Computing, Networking, Storage, and Analysis. We are looking forward to exhibiting VECMA at the upcoming EOSC-hub week in Malaga in April, at the International Conference on Computational Science (ICCS) in Faro in June, at the ISC 2019 in Frankfurt, and at the ECCOMAS Conference on Multi-Uncertainty, Multi-Physics and Multi-Scale Modelling in Porto in July. A full list of events in which we are taking part can be found at <https://www.vecma.eu/events/>.

Through our Partners, we have at our disposal the computational resources of major HPC centres for tests, pre-production, and production runs for the VECMA Applications and as a testing ground for the VECMA VVUQ Toolkit. We actively participate in collaborations with other EC H2020 ventures, most closely with the CompBioMed Centre of Excellence ([www.compbiomed.eu](http://www.compbiomed.eu)).

I feel that this project, having surmounted many of its early challenges to swiftly reach a fully productive phase, is poised to move forward in capable hands. I hope you will enjoy reading this Newsletter and I encourage you to get in touch with us about topics of interest or a potential participation in our endeavour.

## Verified Exascale Computing for Multiscale Applications

Supercomputers run powerful simulations that model some of the greatest challenges scientists are decoding, from achieving nuclear fusion through finding safer drugs to predicting future climate change. The power of accurate prediction requires multiscale calculation of vast amounts of data. But to confirm reliability so researchers and end-users can have confidence in their results, computer simulations have to both correctly represent processes and accurately quantify the uncertainties associated with their calculations.

The VECMA Project aims to raise the bar of confidence in computer simulations. Its purpose is to enable a representative range of diverse multiscale applications to run on current multi-petascale computers and future exascale environments in such a manner that their output is deemed to be of high fidelity and thus “actionable”. That is, the calculations and simulations are certified as validated, verified, and equipped with quantified error bars, so that they may be relied upon to take important decisions in the domains of concern. In short, we develop generic VVUQ approaches for multiscale applications on current tier-0 HPC and emerging exascale

machines.

The central deliverable will be an open source toolkit for multiscale VVUQ based on generic multiscale VV and UQ primitives, to be released in stages over the lifetime of this project, fully tested and evaluated in emerging exascale environments, actively promoted over the lifetime of this project, and made widely available in European HPC centres.

VECMA runs until June 2021 and involves nine academic and industrial partners led by University College London ([cordis.europa.eu/project/rcn/215829/factsheet/en](http://cordis.europa.eu/project/rcn/215829/factsheet/en)).

Our website ([www.vecma.eu](http://www.vecma.eu)) is full of all the latest news and information about VECMA, its Partners, events, publications, and more. We have an active presence and growing following on Twitter ([@VECMA4](https://twitter.com/VECMA4)). We are funded by the European Commission's (EC) Future and Emerging Technologies (FET) programme ([ec.europa.eu/programmes/horizon2020/en/h2020-section/future-and-emerging-technologies](http://ec.europa.eu/programmes/horizon2020/en/h2020-section/future-and-emerging-technologies)) under grant no. 800925.



# Project Implementation: Work Packages, Advisory Boards, Partners, Deliverables

## Work Packages

To achieve our previously described ambitions we have developed a concise work plan consisting of six Work Packages (WP):

**WP1 - Management:** Overseeing project on a day-to-day and strategic basis, including direction of and liaison between Partners, as well as interfacing with the EC and external bodies.

**WP2 - Algorithms & Formalisms:** Theoretical and algorithmic underpinning for VVUQ methods, including the design of the UQPs and VVPs.

**WP3 - Software Implementation:** Implementation of WP2 in the software which comprises the VVUQ Toolkit. Staged re-

lease of these tools over the lifetime of VECMA.

**WP4 - Applications:** Utilization of algorithms, tools, and services developed in WP2, WP3, and WP5 in multiscale applications.

**WP5 - Infrastructure:** Infrastructure of high performance computing environments on which VVUQ-enabled applications will run, in the current multi-petascale and the future exascale eras.

**WP6 - Dissemination & Exploitation:** Dissemination of VECMA, including full scale promotion to accompany the various software releases across a wide range of scientific and social scientific domains.

## Advisory Boards

The **Scientific Advisory Board (SAB)** oversees scientific excellence and outreach, and consists of individuals external to VECMA, internationally recognised, and active in different scientific domains. Current membership comprises:

- **Dr Rob Akers**, UKAEA
- **Prof. Peter Challenor**, University of Exeter
- **Dr Mariano Vazquez**, Barcelona Supercomputing Centre (BSC)
- **Prof. Petros Koumoutsakos**, Eidgenössische Technische Hochschule Zürich (ETH)
- **Prof. Shantenu Jha**, Rutgers University

The **Innovation Advisory Board (IAB)** oversees exploitation, economic, and societal impact, and consists of individuals external to VECMA, representing Government, NGOs, Industry, and Academia. Current membership comprises:

- **Prof. Beniamino Di Martino**, University of Campania
- **Prof. Pär Strand**, Chalmers University of Technology
- **Dr Javier Garcia-Blas**, University Carlos III of Madrid

## Partners



The Centre for Computational Science (CCS) at UCL is an internationally leading centre for computational science research using high performance computing. The CCS is currently comprised of ca 20 members and pursues a diverse range of research unified by common computational approaches, from theory and design of algorithms to implementations and middleware on internationally distributed HPC systems. UCL leads the overall project and will take a substantial role in WP1 and WP4.



UNIVERSITY  
OF AMSTERDAM

The University of Amsterdam (UvA) is an intellectual hub. It collaborates with hundreds of national and

international academic and research institutions, as well as businesses and public institutions. UvA has 7 faculties, 3000 academic staff members, and 30000 students, and is one of Europe's leading research universities. UvA will be WP2 leader and will bring in two applications (in the biomedical domain), and will contribute to multiscale UQ algorithms and UQPs.



Brunel  
University  
London

The Department of Computer Science at Brunel University London is an interdisciplinary centre that includes researchers with a range of backgrounds including computer science, engineering, mathematics, and psychology. UBRU will lead WP3 and provide one deep-track multi-scale application as part of WP4, bring substantial contributions to the specification of the primitives in WP2, and support activities in WP6.



The Leibniz Supercomputing Centre (Leibniz-Rechenzentrum, LRZ) is part of the Bavarian Academy of Sciences and Humanities (BADW) and has been an active player in the area of high performance computing (HPC) for over 20 years. BADW-LRZ operates SuperMUC, a top-level supercomputer with 155,000 x86 cores and a peak performance of over 3 PFlop/s, as well as a number of general purpose and specialized clusters and cloud resources. BADW-LRZ will contribute to the project by providing computing resources and support to project partners.



Bull SAS is the newest member of the Atos family, a leader in digital services with pro forma annual revenue of € 5.6 billion and 96,000 employees in 72 countries. Bull SAS will mainly participate in WP2, WP3, and WP5. Smaller participation is also planned in WP1, WP4, and WP6.



CBK Sci Con Limited is a consultancy that offers technical and management advice to business in e-science domains. CBK sits at the interface between academia and industry, and its main areas of focus include High Performance Computing and Modelling and Simulation across a number of sectors. CBK will lead WP6.



Centrum Wiskunde & Informatica (CWI) is the Netherlands' national research institute for mathematics and computer science. Its strength lies in discovering and developing new ideas that benefit society and the economy. CWI will contribute to WP2 (development of multiscale UQ algorithms and UQPs) and provide one application (Climate) to WP4.



The Max-Planck-Gesellschaft zur Förderung der Wissenschaften e.V (MPG) is represented in this proposal by the Max Planck Institute for Plasma Physics (IPP) which is one of the largest fusion research centres in Europe. IPP will lead Work Package 4 and bring in expertise on aspects of VVUQ.



PSNC leads WP5 and will be responsible for integration of applications, tools, and components from WP2, WP3, and WP4, providing HPC infrastructure for the deployment of the consistent VECMA platform, validation with end-users in real-world scenarios, evaluating the platform against established benchmarks, and providing feedback to all WPs in order to introduce changes in the software components researched and developed.

## Deliverables

### Month 3 (September 2018) public releases:

- **D1.1 Project Handbook:** Describes the organisation and internal procedures of VECMA with regard to day-to-day communication and progress towards timely production of deliverables and within budget.
- **D1.2 Quality Assurance Plan:** Establishes the management infrastructure for efficient and constant monitoring and operation of the day-to-day project activities and deliverables before finalization.
- **D6.1 Detailed Dissemination Action Plan:** Acts as a detailed and comprehensive report on the dissemination actions carried out by the project.

### Month 6 (December 2018) public releases:

- **D1.3 Data Management Plan:** The Data Management Plan acts as a detailed and comprehensive document on the data management plan that is being followed to guide the use of various types of data by the project.
- **D6.2 Innovation Plan:** The Innovation Plan is designed to promote interdisciplinary entrepreneurial opportunities within the research activities of the VECMA project, from invention through to exploitation.

All deliverables at public dissemination level are available at [www.vecma.eu/about/deliverables/](http://www.vecma.eu/about/deliverables/).

## News and Updates

### Multi-Peta/Exascale infrastructure

We are pleased to announce that UCL CCS has been awarded an allocation of 25,000 node hours on Summit, the current number one in the Top 500 list of supercomputers, operating at a peak of ca 200 petaFLOPS (when UK flagship supercomputer ARCHER operates at a peak of 2.55 petaFLOPS). This is the equivalent, to about 15 million core hours on Titan, the former number one in the Top 500, which CCS have been

using for over a year. The access is given inter alia to extend the INSPIRE DOE INCITE project (<https://www.combiomed.eu/department-of-energy-incite-supercomputing-award/>), designed to help determine drugs to administer to cancer patients whose target proteins have developed resistance to the first level drugs they have been treated with.

We are also pleased to announce that our supercomputing Partners have committed the following resources specifically for the VECMA Toolkit & Applications:

- PSNC: minimum 5 million CPU hours on its biggest HPC system, "Eagle", comprising 33,000 cores, with 300TB of

memory, and operating at 1.4 petaFLOPS.

- LRZ: minimum 6 million CPU hours on SuperMUC, a PRACE Tier-0 system, supporting per-node power monitoring and reporting per-job energy-to-solution. Plans for access to SuperMUC NG are under way.
- Atos Bull: 640,000 CPU hours on two supercomputers.

## Toolkit Release: FabSim3 v0.6 & EasyVVUQ v0.1

The VECMA project has hit its milestone of internally releasing a version of the VECMA Toolkit. This internally-released version of the VECMA VVUQ Toolkit comprises two parts, the first release of the EasyVVUQ software and a new release of FabSim3 containing a number of new features. FabSim software facilitates computational research through automation on large-scale and distributed e-infrastructures, while EasyVVUQ brings features related to Validation, Verification, and Uncertainty Quantification in this iteration containing the ba-

sic architecture for creating uncertainty quantification pipelines for analysing simulation output, as well as basic statistical analysis capabilities.

Alpha users are currently testing and evaluating the new software. Mark your calendars for June 2019 where we plan to have the first full public and open release of the VECMA Toolkit.

## FabSim3

As of January 16th 2019, the FabSim3 GitHub repository is entirely public, allowing users to download any version of the code and to read and contribute to the various issues raised in the repository. Switching to open development has several major advantages according to the FabSim3 Development

Team, including continual user feedback and showcasing of this middleware's progress, and saving money on provisioning specific solutions for private repositories. FabSim3 is available at [github.com/djgroen/FabSim3](https://github.com/djgroen/FabSim3).



**HiDALGO**

The new H2020-funded Centre of Excellence project HiDALGO ([hidalgo-project.eu/](http://hidalgo-project.eu/)) on HPC and Big Data Technologies for

Global Systems has kicked off on 1st December 2018 (€8m


total funding). The CoE addresses the importance of decision making in global, multi-dimensional problems, and is anticipated to provide a significant boost to the development of the multiscale migration application at Brunel University.

## Upcoming Events

VECMA Bootcamp  
13-14 February 2019  
Brunel University, London UK

EOSC-hub Week 2019  
10-12 April  
Prague, Czech Republic

All-Hands Meeting  
9-10 May 2019  
Amsterdam, Netherlands



In preparation for first major release of the VECMA toolkit (June 2019), all the project developers are meeting for a concentrated development session. The aim of the Bootcamp is to ensure that all of fast track codes are stable and contain the features needed to support our internal applications. We hope to follow up with a Hackathon later in the year to bring our VVUQ tools to external users applications.



Present and future service providers and users of the European Open Science Cloud (EOSC) come together at this event, during which an open call will be launched. Stronger involvement from the ESFRI communities, the HPC Centers of Excellence and the public sector will be sought through three interactive sessions dedicated to discuss how EOSC can make a difference for these different target groups. Industry



The VECMA All-Hands Meeting (AHM) is our largest meeting, to which all members, Core, and Associate Partners are invited. It occurs on an annual basis, and this year it will be organised at the University of Amsterdam (UvA), a VECMA Core Partner and WP2 leader. Inter-WP meetings of a more in-depth approach will prelude the AHM on 8 May 2019.