



Introducing VECMAtk - Verification, Validation and Uncertainty Quantification for Multiscale and HPC Simulations

Derek Groen^{1(✉)}, Robin A. Richardson², David W. Wright², Vytautas Jancauskas¹⁰, Robert Sinclair², Paul Karlshoefler⁹, Maxime Vassaux², Hamid Arabnejad¹, Tomasz Piontek⁵, Piotr Kopta⁵, Bartosz Bosak⁵, Jalal Lakhilil⁴, Olivier Hoenen⁴, Diana Suleimenova¹, Wouter Edeling⁶, Daan Crommelin^{6,7}, Anna Nikishova⁸, and Peter V. Coveney^{2,3,8}

¹ Department of Computer Science, Brunel University London, London, UK

Derek.Groen@brunel.ac.uk

² Centre for Computational Science, University College London, London, UK

³ Centre for Mathematics and Physics in the Life Sciences and Experimental Biology, University College London, London, UK

⁴ Max-Planck Institute for Plasma Physics - Garching, Munich, Germany

⁵ Poznań Supercomputing and Networking Center, Poznań, Poland

⁶ Centrum Wiskunde & Informatica, Amsterdam, The Netherlands

⁷ Korteweg-de Vries Institute for Mathematics, Amsterdam, The Netherlands

⁸ University of Amsterdam, Amsterdam, The Netherlands

⁹ BULL/ATOS, Paris, France

¹⁰ LRZ, Garching, Germany

Abstract. Multiscale simulations are an essential computational method in a range of research disciplines, and provide unprecedented levels of scientific insight at a tractable cost in terms of effort and compute resources. To provide this, we need such simulations to produce results that are both robust and actionable. The VECMA toolkit (VECMAtk), which is officially released in conjunction with the present paper, establishes a platform to achieve this by exposing patterns for verification, validation and uncertainty quantification (VVUQ). These patterns can be combined to capture complex scenarios, applied to applications in disparate domains, and used to run multiscale simulations on any desktop, cluster or supercomputing platform.

Keywords: Multiscale simulations · Verification · Validation · Uncertainty quantification

1 Introduction

The overarching goal of computational modeling is to provide insight into questions that in the past could only be addressed by costly experimentation, if at all. In order for the results of computational science to impact decision making

outside of basic science, for example in industrial or clinical settings, it is vital that they are accompanied by a robust understanding of their degree of validity. In practice, this can be decomposed into checks of whether the codes employed are solving equations in an accurate manner (*verification*), solving the correct equations to begin with (*validation*), and providing estimates that comprehensively capture uncertainty (*uncertainty quantification*) [1,2]. These processes, collectively known as VVUQ, provide the basis for determining our level of trust in any given model and the results obtained using it [3]. Recent advances in the scale of computational resources available, and the algorithms designed to exploit them mean that it is increasingly possible to conduct the additional sampling required by VVUQ even for highly complex calculations and workflows. The goal of the VECMA project (www.vecma.eu) is to provide an open source toolkit containing a wide range of tools to facilitate the use of VVUQ techniques in multiscale, multiphysics applications. At this initial stage, these range from fusion and advanced materials through climate and forced population displacement, to drug discovery and personalised medicine.

Multiscale computing presents particular difficulties as such applications frequently consist of complex workflows where uncertainty propagates through highly varied components, some of which may only be executed conditionally. Additionally, uncertainties may be associated with the process of transforming the output variables from one scale to another, e.g. coarse to fine scale or vice versa. Although a wide range of toolkits exist to facilitate multiscale computing (see the review by Groen et al. [4]), applying rigorous VVUQ is a major challenge that still needs to be addressed in this domain. The goal of the VECMA toolkit (VECMAtk) is to provide open source tools which implement VVUQ approaches that range from those which treat components or workflows as an immutable ‘black box’ to semi-intrusive methods in which components of the workflow may be replaced by statistically representative, but cheaper, surrogate models [5].

Key to VECMA’s approach is an understanding of the growing size and diversity of available supercomputers as we move to the exascale [6]. For the first time, the vast number of cores available on modern systems make it conceivable for researchers to execute the ensembles necessary to sample phase space for VVUQ analyses concurrently for even very computationally intensive simulations. Moreover, the use of ensembles of simulations offers an efficient way to use large supercomputers. However, effectively running and managing multiscale workflows composed from components with divergent computational requirements (for example with some models capable of exploiting GPUs and others not) represents a key challenge our tools must address. Nonetheless, advances in hardware continue to present opportunities that can be exploited. For example, some intrusive VVUQ methods which make extensive use of machine learning techniques are able to take advantage of hardware designed to accelerate these techniques (such as fast SSD storage). The support for compute intensive applications and workflow management within VVUQ analyses are key factors motivating the development of the VECMAtk.

The philosophy of VECMAtk is that, as far as possible, we want to be able to add VVUQ to existing scientific workflows without changing researchers conception of the problem they are investigating. This goal informs the conception of the toolkit which is designed (i) to be highly modular, (ii) have minimal installation requirements and, (iii) provide control over where (locally or on remote resources) and at what time analysis takes place.

This means that VECMAtk is a collection of elements that can be reused and recombined in different workflows. Our aim is to define stable interfaces, data formats and APIs that facilitate VVUQ in the widest range of applications. The modularity of the toolkit allows us to account for the differing requirements of VVUQ sampling and analysis and provide tools tailored for the needs of both steps of the process. Several software packages or libraries are already available for performing VVUQ (such as OpenTurns [7], UQLab [8], Uncertainty [9], Chaospy [10], etc.), but in many cases these are closed source and none of them provide the separation of concerns needed to allow the analysis of both small local computations and highly compute intensive kernels (potentially using many thousands of cores and GPUs on HPC or cloud resources). We aim to reuse existing tools where appropriate to provide robust and optimised code for sampling and analysis.

To enable researchers to use the toolkit and derive custom approaches, we will provide a range of exemplar workflows and documentation. Essential to the design of these exemplars is the need to lower the barriers to usage and, naturally, to be transferable to a large range of applications with minimal modification.

2 The VECMA Toolkit

The goal of the toolkit is to facilitate the process of researchers mapping the requirements of VVUQ to their specific scientific problem and existing workflow(s). The main factors that shape the development approach of VECMAtk are: (a) the need to fit into existing applications with minimal modification effort, (b) the wide range of target application domains, (c) the flexible and recombinable nature of the toolkit itself, (d) the geographically distributed nature of the users and particularly the developers.

2.1 Development and Prototyping Process

To support these needs and characteristics, we have chosen to adopt an evolutionary prototyping approach. In this approach, a existing application developers initially establish VVUQ techniques using their own scripts, which they share with the VECMAtk developers together with additional needs that they have not been able to easily address themselves. These scripts and requirements, together with existing libraries, form the base from which the development team develops initial prototypes. The prototypes are then tested and refined at frequent intervals, with the user feedback and integration testing guiding further developments. As a result, some prototypes are reduced in scope, simplified, or removed

altogether; and some prototypes are being refined into more advanced, flexible, scalable and robust tools. Regular development meetings within the project help to monitor and disseminate progress as well as providing a venue in which we ensure that all component development teams are following best development practices (for example making use of version control and continuous integration). Although we work closely with a group of application developers, both FabSim3 and EasyVVUQ are publicly available, and anyone can independently install, use and modify these tools to suit their own purposes.

As part of our prototyping process, we identify common workflow patterns and software elements (for example those needed to encode complex parameter distributions) that can be abstracted for re-use in a wide range of application scenarios. We label patterns found in verification and validation contexts *VVPs*, and those associated with uncertainty quantification or sensitivity analysis *UQPs*. The definition of a VVP or UQP should never require the use of any specific execution management platform, as the toolkit is envisioned to provide multiple solutions that facilitate workflows. Examples include diverse sampling algorithms and job types running on heterogeneous resources. Within VECMAtk, we categorize procedures that treat underlying applications as a black box (*non-intrusive*), that account for the coupling mechanisms between sub-models (*semi-intrusive*) or the algorithms of the submodels (*fully intrusive*).

2.2 Key Components

QCG-Broker/Computing: Easy and efficient access to computing power is crucial when a single run of an application is demanding or a large number of application replicas has to be executed to guarantee reliable VVUQ. To fulfil this requirement, VECMAtk uses QCG¹ which provides advanced capabilities for the unified execution of complex jobs across single or multiple computing clusters. The QCG infrastructure, which is presented in Fig. 1, uses the QCG-Broker service to manage execution of computational experiments, e.g. through multi-criteria selection of resources, while several QCG-Computing services offer unified remote access to underlying resources. The QCG services can be accessed with numerous user-level tools [11], of which a few examples are provided in the aforementioned figure.

FabSim3: The combination of different UQPs and VVPs in one application also leads to a cognitively complex workflow structure, where different sets of replicas need to be constructed, organized, executed, analyzed, and actioned upon (i.e. triggering subsequent execution and/or analysis activities). FabSim3 [12]² is a freely available tool that supports the construction, curation and execution of these complex workflows, and allows users to invoke them on remote resources using one-liner commands. In contrast to its direct predecessor FabSim, FabSim3 inherently supports the execution of job ensembles, and provides a plug-in system

¹ <http://www.qoscosgrid.org>.

² <https://github.com/djgroen/FabSim3>.

which allows users to customize the toolkit in a modular and lightweight manner (e.g., evidenced by the minimalist open-source FabDummy plugin). We provide an overview of the FabSim3 architecture in Fig. 3. In the context of VECMA FabSim3 plays a key role in introducing application-specific information in the Execution Layer, and in conveniently combining different UQPs and VVPs.

QCG Pilot Job: A *Pilot Job*, is a container for many subjobs that can be started and managed without having to wait individually for resources to become available. Once the Pilot Job is submitted, it may service a number of defined VVUQ subtasks (as defined by e.g. EasyVVUQ or FabSim3). The QCG Pilot Job mechanism provides two interfaces that may be used interchangeably. The first one allows to specify a file with the description of sub-jobs and execute the scenario in a batch-like mode, conveniently supporting static scenarios. The second interface is offered with the REST API and it can be accessed remotely in a more dynamic way. It will be used to support scenarios where a number of replicas and their complexity dynamically changes at application runtime.

EasyVVUQ is a Python library, developed specifically for VECMA, designed to simplify the implementation of creation of (primarily blackbox) VVUQ workflows involving existing applications. The library is designed around a breakdown of such workflows into four distinct stages; sampling, simulation execution, result aggregation, and analysis. The execution step is deemed beyond the remit of the package (it can be handled for instance by FabSim3 or QCG Client), whilst the other three stages are handled separately. A common data structure, the *Campaign*, which contains information on the application being analyzed alongside the runs mandated by the sampling algorithm being employed, is used to transfer information between each stage. The architecture of EasyVVUQ is shown in Fig. 2.

The user provides a description of the model parameters and how they might vary in the sampling phase of the VVUQ pattern, for example specifying the distribution from which they should be drawn and physically acceptable limits on their value. This is used to define a *Sampler* which populates the Campaign with a set of run specifications based on the parameter description provided by the user. The Sampler may employ one of a range of algorithms such as the Monte Carlo or Quasi Monte Carlo approaches [13]. At this point all of the information is generic in the sense that it is not specific to any application or workflow. The role of the *Encoder* is to create input files which can be used in a specific application. Included in the base application is a simple templating system in which values from the Campaign are substituted into a text input file. For many applications it is envisioned that specific encoders will be needed and the framework of EasyVVUQ means that any class derived from a generic Encoder base class is picked up and may be used. This enables EasyVVUQ to be easily extended for new applications by experienced users.

The simulation input is then used externally to the library to execute the simulations. The role of the *Decoder* is twofold, to record simulation completion in the Campaign and to extract the output information from the simulation runs. Similarly to the *Encoder* the *Decoder* is designed to be user extendable to

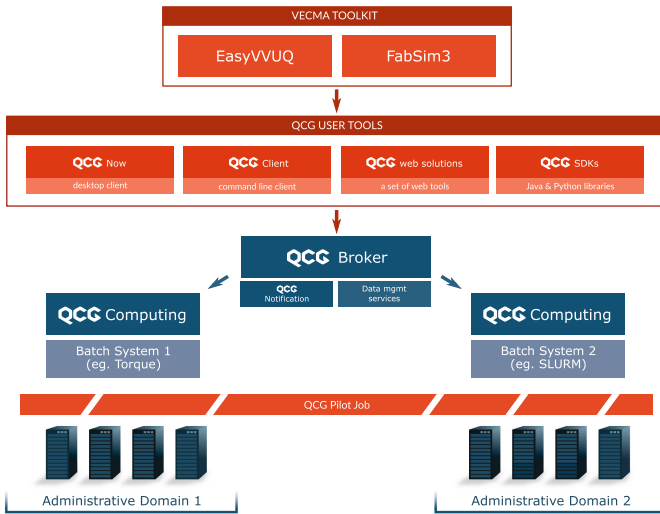


Fig. 1. Simplified overview of QCG usage in VECMA. Jobs requested by the toolkit layer may be farmed out to one or more queues on different computing resources.

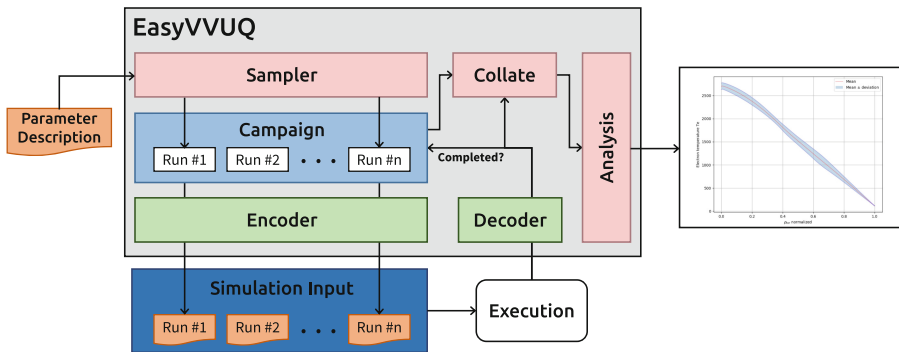


Fig. 2. The architecture of EasyVVUQ. The UQ workflow is split into a sampling, execution and analysis stage, orchestrated via a (persistent) ‘Campaign’ object.

facilitate analysis of a wide range of applications. The Decoder is used in the collation step to provide a combined and generic expression of the simulation output for further analysis (for example the default is to bring together output from all simulation runs in a Pandas dataframe). Following the output collation we provide a range of analysis algorithms which produce the final output data. Whilst the library was originally designed for acyclic ‘blackbox’ VVUQ workflows, development is ongoing to allow the library to be used in more complex patterns.

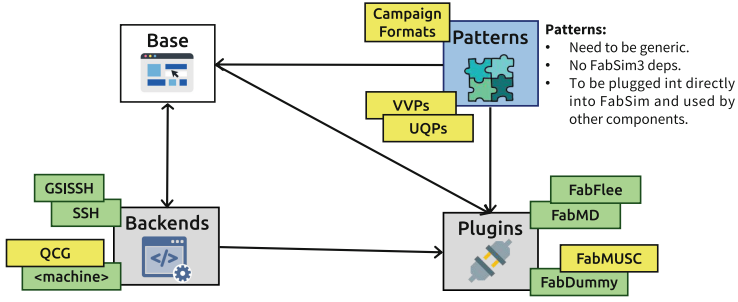


Fig. 3. Essential building blocks in the FabSim3 component of VECMATk, and their interdependencies. Components in yellow are under development as of June 2019

2.3 How the Components Work Together

The components in VECMATk (FabSim3, EasyVVUQ, QCG Pilot Job and other QCG components) can be combined in a variety of ways, enabling users to combine their added values while retaining a limited deployment footprint. At time of writing, we are working on the following combinations:

- FabSim3 has been integrated with QCG-Client to enable job submission to QCG Broker, allowing users to schedule their jobs across multiple remote (QCG-supporting) machines.
- EasyVVUQ can use FabSim3 to facilitate automated execution. Users can convert their EasyVVUQ campaigns to FabSim3 ensembles using a one-liner (`campaign2ensemble`), and FabSim3 output is ordered such that it can be directly moved to EasyVVUQ for further decoding and analysis.
- Integration between QCG Pilot Job and FabSim3 is under way, enabling users to create and manage pilot jobs using FabSim3 automation.
- Integration between QCG Pilot Job and EasyVVUQ is under way, enabling EasyVVUQ users to execute their tasks directly using pilot jobs.

3 Initial Applications

The following section gives an overview of applications that currently use the VECMATk and are guiding the development of new features for the toolkit. We present a detailed look at a fusion calculation and brief description of climate, population displacement, materials, force field, and cardiovascular modeling.

3.1 Fusion Example

Heat and particle transport in a fusion device play a major role in the performance of the thermo-nuclear fusion reaction. Current understanding is that turbulence arising at the micro space and time scales is a key factor for such transport which has effects on much larger scales. Multiscale simulation are

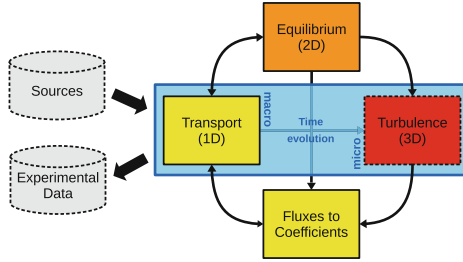


Fig. 4. Schematic view of the targeted fusion application.

developed to bridge such disparate spatiotemporal scales, for instance Luk et al. [14] couple a gyro-fluid 3D turbulence submodel to a 1D transport solver that evolves the temperature profiles over the macro scales. The turbulence submodel provides heat fluxes from which the transport coefficients are derived, but its output is inherently “noisy”. Hence, the calculated profiles are exposed to this noisy signal, and uncertainties will propagate from one submodel to the next. Additional uncertainties come from external sources and from experimental measurements against which the simulation could be validated. This leads to a complex scenario as depicted in Fig. 4.

The goal here is to produce simulated quantities of interest (temperature profile, density, etc.) and their confidence intervals and propagate this additional information through a complex cyclic workflow that involves several components with different properties, computational costs and uncertain inputs. The confidence intervals allow for better validation of the interpretative simulation against experimental results (for existing tokamaks machines), and give insight into the confidence of predictive simulations.

Following previous work on quantifying the propagation of uncertainty through a ‘black-box’ model in fusion plasma [15], we started to apply a polynomial chaos expansion method (PCE) to the cheaper single-scale models from Fig. 4, but taken separately (each one becomes a black-box). The method doesn’t require changes to the underlying model equations and provides a quantitative measure for which combinations of inputs have the most important impact on the results. The PCE coefficients are used to compute statistical metrics that are essential for the basic descriptions of the output distribution. To compute those coefficients we use a quadrature scheme which depends on the polynomial type which itself depends on the probability distribution of uncertain inputs. As a result, integration rules along each axis can be calculated using a tensor product.

Even with a limited number of uncertain parameters resulting from the turbulence code (e.g 8 for a fluid code using a flux-tube approximation), and assuming these parameters are not correlated, this method will necessitate approximately 1.7 millions of runs of *Fluxes to Coefficients* and *Transport* codes if we want to calculate their propagation through these models. As these are the cheapest

codes in the application, this represents only 512 CPU-hours of an embarrassingly parallel job which remains tractable with traditional means. But when the complexity of the models and the number of parameters increase, such quantities of runs with potentially large range of run times becomes very challenging and will require using advanced capabilities of ‘smart’ pilot-job software.

3.2 Variety of Other Applications

Climate. Computational models for atmospheric and oceanic flows are central to climate modeling and weather forecasting. As these models have finite resolutions, they employ simplified representations, so-called parameterizations, to account for the impact of unresolved processes on larger scales. An example is the treatment of atmospheric convection and cloud processes, which are important for the atmospheric circulation and hydrological cycle, but are unresolved in global atmospheric models. These parameterizations are a source of uncertainties: they have parameters that can be difficult to determine, and even their structural form can be uncertain.

A computationally very expensive approach to improve parameterizations and reduce uncertainty is by locally replacing the parameterization with a high-resolution simulation. In [16] this is applied regionally, by nesting the Dutch Atmospheric Large-Eddy Simulation (DALES) model in a selected number of global model columns, replacing the convection parameterization in these columns. The local DALES models run independently from each other and only exchange mean profiles with the global model. While this set-up allows for the use of massively parallel computer systems, running a cloud-resolving simulation in every single model column of a global model remains computationally unfeasible.

Within VECMA we are therefore developing tools for surrogate modeling. More specifically, we aim for statistically representative, data-driven surrogate models that account for the uncertainties in subgrid-scale responses due to their internal nonlinear and possibly chaotic dynamics. The surrogates are to be constructed from a (limited) set of reference data, with the goal of accurately reproducing long-term statistics, in line with the approach from [17, 18].

Forced Displacement. Accurate predictions of forced population displacement can help governments and NGO’s in making decisions as to how to help refugees, and efficiently allocate humanitarian resources to overcome unintended consequences [19]. We enable these simulations by establishing an automated agent-based modeling approach, FabFlee, which is a plugin to FabSim3 that uses the Flee agent-based simulation code. Flee forecasts the distribution of incoming refugees across destination camps [20], while FabFlee provides an environment to run and analyze simulation ensembles under various policy constraints, such as forced redirections, camp and border closures [21]. At time of writing, we are combining FabFlee with EasyVVUQ to more efficiently perform sensitivity analysis for varying agent awareness levels and speed limits of refugee movements [22].

Materials Modeling. Prediction of nanocomposite material properties requires multiscale workflows that capture mechanisms at every characteristic scale of the material, from chemical specificity to engineering testing conditions. For nanocomposite systems, the characteristic time and length of its nanostructure and macrostructure are so far apart, their respective dynamics can be simulated separately. We use DealLAMMPS [23,24], a new program that simulates the nanoscale with LAMMPS molecular dynamics, and the macroscale is simulated using deal.II, a finite element solver. Boundary information is passed from the FEM model to the LAMMPS simulations, the stresses arising from these changes are used to propagate the macroscale model. This workflow creates a vast number of short nanoscale, unpredictable at runtime, simulations at each macroscale timestep. Handling the execution of these task requires automation and coordination between several resources. As just the nanoscale errors due to uncertainty in the boundary condition and even the stress measured for a boundary change are correlated, complex and expensive to calculate. Over several iterations, errors can proliferate and careful consideration and understanding of these is needed from tools that VECMAtk can provide.

Molecular Dynamics Force Fields. Molecular dynamics calculations are used not only in materials modeling but in a wide range of other fields. Choices made in the design of these calculations such as the parameterization of the force field describing chemical components within the system and cut-offs used for long range interactions can have a profound effect on the results obtained and their variability. One field in which this of particular interest in free energy calculations which are increasingly widely used in modern drug design and refinement workflows. The binding affinity calculator (BAC) [25] we have developed automates this class of calculation, from model building through simulations and analysis. In order to understand the impact of forcefield parameter decisions on calculations performed using the BAC we are creating new workflows which incorporate sensitivity analysis through EasyVVUQ. The use of protocols based on ensemble simulations (known as TIES and ESMACS [26]) will give us the ability to adjust the simulation duration and ensemble size in order to robustly determine the uncertainty of results for comparison. This effort builds on previous work which uses Pilot Job manager to handle the execution of ensembles of multiple runs to enable bootstrap error statistics to be applied to calculations for individual protein-ligand pairs.

Cardiovascular Simulation. Haemodynamic simulations provide a non-invasive means of estimating flow rates, pressures and wall shear stresses in the human vasculature. Through MRI and CT scans, patient specific models may be used, with clinical applications such as predicting aneurysm rupture or treatment. We use a 3D lattice-Boltzmann solver, HemeLB [27], to simulate the continuum dynamics of bloodflow through large and highly sparse vascular systems efficiently [28]. A recent validation study focused on HemeLB simulations of a real patient Middle Cerebral Artery (MCA), using transcranial Doppler measurements of the blood velocity profile for comparison, as well as exploring the effects of a change in rheology model or inlet flow rate on the results [29]. We are

now running full 3D simulations of the entire human arterial tree, with an aim of also integrating the venous tree and a cyclic coupling to state-of-the-art heart models. This necessarily introduces an even greater number of input parameters which, when combined with the great computational expense of the submodels, presents a real challenge for validation and verification of our application, and particularly with regards to a computationally feasible sensitivity analysis and uncertainty quantification within such a system.

Biomedical Model. Coronary artery stenosis is a cardiovascular disease of narrowing of the coronary artery due to clustering of fatty plaque. A common treatment is to dent the fatty plaque into the artery wall and to deploy a stent in order to keep the artery open. However, up to 10% cases end up in the re-narrowing of the artery due to an excessive healing response, which is called in-stent restenosis (ISR) [30]. The multiscale in-stent restenosis model simulates this process at different time scales [31].

In [32], uncertainty of the response of the two-dimensional ISR model on the cross-sectional lumen area was estimated and analyzed. Uncertainty quantification showed up to 15% aleatory and about 25% total uncertainty in the model predictions. Additionally, sensitivity analysis identified the endothelium regeneration time as the most influential parameter on the model response. In [33], the semi-intrusive multiscale metamodeling uncertainty quantification method was applied to improve the performance of the uncertainty analysis. Depending on the surrogate used, the simulation time of the semi-intrusive method was up to five times faster than of a Monte Carlo method. In future work, the semi-intrusive metamodeling method will be applied to the three dimensional version of the ISR model employing the VECMAtk, since a black-box approach is not feasible for this application due to high computational demand [34].

4 Roadmap and Release Strategy

For the VECMA toolkit we have adopted a release schedule with two release types: Minor releases are tagged every 3 months, advertised within the project, made public without dedicated advertising, and with a limited amount of additional documentation and examples. Major releases will be made in June 2019, June 2020 and December 2020, and are public and fully advertised. They are accompanied with extensive documentation, examples, training events and dedicated uptake activities. We are at present able to guarantee formal support for the VECMA toolkit up to June 2021, and can future-proof VECMAtk at least until 6 months after the final planned release of the toolkit. The June 2019 VECMAtk release contains FabSim3, EasyVVUQ, and several functionalities provided by QCG. Later releases will likely feature additional components. In-between releases, users will be able to access the latest code for FabSim3 and EasyVVUQ, as we are maintaining an open development environment.

Containerization of VECMAtk. Containerization [35] is a method that allow us to create a virtual machine (VM) in more easily compared to the traditional

virtualization approach. These type of VMs, called *containers*, share the operating system kernel, as well as resources such as storage and networking. Containers matter because they enhance reproducibility: every run using them is guaranteed to have the same settings and configurations. Additionally, containers provide application platform portability, as they can be migrated to other computing environments without requiring code changes. Each container includes the run-time components – such as files, environment variables and libraries – necessary to run the desired software on a single control host, accessing a single kernel. Among the implementations available, we choose to use Docker (docs.docker.com).

In this work, we set up the FabSim GitHub repository with Travis-CI. After each successful test executed by Travis, a docker image is configured and built, and pushed into Docker Hub. For this work, A Docker image is provided as well through the Docker Hub³. And, the docker bundle for easy deployment is available in Docker Hub via: `docker pull vecmafabsim3/fabsimdocker`.

However, although Docker is one of most popular container technologies for software reproducibility, it has low adoption in the HPC world since it requires users with root access to run Docker and execute a containerized applications. To support high performance computing use cases, where users should only have access to their own data, Singularity⁴ can be used as a solution for container system in HPC environment. Singularity containers differ from Docker containers in several important ways, including the handling of namespaces, user privileges, and the images themselves. A singularity image, alongside docker image for this toolkit, is available via: <https://singularity-hub.org/collections/2536>.

5 Conclusions

In this paper we have outlined the design and development process used in the creation of the VECMA toolkit for validation, verification and uncertainty quantification (VVUQ) of multiscale HPC applications. A number of exemplar applications from a diverse range of scientific domains (fusion, climate, population displacement, materials, drug binding affinity, and cardiovascular modeling) are being used to test and guide the development of new features for the toolkit. Through this work we aim to make VVUQ certification of complex, multiscale workflows on high end computing facilities a standard practice.

Acknowledgements. We are grateful to the VECMA consortium, Scientific Advisory Board and the VECMA α users for their constructive discussions and input around this work. We acknowledge funding support from the European Union’s Horizon 2020 research and innovation programme under grant agreement 800925 (VECMA project, www.vecma.eu), and the UK Consortium on Mesoscale Engineering Sciences (UKCOMES) under the UK EPSRC Grant No. EP/L00030X/1.

³ <https://hub.docker.com/r/vecmafabsim3/fabsimdocker>.

⁴ <https://singularity.lbl.gov>.

References

1. Oberkampf, W.L., DeLand, S.M., Rutherford, B.M., Diegert, K.V., Alvin, K.F.: Error and uncertainty in modeling and simulation. *Reliab. Eng. Syst. Saf.* **75**(3), 333–357 (2002)
2. Oberkampf, W.L., Roy, C.J.: *Verification and Validation in Scientific Computing*. Cambridge University Press, Cambridge (2010)
3. Roy, C.J., Oberkampf, W.L.: A comprehensive framework for verification, validation, and uncertainty quantification in scientific computing. *Comput. Methods Appl. Mech. Eng.* **200**(25), 2131–2144 (2011)
4. Groen, D., Knap, J., Neumann, P., Suleimenova, D., Veen, L., Leiter, K.: Mastering the scales: a survey on the benefits of multiscale computing software. *Philos. Trans. Roy. Soc. A* **377**(2142), 20180147 (2019)
5. Nikishova, A., Hoekstra, A.G.: Semi-intrusive uncertainty quantification for multiscale models. arXiv preprint [arXiv:1806.09341](https://arxiv.org/abs/1806.09341) (2018)
6. Alowayyed, S., Groen, D., Coveney, P.V., Hoekstra, A.G.: Multiscale computing in the exascale era. *J. Comput. Sci.* **22**, 15–25 (2017)
7. Baudin, M., Dufloy, A., Iooss, B., Popelin, A.-L.: OpenTURNS: an industrial software for uncertainty quantification in simulation. In: Ghanem, R., Higdon, D., Owhadi, H. (eds.) *Handbook of Uncertainty Quantification*, pp. 2001–2038. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-12385-1_64
8. Marelli, S., Sudret, B.: UQLab: a framework for uncertainty quantification in Matlab. In: *Vulnerability, Uncertainty, and Risk: Quantification, Mitigation, and Management*, pp. 2554–2563 (2014)
9. Tennøe, S., Haldnes, G., Einevoll, G.T.: Uncertainpy: a Python toolbox for uncertainty quantification and sensitivity analysis in computational neuroscience. *bioRxiv* (2018). <https://doi.org/10.1101/274779>
10. Feinberg, J., Langtangen, H.P.: Chaospy: an open source tool for designing methods of uncertainty quantification. *J. Comput. Sci.* **11**, 46–57 (2015)
11. Piontek, T., et al.: Development of science gateways using QCG – lessons learned from the deployment on large scale distributed and HPC infrastructures. *J. Grid Comput.* **14**(4), 559–573 (2016)
12. Groen, D., Bhati, A.P., Suter, J., Hetherington, J., Zasada, S.J., Coveney, P.V.: FabSim: facilitating computational research through automation on large-scale and distributed e-infrastructure. *Comput. Phys. Commun.* **207**, 375–385 (2016)
13. Sobol, I.: On Quasi-Monte Carlo integrations. *Math. Comput. Simul.* **47**(2), 103–112 (1998)
14. Luk, O., Hoenen, O., Bottino, A., Scott, B., Coster, D.: ComPat framework for multiscale simulations applied to fusion plasmas. *Comput. Phys. Commun.* **239**, 126–133 (2019)
15. Preuss, R., von Toussaint, U.: Uncertainty quantification in ion–solid interaction simulations. *Nucl. Instrum. Methods Phys. Res. Sect. B: Beam Interact. Mater. Atoms* **393**, 26–28 (2017)
16. Jansson, F., van den Oord, G., Pelupessy, I., Grönqvist, J., Siebesma, A., Crommelin, D.: Regional superparameterization in a global circulation model using large eddy simulations (2018, in press)
17. Verheul, N., Crommelin, D.: Data-driven stochastic representations of unresolved features in multiscale models. *Commun. Math. Sci.* **14**(5), 1213–1236 (2016)
18. Verheul, N., Viebahn, J., Crommelin, D.: Covariate-based stochastic parameterization of baroclinic ocean eddies. *Math. Clim. Weather Forecast.* **3**(1), 90–117 (2017)

19. Groen, D.: Simulating refugee movements: where would you go? *Procedia Comput. Sci.* **80**, 2251–2255 (2016)
20. Suleimenova, D., Bell, D., Groen, D.: A generalized simulation development approach for predicting refugee destinations. *Sci. Rep.* **7**, 13377 (2017)
21. Suleimenova, D., Bell, D., Groen, D.: Towards an automated framework for agent-based simulation of refugee movements. In: Chan, W.K.V., D'Ambrogio, A., Zacharewicz, G., Mustafee, N., Wainer, G., Page, E., (eds.): *Proceedings of the 2017 Winter Simulation Conference, Las Vegas, Nevada, IEEE*, pp. 1240–1251 (2017)
22. Suleimenova, D., Groen, D.: How policy decisions affect refugee journeys in South Sudan: a study using automated ensemble simulations. *J. Artif. Soc. and Soc. Simul.* (2019, submitted)
23. Vassaux, M., Richardson, R.A., Coveney, P.V.: The heterogeneous multiscale method applied to inelastic polymer mechanics. *Philos. Trans. Roy. Soc. A* **377**, 20180150 (2019)
24. Vassaux, M., Sinclair, R.C., Richardson, R.A., Suter, J.L., Coveney, P.V.: The role of graphene in enhancing the mechanical properties of epoxy resins. *Adv. Theory Simul.* **2**, 1800168 (2019)
25. Sadiq, S.K., Wright, D., Watson, S.J., Zasada, S.J., Stoica, I., Coveney, P.V.: Automated molecular simulation based binding affinity calculator for ligand-bound HIV-1 proteases. *J. Chem. Inf. Model.* **48**(9), 1909–1919 (2008). PMID: 18710212
26. Wan, S., et al.: Rapid and reliable binding affinity prediction of bromodomain inhibitors: a computational study. *J. Chem. Theory Comput.* **13**(2), 784–795 (2017). PMID: 28005370
27. Mazzeo, M.D., Coveney, P.V.: HemeLB: a high performance parallel lattice-Boltzmann code for large scale fluid flow in complex geometries. *Comput. Phys. Commun.* **178**(12), 894–914 (2008)
28. Patronis, A., Richardson, R.A., Schmieschek, S., Wylie, B.J., Nash, R.W., Coveney, P.V.: Modelling patient-specific magnetic drug targeting within the intracranial vasculature. *Front. Physiol.* **9**, 331 (2018)
29. Groen, D., et al.: Validation of patient-specific cerebral blood flow simulation using transcranial Doppler measurements. *Front. Physiol.* **9**, 721 (2018)
30. Fernández-Ruiz, I.: Interventional cardiology: drug-eluting or bare-metal stents? *Nat. Rev. Cardiol.* **13**(11), 631–631 (2016)
31. Caiazzo, A., et al.: A complex automata approach for in-stent restenosis: two-dimensional multiscale modelling and simulations. *J. Comput. Sci.* **2**(1), 9–17 (2011)
32. Nikishova, A., Veen, L., Zun, P., Hoekstra, A.G.: Uncertainty quantification of a multiscale model for in-stent restenosis. *Cardiovasc. Eng. Technol.* **9**(4), 761–774 (2018)
33. Nikishova, A., Veen, L., Zun, P., Hoekstra, A.G.: Semi-intrusive multiscale meta-modeling uncertainty quantification with application to a model of in-stent restenosis. *Philos. Trans. A* **377**(2142), 20180154 (2018)
34. Zun, P.S., Anikina, T., Svitenkov, A., Hoekstra, A.G.: A comparison of fully-coupled 3D in-stent restenosis simulations to in-vivo data. *Front. Physiol.* **8**, 284 (2017)
35. Docker: Docker for the Virtualization Admin. eBook (2016)