Large-scale **EX**ecution for **I**ndustry & **S**ociety

Deliverable D9.8

Impact on Productivity and Business Process Improvement for Earth

Co-funded by the Horizon 2020 Framework Programme of the European Union
Grant Agreement Number 825532
ICT-11-2018-2019 (IA - Innovation Action)

DELIVERABLE ID | TITLE	D9.8 | Impact on Productivity and Business Process Improvement for Earth
RESPONSIBLE AUTHOR	Stéphane Louise (CEA)
WORKPACKAGE ID | TITLE	WP9 | Impact on Productivity and Business Process Improvement for Earth
WORKPACKAGE LEADER	CEA
DATE OF DELIVERY (CONTRACTUAL)	31/12/2021 (M36)
DATE OF DELIVERY (SUBMITTED)	06/01/2022
VERSION | STATUS	V1.4 | Final
TYPE OF DELIVERABLE	R (Report)
DISSEMINATION LEVEL	PU (Public)
AUTHORS (PARTNER)	CEA; AWI; GFZ; ITHACA
INTERNAL REVIEW	Alberto Scionti (LINKS); Fabien Brocheton (NUMTECH)

Project Coordinator: Dr. Jan Martinovič – IT4Innovations, VSB – Technical University of Ostrava
E-mail: jan.martinovic@vsb.cz, **Phone**: +420 597 329 598, **Web**: https://lexis-project.eu

DOCUMENT VERSION TABLE

VERSION	MODIFICATION(S)	DATE	AUTHOR(S)
1.0	Initial release	6/12/2021	Stéphane Louise (CEA), Danijel Schorlemmer (GFZ); Lorenza Bovio (ITHACA); Natalja Rakowsky (AWI)
1.1	Added list of partners, glossary entries, bibliography	13/12/2021	Stéphane Louise (CEA); Natalja Rakowsky (AWI)
1.2	Final check of the deliverable	19/12/2021	Kateřina Slaninová (IT4I)
1.3	2nd final check of the deliverable	28/12/2021	Jan Martinovič (IT4I)
1.4	Update of the deliverable according to the comments from the final check	05/01/2022	Sven Harig, Natalja Rakowsky (AWI); Danijel Schorlemmer, Marius Kriegerowski (GFZ); Stéphane Louise (CEA)

GLOSSARY

ACRONYM	DESCRIPTION
ADAS	Advanced Driver-Assistance Systems
AGPL	GNU Affero General Public License
AWS	Amazon Web Service Inc.
CPS	Cyber-Physical System(s)
EPOS	European Plate Observation System
FESOM2	Finite volumE Sea ice-Ocean Model
FPGA	Field Programmable Gate Arrays
GPL2	GNU General Public Licence Version 2
HPC	High-Performance Computing
I&C	Instrumentation & Control
I/O	Input/Output
MoC	Model of Computation
MPI	Mesage Passing Interface
TOAST	Tsunami Observation And Simulation Terminal, a decision support system for tsunami early warning
VGI	Volunteered Geographic Information

TABLE OF PARTNERS

ACRONYM	PARTNER
Avio Aero	GE AVIO SRL
Atos	BULL SAS
AWI	ALFRED WEGENER INSTITUT HELMHOLTZ ZENTRUM FUR POLAR UND MEERES-FORSCHUNG
BLABS	BAYNCORE LABS LIMITED
CEA	COMMISSARIAT A L ENERGIE ATOMIQUE ET AUX ENERGIES ALTERNATIVES
CIMA	CENTRO INTERNAZIONALE IN MONITORAGGIO AMBIENTALE - FONDAZIONE CIMA
CYC	CYCLOPS LABS GMBH
ECMWF	EUROPEAN CENTRE FOR MEDIUM-RANGE WEATHER FORECASTS
EURAXENT	MARC DERQUENNES
GFZ	HELMHOLTZ ZENTRUM POTSDAM DEUTSCHESGEOFORSCHUNGSZENTRUM GFZ
EURAXENT	MARC DERQUENNES
ICHEC	NATIONAL UNIVERSITY OF IRELAND GALWAY / Irish Centre for High-End Computing
IT4I	VYSOKA SKOLA BANSKA - TECHNICKA UNIVERZITA OSTRAVA / IT4Innovations National Supercomputing Centre
ITHACA	ASSOCIAZIONE ITHACA
LINKS	FONDAZIONE LINKS / ISTITUTO SUPERIORE MARIO BOELLA ISMB
LRZ	BAYERISCHE AKADEMIE DER WISSENSCHAFTEN / Leibniz Rechenzentrum der BAdW
NUM	NUMTECH
O24	OUTPOST 24 FRANCE
TESEO	TESEO SPA TECNOLOGIE E SISTEMI ELETTRONICI ED OTTICI

TABLE OF CONTENTS

LIST OF FIGURES

EXECUTIVE SUMMARY

The impact of the earthquake and tsunami on Productivity and Business Process Improvement within this workpackage can be summarized as:

- A holistic impact on public sector, in the area of disaster response; time reduction for first damage and inundation assessment and a better accuracy of the projected values of those,
- A demonstrator of cyber-physical systems, with HPC in the loop, for industrial control in complex cases,
- The demonstration of Digital twin of country buildings and infrastructure, with a way to automatically build that twin without costly surveys for obtaining the needed detailed information,
- Time reduction for tsunami simulations, allowing faster risk estimations, and impact on similar simulations for industry.

POSITION OF THE DELIVERABLE IN THE WHOLE PROJECT CONTEXT

While the WP6 on Earthquake and Tsunami simulation for emergency reactions is best described as an application workflow for the LEXIS Platform like, e.g. WP7, it is also a demonstrator of technologies and concepts that are nowadays emerging for HPC and Cloud infrastructures, especially with time-constrained and compute intensive applications, which requires novel Models of Computations, accelerated communications and synchronisations.

DESCRIPTION OF THE DELIVERABLE

This deliverable D9.8 provides the current prospects for the Impact on Productivity and Business Process Improvement as a consequence of the work done during the design and implementation of the workflow of WP6. It shows what are the prospects with regards to emergency responses, quick tsunami simulations, and easily available digital twins for building sets, or towns and cities within one country or region.

1 INTRODUCTION

The workflow developed as part of WP6 is a loss, damage and inundation assessment from earthquake events and deducing simulated earthquake damage and projected tsunami related inundation. It is of course of utterly importance to have these kinds of simulation as early as possible as part of emergency response and policy maker's best decisions (Section 2). Nonetheless, the development of this particular workflow also opens prospect toward new fields and new potential productivity and business process improvements such as compute intensive Cyber-Physical Systems (Section 3), digital twin models for buildings within a given city or a given region or country with a very low price tag compared to usual way of doing on-field surveys and studies or building up from data from land registries (Section 4) and tsunami scenario quick simulations (Section 5).

2 COMPLETE PILOT IMPACT

WP6 workflow allows for both early and detailed assessments of a potential threat to populated areas in the event of an earthquake and potentially earthquake-induced tsunami. Making the process from the event of an earthquake occurrence to the simulation of damage assessments and tsunami inundations associated with it in a matter of a few minutes to its first availability to emergency rescues and policy makers is one of the key-points of this workflow. As such it is already groundbreaking and has the potential to save hundreds or thousands of lives would it be put into action by a given country to plan for early warning of population and preparations of emergency services.

Early assessments were primarily provided, in the case of tsunamis, by a network of buoyant networks of captors, which meant it was out of reach of lower income countries because of the sheer cost of this infrastructure and its maintenance. With early/on time computation and simulation, the tsunami assessment would now be reachable for low income countries since a much less expensive array of earthquake detectors and access to cloud computing would be sufficient to allow for early warning. Even for richer countries the potential higher accuracy associated by a multi-factor simulation would help separate minor impact events from meaningful events, thus reducing the number of unnecessary evacuations, if correctly calibrated. This has the potential to both reduce the costs of warning and increase population confidence in the early detection and evacuation alerts.

It also allows theoretically to prepare and handle emergency services and other logistical forces to focus the mains efforts where they would be required the most, facilitating the damage control and emergency response, and preparing for a quicker recovery in the future. Decision makers can orientate emergency response to places where victims are expected to be in the most needed, and where important infrastructures could need emergency help and backup solutions.

The possibility to use the proposed workflow has been investigated in relation to the Copernicus Emergency Management Service, provided by the European Commission. The models integrated in the workflow showed potential in reducing the satellite tasking timing and in generating early possible extent and damage information concerning the impact of the event.

[1], [2] provide earlier research work on the subject of emergency simulation and assessment.

3 IMPACT ON CYBER-PHYSICAL SYSTEMS

Cyber-physical systems (CPS) [3] are typically physical systems, in interaction with the physical world that require time-constrained decision making and real-time adjustments between captors or perceptions or events on one hand and actuators or reactions or information displays on the other hand. Typical examples fall into the field of Instrumentation & Control (I&C) like power plants [4], avionic and aerospace systems [5], or ADAS (Advanced Driver Assistance Systems) for semi-autonomous cars and trucks [6]. Until the 2000s decade, whilst the embedded computing power required to perform the real-time task of piloting (or aiding to pilot) these CPS was increasing steadily, systems were designed to be mostly reliant on their own and avoid having to communicate outside of their limited environment, unless it was not possible to do so: Critical real-time and safety systems especially would have their functions and way of working proven statically (off-line) in a mandatory fashion. Even if the aeronautic system would put important decision on the pilot if one is available (avionics) and ADAS would only be recommendation with regards to the steering or emergency breaking, the low level decisions are done automatically so e.g. the structural integrity of a power plant or a commercial plane can not be jeopardised by an out-of-specification decision of the pilot. For rockets, any out-of-planned trajectory of a non-manned rocket would be quickly resolved by the preemptive destruction of the rocket (Flight termination systems), or mission abort and emergency escape systems for manned missions.

Nonetheless, as many of us have experienced while e.g. driving one's own car, having feedback on the road congestion and road closure may avoid lot of delays and ease the experience of driving. More generally speaking, having some high-level near-real-time information like planned navigation for cars, or power generation planning and weather forecasts for power plants greatly improves the smoothness, the energy efficiency and the life expectancy of CPS.

More recently, concepts like HPC-in-the-loop emerged in CPS [7], because some high-level data are either too large to fit in embedded systems or the content of the data has a short-term value only (like road traffic). Reciprocally, some HPC workload as limited-time value. Weather forecast is one of the early example: Increasing the resolution of the weather simulation greatly improves the accuracy of the forecast, but doing too much means that the result of the computing can be obtained after the time the forecast is no longer a forecast but already something of the past. Therefore there is an optimal time when the simulation in itself has an interest beyond the capability of computers to model the weather: the simulation outcomes should come up with results from the simulation several hours before the potential weather events, otherwise the usefulness of the simulation outcomes are greatly reduced[1].

The case of weather forecast is not so difficult to decide. Lately, HPC simulations with real-time related best utilization have grown in number and complexity. This is one of the first example we treated here with the combined simulation of earthquake and tsunami and the resulting damage assessment within a single (but rather complex) workflow. As such the workflow in WP6 represents a first implementation on a formal model of HPC simulation with deadlines within a non linear workflow. This alone is an achievement and represent an industrial interest because of its applicability in e.g. complex digital twins (as also discussed in section 4) that would allow to extract valued information from the status and health of important infrastructures like large bridges, mines, core conditions of nuclear power plants, etc.

BEYOND THE REAL-TIME SIMULATIONS AND EARLY AWARENESS

Another possible application of delay-bounded computing as implemented in the WP6 workflow, it would be possible to allow, in a cost-effective way, best-effort compute within a cloud infrastructure while a strict quality of service is also maintained. Therefore, one can use the tools developed within the WP6 to add best effort tasks that can be carried out on cores on available compute resources but giving the priority to the foremost tasks that have on-premise computing resources. As soon as the on-promise resources are required, the best effort compute tasks can be shut down immediately by the orchestrator without loosing time waiting for completion of background and best-effort tasks which could be reallocated once the computing resources are available again.

This business model is advocated by Amazon Web Service Inc. (AWS), but lacks the low level support developed in our workflow, at least at the moment of writing this report and within the knowledge of the authors.

4 DIGITAL TWIN IMPACT

The key to useful rapid assessments of damage due to natural hazards is not only the speed of processing but also the knowledge of the building stock in the affected areas. This knowledge is described in exposure models that provide the distribution of different building classes, their value and the people expected inside them. While there are many exposure models developed for many areas of the world, they are usually given a larger-scale aggregates and describe the distribution of building types on the municipal or state level. To make the spatial distribution within the given administrative unit more realistic, satellite-based built-area detection is used to distribute the affected assets. However, because earthquake shaking can vary strongly even on urban scales and tsunamis affect every building differently depending on its exact location, a more precise description of building locations and types is desired.

We developed a digital twin of the global building stock by creating a machinery to process open building data algorithmically to create a building-specific exposure model of the world. Our main data source is OpenStreetMap, a geographic database to which millions of contributors provide so-called Volunteered Geographic Information (VGI). This database is growing by approx. 150,000 buildings every day and currently contains close to half a billion building footprints, resulting in complete or almost complete coverage in many areas or countries of the world, see Figure 1 as an example. This digital twin also includes a plethora of additional information about the overall infrastructure like roads, railroads, hospitals, and schools. The constant processing allows to include each new contribution to OpenStreetMap to be used in the digital twin within a few minutes. There, each building with its associated data will be processed in order to estimate in a probabilistic way the necessary exposure indicators to assess the building type, its value and the number of people inside. This ever growing database, fed by volunteers but also by data donations coming from governmental open-data policies, offers a cost-effective system to global and local exposure data, in particular in areas where cadastral data is not accessible or does not exist in any form useful for exposure modelling. Governments in such areas can use these resources freely and also stimulate the collection of VGI to improve the exposure data to be used by them.

We have demonstrated the ability to process building data at a global scale. The building database (OpenBuildingMap) with currently close to half a billion building footprints is kept up-to-date on a minutely basis. The processing is based on a newly developed engine that is scallable with multiple processing queues and processes building data based on rules. Thus, this engine and the entire processing system can easily be made to work on smaller scales, e. g. on a country level. The open architecture of the processing engine also allows the users to enrich the processing procedures with additional rules.

[1] Forecasting the weather of yesterday is usually considered not useful.

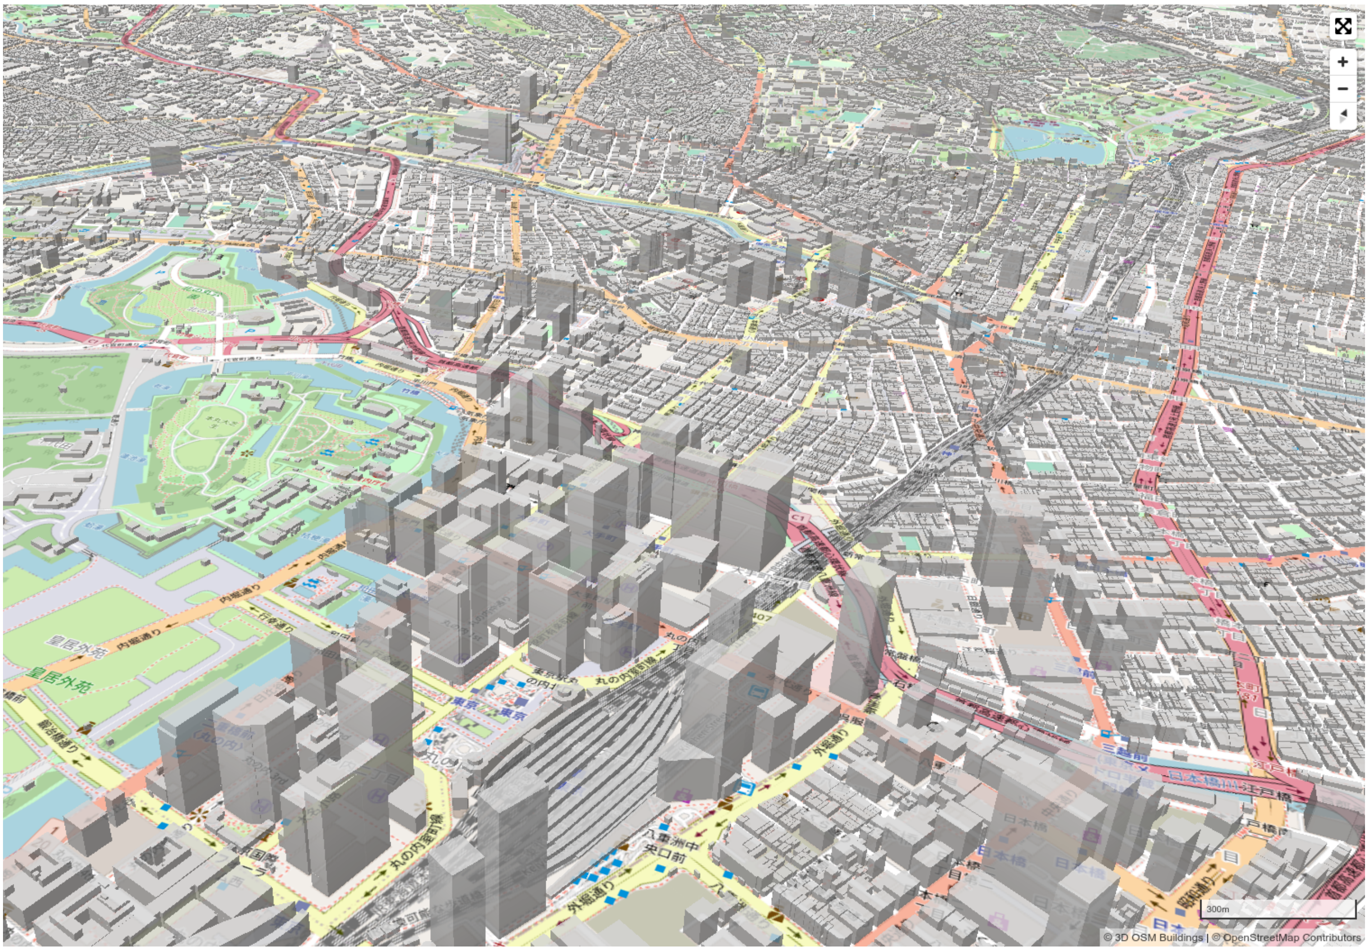

Figure 1: 3D building footprints in the area around the Tokyo railway station and the Imperial Palace, Tokyo, Japan. The map background is provided by OpenStreetMap, the data is from OpenStreetMap and processed in OpenBuildingMap.

To make the exposure model usable quickly, we populate an exposure database with all exposure-model data for fast access of exposure data [8]. This avoids costly queries on the building database and also allows to compensate a-priori for missing data (buildings not contained in the OpenBuildingMap) as the exposure model is provided on tiles of approx. $100 \times 100\,\mathrm{m}$ size. These tiles represent the complete assets of the tile area, either as detailed building information or as summary if building footprints are missing. All codes are freely available under the AGPL license to empower anybody interested in implementing this model and processing engine for their own particular use case. Governmental actors and stakeholders are also invited to use the exposure model directly without setting up an own processing infrastructure. For more details, see Deliverable D6.4 [9].

5 TSUNAMI SIMULATIONS IMPACT

In the early LEXIS phase, we aimed at very fast simulations of the tsunami inundation by determining how coarse the resolution on land may be for a reasonable estimate. The result, approx. 200m horizontal resolution, is helpful for situations with limited compute capacity. For a description of the first test cases, see Deliverable D6.2 [10].

In the major step, the tsunami simulation code was considerably optimised. The most important steps were the switch from double to single precision arithmetic and the implementation of MPI parallelisation. The single precision arithmetic reduces the time to solution on fixed resources by about 30% except for very small setups that fully fit into the cache. The MPI parallelisation allows to employ more than one compute node for a TsunAWI run. As both the strong and weak scaling are very good, we obtain faster time to solution for medium to large setups, and we can now tackle even simulations of a full ocean basin with inundation on all coast in real time, provided sufficient resources are at hand. For example, a 12 hour TsunAWI simulation on a mesh with 11 Million nodes of the whole Indonesian Archipelago, resolving all coasts with 150m or better, takes about 1 hour on one compute node, but only 2min on 40 compute nodes of a recent Linux cluster. Further experiments conducted at Atos with Burst Buffers showed promising acceleration of IO operations and potential for on-datanode data post-processing, see D3.6, section "3.2.3.1 Acceleration of I/O with Help of Smart Burst Buffer" [11]. A work of evaluation of the speed-up that can be achieved by utilizing the Bust-Buffers as accelerators for distributed memory management will be available in the PhD thesis memoir of Erwan Lenormand [12] that will be defended in the second half of January 2022. An article based partially on this work is planned but not yet submitted. The results shown in this PhD Thesis envision a possible speed-up of 2.5 fold on average and up to 1,000% with a careful use of the scheduling and node selection policy.

Improving the time to solution by a very significant factor allows for probabilistic risk assessment (which can require running hundreds of thousands simulations).

With these improvements, TsunAWI will be listed as part of the new Tsunami Thematic Core Service for Europe in the European Plate Observation System (EPOS) research infrastructure. Furthermore, we are in contact with gempa GmbH. The company provides support for SeisComp, a leading software for earthquake warning centers, and distributes TOAST as a decision support system for tsunami warning coupled to SeisComp[2]. In Indonesia, TOAST is connected to a database of pre-computed TsunAWI scenarios. The aim is to add an interface to an online simulation with TsunAWI.

Furthermore, TsunAWI and the user guide is freely available (GPL2 license) and computational meshes for many regions will be added for free download[3].

TsunAWI is the small sibling of the full ocean model FESOM2[4] [13], and a good test bed for new technologies. This time, TsunAWI inherited FESOM2's framework for MPI [14], however, only TsunAWI contains an additional OpenMP layer for parallelisation. FESOM2 will benefit from the expertise in hybrid parallelisation, and we will use TsunAWI as a test bed for new approaches for optimisation techniques in particular for numerical kernels on unstructured meshes. Summing up, the acceleration technology will have an impact on climate change studies with the big brother FESOM and on similar simulations in general.

[2]SeisComp: https://www.gempa.de/products/toast/
[3]Current state with data for the Chilean study area and link to the code: https://doi.org/10.5281/zenodo.5799056
[4]FESOM2: https://fesom.de/

6 CONCLUSION

As has been seen, far from being limited to the most obvious impact of emergency response, this workflow opens up new fields of applications that are worth mentioning with a high potential for governments and businesses. Thank to meaningful reduction in execution time and several levels of simulations, emergency responders can have access to near real-time estimations and would be in a position to focus their efforts were it is the most relevant. Used of time-constrained simulations on the cloud can alleviate beyond first-responder's duties, other fields including enhanced and intelligent automatizing in production processes (also known as Industry 4.0), floats of vehicles, but other application e.g. from the work of GFZ can disrupt the field of building digital twins for cities and towns. As such, this workflow opened up a new pathway toward new fields of applications, which in itself is a major achievement which can be further developed. It is also worth mentioning that most of the code developed within the project and the WP6 is open-source.

REFERENCES

[1] P. Beckman, S. Nadella, N. Trebon, and I. Beschastnikh, "Spruce: A system for supporting urgent high-performance computing," *International Federation for Information Processing Digital Library; Grid-Based Problem Solving Environments;*, vol. 239, 2007-11. DOI: 10.1007/978-0-387-73659-4_16.

[2] K. Yoshimoto, D. Choi, R. Moore, A. Majumdar, and E. Hocks, "Implementations of urgent computing on production hpc systems," *Procedia Computer Science*, vol. 9, pp. 1687–1693, 2012, Proceedings of the International Conference on Computational Science, ICCS 2012, ISSN: 1877-0509. DOI: 10.1016/j.procs.2012.04.186. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1877050912003079.

[3] E. A. Lee, "Cyber physical systems: Design challenges," in *2008 11th IEEE international symposium on object and component-oriented real-time distributed computing (ISORC)*, IEEE, 2008, pp. 363–369.

[4] S. Rudakov and C. E. Dickerson, "Harmonization of ieee 1012 and iec 60880 standards regarding verification and validation of nuclear power plant safety systems software using model-based methodology," *Progress in Nuclear Energy*, vol. 99, pp. 86–95, 2017.

[5] V. Hilderman and T. Baghi, *Avionics certification: a complete guide to DO-178 (software), DO-254 (hardware)*. Avionics Communications, 2007.

[6] S. Fürst, J. Mössinger, S. Bunzel, *et al.*, "Autosar–a worldwide standard is on the road," in *14th International VDI Congress Electronic Systems for Vehicles, Baden-Baden*, vol. 62, 2009, p. 5.

[7] M. T. McMahon, B. M. Wickham, and E. H. Page, "An architecture for low overhead grid, hpc, and cloud-based simulation as a service," in *Simulation for Cyber-Physical Systems Engineering*, Springer, 2020, pp. 321–345.

[8] D. Schorlemmer and T. E. Zadeh. "Exposure and building data." (2022), [Online]. Available: https://doi.org/10.5281/zenodo.5821041.

[9] LEXIS deliverable, *D6.4 pilots results: Improved scenarios measurements and evaluation*.

[10] ——, *D6.2 pilots improvements: Solutions adopted*.

[11] ——, *D3.6 data flow optimisation and data system core*.

[12] E. Lenormand, "Unification des mémoires réparties dans les systèmes hétérogènes (on the unification of distributed memory in heterogeneous systems)," Ph.D. dissertation, Université Paris-Saclay, 2022-01.

[13] S. Danilov, D. Sidorenko, Q. Wang, and T. Jung, "The Finite-volumE Sea ice–Ocean Model (FESOM2)," vol. 10, pp. 765–789, 2017. DOI: 10.5194/gmd-10-765-2017.

[14] N. Koldunov, V. Aizinger, N. Rakowsky, *et al.*, "Scalability and some optimization of the Finite-volumE Sea ice–Ocean Model, Version 2.0 (FESOM2)," vol. 12, pp. 3991–4012, 2019. DOI: 10.5194/gmd-12-3991-2019.