UM-Bridge: From Linking Models and Methods to Turn-Key HPC in the Cloud

Linus Seelinger[1], Anne Reinarz[2], Robert Scheichl[1]

[1]) Heidelberg University, Institute for Applied Mathematics, Heidelberg, Germany
linus.seelinger@iwr.uni-heidelberg.de
[2]) Durham University, Department of Computer Science, Durham, UK

Linking advanced methods to complex models can imply enormous technical effort, even in cases where such a link is algorithmically simple. This is holding back many interesting applications in science and engineering, and hinders widespread adoption of advanced methods.

UM-Bridge [1] is a universal software interface for linking methods and models. Many methods in uncertainty quantification (UQ), optimization, model-based learning, etc., consider a model to be a mathematical function $F : \mathbb{R}^n \to \mathbb{R}^m$ mapping a parameter vector onto a vector of model predictions. Possibly, derivatives of F are required. UM-Bridge makes this universal mathematical "interface" available as an equally universal software interface. This is achieved by introducing a microservice-inspired architecture, linking both sides through network communication (fig. 1).

Figure 1: Microservice inspired architecture linking uncertainty quantification and model software via UM-Bridge.

This approach opens an entire range of possibilities:

- Level 1: Methods and models can be coupled across arbitrary languages and frameworks. Easy to use integrations for a number of languages are available, handling all network communication behind the scenes (table 1).
- Level 2: Through UM-Bridge, models can easily be containerized [2]. We obtain strong separation of concerns between model and method experts, accelerating development. Further, reproducibility of models becomes possible. With contributions from a number of UQ community members, we provide the first library of ready-to-run UQ benchmark problems.
- Level 3: Scaling small-scale applications to HPC-scale cloud systems becomes trivial. Prototype-grade, thread-parallel method codes can control entire clusters of computationally expensive models. Similar support for "traditional" HPC systems is in development.

Linking to more method packages is an ongoing community effort, aiming at providing a single entry point to a wide range of state-of-the-art methods and models. Since UM-Bridge lowers the entry bar to applying uncertainty quantification and data science on existing models, early industry adoption is

already underway. HPC scalability of UM-Bridge on cloud systems is being investigated, with promising early results.

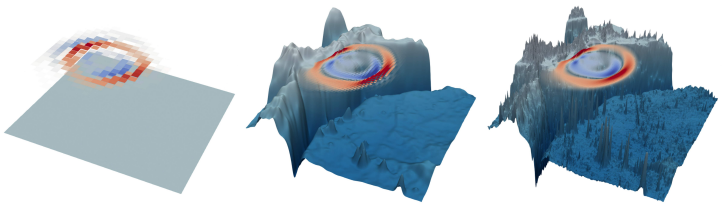

Figure 2: Advanced tsunami simulation of the 2011 Tohoku event. The model is available as a ready-to-use Docker [4] container and can be controlled from any UM-Bridge client. The model hierarchy was initially used for an HPC-scale Bayesian inversion via parallelized multilevel Markov chain Monte Carlo [3].

Language / Framework	Client	Server
Python	✓	✓
C++	✓	✓
MATLAB	✓	✗
R	✓	✗
Julia	planned	✗
MUQ	✓	✓
PyMC	✓	✗
QMCPy	✓	✗
Sparse Grids Matlab Kit	✓	✗
tinyDA	✓	✗

Table 1: Table of client and server support in various languages / frameworks.

References

[1] L. Seelinger, V. Cheng-Seelinger, et al. *UM-Bridge: Uncertainty quantification and modeling bridge*, Journal of Open Source Software, accepted for publication.

[2] Kurtzer, G. M., Sochat, et al. (2017). *Singularity: Scientific containers for mobility of compute*, PloS One.

[3] Seelinger, L., Reinarz, A., et al. (2021). *High performance uncertainty quantification with parallelized multilevel markov chain monte carlo*, SC '21, Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis.

[4] Merkel, D. (2014). *Docker: Lightweight linux containers for consistent development and deployment*, Linux Journal.