PROCEEDINGS OF SCIENCE

Environmental Computing 1.0: The Dawn of a Concept

Matti Heikkurinen[1]
Ludwig-Maximilians-Universität München
Oettingenstraße 67, 80538 München,
Germany
E-mail: heikku@nm.ifi.lmu.de

Michael Schiffers
Ludwig-Maximilians-Universität München
Oettingenstraße 67, 80538 München,
Germany
E-mail: schiffer@nm.ifi.lmu.de

Dieter Kranzlmüller
Leibniz Supercomputing Centre (LRZ) of the Bavarian Academy of Sciences and Humanities
Boltzmannstraße 1, 85748 Garching bei München
Germany
E-mail: Dieter.Kranzlmueller@lrz.de

There is a long history of numerical modelling of various natural phenomena for purposes such as weather prediction or analysis of different earthquake-scenarios. In this paper we present the next logical step: combining multiple models together in a dynamically extensible framework in order to gain a better understanding of the nature and impact of inherently interlinked and dependent environmental phenomena. We call this approach *Environmental Computing*, which encompasses both the link to a broad range of environmental issues that can be approached using the framework model, and the notion that the component models and their features can be evaluated algorithmically to evaluate the accuracy and applicability of the results in different situations.

Reaching this goal requires new technologies, commonly accepted approaches, standards and policies. This multi-pronged approach is necessary, since the combination of different models will bring forth challenges related to the compatibility of the execution environments as well as issues with the syntax and semantics of data. The data challenge applies both to the input of the while model ensemble as well as mechanisms of inter-model data exchange.

To showcase the progress made towards this goal so far, we will present the technical and operational frameworks for the environmental multi-modelling, as well as specific case studies that have acted as proofs-of-concepts or pilot tests establishing the state of the art in this domain. These case studies, together with a short summary of related work, will lay the foundations for a discussion on the methods that could be used to assess the impact and benefits of Environmental Computing in this and other contexts.

PoS(ISGC2015)030

---

[1] Corresponding Author

We will conclude the paper with a discussion related to the potential impact of successful deployments of the environmental computing tools as well as a summary of future research and other next steps.

International Symposium on Grids and Clouds (ISGC) 2015
15 -20 March 2015,
Academia Sinica, Taipei, Taiwan

1. Introduction

The history of numerical modelling of natural phenomena for purposes such as weather prediction or analysis of different earthquake-scenarios is at least as long as one of computer simulations. In this paper we present the next logical step in advancing the accuracy and efficiency of modelling: combining the different models together in a dynamically extensible framework in order to gain a better understanding of the nature and impact of inherently interlinked and dependent environmental phenomena. For example, an earthquake may trigger landslides that in turn will change the probabilities of flooding (either due to waves triggered by the quake or a landslide or due to simultaneous rainfall). Additionally, preparation and response to disaster scenarios requires linking this multi-model system with several other disciplines that are crucial for societal resilience. These range from technical (civil engineering approaches needed for flood defences or increased earthquake resilience) to social sciences (planning and communicating the responses so that they are accepted by the population).

Detailed understanding of environmental phenomena is becoming more and more important due to several coinciding developments. The exposure of the society is changing, as urbanisation is concentrating the population (and the risks associated with environmental phenomena) to an unprecedented degree. At the same time, climate change is changing the weather patterns as well as most likely making extreme weather events more frequent. The intergovernmental agreements negotiated as a response to these scenarios (such as the recently agreed Sendai framework for disaster risk reduction [1]) add obligations and mandates to the governments with regard to civil protection. And finally, both companies and the general public are becoming more and more aware of the potential of the modelling technologies, which increased expectations in terms of early warning and disaster response.

We use the term *Environmental Computing* to denote the technologies, common approaches, standards and policies that will make feasible the new modelling capabilities to address the challenges outlined above. A multi-pronged approach is necessary, since the combination of different models and using these model ensembles (or workflows) as an integral part of decision making will bring forth several challenges. Even in the case of well-understood models, the combination may behave in an unpredictable manner e.g. due to differences in data syntax or semantics. Or the model may work correctly, but will be so inefficient that the results are only relevant for research purposes (i.e. the results arrive only after the event being modelled has already over). Thus, documenting the behaviour of individual models as well as their combinations in a way that non-experts can anticipate the system-level behaviour is crucial.

To lay the groundwork for addressing these issues, standardised technical interfaces and best practice information will need to be collected and presented in a way that is understandable for the developers of the other models and model integrators. Codifying these practices in technical, procedural and organisational standards is equally important for the take-up of the results. Especially in the case of early warning, the stakes can be high – for example, overestimating an environmental risk that leads to an unnecessary evacuation may cause considerable damages. For this reason, the model and workflow metadata should ideally be machine readable, so that its correctness can be (at least partially) determined automatically and the crucial information can be extracted and summarised with minimal human intervention (see Figure 1).

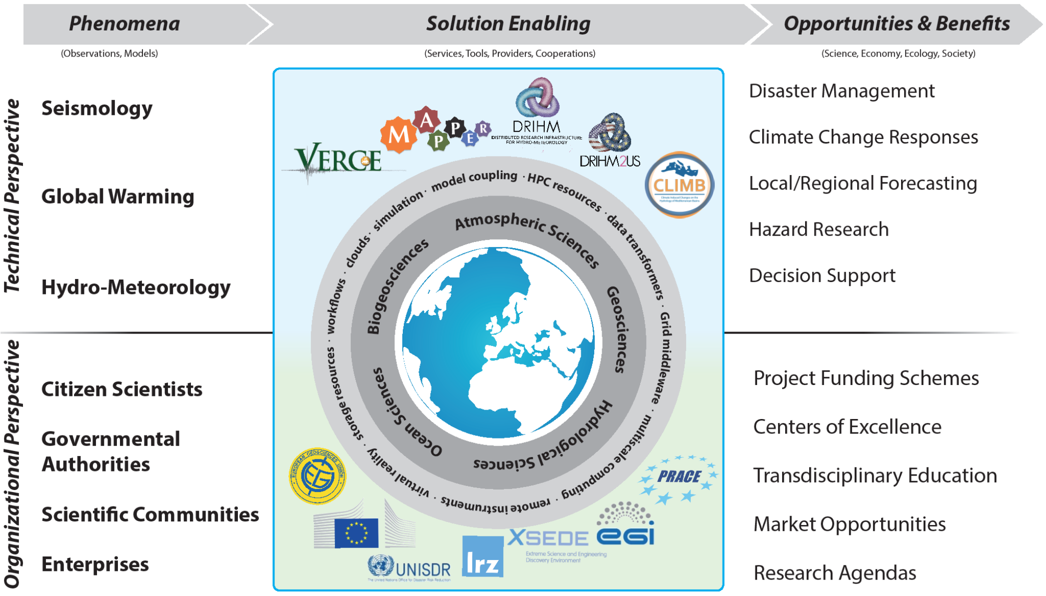

Figure 1: Environmental Computing covers technical and organizational aspects and spans the whole value chain from phenomena observation to applications offering new opportunities for scientists, economy and society

In this paper we will focus on the issues related to environmental modelling in the context of disaster risk reduction and response. While there are numerous other applications where environmental multi-modelling can and will be used, the combination of urbanisation, climate change and political developments make disaster risks and responses particularly timely domain of study. We will present the technical and operational frameworks for the environmental multi-modelling, and then present specific case studies that have acted as proofs-of-concepts or pilot tests establishing the state of the art in this domain. These case studies, together with a short summary of related work, will lay the foundations for a discussion on the methods that could be used to assess the impact and benefits of Environmental Computing in this and other contexts.

The contribution of the paper is threefold: We first outline some of the requirements for a conceptual framework to facilitate solving research questions in the emerging context of Environmental Computing (Section 2). Section 3 discusses examples. Given the framework, the second contribution consists of sketching a methodology to solve issues in Environment Computing (Section 4) using instantiation patterns. The third contribution is unveiled in Section 5 in which we apply the methodology to a hydro-meteorological use case. As such, the application may serve both as a blueprint for similar scenarios and as a "breeding environment" for future extensions. We discuss aspects of socio-economic impacts in Section 6. Finally, Section 7 concludes the paper and sheds light on further work.

2. Fundamental requirements of a technical framework for Environmental Computing

Creating a technical framework for Environmental Computing requires solving challenges related to technical interoperability (both computationally and from a data management perspective) and to domain-independent metadata management of the model components. These two categories are equally important; the technical interoperability will ensure that the software runs and produces results in a way that is predictable and (relatively) efficient. However, without adequate metadata it is impossible to judge whether the results are correct or not. For example, a programme may produce output files successfully, but if the output format changes between programme versions, another model using the data as input may not work correctly anymore. Thus the documentation of syntax and semantics of the output files needs to be reviewed carefully and kept up to date. This requires non-trivial amounts of effort, which brings up important non-functional requirements: for example, the overall execution efficiency of the framework and the increased utility value of the results need to be high enough to justify the investments in the adaptation and documentation of the model components and the workflow tying them together.

In general, the reason that the effort of building such a framework is justified is based on the typical evolutionary paths of every individual component model. Each model solution tends to represent typically decades worth of incremental improvements that allow the software to match the behaviour of a particular sub-system – such as the absorption/evaporation of rain on specific surface types or the propagation of seismic waves in different types of rocks – very accurately over a broad range of situations. This incremental approach means that many of the interfaces, data structures and input/output formats have been developed with the specific application in mind, in a manner that is difficult to separate from the programme logic itself. Hence, the knowledge and in-depth understanding of the particular phenomenon is embodied in the software and the expertise in fine-tuning initial parameters and operational environments. Simply re-implementing the software – even if the resources for it were made available – would lose a lot of this tacit information related to the modelling software.

From the functional perspective, the framework needs to be able to solve (or assist in solving) at least the following issues:

1. Managing dependencies with the execution environments
2. Dynamically linking models into workflows
3. Supporting accessing data – both from external, static sources as well as ingesting data between models (the output of one model serves as input for another model).

The two first requirements are relatively straightforward. The models can be grouped together based on the commonalities of their execution environment requirements, and based on this information a manageable number of “profiles” can be generated and also communicated to developers as suggestions to take into account as mandatory context information. Workflow systems are conceptually mature and several implementations exists that can either be adapted or used as starting points when defining the desired functionality – even under the assumption of dynamics in orchestrating execution chains. However, the final requirement – data access – is considerably less straightforward issue: the file syntax, access protocols and especially semantics of the input and output files can vary in much more diverse and fundamental ways than execution environments or execution order dependencies between the models in a workflow.

On a technical level, the linking of input and output data from different models is accomplished by relying on different file standards, such as WaterML [2] or NetCDF-CF [3] in the Earth Science disciplines. However, these standards may not fully cover issues related to the accuracy of a model in certain parameter values or to any application-specific extensions of the standard. To overcome these challenges, it is necessary to create a metadata framework that is common – and consistently used – across the whole multi-model chain. An optimal metadata framework thus combines several partially conflicting features – it has to be simple, implementable with low initialisation and maintenance efforts, easy to use by non-experts, and comprehensive enough to avoid architectural "overkills".

3. Example frameworks

There are some frameworks already in place that provide parts of the basic functionalities (functional and metadata-related) for Environmental Computing as outlined above. We mention just three. Other examples are reported in [4].

The EU-funded DRIHM project (Distributed Research Infrastructure for Hydro-Meteorology[2] [5] focuses on hydro-meteorological multi-model systems with emphasis on post-event analysis. DRIHM supports linear model chains addressing all three of the functional requirements (execution environment dependencies; workflows; linking of input and output files) as well as a metadata framework called M.A.P. (Metadata, Adaptors, Portability) [6]. M.A.P. is an archetypal example of the Pareto "80-20" rule: the goal is to have sufficient semantic expressiveness to capture most of the common requirements faced by the project and to enable researchers to manually determine whether a particular combination of models would work. This avoids some of the complexity and risks related to developing an automated solution before the procedure being supported has stabilised. While the manual approach is easier for human operators, the drawback is the slightly higher maintenance efforts that are needed in order to keep model descriptions up to date.

The MAPPER (Multiscale Applications on European e-Infrastructures[3]) framework [7] focuses on modelling, predicting and controlling multiscale systems where processes acting at different scales coexist and interact. Typical application areas are *"Urgent Computing"* scenarios [8] where priority driven tasks need to be orchestrated across (super-) computing centres to react in right time. While MAPPER also addresses the same functional requirements as DRIHM, the emphasis is more on orchestrating multi-physics models in the same scientific domain than on data fusion, data ingestions, and metadata management. Consequently, MAPPER supports loosely and tightly coupled simulations by providing the respective frameworks as described by Borgdorff et al. in [7].

The CAPRA-GIS [9] framework focuses on calculating probable losses (in terms of lives lost and direct economic damages) based on statistical analysis of different disaster scenarios. The calculation takes into account the exposed population and the infrastructure at the location where the (simulated) disasters of different types (earthquakes, floods, typhoons etc.) occur, and

---

[2] http://www.drihm.eu/

[3] http://mapper-project.eu

the “robustness” of the infrastructure (e.g., based on civil protection standards and building codes). In terms of functional requirements it is focused on common data exchange standards across different models, mandating presenting the results of the disaster scenarios in AME file format[4] independent of the risk type (flood, earthquake, volcanic activity). Describing the execution environments where the disaster scenarios are modelled falls outside the scope of the software. Similarly, the metadata aspects are mostly limited to compliance with the AME file format. An application is described in [10].

OpenMI (Open Modelling Interface [11]) provides a standard to pass data between models as they run. As such, it specifies (the interfaces of) a framework for model engines to be included in integrated compositions. As a response to the EU Water Framework Directive calls for integrated water management, OpenMI itself was originally developed to consider the interactions of environmental processes, in particular involving water [11], but it has since been realised to be considerably more flexible. It is now considered an interface standard between software components that can be applied to linking any combination of models, databases and associated tools and has been ratified by the Open Geospatial Consortium (OGC)[5]. OpenMI allows two-way exchange of data between compliant components as they run and one-way passing of data from a driving component to a second one, set up only to receive data – a similarity to MAPPER’s loose and tight coupling. FluidEarth [12] is a Windows-based implementation for OpenMI 2.0.

All these frameworks have proven their value in their particular niches. For example, the post-event analysis of the 2014 Genoa flash flood [13] performed on the DRIHM infrastructure demonstrated that the accuracy of the hydro-meteorological predictions could be improved considerably by the use of multi-model approaches and more detailed simulations. However, to accomplishing this requires roughly two orders of magnitude increase in the computing capacity compared to the current, standard operational systems. The MAPPER approach and its application to hydrology have been reported by Belgacem et al. in [14] where they successfully simulated irrigation canals and rivers in 3D. Finally, CAPRA-GIS has enabled a large-scale collaboration between UNISDR [15] and different research partners that has successfully produced the series of Global Assessment Reports [16], which were crucial input for the negotiations leading to the Sendai Framework [1]. Bulatewicz et al. demonstrated the value of OpenMI when coupling existing models into a cohesive group to make them run as a system rather than merging them together into a single ‘super-model’ [17].

At the same time, these frameworks are either relatively unproven for more general use cases (like climate change forecasts) or they are intentionally focusing on solving very specific problems and leaving some of the environmental computing aspects (such as execution environments and model metadata in the case of CAPRA-GIS) outside their scope. However, all of them include similar static structural components (such as model couplers) and dynamic processes (workflow-like processes that either trigger the execution of multi-model ensembles or at least aggregate data into higher level summaries for educated decision making). As such they should be seen as precursor of a more generalised, multi-purpose framework that is

---

[4] https://dl.dropboxusercontent.com/u/10355027/Capra/Introduction%20to%20CAPRA%20AME%20objects.pdf

[5] http://www.opengeospatial.org/

emerging through the more active exchange of best practice information and more ambitious projects in the future.

4. Scheduled and urgent instantiations

When applying a particular Environmental Computing framework, understanding the timeframes involved is the key in achieving a satisfactory end result. Compared to scheduled analysis runs, analysing a specific risk-scenario in near-real time requires different kinds of approaches– from the selection of services providing access to physical resources to approaches to curating the end results of the simulations to testing hypotheses related to the long-term behaviour of different interconnected earth systems. It is usually enough to divide the instantiation approach into two categories: *scheduled* and *urgent* instantiations.

In a *scheduled* instantiation (the standard case) the deadline for producing final results is known and can easily be met with the resources allocated for a task. The deadline itself may vary from minutes (to answer whether a particular observation requires further analysis) to years (publication of a policy document such as a Global Assessment Report [16]). However, if it is possible to identify input data sources, computing and storage resources and hypotheses to test with a particular workflow well in advance, we can treat the instantiation processes in a similar manner. The execution of the workflow is essentially a routine endeavour, with well-known results that may trigger further actions (e.g. extreme weather warning).

In some cases this further action may require triggering other workflows, possibly as an *urgent* instantiation that requires a more dynamic composition of computational tasks and resources in order to meet an unplanned deadline. Typical scenarios triggering an urgent instantiation are related to predicting or reacting to a scenario that is potentially disastrous. In these situations everyone involved in the process must be aware that every minute counts (e.g., to steer and optimise the disaster response in an ongoing crisis), and that meeting hard deadlines may produce outcomes that can be dramatically different both qualitatively and quantitatively when compared to “status quo”, default approaches. For example, being able to give an advance warning of an impending flood even few hours earlier will have a dramatic impact on human impact and can also reduce economic consequences considerably. Similarly, knowing which areas will be worst affected in the coming hours will allow directing equipment and emergency supplies much more efficiently.

The urgent scenario brings up certain additional challenges when compared to scheduled instantiations, both technical and procedural in nature. Freeing up resources (selectively interrupting on-going calculations) and pooling them together quickly enough can present technical challenges. However, frameworks as those reported before may help to overcome (at least some of) the difficulties. On a procedural level, the urgent scenario requires rapid (and stable) communication within and across participating organisations to respond to questions like:

- Who decides that the hazardous event warrants triggering urgent computing procedures?
- Which jobs on which computing resources can be interrupted or postponed?
- How to deal with software licensing issues (e.g., licenses that do not allow remote use)? Under which conditions can exceptions be made?

- Are there any restricting regulations related to use of data (especially personal data)?
- How and when to roll-back which systems in the aftermath of an event?

In major disaster situations these issues can be bypassed as not being relevant for dealing with acute emergency. However, even in these situations it could be beneficial to involve computing centres from countries not directly touched by the emergency (already due to the fact that maintaining computing services is easier in situations where the surrounding infrastructure is intact). On the other hand, bypassing regulations for an event that is thousands of kilometres away may be much more difficult to justify – the sense of urgency ("every minute counts") is much harder to convey to parties not directly influenced by the event. We should also bear in mind that Urgent Computing scenarios may be related to crises that are less obviously urgent than earthquakes or floods. For example, determining the impact of and response to major methane emissions on climate might require computational resources and data sources that are not available as part of the normal climate research. Although the fact that "urgency" may not be strictly tied to any specific deadline helps with the technical setup, it may make the procedural part of the work more challenging.

5. The hydro-meteorological use case

Both instantiation patterns (scheduled and urgent) were successfully applied in the EU-funded DRIHM project to solve numerous challenges related to hydro-meteorology. Hydro-meteorology is a discipline combining aspects from meteorology, hydrology, and hydraulics to produce more accurate models of the impact of extreme weather in terms of flooding and flooding-related phenomena. It is inherently multi-disciplinary and requires multi-model approaches. Until very recently most of the coupling of the models was performed in a manner that required either manual steps in orchestrating component models and transforming the data accordingly for proper ingestion, or the model chain was limited to specific models on each of the stages. In practice, however, it is often desirable to test different meteorological models with different parameters to find the optimal solution for predicting where, when and how much rain will fall in the area being analysed. Similarly, modelling the discharge of water is greatly aided if it is possible to fine-tune the models describing the topology of the area, behaviour of the surface matter (absorption, flow through the matter and surface flows etc.) and other models predicting the behaviour of water as it flows through the catchment area. The impact of the water in the areas prone to flooding is an equally complex modelling challenge – the impact of flood depends on myriad factors ranging from the design of the buildings to specifics of the civil defence procedures (e.g., in terms of installing temporary flood defences or triggering evacuation orders).

The above features of hydro-meteorological modelling already point towards a need for an extensible, modular multi-model workflow system. However, the ability to use "non-traditional" data sources provides additional motivation, which will become more and more important as "Citizen Science" [18], social networks and new Internet-enabled sensors ("Internet of things - IoT") are producing larger amounts of data, accessible through interfaces that support advanced data mining techniques. Already with the amount of weather-related data during the Genoa flash flood of 2011, post-event analysis indicated that including data from sources such as Weather

Underground [19] increased the accuracy of the modelling of this particular extreme event (as presented in [20]). Other projects (such as WeSenseIT [21]) have also performed feasibility studies of including social media data in the analysis of hydro-meteorological phenomena. A practical guide to apply the (DRIHM instantiation of the) Environmental Computing framework is given in [22].

6. Socio-economic impact

If it difficult to overstate the socio-economic importance of the developments in the area of environmental computing, even if we concentrate only on disaster risk reduction. On global level, (as reported in the 2015 UNISDR Global Assessment Report [16]) the *average* economic losses due to disasters such as earthquakes, tsunamis, cyclones and flooding are between 250 and 300 billion dollars. In addition to major disaster scenarios outlined above, so-called extensive risks (minor but recurrent disaster risks) tend to burden low and middle-income countries disproportionately. In addition to considerable losses (estimated at 94b$ in the last decade), these risks are responsible for most of the disaster morbidity and displacement and have a serious detrimental impact on the social and economic development. Hence even a moderate improvement in the accuracy and speed of modelling disaster risks will have a considerable positive humanitarian and economic impact.

It should be noted that the above indicators deal mostly with the direct impacts of different disaster scenarios. A foresight study commissioned by the UK government that concluded its work in 2012 [23], outlined some of the indirect impacts – such as economic “contagion” effects through globalisation, long-term impact of reduced saving/investment incentives and the life-long impact of malnutrition in children during the critical times in development. These factors provide even more profound indicators that show that the investments in environmental computing infrastructure and services is crucial for stability, sustainability and development – both in the industrialised high-income countries as well as in the low to middle-income ones.

With regard to these investments, it is important to keep in mind that concentrated efforts and considerable long-term financial commitments are needed. The Foresight report [23] mentions that high-resolution, multi-model forecasts will require access to supercomputers in the exaflop range. This level of computing power will most likely require pooling international resources and expertise already due to financial constraints and availability of technical competence. Equally important are the organisational and policy developments that allow efficient, cross-border use of data sources – not only in the acute emergency situations, but also for the disaster risk reduction activities.

7. Conclusion and further work

To conclude, the three multi-model frameworks presented in this paper demonstrate both the feasibility and demand of Environmental Computing as a concept. The frameworks present a set of novel features that – while perhaps falling short of the fully automated, “plug and play” multi-model vision – demonstrate the added value of treating models, their execution environments, related metadata and workflows tying them together as entities that can be managed, selected and activated based on logical algorithms (or at least heuristics steering the decision making of the model users). For example in the case of DRIHM, the M.A.P approach and the workflow system that it enabled made a "single click" execution of a complex workflow possible, whereas previously any change of a model component would require at least a days worth of adaptation, testing and verification work. The increased operational efficiency released some of the effort into experimentation, which in turn led into successful tests of the citizen science concepts with considerable potential for further research and development.

However, until the discipline has a clearly identified name (such as "Environmental Computing" put forth in this paper), there is a real risk that these success stories are seen as being relevant only to their specific sub-discipline. However, as the case of CAPRA-GIS illustrates, in the high-level decision making it is crucial to bring together scenarios and impact assessments using a very broad range of tools and approaches. We believe that by finding a common term for discipline that is dealing with linking environmental models together and scaling them up both in terms of the maturity of the solution (e.g. standardised metadata, common approaches to descriptions of execution environments and testing procedures) as well as their ability to use top-tier computational services (up to supercomputers) will greatly speed up the overall maturing of the multi-modelling solutions. As discussed in Section 6, the socio-economic impact of these new capabilities is considerable, even with very modest assumptions related to their applicability and initial effectiveness.

Combining the modular multi-model capabilities (such as ones demonstrated by the DRIHM approach) with the technical and procedural support for urgent computing scenarios (such as illustrated by the MAPPER approach) would have a considerable impact on the speed and accuracy of the short-term response (immediately before, during and after the disaster). However, integrating these capabilities with tools that support long-term planning and policy formation activities (such as CAPRA-GIS) might have even larger impact by providing more detailed and accurate overview of the global risks and possible responses would allow targeting civil protection investments and policies to areas where they will bring most benefits.

To realise this vision, we need to continue aligning the activities on technical, procedural and policy levels to ensure efficient sharing of best practices and generalisation of the Environmental Computing frameworks so that they can support decision making more effectively. Launching the term "Environmental Computing" so that it is recognised by all the key stakeholders is an important first step towards structuring this new discipline.

References

[1] UN General Assembly: Resolution 69/283, adopted by the General Assembly on 3 June 2015, Sixty-ninth session of the UN General Assembly, Sendai Framework for Disaster Risk Reduction 2015 -2030, http://www.preventionweb.net/files/43291_sendaiframeworkfordrren.pdf, accessed 10th July 2015

[2] Open Geospatial Consortim (OGC), WaterML 2.0 Standard, http://www.opengeospatial.org/standards/waterml, accessed 10th Juy 2015

[3] UNIDATA, Network Common Data Form (NetCDF), http://www.unidata.ucar.edu/software/netcdf/, accessed 10th Juy 2015

[4] Joern Birkmann, Stefan Kienberger and David Alexander (eds.): Assessment of Vulnerability to Natural Hazards: A European Perspective, Elsevier 2014, ISBN: 978-0-12-410528-7

[5] M. Schiffers, D. Kranzlmüller, A. Clematis, D. D'Agostino, A. Galizia, A. Quarati, A. Parodi, N. Rebora, M. Morando, N. Rebora, E. Trasforini, L. Molini, F. Siccardi, G. C. Craig, A. Tafferner: Towards a Grid Infrastructure for Hydro-Meteorological Research, Computer Science, Vol. 12, pp. 45-62, 2011

[6] DRIHM: Deliverable D6.2 (Report on application services delivery, 2015), http://www.drihm.eu/images/Deliverable/last/drihm-dwp6.2-20150228-4.0-HRW-Report_on_application_services_delivery.pdf, accessed 10th July 2015

[7] Borgdorff, Joris; Falcone, Jean-Luc; Lorenz, Eric; Bona-Casas, Carles, Chopard, Bastien; Hoekstra, Alfons G.: Foundations of distributed multiscale computing: Formalization, specification, and analysis. In: Journal of Parallel and Distributed Computing 73 (2013), Nr. 4, S. 465–483

[8] N. Trebon: Enabling Urgent Computing within the Existing Distributed Computing Infrastructure, Ph.D. dissertation, University of Chicago, 2011.

[9] CAPRA Probabilistic Risk Assessment Program, CAPRA-GIS, http://www.ecapra.org/capra-gis, accessed 10th July 2015

[10] Marulanda, Mabel, Carreño, Martha, Cardona, Omar, Ordaz, Mario and Barbat, Alex: Probabilistic earthquake risk assessment using CAPRA: Application to the city of Barcelona, Spain, Natural Hazards, 69 (2013), issue 1, p. 59-84

[11] Gregersen, J. B., Gijsbers, P. J. A. and Westen, S. J. P.: OpenMI: Open modelling interface, Journal of Hydroinformatics, 9.3 (2007), 175-191

[12] Harpham, Q. K., Cleverley, P., Kelly, D: The FluidEarth 2 Implementation of OpenMI 2.0, Journal of Hydroinformatics 16.4 (2014), 890-906

[13] A. Parodi: The Genoa 2014 event, Presentation at the National Committee of the Italian Civil Protection Department, 23rd February 2015, Rome Italy, http://www.drihm.eu/images/ppt/drihm_drihm2us_genoa_2014_post_event_analysis_commissione_grandi_rischi_SHORT_ENGLISH.pdf, accessed 10th July 2015

[14] M.B. Belgacem, B. Chopard, J. Borgdorff, M. Mamonski, K. Rycerz, D. Harezlak: Distributed Multiscale Computations Using the MAPPER Framework, in Proceedings of the 2013 International Conference on Computational Science, Vol. 18, pp. 1106-1115.

[15] The United Nations Office for Disaster Risk Reduction (UNISDR), http://www.unisdr.org/, accessed 10th July 2015

[16] UNISDR: Global Assessment Report on Disaster Risk Reduction (GAR) 2015, http://www.preventionweb.net/english/hyogo/gar/2015/en/home/index.html, accessed 10th July 2015

[17] Bulatewicz T., Yang X., Peterson J.M., Staggenborg S., Welch S.M., Steward D.R.: Accessible integration of agriculture, groundwater, and economic models using the Open Modeling Interface (OpenMI): methodology and initial results. Hydrology and Earth System Sciences, 14 (2010), 521-534

[18] European Commission: Green Paper on Citizen Science - Citizen Science for Europe: Towards a better society of empowered citizens and enhanced research, Brussels, Belgium, 2013.

[19] Weather Underground, http://www.wunderground.com/, accessed 10th July 2015

[20] T. Bedrina and M. Heikkurinen: Citizen-scientists' weather observations – importance for society, http://www.drihm.eu/images/documents/CitizenScientist/Citizen-scientists%20weather%20observations%20-%20DRIHM%20vision.pdf, accessed 10th July 2015

[21] WeSenseIt project, http://wesenseit.eu/, accessed 10th July 2015

[22] E. Danovaro, L. Roverelli, G. Zereik, A. Galizia, D. D'Agostino, A. Quarati, A. Clematis, F. Delogu, E. Fiori, A. Parodi, C. Straube, N. Felde, Q. Harpham, B. Jagers, L. Garrote, L. Dekic, M. Ivkovic, E. Richard, O. Caumont: Setup an hydro-meteo experiment in minutes: the DRIHM e-infrastructure for hydro-meteorology research, in IEEE Proceedings of the 10th IEEE International Conference on e-Science, pp. 47-54, 2014.

[23] The Government Office for Science: Foresight Reducing Risks of Future Disasters: Priorities for Decision Makers, Final Project Report, London 2012, https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/286476/12-1289-reducing-risks-of-future-disasters-report.pdf, accessed 10th July 2015