

# 7<sup>th</sup> BSC Severo Ochoa Doctoral Symposium

24<sup>th</sup> Spring 2020 2018

## Book of Abstracts



**Barcelona  
Supercomputing  
Center**  
*Centro Nacional de Supercomputación*



**EXCELENCIA  
SEVERO  
OCHOA**



*Book of Abstracts*  
7th BSC Severo Ochoa Doctoral Symposium

*Editor*  
Carolina Olmo

*Cover*  
Design based on artwork created by macrovector.com

*This is an open access book registered at UPC Commons*  
(upcommons.upc.edu) under a Creative Commons license to protect its  
contents and increase its visibility.

*This book is available at*  
<https://www.bsc.es/education/predoctoral-phd/doctoral-symposium/7th-bsc-so-doctoral-symposium>

*published by*  
Barcelona Supercomputing Center

*supported by*  
The “Severo Ochoa Centres of Excellence” programme

7th Edition, May 2020



# ACKNOWLEDGEMENTS

The BSC Education & Training team gratefully acknowledges all the PhD candidates, Postdoc researchers, advisors and experts for contributing to this Book of Abstracts for the 7th BSC Severo Ochoa Doctoral Symposium 2020. The contribution is very significant because it happens in the dramatic context caused by the Covid19.

BSC Education & Training team  
education@bsc.es



## EDITORIAL COMMENT

We are proud to present the Book of Abstracts for the 7th BSC Severo Ochoa Doctoral Symposium.

During more than ten years, the Barcelona Supercomputing Center has been receiving undergraduate, master and PhD students, and providing them training and skills to develop a successful career. Many of those students are now researchers and experts at BSC and in other international research institutions.

In fact, the number of students has never decreased. On the contrary, their number and research areas have grown and we noticed that these highly qualified students, especially the PhD candidates, needed a forum to present their findings and fruitfully exchange ideas. As a result, in 2014, the first BSC Doctoral Symposium was born.

In this 7th edition of the BSC Severo Ochoa Doctoral Symposium the context has dramatically changed due to the Covid-19 crisis that has made it impossible for it to take place in the usual format. However, we did not want to miss the chance of giving visibility to our PhD student's research and thanks to their enthusiasm even in such a difficult situation the Symposium has taken a digital form. This Book of Abstracts is the result of their contributions.





# WELCOME ADDRESS

I am delighted to welcome all the PhD students, Postdoc researchers, advisors, experts to the 7th BSC Severo Ochoa Doctoral Symposium.

This 7<sup>th</sup> edition of the BSC Severo Ochoa Doctoral Symposium has substantially changed given the current Covid-19 circumstances. Nevertheless, the goal of the occasion continues to be providing a framework to share research results of the projects developed by PhD thesis that use High Performance Computing in some degree.

The symposium was conceived in the framework of the Severo Ochoa Program at BSC, following the project aims regarding the talent development and knowledge sharing. Keeping that in mind, the symposium provides an interactive forum for PhD students considering both the ones just beginning their research and others who have developed their research activities during several years.

As a consequence, I highly appreciate the support provided by BSC and the Severo Ochoa Center of Excellence Programme that make possible to celebrate this event.

I must add that I am very grateful to the BSC directors for supporting the symposium, to the group leaders and to the advisors for encouraging the participation of the students in the event.

And last but not least, I would like to thank all PhD students and Postdoc researchers for their papers and effort. I wish you all the best for your career and I really hope you enjoy this great opportunity to meet other colleagues and share your experiences.

Dr. Maria Ribera Sancho  
Manager of BSC Education & Training



The background of the page is filled with a variety of light gray line-art icons. These icons include lightbulbs (some with radiating lines to indicate they are lit), puzzle pieces, globes, and stylized human figures with their arms raised in a gesture of triumph or excitement. The icons are scattered across the entire page, creating a dense and thematic pattern.

# Abstracts

---

# Model validation and data insertion with FALL3D-8.0: exploiting geostationary satellite retrievals of volcanic ash and SO<sub>2</sub>

Andrew Prata, Leonardo Mingari, Arnau Folch  
 \*Barcelona Supercomputing Center, Barcelona, Spain  
 E-mail: {andrew.prata, leonardo.mingari, arnau.folch}@bsc.es

**Keywords**—Satellite retrievals, volcanic ash, volcanic SO<sub>2</sub>, dispersion modelling, validation, FALL3D

## I. EXTENDED ABSTRACT

The new version of FALL3D has recently been released with several new features and improvements in model physics, solving algorithms, code accuracy and performance [1]. Among the new features are a data insertion scheme and the ability to simulate volcanic SO<sub>2</sub> clouds. The data insertion scheme enables users to initialise model runs from satellite retrievals. This modelling approach is useful for removing uncertainties associated with source term parameters such as the mass flow rate, plume height, source duration and start time. Here we demonstrate and validate the new data insertion scheme in *FALL3D-8.0* using geostationary satellite retrievals of volcanic ash and SO<sub>2</sub>.

### A. Satellite retrievals

1) *Volcanic ash*: The ash detection scheme presented here exploits the reverse absorption signature between 11 and 12  $\mu\text{m}$  and is based on applying successive masks that flag pixels as ‘ash-affected’ before attempting a subsequent quantitative ash retrieval. We use the June 2011 eruption of Puyehue-Cordón Caulle (Chile) as a case study and apply the ash retrieval to SEVIRI (Meteosat-9) measurements. Once pixels have been identified as being ‘ash-affected’ we apply a Look-up Table (LuT) approach [2] to retrieve volcanic ash optical depth ( $\tau$ ), effective radius ( $r_e$ ; in  $\mu\text{m}$ ), and column mass loading ( $m_l$ ; in  $\text{g m}^{-2}$ ). The temperature difference model employed here is based on the forward model developed by [3] and [4]. Uncertainties using this method are estimated to be up to 50% [4], [5].

TABLE I. SUMMARY OF THE SAL AND FMS VALIDATION SCORES.

Validation metrics	S	A	L	SAL	FMS
2011 Cordón Caulle					
24 h	-1.00	-0.22	0.08	1.30	0.42
48 h	-0.83	0.08	0.32	1.24	0.14
72 h	0.46	0.89	0.36	1.71	0.10
2019 Raikoke					
24 h	-0.79	-0.61	0.05	1.46	0.23
48 h	-0.93	-0.91	0.03	1.87	0.20

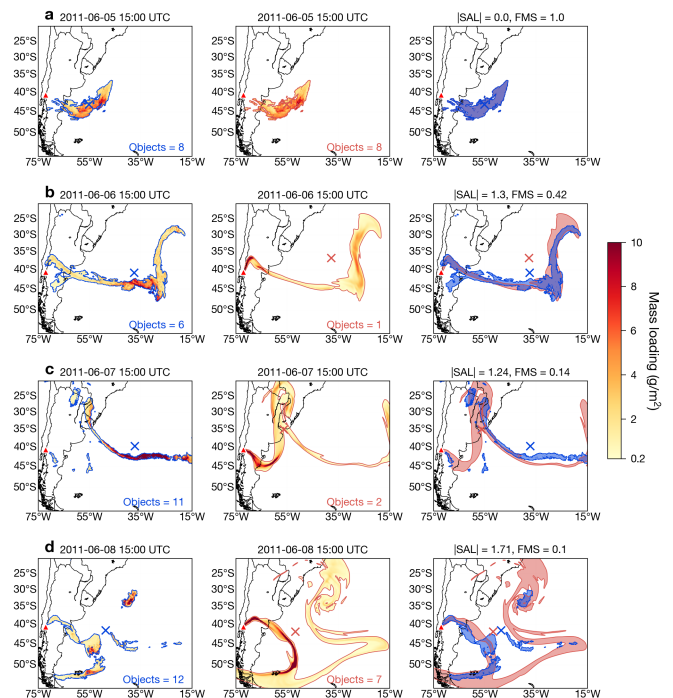


Fig. 1. *FALL3D-8.0* validation of fine ash mass loading using SEVIRI mass loading retrievals. Left panels show satellite retrievals. Middle panels show *FALL3D-8.0* ash simulations. Right panels show spatial overlap of model vs. observed fields.

2) *Volcanic SO<sub>2</sub>*: We apply a three-channel technique to IR geostationary satellite measurements to retrieve total SO<sub>2</sub> column densities in Dobson Units (DU) [6]. This retrieval exploits the SO<sub>2</sub> absorption feature near 7.3  $\mu\text{m}$ . We use the June 2019 eruption of Raikoke (Russia) as a case study and apply the retrieval to AHI (Himawari-8) measurements. To determine whether there is an SO<sub>2</sub> signal in the data, we first construct a synthetic 7.3  $\mu\text{m}$  brightness temperature by interpolating from 6.9 to 11.2  $\mu\text{m}$  in the radiance space and then converting to brightness temperature via the Planck function [6]. One can identify SO<sub>2</sub> clouds by taking the difference between these two variables:

$$\Delta T_{SO_2} = T_{BC}^{7.3} - T_B^{7.3} \quad (1)$$

The  $\Delta T_{SO_2}$  calculated via Eq. (1) is a function of the total column density of SO<sub>2</sub>.

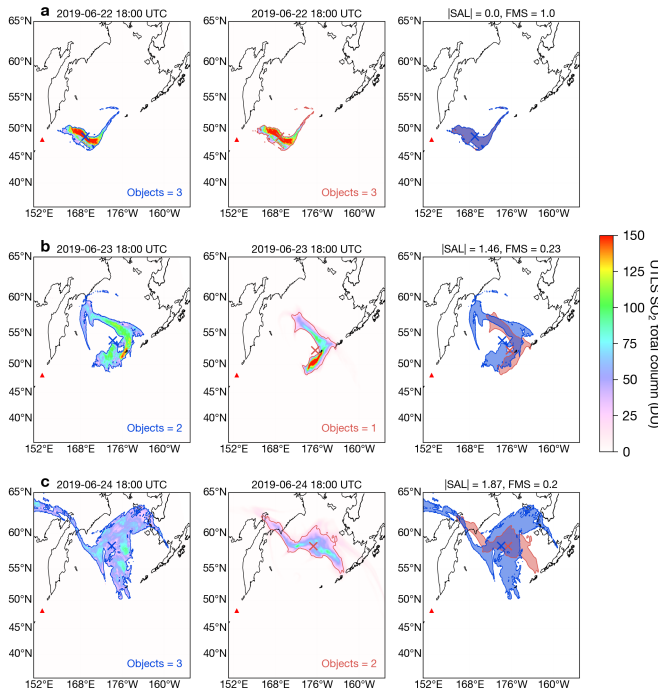


Fig. 2. Same as Fig. 1 but for the Raikoke case study and AHI upper-troposphere lower-stratosphere (UTLS) total column burdens retrievals (DU).

The  $\text{SO}_2$  retrieval is based on constructing this function from offline radiative transfer calculations.

### B. Validation Metrics

We use the Structure, Amplitude and Location (SAL) metric [7] to quantitatively compare satellite retrievals of volcanic ash and  $\text{SO}_2$  to corresponding FALL3D simulations. As in [8] and [9], we also use the Figure of Merit in Space (FMS) score as a complement to SAL for comparing the spatial coverage of observed vs. modelled fields. A detailed mathematical description of the SAL metrics is presented in [7]. SAL varies from 0 (best agreement) to 6 (worst agreement).

### C. Results

1) *2011 Cordón Caulle*: Figure 1 shows how the satellite retrievals and the model simulations compare using data insertion. FALL3D accurately represents the spatial structure of the satellite retrievals with a SAL score of 1.3 and FMS of 0.42 (Fig. 1b) after 24 hours. After 48 hours, the SAL score is 0.77 and FMS is 0.14 (Fig. 1c; Table I). The main difference between the model and observations at this time is in the centres of mass ( $L = 0.32$ ). This is due to a second input of mass used in the Cordón Caulle simulations in addition to the large masses retrieved from the satellite near the centre of the domain (near  $43^\circ\text{S}$ ,  $35^\circ\text{W}$ ). The satellite is likely over-estimating mass in this part of the ash cloud because of the underlying meteorological cloud layer that has not been accounted for in the radiative transfer modelling.

2) *2019 Raikoke*: Figure 2 shows the satellite retrievals and model simulations for the Raikoke case study. Over the first 24 hours the SAL score increases from 0 to 1.46 while the FMS decreases from 1 to 0.23 (Fig. 2b). The SAL score is

mainly affected by the S and A scores whereas the L score is low (0.05) indicating the FALL3D is able to track the centre of mass of  $\text{SO}_2$  very well when initialised with satellite retrievals.

### D. Conclusions

In general FALL3D-8.0 is able to reproduce observations with a high degree of accuracy when initialised using the new data insertion scheme. Both simulations for  $\text{SO}_2$  and ash maintained SAL scores below 2 out to 48 hours after data insertion.

## II. ACKNOWLEDGMENT

This work has been supported by funding from the European Unions Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement H2020-MSCA-COFUND-2016-754433.

## REFERENCES

- [1] A. Folch *et al.*, “FALL3D-8.0: a computational model for atmospheric transport and deposition of particles, aerosols and radionuclides – Part 1: Model physics and numerics,” *Geoscientific Model Development*, vol. 13, no. 3, pp. 1431–1458, Mar. 2020. [Online]. Available: <https://www.geosci-model-dev.net/13/1431/2020/>
- [2] A. J. Prata and A. T. Prata, “Eyjafjallajökull volcanic ash concentrations determined using spin enhanced visible and infrared imager measurements,” *Journal of Geophysical Research: Atmospheres*, vol. 117, 2012.
- [3] A. J. Prata, “Infrared radiative transfer calculations for volcanic ash clouds,” *Geophysical Research Letters*, vol. 16, no. 11, pp. 1293–1296, Nov. 1989.
- [4] S. Wen and W. I. Rose, “Retrieval of sizes and total masses of particles in volcanic clouds using AVHRR bands 4 and 5,” *Journal of Geophysical Research*, vol. 99, p. 5421, 1994. [Online]. Available: <http://doi.wiley.com/10.1029/93JD03340>
- [5] S. Corradini *et al.*, “Mt. etna tropospheric ash retrieval and sensitivity analysis using moderate resolution imaging spectroradiometer measurements,” *Journal of Applied Remote Sensing*, vol. 2, no. 1, p. 023550, 2008.
- [6] A. Prata *et al.*, “Global, long-term sulphur dioxide measurements from TOVS data: A new tool for studying explosive volcanism and climate,” *Geophysical Monograph - American Geophysical Union*, vol. 139, pp. 75–92, 2003.
- [7] H. Wernli *et al.*, “SALA Novel Quality Measure for the Verification of Quantitative Precipitation Forecasts,” *Monthly Weather Review*, vol. 136, no. 11, pp. 4470–4487, Nov. 2008.
- [8] K. L. Wilkins *et al.*, “Using data insertion with the NAME model to simulate the 8 May 2010 Eyjafjallajökull volcanic ash cloud: ASH CLOUD MODELING USING DATA INSERTION,” *Journal of Geophysical Research: Atmospheres*, vol. 121, no. 1, pp. 306–323, 2016.
- [9] A. Marti and A. Folch, “Volcanic ash modeling with the NMMB-MONARCH-ASH model: quantification of offline modeling errors,” *Atmospheric Chemistry and Physics*, vol. 18, no. 6, pp. 4019–4038, 2018.



**Dr Andrew Prata** received his PhD in Atmospheric Physics at Monash University (Australia) in 2017 after completing an Honours Degree at Monash in 2012 and a BSc in Atmosphere and Ocean Science at the University of Melbourne and the University of California Los Angeles (UCLA) in 2010. Prior to completing his PhD, he completed an internship at NASA JPL (USA) and worked as a Research Assistant at the Australian Bureau of Meteorology. In 2018, he completed a postdoc in the Meteorology Department at the University of Reading (UK). He is currently appointed as a Postdoctoral Research Fellow within the Computer Applications in Science and Engineering (CASE) group at the Barcelona Supercomputing Center, Spain.

# dislib: Large Scale High Performance Machine Learning in Python

Javier Álvarez Cid-Fuentes, Rosa M. Badia  
Barcelona Supercomputing Center, Barcelona, Spain  
E-mail: {javier.alvarez, rosa.m.badia}@bsc.es

*Keywords*—*machine learning, python, distributed computing, high performance computing*

## I. EXTENDED ABSTRACT

In recent years, machine learning has proven to be an extremely useful tool for extracting knowledge from data. This can be leveraged in numerous research areas, such as genomics, earth sciences, and astrophysics, to gain valuable insight. At the same time, Python has become one of the most popular programming languages among researchers due to its high productivity and rich ecosystem. Unfortunately, existing machine learning libraries for Python do not scale to large data sets, are hard to use by non-experts, and are difficult to set up in high performance computing clusters. These limitations have prevented scientists from exploiting the full potential of machine learning in their research. In this work, we present dislib [1], a distributed machine learning library on top of PyCOMPSs programming model [2] that addresses the issues of other similar existing libraries.

### A. Overview

As said before, dislib is built on top of PyCOMPSs programming model. PyCOMPSs is a task-based programming model that makes the development of parallel and distributed Python applications easier. PyCOMPSs consists of two main parts: programming model and runtime. The programming model provides a series of simple annotations that developers can use to define potential parallelism in their applications. The runtime analyzes these annotations at execution time, and distributes the computation automatically among the available resources. In essence, dislib is a collection of PyCOMPSs applications that exposes two main interfaces to developers: a distributed data structure called distributed array (ds-array), and an estimator-based API. A ds-array is a 2-dimensional matrix divided in blocks that are stored across different computers. Ds-arrays offer a similar API to NumPy [3], which is one of the most popular numerical libraries for Python. The difference between NumPy arrays and ds-arrays is that ds-arrays are internally parallelized and use distributed memory. This means that ds-arrays can store and process much larger data than NumPy arrays. Implementing a NumPy-like API makes ds-arrays easy to use and intuitive to developers already familiar with NumPy. It also allows to parallelize NumPy codes by just replacing NumPy arrays with ds-arrays.

Distributed arrays serve as a building block for machine learning algorithms included in dislib. These algorithms are

exposed through an estimator-based API inspired by scikit-learn [4], a widely used machine learning library for Python. The typical dislib application consists of the following steps:

- 1) Load data into a ds-array
- 2) Instantiate an estimator object with parameters
- 3) Fit the estimator with the loaded data
- 4) Retrieve information from the estimator or make predictions on new data

Each machine learning model in dislib is enclosed in a class implementing the estimator interface, where an estimator is anything that learns from data that implements two main methods: `fit` and `predict`. Using the estimator abstraction provides a unified API across different algorithms, which reduces the complexity of dislib, and facilitates the implementation of high level applications and utilities that can use different algorithms in an interchangeable manner. Distributed arrays and dislib estimators allow non-expert developers to easily create distributed machine learning applications using regular sequential code. Figure 1 shows an example application that runs the K-means clustering algorithm [5].

```
1 from dislib.cluster import KMeans
2 from dislib.data import load_txt_file
3
4 x = load_txt_file("data.csv", block_size=...)
5 kmeans = KMeans(n_clusters=10)
6 kmeans.fit(x)
7 print(kmeans.centers)
```


Figure 1: dislib code.

### B. Performance evaluation


We compare dislib’s performance to two similar libraries: MLlib and Dask-ML. We use these libraries because they share many characteristics with dislib, including an interface based on estimators, automatic data management and distribution, and support for HPC clusters. We use K-means as a benchmark, which is a popular unsupervised learning algorithm, and we run our experiments on the MareNostrum 4 supercomputer<sup>1</sup>.

Figure 2 shows execution times and speedup of the K-means algorithm in MLlib, Dask-ML, and dislib. We experiment with two different granularities: 50 features and 50

<sup>1</sup><https://www.bsc.es/marenostrum>



(a) 1 billion samples with low granularity (50 features and 50 clusters).



(b) 500 million samples with high granularity (100 features and 500 clusters).

Figure 2: Execution times and speedup of K-means.

clusters (low granularity), and 100 features and 500 clusters (high granularity). We run the algorithm for 5 iterations with randomly generated data divided in 1,536 partitions in all cases. Missing points in the plots mean that the execution failed due to memory issues. In the case of dislib, we run two sets of experiments: using 48 and 12 processes per node (labeled in Figure 2 as *dislib* and *dislib-12* respectively). This is to get a better comparison against Dask-ML, which can only use 12 processes per node due to memory limitations.

We see that dislib outperforms MLib and Dask-ML both in terms of execution time and data size. More precisely, dislib can be up to 4 times faster than MLib and Dask-ML. When running with 192 and 384 cores with 500 million samples, Dask-ML achieves similar performance to dislib with 12 processes per node. This means that some of the difference in performance between dislib and Dask-ML is due to dislib being able to utilize all the available cores. However, when running with 768 and 1,536 cores, dislib with 12 processes per node outperforms Dask-ML. This means that dislib scales better in terms of the number of nodes.

Our evaluation shows that dislib greatly outperforms MLib and Dask-ML, both in execution time and ability to process large data sets. Moreover, running dislib using HPC queue systems like Slurm is completely automatic, whereas MLib and Dask-ML require writing custom deployment scripts, and fiddling with configuration parameters to obtain good performance. Moreover, Dask-ML provides limited code portability as the IP of the Dask scheduler needs to be defined in the source code of the application. In terms of memory management, our experiments show that both Dask-ML and MLib often raise out-of-memory errors when processing large data sets, while dislib is much more robust. Dask-ML can handle larger data sets than MLib, but cannot use all the available computational resources. Unlike dislib, both MLib and Dask-ML required manual configuration of the number of workers per node, and the amount of CPU and memory per worker. This makes MLib and Dask-ML less accessible to non-experts, as finding the optimal configuration for a particular cluster

is a difficult task that requires an extensive trial and error process. Moreover, this means that the performance of MLib and Dask-ML is limited by the ability of the user to tune the configuration parameters.

## II. ACKNOWLEDGMENTS

This work has received funding from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement H2020-MSCA-COFUND-2016-754433. The research leading to these results has also received funding from the collaboration between Fujitsu and BSC (Script Language Platform).

## REFERENCES

- [1] J. Álvarez Cid-Fuentes, S. Solà, P. Álvarez, A. Castro-Ginard, and R. M. Badia, “dislib: Large Scale High Performance Machine Learning in Python,” in *Proceedings of the 15th International Conference on eScience*, 2019, pp. 96–105.
- [2] F. Lordan, E. Tejedor, J. Ejarque, and et al., “ServiceSs: an interoperable programming framework for the Cloud,” *Journal of Grid Computing*, vol. 12, no. 1, pp. 67–91, 2014.
- [3] S. van der Walt, S. C. Colbert, and G. Varoquaux, “The NumPy Array: A Structure for Efficient Numerical Computation,” *Computing in Science & Engineering*, vol. 13, no. 2, p. 22, 2011.
- [4] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and Édouard Duchesnay, “Scikit-learn: Machine Learning in Python,” *Journal of Machine Learning Research*, vol. 12, no. Oct, pp. 2825–2830, 2011.
- [5] J. A. Hartigan and M. A. Wong, “Algorithm AS 136: A K-Means Clustering Algorithm,” *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, vol. 28, no. 1, pp. 100–108, 1979.



**Javier Álvarez Cid-Fuentes** is a researcher at the Workflows and Distributed Computing group of the Barcelona Supercomputing Center. His research interests include large scale distributed machine learning and parallel programming models for distributed infrastructures. Álvarez Cid-Fuentes received a Ph.D. in computer science from the University of Adelaide in 2018.

# Decadal Climate Prediction at the BSC

Simon Wild<sup>#1</sup> and Roberto Bilbao<sup>#</sup>

<sup>#</sup>Barcelona Supercomputing Center

<sup>1</sup>simon.wild@bsc.es

## Decadal Prediction, Climate Prediction, Climate Variability

### A. Introduction

Initialised decadal climate predictions have been made available for users as a potential source of near-term climate information with the aim of supporting climate-related decisions in key economic and societal sectors such as energy, agriculture and insurance. Observed climate variability on the decadal timescale can be described as the superimposition of an anthropogenically-driven trend on natural variability of the climate system. The trend can be considered to be driven by changes in anthropogenic emissions (mainly greenhouse gases and anthropogenic aerosols). Natural variability is generated internally by interactions between and within different components of the climate system (atmosphere, ocean and sea ice) or by external factors such as volcanic eruptions and solar activity. The variability modes on timescales of several years and longer can then provide a source of potential predictability and thus lead to skill of decadal predictions. In this context, there is a growing interest from many stakeholders for climate services on 1-10 year timescales, but some efforts are still needed from the climate science community to assess the forecast quality on such timescales.

### B. Study Overview

This study will provide a quality assessment of the recently produced retrospective decadal forecasts (hindcasts) at the Barcelona Supercomputing Center (BSC) using the EC-Earth coupled global climate model. This set of ensemble hindcasts is BSC's contribution to the Decadal Climate Prediction Project (DCPP) Component A, comprising 10 members full-field initialised yearly on 1st November from 1960 to 2018 and covering 11 forecast years. In plain language we will assess the question whether our decadal predictions are any good or more precisely:

- Are the initialized decadal predictions better than the non-initialized simulations?
- Can we identify reasons if the answer was “no” for the questions above for a given geographical reason?

### C. Results

We will show typical measures for the quality of deterministic and probabilistic hindcasts, such as anomaly correlation coefficients of standard variables such as near surface temperatures, verified against observations and reanalysis products. In this context the effect on verification measures by mimicking observational products for near surface temperatures, combining sea surface temperatures with near surface air temperatures over land, will be discussed. We will also assess the impact of initialization on forecast reliability quantifying the statistical relationship between the predicted probabilities and the observed relative frequency of an event.

Initialized hindcasts will further be compared to an ensemble of non-initialized historical simulations assessing the potential added value of initializing the model towards the observed climate state in near-term climate predictions (one example is figure 1). Over few ocean regions initialized

predictions show significant skill beyond the relation to trends due to external forcing. We will also discuss the possibility of an overestimation of the added value due to different behaviour from the observed climate in non-initialized historical simulations.

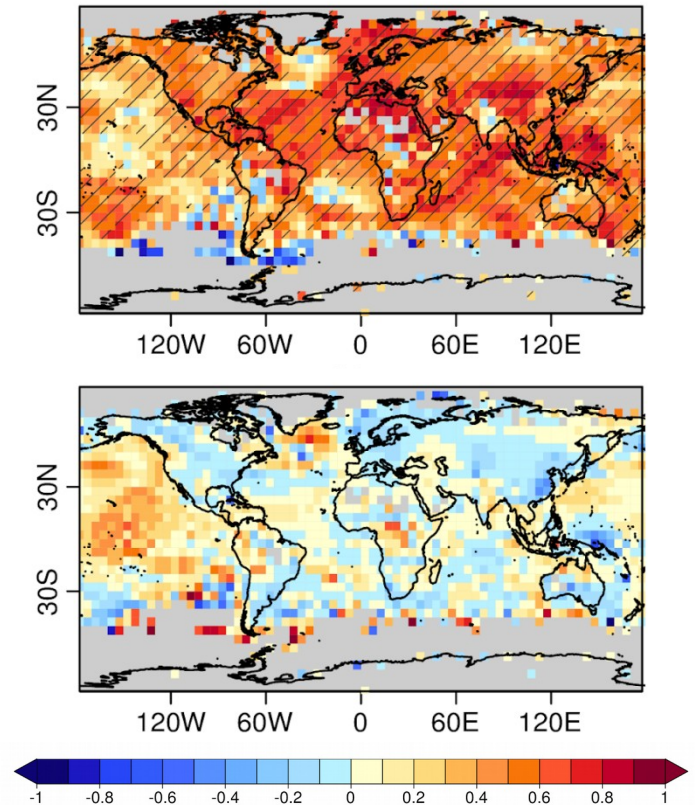


Fig. 1 Top: Anomaly correlation coefficient (ACC) between first forecast year of ensemble mean of initialized decadal predictions and observations for surface temperature and 1961-2018. Hatching indicates significant correlation (at the 95% level). Bottom: Difference between ACC of initialized and non-initialized simulations

### D. Summary

The recently produced decadal predictions at the BSC show added value of initialization for some standard variables for some regions especially for early forecast years, e.g. for surface temperature over the Pacific in forecast year 1 (figure 1, bottom). Our hindcasts show however barely any benefit from initialization beyond the first forecast year. The ocean dynamics in the North Atlantic region in response to initialization can be shown to play a crucial role why the added value of initialization is limited.

### E. ACKNOWLEDGEMENTS

We are grateful for the contributions of Juan C Acosta-Navarro<sup>#</sup>, Arthur Amaral-Ramos<sup>#</sup>, Pierre-Antoine Bretonnière<sup>#</sup>, Ruben Cruz-García<sup>#</sup>, Francisco J Doblas-Reyes<sup>#</sup>, Markus G Donat<sup>#</sup>, Pablo Echevarría<sup>#</sup>, Eduardo Moreno-Chamarro<sup>#</sup>, Pablo Ortega<sup>#</sup>, Yohan Ruprich-Robert<sup>#</sup>, Valentina Sicardi<sup>#</sup>, Etienne Tourigny<sup>#</sup>, Javier Vegas<sup>#</sup>, Deborah Verfaillie<sup>#</sup> who are all co-authors of this study.



This project received funding from the European Union's research and innovation programme under the Marie Skłodowska-Curie grant agreement no H2020-MSCA-COFUND-2016-75433 and the EUCP project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement no 776613.

## **Author biography**

**Simon Wild** was born near Munich, Germany, in 1985. He received his B.Sc in Meteorology in 2009 and subsequently his M.Sc in Meteorology in 2012 at the Freie Universitaet Berlin, Germany. Simon worked as a climate scientist at the School of Geography, Earth and Environmental Sciences at the University of Birmingham, UK from 2011 until 2018 where he also received his PhD (entitled *North Atlantic Winter Wind Storm Variability across different Time Scales*) in 2018. He then joined the Climate Prediction Group at the Barcelona Supercomputing Center where he has been working since.

# Assessment of the Impact of the Covid-19 Lockdown on Air Pollution over Spain using Machine Learning

Hervé Petetin<sup>\*1</sup>, Dene Bowdalo<sup>\*2</sup>, Albert Soret<sup>\*3</sup>

<sup>\*</sup> *Barcelona Supercomputing Center (BSC)*

{<sup>1</sup>herve.petetin, <sup>2</sup>dene.bowdalo@bsc.es, <sup>3</sup>albert.soret}@bsc.es

**Keywords**— air quality, machine learning, covid-19

## EXTENDED ABSTRACT

The rapid spread of the Covid-19 pandemic over Spain recently forced the Spanish authorities to implement drastic measures of social distancing through a strict lockdown of the population starting on March 14<sup>th</sup>. As hospitalizations were still strongly increasing, a second and more stringent phase of the lockdown was implemented from the March 30<sup>th</sup> to April 9<sup>th</sup> with workers from all non-essential economical activities forced to stay at home.

This situation has impacted numerous activity sectors, including road transport, air traffic and part of the industries. As a consequence, air pollutant emissions have been greatly reduced. Although such a large change of emission forcing is expected to reduce the air pollutant concentrations in Spanish urban areas, the extent of such reductions remains uncertain. Key to this uncertainty are the highly variable meteorological conditions that can either attenuate or amplify changes of air pollution concentration originally driven by changes of emissions. Thus, assessing the impact of the Covid-19 lockdown solely based on the analysis of the concentration time series can often be misleading since at least part of the variability is driven by the meteorology.

In this study, we explore the use of machine learning algorithms for estimating the business-as-usual NO<sub>2</sub> concentrations that would have been observed without the Covid-19 lockdown based on ERA5 meteorological data and additional time features. Trained on past data, these ML models can learn the complex relationships between meteorology and NO<sub>2</sub> concentrations, indirectly assuming an average emission forcing. By using these ML models to predict the NO<sub>2</sub> concentrations under the current situation (with very different emission forcing), we expect the discrepancies between predictions and observations to be related to a large extent to the reduction of emissions induced by the lockdown regardless of the meteorological conditions.

In this study, the reduction of NO<sub>2</sub> pollution is investigated in all 50 provinces of Spain.

## A. Data

We selected one air quality (AQ) surface monitoring stations in each of the 50 Spanish provinces (Figure 1). Stations are selected according to several criteria of interest, including high population density and high data availability. NO<sub>2</sub> measurements are extracted from the BSC in-house GHOST (Globally Harmonized Observational Surface Treatment) project (Bowdalo et al., in preparation). GHOST gathers numerous publicly available atmospheric datasets together with harmonized detailed metadata and quality assurance (QA) flags. It aims at facilitating a greater quality of data analysis in the atmospheric chemistry community. In this study, we used the EEA AQ eReporting dataset, and applied a

careful QA screening to the entire NO<sub>2</sub> dataset in order to remove unrealistic or very suspicious values.

Meteorological data are taken from ERA5 reanalysis dataset, whose spatial resolution is 31 km globally. We consider the following parameters at the daily scale: daily mean, minimum and maximum 2-m temperature, surface wind speed, 10-m zonal and meridian wind speed components normalized by 10-m wind speed (in order to provide information solely on wind direction), surface solar radiation downward, surface pressure and total cloud cover.

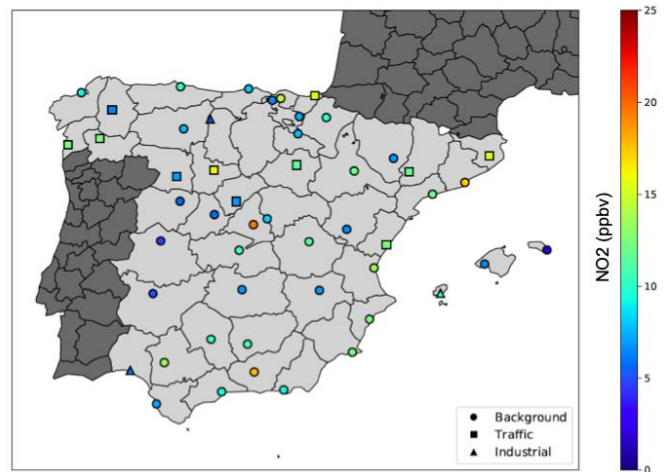


Fig. 1. Climatological (01/01/2013-06/04/2020) NO<sub>2</sub> mixing ratio at the stations selected in each of the 50 provinces of Spain (Canarias islands not shown).

## B. Methodology

In this study, we used the Gradient Boosting Machine (G) machine learning model. One ML model is developed at each of the 50 stations. The target is the daily mean NO<sub>2</sub> mixing ratio. Following the studies of Grange et al. (2018) and Grange and Carslaw (2019), we considered a set of features including the previously mentioned meteorological variables and several time features, namely the Julian date (as the seasonal term) and the weekday.

The complexity of using ML models for predicting pollutant concentrations in this case relies in its typical non-stationarity, as emissions often following trends due for example to the implementation of more restrictive regulation. In order to limit this potential issue, we trained the ML models only over the 3 years from 2017 to 2019, assuming that trends over that short period are relatively low compared to the variability induced by the meteorology, and let the year 2020 for testing. In this way, the performance of the ML model can be assessed looking at the results in 2020 before the lockdown (not after since emission forcing is different). To estimate the uncertainty of our ML models, we replicated our approach over the past years: training over 2013-2015 and calculation

of the uncertainty over January-April 2016; then idem for 2014-2016, 2015-2017 and 2016-2018.

### C. Results

Preliminary results indicate that the performance of the ML models is quite heterogeneous among the different regions. In many regions, the discrepancies with observed concentrations (before lockdown) remain within the uncertainty range. However, some bias is often highlight, in some cases larger than the uncertainty. Discrepancies may be due to deficiencies in the ML model and/or changes of emissions in the vicinity of the station. However, in most regions, the predictions followed reasonably well the variability of the observed concentrations, which suggests that the ML models are capturing well the influence of the meteorological conditions.




Fig. 2 Seven-days running average of NO<sub>2</sub> mixing ratios at Madrid (top panel) and Barcelona (bottom panel) from January to present day. Observed concentrations are shown in blue, predicted business-as-usual (BAU) NO<sub>2</sub> concentrations are shown in red, with a 95% confidence interval. The NO<sub>2</sub> concentrations from all previous years since 2013 are shown in grey. The lockdown period starts after the vertical black bar. Note that ML predictions were not possible after late March due to not yet available ERA5 data.

First results in Madrid and Barcelona regions are shown in Figure 2. In Madrid, before lockdown, the ML model performs quite well but with a moderate positive bias. It is not able to capture all the events, and the errors found during the first weeks of February are found to exceed the uncertainty range estimated from previous years. Including ERA5 information on the planetary boundary layer height and/or atmospheric stability may help solving this issue. After the lockdown, observed NO<sub>2</sub> mixing ratios quickly decrease, down to values around 8 ppbv, that have never been observed at this period of the year since 2013. The ML model predicts business-as-usual NO<sub>2</sub> mixing ratios of about 15 ppbv over that period. The difference with observed concentrations is found to be statistically significant at a 95% confidence level, with a best estimate of -47%.

Similar results are obtained in Barcelona where the ML model performs better as the observations are within the predicted uncertainty range before lockdown. The observed NO<sub>2</sub> concentrations dropped down to 8 ppbv while ML predictions exceed 16 ppbv, which means a reduction of 50%, in agreement with what is observed in Madrid.

In some other regions, results are more complex to interpret given the potential deficiencies of the ML models. Discrepancies between observations and business-as-usual predictions during the lockdown often remain within the uncertainty interval of the ML model, which prevents from concluding on statistically significant reductions of NO<sub>2</sub> pollution. Including more recent data may help to detect more

significant changes.

### D. Conclusion and Perspectives

Further efforts are obviously required for analyzing the results obtained in the different regions of Spain but these preliminary results suggest that the methodology is suitable, at least in some regions, for identifying and quantifying the reduction of NO<sub>2</sub> pollution that is not due to meteorological conditions. It is worth noting that part of the discrepancies previously highlighted may be due not only to changes in the emissions but also to changes in the chemistry, although we assume this impact is lower.

Processing the most recent data, during the second phase of the lockdown, should soon provide us a more complete picture of the impact of these reduced emissions on the NO<sub>2</sub> levels in Spanish urban areas. Including additional ERA5 variables may also help to improve the skills of the ML models, which is required for confirming the statistical significance of NO<sub>2</sub> reduction in some regions.

### E. ACKNOWLEDGEMENTS

This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement H2020-MSCA-COFUND-2016-754433.

### References

- [1] Grange, S. K. and Carslaw, D. C.: Using meteorological normalisation to detect interventions in air quality time series, *Science140f The Total Environment*, 653, 578–588, doi:10.1016/j.scitotenv.2018.10.344, 2019.
- [2] Grange, S. K., Carslaw, D. C., Lewis, A. C., Boleti, E., and Hueglin, C.: Random forest meteorological normalisation models forSwiss PM 10 trend analysis, *Atmospheric Chemistry and Physics*, 18, 6223–6239, doi:10.5194/acp-18-6223-2018, 2018

### Author biography



**Hervé Petetin** was born in Caen, France in 1986. He holds an engineering diploma from the Ecole Centrale de Lille, France, a M.Sc. in Mechanics and fluid dynamics from the University of Science and Technology Lille, a M.Sc. in Atmospheric physics and chemistry from the University of Paris Est Creteil, France, and a Ph.D. in Atmospheric physics and chemistry from the University of Paris Diderot, France.

His research has first focused on the study of air quality and more specifically the aerosol pollution in large megacities, using regional chemistry-transport models and in-situ observations. As a member of the IAGOS European Research Infrastructure, he then investigated during several years the variability and trends of two important intermediate-lifetime gaseous compounds, namely ozone and carbon monoxide, in the troposphere based on airborne in-situ observations provided by the IAGOS fleet.

In 2018, he obtained a postdoctoral funding at the BSC from the STARS program (Marie-Sklodowska-Curie Action

COFUND program) for exploring the use of machine learning in air quality (including the correction of air quality forecasts).

# Exploring Gender Differences in Bioinformatics Research

Jorge Saldivar<sup>1</sup>, Nataly Buslón<sup>2</sup>, María José Rementería<sup>3</sup>

*Social Links Analytics Group, Life Science Department, Barcelona Supercomputing Center (BSC), Barcelona, Spain*

<sup>1</sup>jorge.saldivar@bsc.es, <sup>2</sup>nataly.buslon@bsc.es, <sup>3</sup>maria.rementeria@bsc.es

**Keywords**— gender inequality, bioinformatics, science of science

## EXTENDED ABSTRACT

Women’s underrepresentation in science is an unfair situation apparent in different research fields. This article examines the problem of gender bias in the field of Bioinformatics. Through a quantitative study of more than 40K articles, we discovered that even when the difference between women and men representation in academic publications has been slightly reduced in the last years, men still dominate the scientific activity. We also found that publications led by female researchers are less recognized in terms of citations than articles led by male scholars.

### A. Introduction

The underrepresentation of women in the scientific community is noticeable across research fields [1]. Women scientists publish less than men and have fewer opportunities than men to participate in international collaborations. In developed countries, articles with women in leading authorship positions receive fewer citations than those with men in the same positions [2]. By leveraging quantitative methods, this article explores aspects of gender differences in bioinformatics research with the aim to unveil unfair situations of gender inequalities as well as study whether already published gender biases in other fields of research are also present in bioinformatics.

### B. Methods

Five representative journals in the field of Bioinformatics were chosen for the study, namely Oxford Bioinformatics, Plos Computational Biology, Nucleic Acids Research, BMC Bioinformatics, and BMC Genomics. Meta-data of articles (e.g., title, author names, publication year, abstract, number of citations, PubMed identifier) published in these journals were extracted from Scopus—one of today’s most complete repositories of scientific manuscripts. The data collection process was run in August 2019.

The search engine of Scopus was queried to find articles published between 2005 and 2017 in the selected journals. The lower-bound limit in the range of years was chosen because the journal Plos Computational Biology appeared in 2005. Also, we decided to limit our search to 2017 because we wanted to leave enough time (one full year) to ensure that Scopus already indexed the majority, if not all, of articles of the selected journals by the moment of the search.

In total, meta-data of 47,427 articles and their corresponding authors were collected. After getting rid of duplicates (194) and entries without DOI (401), 46,832 records were considered for further analysis. Table 1 summarizes the data collected per journal.

The gender of the 143,960 authors who produced the 46,832 articles was inferred using the services NamSor and gender-guesser. We hit the API of NamSor asking for the

gender of authors by providing the authors’ first and last names. NamSor API returns male, female, or unknown. In case NamSor fails to identify the gender, the python package gender-guesser was employed as a backup to find out the gender of authors. The gender of 27,705 authors (19%) could not be detected by either any of the two gender identification tools. In 10% of the cases, gender could not be identified because we could not get the author’s name. For the rest, we empirically found that, as also reported in [3], Asian names are challenging for gender identification services. After removing authors with unidentified gender, 116,255 authors left in the dataset.

TABLE I  
SUMMARY OF THE DATA COLLECTED PER JOURNAL

Journal	ISSN	No. of articles
Oxford Bioinformatics	1460-2059	8,466
Plos Computational Biology	1553-734X	5,121
Nucleic Acids Research	1362-4962	15,489
BMC Bioinformatics	1471-2105	7,703
BMC Genomics	1471-2164	10,053
Total articles		46,832

We found that 1,382 of the 46,832 articles either do not provide information about their authors or all of their authors have unknown gender. These records were removed from the dataset, remaining 45,450 articles for the analysis.

### C. Results

A mean of 4.8 of people authored articles published in the journals of interest, and this average shows a moderate increase from 4.1 in 2005 to 5.1 in 2017. As expected and in line with several previous research [4][5], the majority of the 116,255 authors are male (64.7%), while the rest 35.2% are female.

#### a. Gender distribution

In 60% of the articles (29,926), male researchers dominate the list of authors. More than one-third of the articles were authored only by men (15,583) while articles published exclusively by women represent 3.91% of the total. Female researchers lead 26.26% of articles, whereas males appear as first authors in almost 60% of articles. Percentages do not sum up to 100 because, due to limitations in the method used to infer authors’ gender, we do not have gender information about all first authors. The underrepresentation of female researchers is even more noticeable when analyzing the last

authors of articles. In this case, the proportion of female researchers that participate as last authors is less than 20%.

Results show that gender differences gradually decrease over time. Figure 1 shows how the representation of female authors within the total authors increases from 2005 to 2017. The proportion of total authors passed from .27 in 2005 to .32 in 2017. The different actions taken by governments, research centers, universities, and civil institutions to promote the participation of women in science might have impacted the increment. However, male researchers are still clearly dominating the group of authors.




Fig. 1 Evolution of the proportion of male and female authors over time. Women proportion is shown in purple, whereas men proportion is depicted in green.

#### b. Productivity and citations

In general, the 45,450 articles received a median of 28 citations. The top ten most-cited authors are all males, and they accounted for almost 2% of the total citations (223,053 out of 13,368,610).

The average citation by paper within the group of male authors is 46.2, while in female researchers is 43.12. Articles authored exclusively by women received on average fewer citations (33) than articles in which only men participated (54), and the difference is statistically significant ( $p$ -value < .001,  $\alpha$ =.05). We found a similar situation when analyzing articles led by a female researcher and articles in which a woman appears as the last author. The difference between the average citations in both cases showed to be statistically significant ( $p$ -value < .001,  $\alpha$ =.05).

Authors have, in general, a mean of productivity of 1.89 articles. In the case of male researchers, they have productivity of 2.05 articles while females 1.60 articles. As in the case of citations, the top-ten most productive authors are all male. The first woman appears in the position 13th in the productivity ranking (i.e., sorting authors by their number of articles in descending order).

Two females are part of the list of the top ten authors that participated in articles as the last author. As found in previous research [6], articles with females as last authors are less frequently cited (mean=38.50) than articles with males as last authors (mean=50.64), and the difference is statistically significant ( $p$ -value < .001,  $\alpha$ =.05). Contrary to what Bendels et al. [6] have found, the difference does not get larger as more people authored articles. Similarly, we discovered that articles led by female researchers got significantly fewer citations on average than articles with male scholars as first authors ( $p$ -value < .001,  $\alpha$ =.05). Besides, our analyses show that articles authored only by women

received significantly fewer citations on average than articles published only by male authors ( $p$ -value < .001,  $\alpha$ =.05).

#### D. Conclusion and Future Enhancement

Although the last years show a reduction in the inequalities between men and women in scientific production, our results confirm that still there is a noticeable gender difference in favor of male researchers. In future work, we will deep dive into the presented results to further explore other dimensions of the gender bias in bioinformatics.

#### E. ACKNOWLEDGEMENTS

European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement H2020-MSCA-COFUND-2016-754433.

#### References

- [1] Bonham, Kevin S., and Melanie I. Stefan. "Women are underrepresented in computational biology: An analysis of the scholarly literature in biology, computer science and computational biology." *PLoS computational biology* 13, no. 10 (2017).
- [2] L. Vincent, C. Ni, Y. Gingras, B. Cronin, and C. R. Sugimoto. "Bibliometrics: Global gender disparities in science." *Nature News* 504, no. 7479 (2013): 211.
- [3] K. Fariba, C. Wagner, F. Lemmerich, M. Jadidi, and M. Strohmaier. "Inferring gender from names on the web: A comparative evaluation of gender detection methods." In *Proceedings of the 25th International Conference Companion on World Wide Web*, pp. 53-54. 2016.
- [4] H. Luke, D. Stuart-Fox, and C. E. Hauser. "The gender gap in science: How long until women are equally represented?." *PLoS biology* 16, no. 4 (2018): e2004956.
- [5] H. Junming, A. Gates, R. Sinatra, and A. Barabási. "Historical comparison of gender inequality in scientific careers across countries and disciplines." *Proceedings of the National Academy of Sciences* 117, no. 9 (2020): 4609-4616.
- [6] B. Michael, R. Müller, D. Brueggmann, and D. Groneberg. "Gender disparities in high-quality research revealed by Nature Index journals." *PloS one* 13, no. 1 (2018).

#### Author biography



**Jorge Saldivar** is a doctor in Information and Communication Technologies by the University of Trento, Italy. He is a postdoctoral researcher in the Social Link Analytics group at the BSC. Jorge has a broad interest across data science, especially in research that brings together quantitative and qualitative methods. Previously, he was a Data Scientist Fellow at the University of Chicago, USA, and worked as a Postdoctoral Research Fellow in the Department of Electronics and Informatics at the Catholic University of Asuncion, Paraguay.

# A compromise Archive Platform for Monitoring Infrastructures

Carlos Garcia Calatrava\*<sup>†</sup>, Fernando Cucchietti\*, Yolanda Becerra\*<sup>†</sup>

\*Barcelona Supercomputing Center, Barcelona, Spain

<sup>†</sup>Universitat Politècnica de Catalunya, Barcelona, Spain

E-mail: {carlos.garcia, fernando.cucchietti, yolanda.becerra}@bsc.es

**Keywords**—*Big Data, Stream processing, Polyglot Persistence, Monitoring, Archival, IoT.*

## I. EXTENDED ABSTRACT

The great advancement in the technological field has led to an explosion in the amount of generated data. Many different sectors have understood the opportunity that acquiring, storing, and analyzing further information means, which has led to a broad proliferation of measurement devices. Those *sensors'* typical job is to monitor the state of the enterprise *ecosystem*, which can range from a traditional factory, to a commercial mall, or even to the largest experiment on Earth[1].

Big enterprises (BEs) are building their own big data architectures, usually made out of a combination of several state-of-the-art technologies. Finding new interesting data to measure, store and analyze, has become a daily process in the industrial field.

However, small and medium-sized enterprises (SMEs) usually lack the resources needed to build those data handling architectures, not just in terms of hardware resources, but also in terms of contracting personnel who can master all those rapidly evolving technologies.

Our research tries to adapt two world-wide-used technologies into a single but elastic and moldable one, by tuning them, to offer an alternative and efficient solution for this very specific, but common, scenario.

### A. Introduction: A historical approach

#### From one-size-fits-all to one-size-for-each

Traditionally, databases have been considered a passive asset: OLTP systems ingested structured data, in order to facilitate daily operations, and the relational model was considered, *de facto*, the standard model.

As soon as hardware evolved, organizations realized the real potential of data and several different technologies emerged, improving the handling and storage of the data in a wide range of scenarios.

In not many years, databases moved from one-size-fits-all[2] to one-size-for-each, where each scenario had a very specific and efficient data model, and each data model had a plethora of different databases to choose from. For example, Graph databases enabled the full potential of social

networks and key-value databases became crucial in huge online marketplaces[3].

#### Polyglot Persistence

Alongside the increasing amount of new Data Analysis and Machine Learning algorithms, BEs encountered a new problem: Systems were evolving and, sometimes, choosing one single data model was not enough. Thus, BEs started to put in practice polyglot persistence[4]: Multiple data storage technologies were incorporated, chosen based on the way data was used by individual applications. Hence, applications started to benefit from different data models at the same time.

### B. Monitoring Infrastructures and archive databases

In this research, we understand as Monitoring Infrastructure a set of devices, usually called sensors, where each supervises the state of a specific asset alongside time. The global reporting of those sensors is able to describe the state of the whole system, at a given point in time. Archive databases are meant to store the data that a Monitoring Infrastructure produces. The data obtained during this process is crucial for performing tasks such as predictive maintenance or forecasting.

Unlike BEs, SMEs are hardly ever able to implement Polyglot Persistence: Having an expert in each database technology is difficult. Moreover, some companies cannot afford to keep the historical data forever, not just in a single data model, as it becomes huge. Thus, they usually implement a *sliding window* technique, where only a fixed amount of data is kept: As new data is stored, the oldest data is removed. This generates two main problems: Firstly, data is stored just in a single data model, and analyzing it becomes really slow. Secondly, data is being discarded, so important information is lost.

Finally, for preventing data loss, databases are commonly replicated in different machines, which ends up in an impressive consumption of hardware resources.

### C. Soft Goals

The resulting platform should meet the following soft requirements:

- Allow the ingestion, in real time, of monitoring data
- Use as few different technologies as possible, while enabling fast analysis
- Minimize the amount of storage needed

#### D. Proposed solution

The platform was built with just two different well-known big data technologies. Also, a Java/Python API is provided in order to interact with the platform.

##### Apache Kafka

Apache Kafka is a distributed streaming platform, used for building real-time data pipelines and streaming apps[5].

In this scenario, Kafka helps in two different actions: It facilitates the stream data ingestion and, also, thanks to a specifically designed Java consumer, compacts and re-structures the data before sending it to the database.

##### MongoDB

MongoDB is the most popular NoSQL database[6]. It follows a schema-less design: Data Architects define a basic schema, but altering it or adding new fields does not enforce any global-schema modification, which allows a great flexibility. Hence, each record is, theoretically, schema-independent from the other ones.

In this platform, MongoDB was tuned for handling a columnar schema-full data model, apart from the standard document schema-less data model. Thanks to this modification, MongoDB behaves like a Polyglot Persistence system itself, but just using a single database technology.

This not only reduces the number of used technologies and needed experts, but also allows Polyglot Persistence Intra Communication: In traditional polyglot persistence systems, each data-model is held in a different technology, making the communication between them fairly difficult. By placing different data models under the same logical database, it is even possible to perform queries that benefit from different data models at the same time.

Furthermore, it also tackles the problem of database replication. Although both data models are in the very same logical database, they can be placed in different physical databases, by sharding. Thus, a data model can serve as a replica for the other one, and vice-versa: If a machine fails, its data can be recovered from its *sibling* database.

#### E. Experiment design

In the scope of this extended abstract, the experiment set up consisted in feeding the archiving platform with a 4-month pseudo-streaming factory simulation, that incorporates 3000 different sensors, producing data every 30 seconds. More than 1000 Million records were generated. The system was limited to 4 GB of RAM, and data was stored in a traditional HDD using the ZSTD compression technique.

The performance metrics were obtained in independent executions, by clearing RAM and cache. Each execution was performed 3 times, keeping the AVG.

The executions consisted in running three common queries. Each query asked for the data of 10 random sensors, either for 500 random timestamps or for a 3-month time range. Also, this 3-month time-sequential query was divided into two different queries: The first one obtained the RAW values, and the other one *resampled* them, obtaining the mean value per sensor and hour.




Fig. 1. Query execution time comparison for single vs multi-schema database.

#### F. Evaluation

As seen in Figure 1, by adding a columnar schema-full storage to the database, time-sequential queries reduce their execution time drastically, as this model captures the relationship between same-sensor measurements.

However, for Random Access queries, the document schema-less design continues being preferred, as it captures the relationship between same-time measurements.

Finally, regarding the disk consumption, the document schema-less design consumed 15.8GB (2-replica + arbiter), while the double schema design needed 12.8GB.

#### G. Conclusion and future work

The built Archiving Platform keeps a compromise between needed resources and efficiency, while providing an easy-to-use API. It has shown to speed up querying while reducing disk usage, in comparison to traditional schema-less designs.

As future work, it will be important to implement queries that benefit from both data models at the same time, and to add an automatic data recovery mechanism between data models, thus allowing a fully functional replica set.

#### REFERENCES

- [1] P. Golonka *et al.*, "Future archiver for cern scada systems," in *ICALEPCS '17*.
- [2] M. Stonebraker, "Technical perspective: One size fits all: An idea whose time has come and gone," 2008.
- [3] (2020) Amazon key-value usage website. [Online]. Available: <https://aws.amazon.com/en/nosql/key-value/>
- [4] M. Fowler and P. Safalage. (2012) The future is polyglot persistence. [Online]. Available: <https://martinfowler.com/articles/nosql-intro-original.pdf>
- [5] (2020) Apache kafka website. [Online]. Available: <https://kafka.apache.org/>
- [6] (2020) Database engines website. [Online]. Available: <https://db-engines.com/en/ranking>



**Carlos Garcia Calatrava** received his BSc degree in Informatics Engineering from the Facultat d'Informàtica de Barcelona (FIB), Universitat Politècnica de Catalunya (UPC), Spain, in 2016. Two years later, he completed his MSc degree in Innovation and Research in Informatics from the very same university. Regarding his professional background, Carlos has worked in the European Organization for Nuclear Research (CERN, Switzerland) and in the National Institute of Informatics (NII, Japan), among others. Currently, he is working at CASE-BSC, while

pursuing a PhD in Computer Architecture.



# Towards Integrated Urban Simulations

Irene Meta<sup>\*†</sup>, Fernando Cucchiatti<sup>\*</sup>,

<sup>\*</sup>Barcelona Supercomputing Center (BSC), Barcelona, Spain

<sup>†</sup>Universitat Internacional de Catalunya (UIC), Barcelona, Spain

E-mail: {irene.meta, fernando.cucchiatti}@bsc.es

**Keywords**—urban simulations, holistic approach, urban areas, future city.

## I. EXTENDED ABSTRACT

More than half of the world population lives in urban areas. Urbanites are estimated to grow up to 68% of the population by 2050 [1]. This rapid growth requires new contributions from researchers and policy-makers to the development of the future city. Again, understanding how the city will grow is a crucial step in guiding this process towards the best outcome.

Cities are highly complex systems that traditional urban dynamic simulations cannot grasp in their totality, if solved only in a lightly coupled way. In addition, a model is useful only if it can be used in the planning and management practice [2]. It's true that, driven by the urge to improve their models, different sectors are developing multi-layered integrated simulations. Nevertheless, a wider scope of considering the city in its holistic behaviour is missing. Indeed, management, social, and technical barriers restrain the adoption of integrated models, such as 'model complexity, user friendliness, administrative fragmentation and communication' [3].

Simulation models are the test field where to evaluate the mechanisms of how urban systems work, according to different settings and changing conditions [4]. They allow impact assessment for new urban policies and development goals, as well as the evaluation of ongoing ones [5]. With virtual models, planners can test and evaluate changes that, otherwise, are (i) costly, (ii) unethical when people are implied and (iii) whose effects may delay appearing.

The first stage of this research consists of studying the progress made towards integration in each urban system (that we consider 'layers'). To do that, we need to (i) identify these and (ii) develop a methodology to compare integrated models belonging to the identified urban systems.

To identify the layers the city is made of, we start from the City Anatomy [6] research. Indeed, there is no common definition of city, that evolves around political, economic, and social ideas [6]. In Ref. [6] the city is considered 'a system of systems and interactions that foster emergent human behaviour. The city is compared to an organism: the capability of working together and in coordination leads to more than the sum of the individual systems (or layers).

Starting from the Anatomy described in [6], the urban systems are translated into nine layers: the six 'infrastructures' that enable flows from/to and within the city ('Communication Network', 'Water Cycle', 'Energy Cycle', 'Matter Cycle', 'Mobility Network' and 'Nature'), the 'Built Domain', the 'Society', and the 'Environment'.

After having defined the layers, we define a methodology to categorize and compare existing models. These two steps are recursive: in defining the categories, we gain new insights to better translate the urban systems from [6] into layers.

In the process of comparing we consider both Deterministic and Statistical models, and we exclude virtual three-dimensional models that offer only a visualization of the planning distribution.

Deterministic models are 'theory-based' [4] (or 'secure' models according to [7]), where the relationship between model and theory is recursive: models test the theory they are based on and are used to define a better one [4]. Statistical models are 'black-box' or 'theory-laden' [4] models (see 'insecure' models [7]) where the results obtained are more important than the behavior of the simulation algorithms, that may be not fully understood.

We collect examples for every and each layer and according to different categories, which take into account the subject that is modelled, whether the model aims to develop a theory or to exploit one, which simulation technique is used and to which purpose, the temporal and spatial capabilities of the simulation, the data requirements and availability, and the integration and interaction processes [4], [8], [3].

As a result, we conduct an extended review according to the nine layers. The definition and gathering of these categories should help identify which are the open problems yet to solve in the simulation field. From there, we're planning the future development of sustainable mobility simulation.

## ACKNOWLEDGMENT

The references analysis was partially developed in collaboration with IAAC - Valldaura Labs (Institute of Advanced Architecture of Catalonia). We want to thank Vicente Guallart and Dr. Diego Navarro for the insightful discussion.

## REFERENCES

- [1] P. D. United Nations, Department of Economic and Social Affairs, "World Urbanization Prospects: Key Facts of the 2018 Revision."
- [2] E. Miller, "Integrated urban modeling: Past, present, and future," vol. 11, pp. 387–399.
- [3] P. M. Bach *et al.*, "A critical review of integrated urban water modelling – Urban drainage and beyond," vol. 54, pp. 88–107. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1364815213003216>
- [4] M. D. Kilbridge *et al.*, "A Conceptual Framework for Urban Planning Models," vol. 15, pp. B–246. [Online]. Available: <https://pubsonline.informs.org/doi/abs/10.1287/mnsc.15.6.b246>
- [5] V. Couture *et al.*, "Income Growth and the Distributional Effects of Urban Spatial Sorting," p. 59.
- [6] Guallart *et al.*, "City Anatomy: A Framework to support City Governance, Evaluation and Transformation."

- [7] S. Bullock, "Levins and the Lure of Artificial Worlds."
- [8] S. Berling-Wolff and J. Wu, "Modeling urban landscape dynamics: A review," vol. 19, pp. 119–129. [Online]. Available: <https://doi.org/10.1111/j.1440-1703.2003.00611.x>



**Irene Meta** Graduated with honour in 2016, holds a Master in Architecture from the Polytechnic School of Genoa (Italy). She receives a mobility grant to focus on Big Data and Smart City technologies for mobility applications at BSC (Barcelona Supercomputing Center – Centro Nacional de Supercomputación). She later joins the IAAC (Institute of Advanced Architecture of Catalunya) - Valldaura Labs Group, accountable for the cooperation between IAAC and BSC in developing a research project on Urban Simulation Models. She is now a PhD candidate at BSC and UIC (Universitat Internacional de Catalunya). Her doctoral research focuses on the crucial role of big data in urban planning and management and on their potential use for decision-making process and exploratory purpose.

# Natural disasters, remote sensing, and synthetic controls

Feliu Serra-Burriel<sup>\*†</sup>, Pedro Delicado<sup>†</sup>, Fernando Cucchietti<sup>\*</sup>

<sup>\*</sup>Barcelona Supercomputing Center, Barcelona, Spain

<sup>†</sup>Universitat Politècnica de Catalunya, Barcelona, Spain

E-mail: {feliu.serra, fernando.cucchietti}@bsc.es

pedro.delicado@upc.edu

**Keywords:** Natural disasters, Wildfires, Synthetic controls, Remote Sensing, Landsat.

## I. EXTENDED ABSTRACT

Satellite imagery has been used for decades to study changes on Earth’s surface and understand the mechanisms that have shaped it as we know it today. Moreover, substantial improvements in computing power and the increase of data available in recent years have boosted interest for this kind of research. Pixel-based composites of large areas are easily accessible today thanks to the Google Earth Engine platform[1]. These are being used to study the evolution of different ecosystems such as forests[2], as well as the frequency of wildfires. Furthermore, technological advances over the last decades have enabled to precisely monitor variations in extreme weather events[3]. These weather phenomena seem to be larger now in quantity and size due to the increase of climate volatility[4].

The consequences of natural hazards have been mostly studied by comparing pre- and post-disaster conditions, or simple pair-wise comparisons between affected and non-affected areas, rendering inaccurate estimates[5]. We are interested in developing a system that, by means of a synthetic control approach, will enable us to causally evaluate the effects of disturbances over areas of interest using satellite imagery.

Resilience is another field of interest for the research community. The decrease in resilience of regions that are recurrently hit by these events might end up making certain places inhabitable. For example, extreme weather events already have their toll on life expectancy in the US[6]. Hence, large migrations may follow as a result in the long term.

### A. Costs of natural disturbances

Natural disasters have become one of the main concerns of many developed and developing countries that are more prone to be affected by these. Specifically, Western US, Australia and Brazil are some of the regions that have experienced large wildfires in recent years. However, wildfires have historically played a fundamental role in regulating the cycles of wildlife around the globe[7].

It is hard to quantify or measure the extent, severity and the impact of these natural disasters as well as the economic effects of these events. Different sources of data can provide different insights of these events. Moreover, the large costs and low precision of human-made reports on observed wildfires

calls for an accurate estimation of their effects. Aftermath costs calculated by insurers are an under-estimation of the real damage of wildfires. Sanitary and labor expenses, as well as other types of effects are usually not taken into account. There is acute need for quantifying the morbidity effects and medical costs in order to estimate the true economic impact of wildfires.

With regard to the specific case of ecosystem disturbances provoked by wildfires, previous research has focused on the analysis of satellite imagery [8] to evaluate the evolution of forest ecosystems, both before and after wildfires occur[9]. In addition, past studies have centered on the simulation of different temperature and climate scenarios and on the varying degrees of wildfire risks in California[10]. On the other hand, the aim of this study is to estimate the short and long term causal effects of wildfires.

### B. Experimental Environment

Open access data to more than three decades of satellite imagery is available through the Landsat archive[11], as well as many other remote sensing data resources. This spatio-temporal data allows for a longitudinal analysis of these natural hazards before and after they take place. Figure 1 shows the state of vegetation, before, during, and after a large wildfire.

We can obtain proxies of vegetation health by observing various Landsat Surface Reflectance-Derived Spectral Indices (LSR-DSI), such as Normalized Difference Vegetation Index (NDVI), Normalized Burned Ratio (NBR) and other composite. Other sources of data such as satellite nightlight imagery can be good proxies of the socio-economic statuses of societies. The increase in size and frequency of these natural disasters urges the use of new methods in order to rapidly evaluate their extent. By controlling for spatial attributes, incident risk and land use we intend to thoroughly examine these effects.

The methodology proposed in this extended abstract is known as synthetic control. This approach is drawn from social and medical sciences literature[12]. It consists on the creation of a synthetic counterfactual for each treated unit, in order to assess the longitudinal effects of the event of interest. By means of using synthetic controls we intend to study the average causal effect of natural disasters on different regions



Fig. 1. Sample images capturing the Thomas Fire in California in December 2017. The images show the NDVI on surface reflectance over three Landsat 8 images. The images show the vegetation previous, during and after the wildfire. The second image depicts some smoke, meaning that the wildfire was still being controlled.

of the planet. In order to obtain estimates of the effects, we need to observe the characteristics of the affected areas, as well as their respective counterfactual regions. These are the areas that were found to be most similar to their counterparts before the impacts, but were unaffected by them.

### C. Discussion

In this extended abstract we propose the use of remote sensing data, jointly with the use of synthetic controls to estimate spatio-temporal effects of natural disturbance events. By combining literature from econometrics, remote sensing, political science, and the vast amount of data at our disposal, in addition to adequate pre-processing, we can provide significant insights on the effects of these events and enable improvements in the field of mitigation planning, as well as in more efficient regeneration management.

### II. ACKNOWLEDGMENT

We would like to thank Andrew Prata for his scientific acumen, Guillermo Marín for his suggestions on how to improve

the plot, Miquel Serra-Burriel for the insightful discussions and proofreading, and Victor Paradis for his proofreading and editing.

### REFERENCES

- [1] N. Gorelick, M. Hancher, M. Dixon, S. Ilyushchenko, D. Thau, and R. Moore, "Google Earth Engine: Planetary-scale geospatial analysis for everyone," *Remote Sensing of Environment*, vol. 202, pp. 18–27, 2017. [Online]. Available: <https://doi.org/10.1016/j.rse.2017.06.031>
- [2] J. C. White, M. A. Wulder, G. W. Hobart, J. E. Luther, T. Hermosilla, P. Griffiths, N. C. Coops, R. J. Hall, P. Hostert, A. Dyk, and L. Guindon, "Pixel-based image compositing for large-area dense time series applications and science," *Canadian Journal of Remote Sensing*, vol. 40, no. 3, pp. 192–212, 2014. [Online]. Available: <https://doi.org/10.1080/07038992.2014.945827>
- [3] D. R. Easterling, J. L. Evans, P. Y. Groisman, T. R. Karl, K. E. Kunkel, and P. Ambenje, "Observed variability and trends in extreme climate events: A brief review," *Bulletin of the American Meteorological Society*, vol. 81, no. 3, pp. 417–425, 2000.
- [4] M. M. Q. Mirza, "Climate change and extreme weather events: can developing countries adapt?" *Climate Policy*, vol. 3, no. 3, pp. 233–248, 2003. [Online]. Available: <https://www.tandfonline.com/doi/abs/10.3763/cpol.2003.0330>
- [5] D. K. Bolton, N. C. Coops, and M. A. Wulder, "Characterizing residual structure and forest recovery following high-severity fire in the western boreal of Canada using Landsat time-series and airborne lidar data," *Remote Sensing of Environment*, vol. 163, pp. 48–60, 2015. [Online]. Available: <http://dx.doi.org/10.1016/j.rse.2015.03.004>
- [6] O. Deschênes and E. Moretti, "Extreme Weather Events, Mortality, and Migration," *The Review of Economics and Statistics*, vol. 91, no. 4, pp. 659–681, 2009. [Online]. Available: <https://doi.org/10.1162/rest.91.4.659>
- [7] J. G. Pausas and J. E. Keeley, "Wildfires as an ecosystem service," *Frontiers in Ecology and the Environment*, vol. 17, no. 5, pp. 289–295, 2019.
- [8] A. Banskota, N. Kayastha, M. J. Falkowski, M. A. Wulder, R. E. Froese, and J. C. White, "Forest monitoring using landsat time series data: A review," *Canadian Journal of Remote Sensing*, vol. 40, no. 5, pp. 362–384, 2014. [Online]. Available: <https://doi.org/10.1080/07038992.2014.987376>
- [9] G. Ireland and G. P. Petropoulos, "Exploring the relationships between post-fire vegetation regeneration dynamics, topography and burn severity: A case study from the montane cordillera ecozones of western Canada," *Applied Geography*, vol. 56, pp. 232–248, 2015.
- [10] B. B. Westerling, A.L., "Climate change and wildfire in California," *Climatic Change* 87, p. 231–249, 2008.
- [11] L. Lymburner, E. Botha, E. Hestir, J. Anstee, S. Sagar, A. Dekker, and T. Malthus, "Landsat 8: Providing continuity and increased precision for measuring multi-decadal time series of total suspended matter," *Remote Sensing of Environment*, vol. 185, pp. 108–118, 2016. [Online]. Available: <http://dx.doi.org/10.1016/j.rse.2016.04.011>
- [12] A. Abadie, A. Diamond, Hainmueller, and Jens, "Synthetic control methods for comparative case studies: Estimating the effect of California's Tobacco control program," *Journal of the American Statistical Association*, vol. 105, no. 490, pp. 493–505, 2010.



**Feliu Serra-Burriel** received his BSc degree in Economics at Pompeu Fabra University (UPF), Barcelona in 2016. Then, he completed his MSc degree, research programme in Statistics from London School of Economics, London in 2017. The following year, he worked at UPF as a research assistant for professor Christian Fons-Rosen. Since 2018, he has been with the Visualization group of Barcelona Supercomputing Center (BSC) and he also started as a PhD student at the department of Statistics and Operation Research of Universitat Politècnica de Catalunya (UPC), Spain.

# Specific-purpose globalizations for Newton’s method: anisotropic optimization of curved meshes

Guillermo Aparicio-Estremis, Abel Gargallo-Peiró, Xevi Roca  
Barcelona Supercomputing Center (BSC), Barcelona, Spain  
E-mail: {guillermo.aparicio, abel.gargallo, xevi.roca}@bsc.es

*Abstract*—We derive an optimization method to adapt straight-edged and curved piece-wise polynomial meshes to the stretching and alignment of a target metric. Two globalization strategies for the optimization method are proposed: backtracking line search and restricted trust region. To compare both globalization approaches, we derive a specific-purpose implementation of Newton’s method for each globalization. To propose these two new implementations, we present different emulation methods to interchange between both approaches their non-shared globalization features. Once the number of non-linear iterations is comparable, we have been able to improve the inexact Newton implementation, with both globalization methods, to reduce one order of magnitude the total number of sparse matrix-vector products.

## I. INTRODUCTION

The optimization of an objective function that measures a global mesh quantity by changing the coordinates of the mesh nodes has been used in several meshing methods [1]. To optimize the objective function, we can iteratively modify the coordinates of either one free node (local) or all the free nodes (global) per non-linear iteration by using either gradient-based (first-order) or Hessian-based (second-order) optimization methods. Existent literature shares a common conclusion. That is, when highly optimized and accurate meshes are required, especially in isotropic meshes featuring high gradations of the element size, a specific-purpose global feasible Newton method outperforms one-node optimization methods.

Although the existent literature does not deal with piece-wise polynomial meshes and highly stretched elements, it provides valuable knowledge to devise our specific-purpose global optimization method. Our goal is to derive an optimization method to adapt straight-edged and curved piece-wise polynomial meshes to the stretching and alignment of a target metric. The objective function has been proposed before [2]. However, a specific purpose solver where an initial mesh configuration, with nodal coordinates not in the convergence basin of a local optimum, is efficiently driven to an optimal configuration has not been proposed. Accordingly, we are especially interested in a globalized solver that can deal with inaccurate initial approximations with low and high polynomial degrees and targets metrics featuring highly varying size gradations, stretching ratios, and principal directions. Thus, the main research question of this work arises whether it is better to use a backtracking line search or a restricted trust region.

To compare both globalization approaches for the adaption problem to a target metric problem, we derive two specific-purpose implementations of Newton’s method equipped with

backtracking line search (BLS) and restricted trust region (RTR) [3], [4].

## II. RESULTS

As a test example we consider the hexahedral domain  $\Omega := [-\frac{1}{2}, \frac{1}{2}]^3$  equipped with the metric  $\mathbf{M}$  given by

$$\mathbf{M} = \nabla\varphi^T \cdot \mathbf{D} \cdot \nabla\varphi,$$

where

$$\varphi(x, y) := \left( x, \frac{10y - \cos(2\pi x)}{\sqrt{100 + 4\pi^2}} \right)$$

and where  $\mathbf{D}$  is the three-dimensional boundary layer metric with its axes aligned according to the Cartesian axes with a stretching in the  $z$ -direction at the plane  $z = 0$ , that is,

$$\mathbf{D} := \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1/h(z)^2 \end{pmatrix}.$$

The function  $h$  is given by

$$h(x) := h_{\min} + \gamma|x|,$$

with  $h_{\min} = 2 \cdot 10^{-2}$  and  $\gamma = 2$ .

We have generated initial isotropic straight-sided tetrahedral meshes with the same resolution, 1577 nodes. The meshes are of polynomial degree 1, 2, and 4 and composed by 7296, 912 and 114 elements respectively. Perspectives of the mesh are shown in Figures 1(a), 1(b) and 1(c). They are colored according to the point-wise quality measure. The optimized meshes are presented in Figures 1(d), 1(e) and 1(f). We observe that the elements away from the anisotropic region are enlarged whereas the elements lying in the anisotropic region are compressed. In the optimized mesh the minimum is improved and the standard deviation of the element qualities is reduced when compared with the initial configuration.

To optimize the objective function presented in [2] we have applied, for each globalization strategy (BLS and RTR), the standard and proposed inexact Newton optimization methods. The optimization results are shown in Table I. We observe that, for the proposed optimization methods, the total amount of matrix vector products and the number of non-linear iterations have been reduced. Furthermore, we observe that the optimization results for the BLS and RTR strategy are comparable.

TABLE I. NON-LINEAR ITERATIONS, GLOBALIZATION ITERATIONS AND MATRIX-VECTOR PRODUCTS FOR STANDARD AND PROPOSED BACKTRACKING LINE-SEARCH (BLS) AND RESTRICTED TRUST-REGION (RTR) STRATEGIES.

Globalization strategy	Mesh degree	Non-linear iterations		Globalization iterations		Matrix-vector products	
		Standard	Proposed	Standard	Proposed	Standard	Proposed
BLS	1	52	32	83	50	5314	231
BLS	2	75	54	349	136	14043	572
BLS	4	91	84	445	163	16072	733
BLS	8	277	141	1211	319	18773	2424
RTR	1	251	35	0	72	22162	277
RTR	2	272	52	0	122	13061	471
RTR	4	305	80	0	157	15623	742
RTR	8	397	110	0	233	29872	1758




Fig. 1. Slices of tetrahedral meshes of polynomial degree 1, 2, and 4 in columns. Initial straight-sided isotropic meshes and optimized meshes from initial meshes in rows. Meshes are colored according to the pointwise quality measure presented in [2].

### III. CONCLUSIONS

For r-adaption problems, standard implementations of BLS and RTR are not comparable while our new specific-purpose implementations are. To propose these two new implementations of BLS and RTR, we propose different emulation methods to interchange between both approaches their non-shared globalization features. Once the number of non-linear iterations is comparable, we have been able to improve the inexact Newton implementation, with both globalization methods, to reduce one order of magnitude the total number of sparse matrix-vector products.

### IV. ACKNOWLEDGMENT

This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme under grant agreement No 715546. This work has also received funding from the Generalitat de Catalunya under grant number 2017 SGR 1731. The work of X. Roca has been partially supported by the Spanish Ministerio de Economía y Competitividad under the personal grant agreement RYC-2015-01633.

### REFERENCES

[1] L. F. Diachin, P. Knupp, T. Munson, and S. Shontz, “A comparison of inexact newton and coordinate descent meshoptimization techniques,” Lawrence Livermore National Lab.(LLNL), Livermore, CA (United States), Tech. Rep., 2004.

[2] G. Aparicio-Estrems, A. Gargallo-Peiró, and X. Roca, “Defining a stretching and alignment aware quality measure for linear and curved 2d meshes,” in *International Meshing Roundtable*. Springer, 2018, pp. 37–55.

[3] J. Nocedal and S. Wright, *Numerical optimization*. Springer Science & Business Media, 2006.

[4] J. P. Bulteau and J. P. Vial, “A restricted trust region algorithm for unconstrained optimization,” *Journal of Optimization Theory and Applications*, vol. 47, no. 4, pp. 413–435, Dec 1985. [Online]. Available: <https://doi.org/10.1007/BF00942189>



**Guillermo Aparicio-Estrems** received his BSc degree in Mathematics from Universitat Politècnica de Catalunya (UPC), Spain in 2016. The following year, he completed his MSc degree in Applied Mathematics from Universitat Politècnica de Catalunya (UPC), Spain in 2017. Since 2017, he has been with the Geometry and Meshing for simulations group of Barcelona Supercomputing Center (BSC) as a PhD student of Universitat Politècnica de Catalunya (UPC), Spain.

# Visualization of pentatopic meshes

Guillem Belda-Ferrín\*, Eloi Ruiz-Gironés\*, Xevi Roca \*

\*Barcelona Supercomputing Center, Barcelona, Spain

E-mail: {guillem.belda, eloi.ruizgirones, xevi.roca}@bsc.es

**Abstract**—We propose a simple tool to visualize 4D unstructured pentatopic meshes. The method slices unstructured 4D pentatopic meshes (fields) with an arbitrary 3D hyperplane and obtains a conformal 3D unstructured tetrahedral representation of the mesh (field) slice ready to explore with standard 3D visualization tools. The results show that the method is suitable to visually explore 4D unstructured meshes. This capability has facilitated devising our 4D bisection method, and thus, we think it might be useful when devising new 4D meshing methods. Furthermore, it allows visualizing 4D scalar fields, which is a crucial feature for our space-time applications.

**Keywords**—Visualization, 4D mesh, pentatopic mesh

## I. EXTENDED ABSTRACT

In the last years, there has been an emerging interest to generate [1], [2], refine [3], [4], [5], [6], [7], [8], and adapt [9] 4D meshes. In our case, the need to simulate unsteady problems using full space-time (3D space + 1D time) discretizations with unstructured methods prompts our interest. We think that one key issue that is hampering to devise, check, and illustrate new 4D meshing approaches is the lack of a natural approach to visualize 4D meshes. We aim to provide a preliminary solution to this issue, by providing a simple tool to visualize 4D unstructured pentatopic meshes.

Our method is devised to exploit existent 3D visualization software which provides mature user interfaces to interact in real-time with 2D projections of the 3D meshes. To exploit these interfaces, we propose to slice unstructured 4D pentatopic meshes (fields) with an arbitrary 3D hyperplane and obtain a conformal 3D unstructured tetrahedral representation of the mesh (field) slice that is ready to be read with standard 3D visualization tools. Recently, a method to visualize 4D pentatopic meshes has been outlined in [9]. The main difference with our approach is that in [9], the resulting 3D visualization mesh is composed of polyhedra instead of tetrahedra.

We devise the method as follows. First, given a 4D unstructured pentatopic mesh and a hyperplane, we compute element-by-element the intersection of the hyperplane with the element edges. We only consider the intersections that lead to a 4D point, and we ignore all the singular intersections leading to either the null set or the full edge. Then, for the current element, we check if the resulting set of selected points generates a 3D polyhedron. If the points in the hyperplane define a volume, we compute the coordinates of the 4D points expressed in terms of a 3D orthonormal base of the hyperplane. Then, we compute a Delaunay tetrahedralization of the resulting 3D points. Finally, we store a link between the intersection edge points and a unique edge identifier. If the intersection points do not define a volume, we continue with the following element. Using the link between the intersection edge points and the unique edge identifiers, we can now

merge all the local Delaunay tetrahedralizations to obtain a 3D conformal unstructured tetrahedral mesh, without duplicates of the edge points, that represents the 3D slice of the 4D mesh.

## II. RESULTS

We consider the gravitational potential of two masses, defined by Equation (1), such that the positions of the masses evolve in time.

$$V(\vec{x}, t) = -G \left( \frac{m_1}{\|\vec{x} - \vec{p}_1(t)\|} + \frac{m_2}{\|\vec{x} - \vec{p}_2(t)\|} \right) \quad (1)$$

Initially, at time  $t = 0$  the two masses are located at different points of the  $z$ -axis. Then, the masses will move along the  $z$ -axis with constant velocity and opposite direction until they arrive at the same position at  $t = 1$ , see Equation (2).

$$\begin{aligned} \vec{p}_1(t) &= \vec{p}_1 + (0, 0, vt), \quad \vec{p}_1 \in \mathbb{R}^3 \\ \vec{p}_2(t) &= \vec{p}_2 - (0, 0, vt), \quad \vec{p}_2 \in \mathbb{R}^3 \end{aligned} \quad (2)$$

Hence, the isosurface of the gravitational potential will evolve from two connected components merging into a single component. For a fixed isovalue the gravitational potential defines an implicit 3D manifold embedded into a 4D space (3D space + 1D time). We adapt a 4D pentatopic mesh [7] of a hypercube to capture the 3D embedded manifold defined by the isovalue  $V(\vec{x}, t) \approx -10$ . Then, we obtain a 4D pentatopic mesh composed of 16798112 pentatopes and 879778 nodes. To select the elements to refine first, we compute the elements that intersect the manifold. Then, from those elements, we compute the manifold curvature of each one using the Hessian of  $V(\vec{x}, t)$  and we select to refine the 10% of the elements with higher curvature. Therefore, the obtained mesh will be refined close to the isosurface and, in particular, where the curvature is higher. The elements near to the isosurface will be smaller than the far one, which will be coarser.

Figure 1 shows three slices at different times. Each time slice is represented using a 3D mesh in the  $(z, x, y)$  space. The figures on the left, Figures 1(a), 1(c) and 1(e), illustrate the adapted mesh at times  $t = 0$ ,  $t = 0.5$  and  $t = 1$  respectively. In the right hand side, Figures 1(b), 1(d) and 1(f), illustrate the adapted mesh with the associated isosurface at times  $t = 0$ ,  $t = 0.5$  and  $t = 1$  respectively. As we expected, the mesh is more refined close to the isosurface and is coarser far from it. We are showing three different slices in time, but since the presented visualization is a post-process algorithm and we are working with a single 4D pentatopic mesh, we can make as many slices as we require. Figure 2 illustrates a slice with the hyperplane  $x = 0.5$ . We obtain a 3D space-time configuration in which vertical axis is the time axis. We can see the two initial connected components at the bottom of the mesh, and how they merge to a single connected component.

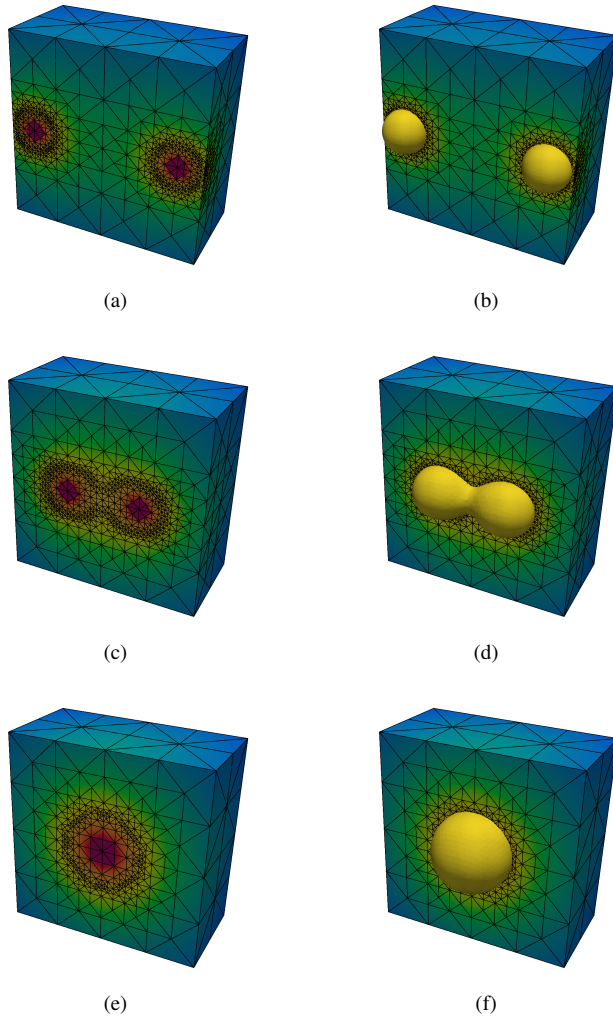


Fig. 1. Different slices in time of an adapted 4D pentatopic mesh are presented. Figures (a) and (b) corresponds to the time slice  $t = 0.0$ . Figures (c) and (d) corresponds to the time slice  $t = 0.5$ . Finally, Figures (e) and (f) corresponds to the time slice  $t = 1.0$ .




Fig. 2. A slice of an adapted 4D pentatopic mesh with the hyperplane  $x = 0.5$  is presented. We obtain the 3D space-time mesh  $(z, y, t)$ , where we can see the time evolution of the isosurface defined by the gravitational potential.

### III. CONCLUSIONS

Our preliminary results show that our approach is suitable to visually explore 4D unstructured meshes with the help of

3D visualization interactive packages This capability has facilitated developing, debugging, and checking our 4D bisection method, and thus, we think it might be useful when devising new 4D meshing methods. Furthermore, it allows visualizing 4D scalar fields, which is a crucial feature for our space-time applications.

### IV. ACKNOWLEDGMENT

This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme under grant agreement No 715546. This work has also received funding from the Generalitat de Catalunya under grant number 2017 SGR 1731. The work of X. Roca has been partially supported by the Spanish Ministerio de Economía y Competitividad under the personal grant agreement RYC-2015-01633. We would like to thank Dr. Abel Gargallo-Peiró for the fruitful conversations and comments.

### REFERENCES

- [1] P. Foteinos and N. Chrisochoides, "4d space-time delaunay meshing for medical images," *Engineering with Computers*, vol. 31, no. 3, pp. 499–511, 2015.
- [2] E. Karabelas and M. Neumüller, "Generating admissible space-time meshes for moving domains in  $d + 1$ -dimensions," *arXiv preprint arXiv:1505.03973*, 2015.
- [3] J. M. Maubach, "Local bisection refinement for  $n$ -simplicial grids generated by reflection," *SIAM Journal on Scientific Computing*, vol. 16, no. 1, pp. 210–227, 1995.
- [4] C. T. Traxler, "An algorithm for adaptive mesh refinement in  $n$  dimensions," *Computing*, vol. 59, no. 2, pp. 115–137, 1997.
- [5] R. Stevenson, "The completion of locally refined simplicial partitions created by bisection," *Mathematics of computation*, vol. 77, no. 261, pp. 227–241, 2008.
- [6] M. Alkämper, F. Gaspoz, and R. Klöforn, "A weak compatibility condition for newest vertex bisection in any dimension," *SIAM Journal on Scientific Computing*, vol. 40, no. 6, pp. A3853–A3872, 2018.
- [7] G. Belda-Ferrín, A. Gargallo-Peiró, and X. Roca, "Local bisection for conformal refinement of unstructured 4d simplicial meshes," in *International Meshing Roundtable*. Springer, 2018, pp. 229–247.
- [8] J. Grande, "Red–green refinement of simplicial meshes in  $d$  dimensions," *Mathematics of Computation*, vol. 88, no. 316, pp. 751–782, 2019.
- [9] P. C. Caplan, "Four-dimensional anisotropic mesh adaptation for space-time numerical simulations," Ph.D. dissertation, Aero/Astro Dpt., MIT, 2019.



**Guillem Belda-Ferrín** received his BSc degree in Mathematics from Universitat Politècnica de Catalunya (UPC), Barcelona in 2014. He completed his MSc degree in Advanced Mathematics and Mathematical Engineering from UPC in 2015. Since 2017, he has been with the Geometry and Meshing for simulations group of Barcelona Supercomputing Center (BSC) as a PhD candidate of the Applied Mathematics doctorate program of UPC, Spain.



# Subdividing linear and curved meshes preserving sharp features of a model

Albert Jiménez-Ramos\*, Abel Gargallo-Peiró\*, Xevi Roca\*

\*Barcelona Supercomputing Center (BSC), Barcelona, Spain

E-mail: {albert.jimenez, abel.gargallo, xevi.roca}@bsc.es

**Abstract**—To provide straight-edged and curved piece-wise polynomial meshes that target a unique smooth geometry while preserving the sharp features and smooth regions of the model, we propose a new fast curving method based on hierarchical subdivision and blending. There is no need for underlying target geometry. It is only required a straight-edged mesh with boundary entities marked to characterize the geometry features, and a list of features to recast through a unique sharp-to-smooth modeling capability. The examples show that the method is well-suited to curve large quadratic and quartic meshes in low-memory configurations.

**Keywords**—*mesh curving, surrogate geometry, geometry modeling, subdivision, blending*

## I. EXTENDED ABSTRACT

### A. Introduction

In flow simulations for wind energy, transport of pollutants, and bio-engineering the boundary of the computational domain is usually represented by a straight-edged mesh obtained by sampling real data. This straight-edged mesh approximates the geometry, at different scales, corresponding to the viscous surfaces to analyze such as topography, urban areas and human organs. The mesh also presents a series of sharp features, vertices, polylines bounded by vertices, that the method should preserve, and that bound the smooth regions of the computational model.

The resolution to approximate the geometry could be insufficient for the required flow analysis, and thus, additional refinement of the boundary mesh would be required. However, a standard refinement approach, when no target geometry is available, could be inadequate for flow simulation in a twofold way. First, the refined mesh might reproduce precisely the geometry of the first straight-edged mesh and thus, introduce artificial flow artifacts close to initially non-smooth features that should be smooth. Second, the refined mesh might target a smooth surface geometry, implicitly determined by the initial straight-edged mesh, but without adequately respecting after successive refinement the sharp features, curves, and vertices, of the computational model. Ideally, vertices should remain fixed, and polylines should target a smooth limit curve.

Solving these issues is essential for those flow analyses that start from a mesh obtained by sampling real data where the computational model presents smooth regions bounded by sharp features. Even they can be useful in aeronautical applications where only legacy data, in a format of vertices, and polyline and surface meshes, is available. In some applications, practitioners might also need, a non-standard but flexible *sharp-to-smooth modeling* capability, to remove some sharp

features ensuring that surrounding regions become smooth along with the removed feature.

Intending to provide piece-wise linear meshes or curved piece-wise polynomial meshes that target a unique smooth geometry while preserving the sharp features of the model, our contribution is to propose a new fast curving method based on hierarchical subdivision and blending with sharp-to-smooth modeling capabilities. Our approach only needs an initial straight-edged mesh with boundary triangles marked with surface identifiers, and a list of features to recast. There is no need for underlying target geometry. The goal of the method is to obtain a volume mesh of the flow domain that under successive refinement leads to smooth regions bounded by the sharp features determined after recasting. The recasting operation is devised to implement a sharp-to-smooth modeling capability. We favored a fast and explicit curving method, based on subdivision and blending, to an implicit approach formulation that features validity guarantees or untangling capabilities, based on boundary curving and optimization, but slower and more memory demanding. This favoring is so since the appearance of invalid elements is small compared with the scale of the generated meshes, and fast local untangling can repair those invalid elements. This work details our mesh curving methodology and illustrates its application with the included numerical examples.

### B. Example

The curved high-order mesh generation procedure proposed in this work is composed of four main steps:

- 1) **Sharp-to-smooth modeling.** We recast some of the feature entities present in the original model, and thus provide a new model improving the smoothness of the surrogate geometry. Each feature vertex (node of the mesh), curve (set of edges of the mesh) and surface (set of triangles of the mesh) is associated with a unique identifier. Therefore, to recast a feature, it is enough to know its identifier. The initial model, [Fig. 1\(a\)](#), contains some vertex and curve features, such as the leading edge or the curve along the fuselage, that have been recast in the final model, [Fig. 1\(b\)](#), to obtain a smoother virtual surface. We highlight that the presence of a non-desired geometry feature has an impact on the smoothness of the generated mesh. We illustrate the continuity of the normal vectors using a zebra mapping. For the initial model, the normals are discontinuous, see [Fig. 1\(c\)](#), and therefore, the mesh is  $C^0$ -continuous;




Fig. 1: (a) Marks on boundary entities of the initial model: curve and vertex features. (b) Marks of the final model: curve and vertex features recast (gray) and preserved (black). Zebra mapping showing isophotes along the leading edge for the (c) initial model, and (d) the final model.

while in the final model the normals are continuous, see Fig. 1(d), so the mesh is  $\mathcal{C}^1$ -continuous.

- 2) **Approximate a surrogate boundary.** Given a linear tetrahedral mesh, we extract its boundary. The boundary is a linear triangular mesh with its entities marked, and by means of the hierarchical subdivision of its elements we generate a curved high-order triangular surface mesh. The curved surface mesh approximates a surrogate boundary composed of feature surfaces with an interior that is  $\mathcal{C}^1$ -continuous and  $\mathcal{C}^2$ -continuous almost everywhere [1], [2]. This surrogate is determined by the subdivision of the curves and surfaces, and preserves the sharp features and smooth regions marked on the boundary of the initial volume mesh.
- 3) **Accommodate the curvature of the boundary.** We accommodate the curvature of the curved surface mesh to the boundary volume elements using an explicit hierarchical blending [3]. In this example, the blending reduces from 358 to 24 the number of invalid elements, that is, in 237 seconds a 93% of the invalid elements have been untangled.
- 4) **Local untangling.** If necessary, we optimize the inverted elements locally following the approach detailed in [4]. Since the mesh after the previous step is close to be optimal, it is a good initial condition for the implicit optimization and in 60 seconds the mesh becomes valid achieving a minimum quality of 0.7. The final quartic mesh is shown in Fig. 2.

### C. Conclusion

The obtained results show that we can generate, from an initial straight-edged mesh, successively refined piecewise linear, quadratic and quartic meshes, that target smooth curves and surfaces, while preserving the initially marked sharp features and smooth regions. The interior of the obtained limit curves is of class  $\mathcal{C}^2$ , and the interior of the surfaces is at least  $\mathcal{C}^1$ -continuous, being of class  $\mathcal{C}^2$  when the surface mesh is structured. For manifolds with boundaries, only a straight-edged mesh with boundary triangles marked with




Fig. 2: Generated valid curved tetrahedral mesh of polynomial degree  $p = 4$  composed of  $1.8 \cdot 10^7$  nodes and  $1.7 \cdot 10^6$  elements.

surface identifiers is required. Then, the method automatically computes the boundary curves and vertices from the triangle marks. A unique sharp-to-smooth modeling capability - not fully available in standard CAD packages - allows removing sharp features, *i.e.* vertices and curves, and smoothly merge the incident entities, *i.e.* curves and surfaces. The resulting surrogate geometry features  $\mathcal{C}^1$ -continuity along the merging region. The proportion of invalid elements is small compared to the size of the meshes, and thus, we showed that it can be fixed using local untangling and curving without the need for a global solver.

## II. ACKNOWLEDGMENT

This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme under grant agreement No 715546. This work has also received funding from the Generalitat de Catalunya under grant number 2017 SGR 1731. The work of the third author has been partially supported by the Spanish Ministerio de Economía y Competitividad under the personal grant agreement RYC-2015-01633. Special thanks to Eloi Ruiz-Gironés.

## REFERENCES

- [1] J. M. Lane and R. F. Riesenfeld, "A theoretical development for the computer generation and display of piecewise polynomial surfaces," *IEEE TPAMI*, no. 1, pp. 35–46, 1980.
- [2] C. Loop, "Smooth subdivision surfaces based on triangles," Master's thesis, The University of Utah, Masters Thesis, 1987.
- [3] A. Perronnet, "Interpolation transfinie sur le triangle, le tétraèdre et le pentaèdre. application à la création de maillages et à la condition de dirichlet," *CR ACAD SCI I-MATH*, vol. 326, no. 1, pp. 117–122, 1998.
- [4] A. Gargallo-Peiró, X. Roca, J. Peraire, and J. Sarrate, "Optimization of a regularized distortion measure to generate curved high-order unstructured tetrahedral meshes," *INME*, vol. 103, no. 5, pp. 342–363, 2015.



**Albert Jiménez-Ramos** received his BSc degree in Mathematics from Universitat Politècnica de Catalunya (UPC), Barcelona in 2017. He completed his MSc degree in Advanced Mathematics and Mathematical Engineering from UPC in 2018. Since 2018, he has been a member of the Geometry and Meshing for simulations group of Barcelona Supercomputing Center (BSC) as a PhD student of the Applied Mathematics doctorate program of UPC, Spain.

# Computing Worst-Case Contention Delays for Networks on Chip

Jordi Cardona<sup>\*†</sup>, Carles Hernandez<sup>\*‡</sup>, Jaume Abella<sup>\*</sup>

<sup>\*</sup>Barcelona Supercomputing Center, Barcelona, Spain

<sup>†</sup>Universitat Politècnica de Catalunya, Barcelona, Spain

<sup>‡</sup>Universitat Politècnica de València, València, Spain

E-mail: {jordi.cardona, carles.hernandez, jaume.abella}@bsc.es

**Keywords**—NoC, Mesh, WCET, ILP; Contention

## I. EXTENDED ABSTRACT

### A. Introduction

Computing performance needs in domains such as automotive, avionics, railway, and space are on the rise. This is fueled by the trend towards implementing an increasing number of product functionalities in software that ends up managing huge amounts of data and implementing complex artificial-intelligence functionalities [1], [2].

Manycores are able to satisfy, in a cost-efficient manner, the computing needs of embedded real-time industry [3], [4]. In this line, building as much as possible on manycore solutions deployed in the high-performance (mainstream) market [5], [6], contributes to further reduce costs and increase availability. However, commercial off the shelf (COTS) manycores bring several challenges for their adoption in the critical embedded market. One of those is deriving timing bounds to tasks' execution times as part of the overall timing validation and verification processes [7]. In particular, the network-on-chip (NoC) has been shown to be the main resource in which contention arises, and hence hampers deriving tight bounds to the timing of tasks [8].

For widely-used wormhole NoCs (wNoCs) [6], [5], several proposals show how to compute latency upperbounds to the different flows communicating on the manycore [9], [10] under some restrictions, e.g. deterministic routing. Unfortunately, WCET estimates computed with wNoCs are generally pessimistic when – as required to achieve composable estimates – no restrictions are imposed on when and where interference occurs in the wNoC. Interestingly, wNoCs offer several software-controllable parameters that allow to optimize (reduce) the worst-case contention delay (WCD) that packets crossing can suffer. These include mapping, routing, and allocation of weights (referred to as walloc) to arbitration policies in each router. NoC contention optimization solutions have been proposed for mapping [11], [12] and combining routing and mapping [13], [14]. Additionally, optimal allocation of weights to achieve fair bandwidth balancing have been also proposed for TDMA [15] and wNoCs [16]. In general, those solutions do not tackle all parameters at once, which leads to globally suboptimal solutions.

Overall, reducing the WCD in NoCs is indeed a multidimensional problem and, to make things worse, strong dependencies exist between the different parameters. For instance,

the impact of routing in WCD is heavily affected by the mapping of tasks to cores.

Despite the inter-dependences among these parameters, to our knowledge, no previous work proposes an integral solution to the problem of WCD reduction simultaneously optimizing mapping, routing and walloc.

### B. Contribution

In this study, we cover this gap by proposing a wNoC ILP- and stochastic-based optimization framework that minimizes WCD estimates (and hence WCET estimates) of applications running in the wNoC-connected manycores.

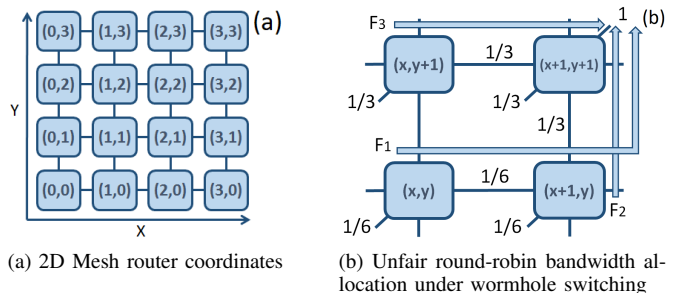


Fig. 1: Mesh manycore system and how mapping, routing and bandwidth allocation parameters have impact in the WCD

We focus the target of this study in systems build with 2D mesh NoC topologies like the ones showed in Figure 1, even though the same analysis can be done to other topologies. In Figure 1(a), we show a block diagram of a 2D mesh multicore system where each node coordinate represents a router that has attached a processor and/or a memory element and it is connected to the other routers forming a mesh topology.

Our target is to create a framework that given some tasks that are going to run on the top of a mesh multicore system, it optimizes the mapping, routing and bandwidth allocation configuration of the mesh all at the same time so as to reduce the WCET estimates of these tasks. One of the possible NoC configuration outputs after running the optimizing framework is shown in Figure 1(b). In this figure, we show a feasible NoC solution for the three main NoC parameters object of these study: mapping assigned as first come first serve (FCFS), XY routing algorithm and round-robin bandwidth arbitration).

In this study:

- 1) We analyze the main wNoC parameters that cause variability in WCD (mapping, routing, and walloc) separately and how they relate to each other. We propose a particular WCD-centric abstraction to address the main sources of jitter in a wNoC, namely: placement of tasks (threads) to cores, routing, and weighted bandwidth allocation (walloc).
- 2) We show that reducing WCD is a multidimensional problem that we decompose into a stochastic-based optimization and an ILP formulation. The former covers the optimization of the routing, whereas the latter optimizes mapping and walloc.
- 3) We compare the effectiveness of the ILP method with respect to hand-made setups and other approaches that optimize a subset of the parameters. Our results confirm that our multidimensional optimization approach achieves performance guarantees that outperform the other ones evaluated. We also show that optimizing virtual-channel (VC) allocation provides a subset of the configurations obtained with walloc, so that optimizing walloc makes VC not to provide any additional advantage.

We focus on high-performance wormhole NoCs in which time-predictability is achieved by leveraging an optimal configuration of parameters. This includes features like arbitration and routing already configurable from software in existing real wNoC designs. Weight allocation, while to our knowledge it has not been implemented in commercial NoCs yet, it is widely used in high-performance routers for off-chip wormhole networks [17]. Given that the implementation cost of weighted arbitration is quite low [16], they can be included with low cost in high-performance on-chip designs. Moreover, modifications required to implement weighted arbitration are local in contrast to hardware proposals that require global changes like, new signals among routers and nodes, different flow-control, global clocks or the like.

## REFERENCES

- [1] "Intel GO Automated Driving Solution Product Brief," <https://www.intel.es/content/dam/www/public/us/en/documents/platform-briefs/go-automated-accelerated-product-brief.pdf>.
- [2] M. Girone, "Computing Challenges at the Large Hadron Collider (LHC)," *Keynote at the HiPEAC Conference 2018*, 2018.
- [3] Kalray MPPA 256 Many-Core Processor, <http://www.kalray.eu/products/mppa-manycore>.
- [4] R. Ginosar, P. Aviely, T. Israeli, and H. Meirov, "RC64: High performance rad-hard manycore," in *2016 IEEE Aerospace Conference*, March 2016, pp. 1–9.
- [5] Tiler, *TILE-Gx Processors Family* <http://www.tiler.com/products/TILE-Gx.php>.
- [6] S. Ramos and T. Hoefler, "Capability models for manycore memory systems: A case-study with Xeon Phi KNL," in *2017 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, May 2017, pp. 297–306.
- [7] M. Paulitsch, O. M. Duarte, H. Karray, K. Mueller, D. Münch, and J. Nowotzsch, "Mixed-criticality embedded systems - A balance ensuring partitioning and performance," in *2015 Euromicro Conference on Digital System Design, DSD 2015, Madeira, Portugal, August 26-28, 2015*, 2015, pp. 453–461. [Online]. Available: <https://doi.org/10.1109/DSD.2015.100>

- [8] M. Panic, E. Quinones, P. G. Zavkov, C. Hernandez, J. Abella, and F. J. Cazorla, "Parallel many-core avionics systems," in *2014 International Conference on Embedded Software (EMSOFT)*, Oct 2014, pp. 1–10.
- [9] M. Panic, C. Hernandez, E. Quinones, J. Abella, and F. J. Cazorla, "Modeling high-performance wormhole NoCs for critical real-time embedded systems," in *2016 IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS)*, April 2016, pp. 1–12.
- [10] D. Rahmati, S. Murali, L. Benini, F. Angiolini, G. D. Micheli, and H. Sarbazi-Azad, "Computing accurate performance bounds for best effort networks-on-chip," *IEEE Transactions on Computers*, vol. 62, no. 3, pp. 452–467, March 2013.
- [11] C.-L. Chou and R. Marculescu, "Contention-aware application mapping for network-on-chip communication architectures," in *2008 IEEE International Conference on Computer Design*, Oct 2008, pp. 164–169.
- [12] C. Zimmer and F. Mueller, "Low contention mapping of real-time tasks onto TilePro 64 core processors," in *2012 IEEE 18th Real Time and Embedded Technology and Applications Symposium*, April 2012, pp. 131–140.
- [13] L. Yang, W. Liu, P. Chen, N. Guan, and M. Li, "Task mapping on SMART NoC: Contention matters, not the distance," in *2017 54th ACM/EDAC/IEEE Design Automation Conference (DAC)*, June 2017, pp. 1–6.
- [14] H. Yu, Y. Ha, and B. Veeravalli, "Communication-aware application mapping and scheduling for NoC-based MPSoCs," in *Proceedings of 2010 IEEE International Symposium on Circuits and Systems*, May 2010, pp. 3232–3235.
- [15] M. Shekhar, H. Ramaprasad, and F. Mueller, "Network-on-chip aware scheduling of hard-real-time tasks," in *Proceedings of the 9th IEEE International Symposium on Industrial Embedded Systems (SIES 2014)*, June 2014, pp. 141–150.
- [16] M. Pani, C. Hernandez, J. Abella, A. Roca, E. Quiones, and F. J. Cazorla, "Improving performance guarantees in wormhole mesh NoC designs," in *2016 Design, Automation Test in Europe Conference Exhibition (DATE)*, March 2016, pp. 1485–1488.
- [17] E. Z. Diego Crupnicoff, Sujal Das, *Deploying Quality of Service and Congestion Control in InfiniBand-based Data Center Networks*, Mellanox Technologies, 2005.



**Jordi Cardona** is a PhD. Student for the CAOS group at BSC. He obtained his M.S. degree in 2018 and graduated in Informatics Engineering in 2016, both titles obtained from the Universitat Politècnica de Catalunya. He enrolled BSC in 2016 where he started working on the analysis of COTS networks on chip for real-time multi-core systems during his master thesis and his current research focuses on monitor and control contention in shared resources of critical real-time systems.

# Modelling and predicting extreme behavior in Critical Real-Time systems with Advanced Statistics

Sergi Vilardell\*<sup>†</sup>, Isabel Serra\*<sup>†</sup>, Francisco J. Cazorla\*,

\*Barcelona Supercomputing Center, Barcelona, Spain

<sup>†</sup> Centre de Recerca Matemàtica, Barcelona, Spain

E-mail: {sergi.vilardell, francisco.cazorla}@bsc.es

iserra@crm.cat

**Keywords**—*CRTES, MBPTA, Extreme Value Theory.*

## I. EXTENDED ABSTRACT

In the last decade, the market for Critical Real-Time Embedded Systems (CRTES) has increased significantly. According to Global Markets Insight [1], the embedded systems market will reach a total size of US \$258 billion in 2023 at an average annual growth rate of 5.6%. Their extensive use in domains such as automotive, aerospace and avionics industry demands ever increasing performance requirements [2]. To satisfy those requirements the CRTES industry has implemented more complex processors, a higher number of memory modules, and accelerators units. Thus the demanding performance requirements have led to a merge of CRTES with High Performance systems. All of these industries work within the framework of CRTES, which puts several restrictions in their design and implementation. Real Time systems require to deliver a response to an event in a restricted time frame or deadline. Real-time systems where missing a deadline provokes a total system failure (hard real-time systems) need satisfy certain guidelines and standards to show that they comply with test for functional and timing behaviour. These standards change depending on the industry, for instance the automotive industry follows ISO 26262 [3] and the aerospace industry follows DO-178C [4]. Researches have developed techniques to analyse the timing correctness in a CRTES. Here, we will expose how they perform on the estimation of the Worst-Case Execution Time (WCET). The WCET is the maximum time that a particular software takes to execute. Estimating its value is crucial from a timing analysis point of view. However there is still not a generalised precise and safe method to produce estimates of WCET [5]. In the CRTES the estimations of the WCET cannot be lower than the true WCET, as they are deemed unsafe; but they cannot exceed it by a significant margin, as they will be deemed pessimistic and impractical.

There are two main frameworks for Timing Analysis. On the one hand there is Static Timing Analysis, which aims at developing an analytical model of the hardware in order to compute the timing of a program. The complexity of the hardware nowadays has increased and Real-Time systems in these days have operating manuals of the order of thousands of pages. If one assumes that they contain complete trustworthy information, it is still a herculean task. On the other hand researches have resorted to a measurement-based approach, called Measurement Based Timing Analysis (MBTA). Here,

the timing analysis is deduced from traces of the execution time. The WCET is deduced from the empirical distribution of the data. In fact the commonly used reference value to obtain is the probabilistic WCET (pWCET). In probabilistic analysis one wants to obtain the probability of exceeding a certain value, and the tool to obtain it lies in Extreme Value Theory [6]. Thus, the pWCET is not a single value like the WCET. Instead it is a distribution function that computes the probability of exceeding high quantiles. This encapsulates the essence of Measurement Based Probabilistic Timing Analysis (MBPTA) [7], [8], [9]. It separates from MBTA in that the empirical distribution of the execution time of multiple experiments, is used to compute a probabilistic WCET (pWCET). This methodology has received the support to be compliant with the safety standards. The aim here is not to estimate the exact WCET, but to compute a distribution that estimates the high quantiles of the empirical distribution of the experiments. All techniques, including the static one need to satisfy two properties. First and foremost, the estimated WCET cannot be lower than the true WCET. Hard real-time systems are design to meet the deadline, otherwise a failure of the system occurs. It is not safe to obtain an optimistic estimation of the maximum delay possible as it leads to potentially dangerous situations. Secondly, the estimated WCET cannot be too far from the true WCET. Determining a resource budget for the software is a big part of critical systems, otherwise it could lead to inefficient and expensive systems.

Extreme Value Theory (EVT) has been the theoretical framework to work with MBPTA and specifically in estimating pWCET. EVT deals with the extreme deviations of the data and provides with tools to contextualize and estimate their behaviour. The tail, the extreme ends of the distribution are characterised based on the value of the extreme value index ( $evi$ ). Tails lighter than exponential ones (so with  $evi < 0$ ) can deliver tighter bounds, as discussed in [10]. Yet, in the context of EVT, either GEV or GPD, distributions with  $evi < 0$  have a compact support, i.e. they have an absolute maximum value that cannot be exceeded. Hence, light tails in the case of EVT have an intrinsic risk of delivering optimistic tail distributions. As we did in our work in [11], we overcame the limitation of the data and delivering a practical solution to obtain pWCET estimates tighter than those of exponential tails while preserving reliability. We did so by complementing EVT with survivability analysis as the theoretical ground for our hypothesis. One of the traits of CRTES is that programs need to finish. They have a maximum budget for time that cannot be exceeded for energy and safety reasons. This translated

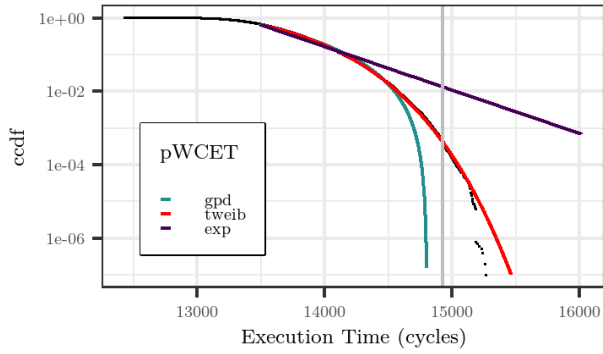


Fig. 2: Whole data of  $n = 10^7$  of railway case study data with different pWCET fittings

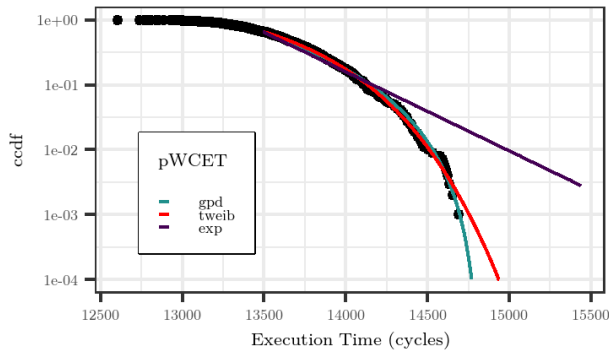


Fig. 1: Sample of  $n = 1000$  of railway case study data with different pWCET fittings

into probability language means that as the time passes when a program is executing, the probability for it to finish gets smaller. Now naturally, these kinds of probability distributions are under the label of light tails, but as we mentioned they deliver optimistic tail distributions. In risk analysis, there is also a property that describes the CRTES program behaviour, and that is Increasing Hazard Rate (IHR). We worked searching by an alternative solution to light tails within the risk analysis domain, and drew an equivalence between IHR and non-heavy tails. From there we derived the next theorem. In order to use IHR distributions for pWCET estimation, we build upon the following theorem proven in [12] and [13]: *Theorem*. Given a non-negative random variable  $X$ , with  $f$  and  $F$  the pdf and cdf, respectively (where  $H(x) = -\log(1 - F(x))$ ,  $x \in \text{support}(X)$ ),

$$\log(f) \text{ concave} \Rightarrow X \text{ IHR} \Leftrightarrow H \text{ convex} \quad (1)$$

From here, we found that a function that satisfies all these properties is the *tailW*. The *tailW* law is constructed using the excess probability function. Thus, the cdf is  $F(x, \alpha, \beta, \nu) = 1 - \exp(-\alpha(x + \nu)^\beta + \alpha\nu^\beta)$  for  $x \geq 0$ ,  $\alpha > 0$ ,  $\beta \geq 0$  and  $\nu > 0$ . Now we can see in Figure 1 we see the complementary cumulative distribution function of the execution times of a program, which represents the probability of having another value bigger than a given execution time. We see how the *tailW* performs better than the exponential function for a small

sample of  $n = 1000$ , and slightly worse than the gpd with light tails. This sample was drawn from a bigger railway case study data of  $n = 10^7$ . Now, if we take the models with the parameters resulting from fitting the sample of  $n = 1000$ , and use them to predict the whole data, we will see how *tailW* performs against the others. In Figure 2 we see how *tailW* is close to the real values of the distribution, while not being overly pessimistic as the exponential function, but also not optimistic and falling behind the data as the gpd with light tails.

**Acknowledgment.** I want to thank Jaume Abella from BSC for being, alongside the authors on the title, an equal contributor to this work.

## REFERENCES

- [1] Global Market Insights, "Embedded system market," 2016.
- [2] M. Duranton *et al.*, "HiPEAC Vision 2019," *HiPEAC 2019*, 2019.
- [3] International Organization for Standardization, "ISO 26262 Road vehicles - Functional Safety," 2018.
- [4] B. Brosgol, "Do-178c: The next avionics safety standard," ser. SIGAda '11. New York, NY, USA: Association for Computing Machinery, 2011, p. 5–6. [Online]. Available: <https://doi.org/10.1145/2070337.2070341>
- [5] J. Abella *et al.*, "Wcet analysis methods: Pitfalls and challenges on their trustworthiness," in *10th IEEE International Symposium on Industrial Embedded Systems (SIES)*, June 2015, pp. 1–10.
- [6] S. Coles, *An introduction to statistical modeling of extreme values*, ser. Springer Series in Statistics. London: Springer-Verlag, 2001.
- [7] F. J. Cazorla *et al.*, "Proartis: Probabilistically analyzable real-time systems," *ACM Trans. Embed. Comput. Syst.*, vol. 12, no. 2s, May 2013. [Online]. Available: <https://doi.org/10.1145/2465787.2465796>
- [8] L. Kosmidis *et al.*, "Achieving timing composability with measurement-based probabilistic timing analysis," in *16th IEEE International Symposium on Object/component/service-oriented Real-time distributed Computing (ISORC 2013)*, June 2013, pp. 1–8.
- [9] F. Wartel *et al.*, "Measurement-based probabilistic timing analysis: Lessons from an integrated-modular avionics case study," in *2013 8th IEEE International Symposium on Industrial Embedded Systems (SIES)*, June 2013, pp. 241–248.
- [10] J. Abella *et al.*, "Measurement-based worst-case execution time estimation using the coefficient of variation," *ACM Trans. Des. Autom. Electron. Syst.*, vol. 22, no. 4, pp. 72:1–72:29, Jun. 2017. [Online]. Available: <http://doi.acm.org/10.1145/3065924>
- [11] S. Vilardell *et al.*, "Software timing analysis for complex hardware with survivability and risk analysis," in *2019 IEEE 37th International Conference on Computer Design (ICCD)*, 2019, pp. 227–236.
- [12] D. Cox and D. Oakes, *Analysis of Survival Data*, ser. Monographs on Statistics and Applied Probability. Chapman and Hall, 1984.
- [13] M. L. Hazelton, "Assessing log-concavity of multivariate densities," *Statistics and Probability Letters*, vol. 81, no. 1, pp. 121 – 125, 2011. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0167715210002774>



**Sergi Vilardell** received his BSc degree in Physics from Universitat Autònoma de Barcelona (UAB), in 2016. The following year, he completed his MSc degree in Modelling for Science and Engineering from UAB in 2017. Since then, he has been working with Isabel Serra from Centre de Recerca Matemàtica (CRM) in Extreme Value Theory problems. Since 2018 he works with the Computer Architecture and Operating Systems (CAOS) group at Barcelona Supercomputer Center (BSC), as well as a PhD student at the department of computer architecture of Universitat Politècnica de Catalunya (UPC), Spain.

# OpenMP static TDG runtime implementation and its usage in Heterogeneous Computing

Chenle YU<sup>\*†</sup>, Sara Royuela<sup>\*</sup>, Eduardo Quiñones<sup>\*</sup>

<sup>\*</sup>Barcelona Supercomputing Center, Barcelona, Spain

<sup>†</sup>Universitat Politècnica de Catalunya, Barcelona, Spain

E-mail: {chenle.yu, sara.royuela, eduardo.quinones}@bsc.es

**Keywords**— *OpenMP, CUDA, CUDA graph.*

## I. EXTENDED ABSTRACT

OpenMP being the standard to use in shared memory parallel programming, it offers the possibility to parallelize sequential program with accelerators by using target directive. However, CUDA Graph as a new, efficient feature is not supported yet. In this work, we present an automatic transformation of OpenMP TDG to CUDA Graph, increasing the programmability of the latter.

### A. Introduction

With the ever growing number of cores on modern CPUs, applications are more likely to be designed to have high scalability, i.e. to have better utilization of the computation power. This is true in Safety-Critical Embedded System (such as ADAS, Advanced Driver Assistance System), High Performance Computing domain and many other computer related areas.

However, efficiently parallelizing an application could be challenging, since there are different existing thread APIs in different programming languages (e.g. Java thread), or in different standards, as pthread for POSIX. In order to render this process easier, Parallel Programming Models have been introduced. Among them, OpenMP is considered as the standard model in shared-memory platform, it has been widely used in the past decades thanks to its performance and programmability. Besides, CUDA as the standard programming model to use if we aim to fully exploit Nvidia cards' performance, although it has a steep learning curve if one is not familiar with thread/block/grid concept of it.

Interestingly, Nvidia introduced a novel task execution model named CUDA Graph [1], which has numerous similarities with Task Dependency Graph (TDG) used in OpenMP, e.g., this feature allows user to define an execution graph once, and reuse it multiple times in the future, the same execution pattern is under consideration to add into the next OpenMP specification. Consequently, having CUDA Graph supported in OpenMP could be used for the next OpenMP specification implementation, and also for increasing the programmability of CUDA Graph. Thus, we hereby present our research aiming to automatically transform from OpenMP tasking program to a CUDA Graph application.

### B. Transforming OpenMP to CUDA Graph

OpenMP 4.0 has defined *target* directive, which allows users to offload the associated task to an accelerator device (e.g. a GPU). However, it does not support CUDA Graph. In order to deploy CUDA Graph in OpenMP, we proceed as the following:

- build a Task Dependency Graph (TDG) from the source OpenMP code
- use this TDG to build CUDA Graph

The first step is accomplished in a previous work by Royuela S. [2], where the author builds the TDG at compile time (referred as static TDG) with a source-to-source compiler: Mercurium[3]. This process requires the knowledge of all task related information at compile time, such as data to consume by tasks, number of loop iterations, etc. to generate the TDG. Although the framework slows down the compilation phase of the application, it parses all *task* constructs and their corresponding *depend* clauses to generate intermediate files, containing execution order information of the OpenMP program. Based on this work, we will see that an implementation of CUDA Graph in OpenMP is possible.

Generating CUDA graph from the intermediate files is done with CUDA Graph API. Host function nodes are created and inserted into the graph through *cudaGraphAddHostNode* function call, while *cudaGraphAddKernelNode* is used for kernel function node. Dependencies among tasks are managed via arrays of *CudaGraphNode\_t*. Finally, the CUDA Graph, mapped from static TDG, is actually instantiated by calling *cudaGraphInstantiate*, and the function *cudaGraphLaunch* is used to launch the executable graph on a specified stream.

Currently, the automatic transformation of OpenMP task program to CUDA Graph has been tested over different benchmarks that have numerous iterations. This is done because CUDA Graph was initially introduced to reduce the host-device communications overhead, the performance gain is supposed to be greater when the number of iterations increases. Secondly, cases where applications run repeatedly before ended by the user are omnipresent, especially in real-time systems. Hence, such comparison is valuable and meaningful. In figure. 1, we show the evolution of *Cholesky* decomposition execution time based on the number of iterations, over a matrix of 4000 x 4000 elements (decomposed to 1540

tasks, or `cudaGraphNode_t`). The execution time of *Cholesky* decomposition using CUDA Graph is roughly 10 times shorter than the original OpenMP program, either for 1 iteration or repetitive execution, as shown in the chart.

In addition, manually writing this example with CUDA Graph is nearly impossible: every task node operates on different data block, and needs almost 10 lines to create the node, set up the correct argument, etc., resulting to a program having more than 15 000 lines. However, such example only had 25 additional lines in our OpenMP program.




Fig. 1. Execution time evolution of Cholesky decomposition w.r.t the number of iterations

### C. Further discussion

As mentioned in section I-A, next OpenMP specification is about to include reusable task graph. The transforming of OpenMP TDG to CUDA Graph fulfills such execution paradigm with the use of GPU accelerator. Regarding homogeneous programming, one implementation of such execution pattern could be: i) if all task information are known at

compile time, generate the static TDG, ii) storage of the graph in memory, and execute it as many times as necessary. Based on our current work, this implementation should be straightforward, and it is in progress.

### D. Conclusion


We have presented our current research work based on the concept of static TDG in OpenMP. Regarding the use of CUDA Graph, we drastically increased its programmability and the performance results satisfied our expectations. In addition, we discussed about how this implementation could be used in the reusable task graph of next OpenMP specification.

## II. ACKNOWLEDGMENT

The realization of this work was possible thanks to the tremendous effort of my PhD advisors: Sara Royuela, Eduardo Quiñones, and Xavier Martorell. This work is part of Ampere european project.

## REFERENCES

- [1] Nvidia, "CUDA 10 Features Revealed: Turing, CUDA Graphs, and More," 2018, <https://devblogs.nvidia.com/cuda-10-features-revealed/>.
- [2] S. Royuela Alcazar, "High-level Compiler Analysis for OpenMP," Ph.D. dissertation, 2018.
- [3] Barcelona Supercomputing Center, "Mercurium," 2019, [pm.bsc.es/mcxx](http://pm.bsc.es/mcxx).



**Chenle** majored in HPC and Cryptography at Sorbonne University, France, during his M.S. He joined BSC in February 2019 as a Ph.D student. His research is focused on HPC, including OpenMP runtime implementation and Heterogeneous Computing.



# Enhanced Performance Using Hybrid Programming Models of Task-based Workflows and MPI

Hatem Elshazly<sup>\*†</sup>, Rosa M. Badia<sup>\*†</sup>

<sup>\*</sup>Barcelona Supercomputing Center, Barcelona, Spain

<sup>†</sup>Universitat Politècnica de Catalunya, Barcelona, Spain

E-mail: {hatem.elshazly, rosa.m.badia}@bsc.es

**Index Terms**—Hybrid Programming Models, MPI, Task-based Parallel Programming Models, Performance, Productivity, High Performance Computing

## I. EXTENDED ABSTRACT

While MPI [1] + X (where X is another parallel programming model) has been proposed and used by the community, we propose a hybrid programming model that combines task-based model + MPI. Task-based workflows offer the necessary abstraction to simplify the application development for large scale execution, and supporting tasks that launch MPI executions enables to exploit the performance capabilities of many-core systems. Hence, application programmers can get the maximum performance out of the underlying systems without compromising the programmability of the application.

We present an extension to PyCOMPSs framework [2], a task-based parallel programming model for the execution of Python applications. Throughout this paper, we name the tasks that natively execute MPI code as *Native MPI Tasks*, as opposed to tasks that call external MPI binaries. Having *Native MPI* tasks as part of the programming model means that in the same source file users can have two types of task: tasks that execute MPI code and other tasks that execute non-MPI code. PyCOMPSs organizes the tasks in Directed Acyclic Graph (DAG) and manages their scheduling and execution, hence users can focus only on the logic of the task.

### A. Native MPI in PyCOMPSs

Tasks are defined in PyCOMPSs by annotating application’s method with Python decorators. Through the `@task` annotation, developers indicate that a function in the code becomes a task. Following the same approach, a method is declared as *Native MPI* task by means of the `@mpi` decorator. The number of MPI processes per *Native MPI* task can be specified using `@constraints` decorator as shown in the sample code snippet in Figure 1.

PyCOMPSs runtime will manage the input and output data of *Native MPI* tasks like any non-MPI task in a completely transparent manner to the user. The runtime will ensure that all the processes in the MPI environment have access to all the input data of the task. The return output of a *Native MPI* task – if any – is a list containing the output of all the MPI processes invoked for the task.

```
@constraints(computingUnits=4)
@mpi(runner='mpirun', computingNodes=1)
@task(returns=int)
def return_ranks(random_num):
    from mpi4py import MPI
    rank = MPI.COMM_WORLD.rank
    return rank*random_num
```

Fig. 1. Simple *Native MPI* task in PyCOMPSs. `return_ranks` task will be executed by 4 MPI processes as specified in `computingUnits` on 1 node. It returns a list of each MPI rank multiplied by the `random_num` input value.

Similar to non-MPI PyCOMPSs tasks, the execution details of *Native MPI* tasks are completely abstracted from the runtime; the MPI environment is encapsulated within the *Native MPI* task that launched it. Thus, one workflow can have multiple *Native MPI* tasks, each with different configuration parameters (i.e., number of computing nodes and MPI processes) and combine them with other tasks in the task execution graph.

PyCOMPSs runtime launches special Python worker processes for *Native MPI* tasks at the time of the task execution to launch the MPI environment and manage the task execution. If two *Native MPI* tasks are scheduled for execution at the same time, the runtime launches an exclusive MPI worker for each of them. Hence, each of the tasks will have its own isolated execution environment.

### B. Evaluation

In this section, we evaluate performance benefits and trade-offs of using *Native MPI* tasks in PyCOMPSs. Experiments were conducted on the MareNostrum4 supercomputer; which includes a set of high-memory computing nodes with 48 cores and 370 GB of memory each. Each experiment was run multiple times: using sequential implementation of the targeted tasks and a parallel implementation with an increasing number of MPI processes (2, 4 and 8). In all experiments, the sequential implementation of the task is used as the baseline.

For the purpose of this evaluation, we developed an application that calculates the term frequency (TF-IDF) of a web archive file. We used an input web archive file of a total size of 186 Gbytes. The application consists of a reading task which reads a record from the file and a compute task that calculates TF-IDF. The total number of tasks for this application is 1440 tasks; 720 read tasks and 720 corresponding compute tasks.

Figure 2 shows the performance results of the application. As shown in Figure 2(a) the average time per compute task decreases while increasing the number of MPI processes per compute task. Using 8 MPI processes per compute task, we obtained up to 7x speedup in the average time per compute task. In addition to that, as shown in Figure 2(b), the performance improvement per compute task is reflected as up to 3x speedup improvement in the total execution time.

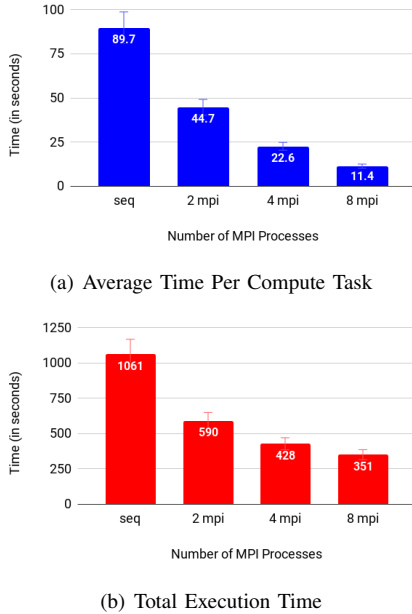


Fig. 2. Performance Results for Web Archive Analysis Application

To further understand the performance and behaviour of *Native Python MPI* tasks in PyCOMPSs, several experiments were conducted on the Web Archive Analysis. Each experiment was launched multiple times with a sequential implementation task and then a parallel *Native MPI* task implementation with different numbers of MPI processes (2, 4, 8, 16 and 48) on different number of nodes (4, 8 and 12).

As shown in Figure 3, as the number of nodes increases, task parallelism increases so the total execution time of both applications improves. For a specific number of nodes, total execution time decreases until it reaches a point after which it starts to increase as the number of MPI processes per *Native MPI* task increases. This point is 8 MPI processes for 4, 8 nodes and 16 MPI processes for 12 nodes. This is because *Native Python MPI* tasks use the `@constraint` decorator of PyCOMPSs to specify the number of MPI processes per task. Increasing the number of MPI processes per task (i.e. increasing task constraints) decreases task parallelism. This effect is mitigated as the number of resources increases because there are enough resources to maintain the same level or allow for more task parallelism. This can be noted in Figure I-B where for 4 and 8 nodes the total execution time degrades at 8 MPI processes but when the number of nodes is increased to 12, this point shifts to 16 MPI processes.




Fig. 3. Scalability Results

### C. Conclusion

Enabling the execution of MPI code natively in PyCOMPSs tasks offers great benefits in terms of both programmability and performance for Python applications. However, a tradeoff arises between MPI parallelism per task and task parallelism that may negatively affect the total time of the application. As future work, we plan to improve the scheduling of tasks to better utilize the underlying infrastructure.

### II. ACKNOWLEDGMENT

This work has been published in the 28th Euromicro International Conference on Parallel, Distributed and Network-based Processing (PDP), 2020 [3].

### REFERENCES

- [1] W. Gropp, E. Lusk, and A. Skjellum, *Using MPI: portable parallel programming with the message-passing interface*. MIT press, 1999, vol. 1.
- [2] E. Tejedor, Y. Becerra, G. Alomar, A. Queralt, R. M. Badia, J. Torres, T. Cortes, and J. Labarta, "Pycomps: Parallel computational workflows in python," *International Journal of High Performance Computing Applications*, 2015.
- [3] H. Elshazly, F. Lordan, J. Ejarque, and R. M. Badia, "Performance Meets Programmability: Enabling Native Python MPI In PyCOMPSs," in *28th Euromicro International Conference on Parallel, Distributed and Network-based Processing*, 2020.

**Hatem Elshazly** received his BSc degree in Computer Science from Cairo University, Egypt in 2012. The following year, he joined as a Masters student and research software engineer at Nile University, Egypt. He completed his MSc degree in Optimizing Data Intensive Applications in 2016. Since 2018, he has been with the workflows and distributed computing group of Barcelona Supercomputing Center (BSC) as well as a PhD student at the department of computer architecture of Universitat Politècnica de Catalunya (UPC), Spain.



# The atmospheric iron cycle in EC-Earth

Elisa Bergas-Massó\*, María Gonçalves-Ageitos\*†, Carlos Pérez García-Pando\*§

\*Barcelona Supercomputing Center (BSC), Barcelona, Spain

†Universitat Politècnica de Catalunya (UPC), Barcelona, Spain

§ICREA, Pg. Lluís Companys 23, 08010 Barcelona, Spain.

E-mail: elisa.bergas@bsc.es

**Keywords**—Iron cycle, Mineralogy, Climate, Earth System Model, EC-Earth3.

## I. EXTENDED ABSTRACT

The ocean is known to act as an atmospheric carbon dioxide ( $CO_2$ ) sink. About a quarter of the  $CO_2$  emitted to the atmosphere since the industrial revolution, has been captured by the ocean [1]. The capacity of the ocean to capture  $CO_2$  highly depends on ocean productivity which relies upon bio-available iron (Fe) for photosynthesis, respiration and nitrogen fixation [2]. Fe is in fact considered to be the limiting nutrient in some remote regions of the ocean known as high-nutrient low-chlorophyll (HNLC) [3]. Understanding and constraining the bio-available iron supply to the ocean is thus fundamental to be able to project future climate.

Fe supply reaches the oceans mainly from rivers as suspended sediment. However, fluvial and glacial particulate Fe is restricted to near-coastal areas. Therefore, the dominant input of iron to open ocean surface is the deposition of atmospheric mineral dust emitted from arid and semiarid areas of the world. Another contributor to atmospheric Fe supply that is not always accounted for in models, is combustion, which main sources are anthropogenic combustion and biomass burning.

Just a fraction of the deposited Fe over ocean can be used by marine biota as nutrient (bio-available). The assumption that soluble Fe can be considered as bio-available will be used here [4]. Freshly emitted Fe-dust is known to be mainly insoluble. Observations, modelling and laboratory studies suggest that the solubility of Fe-dust increases downwind of the sources due to different processes [5] [6]. On the other hand, although the total burden of emitted combustion Fe is known to be smaller than Fe-dust, combustion Fe at emission may be more soluble [7].

The dust Fe content, speciation and ability to dissolve depends upon mineralogy. Earth System Models (ESMs) typically assume that dust aerosols have a globally uniform composition, neglecting the known local and regional variations in the mineralogical composition of the sources. In fact, dust aerosols are a mixture of different minerals, whose relative abundances, particle size distribution (PSD), shape, surface topography and mixing state influence their effect upon climate.

In this work we aim to improve our understanding of the atmospheric delivery of bio-available Fe to the ocean by implementing and constraining Fe emissions and a Fe solubilization atmospheric mechanism in a state-of-the art ESM, EC-Earth. EC-Earth is collaboratively developed by European research

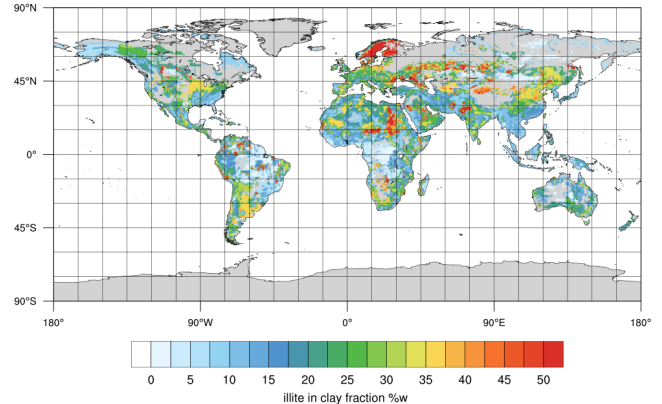


Fig. 1. Illite in clay fraction [%] in Claquin mineralogical dataset.

centers from 10 different countries including the Barcelona Supercomputing Center (BSC) [8]. All our developments will be implemented in the atmospheric chemistry module of EC-Earth: the Tracer Model 5 (TM5) [9].

## A. Iron emissions

Relying on the hypothesis that both dust and combustion aerosols are the major sources of bio-available Fe deposited over open ocean surface, emissions coming from those sources need to be properly assessed in the model.

1) *Fe-dust emissions*: We will implement a soil mineralogical dataset (see Figure 1) [10] [11] and improve and apply an extended version of brittle fragmentation theory (BFT) [12] to represent the emitted PSD of each mineral in each grid cell of the model. Dust Fe content will be then calculated based on standard chemical formulas for each mineral [13] (Nickovic et al. (2013))

2) *Fe-combustion emissions*: Current emission inventories applicable to present and future time periods do not provide explicitly iron species. A common assumption is used here in which iron emissions are estimated based on other combustion species as black carbon (BC) or organic carbon (OC).

In this work scaling factors, which are based on experimental data [14], of total Fe emissions to those of BC (Fe:BC) and OC (Fe:OC) for each of the emission sectors in the IPCC-AR5 inventory [15] are used.

## B. Atmospheric processing of Fe

TM5 simulates aqueous-phase chemistry in aerosol water and cloud droplets as described in [16]. The atmospheric

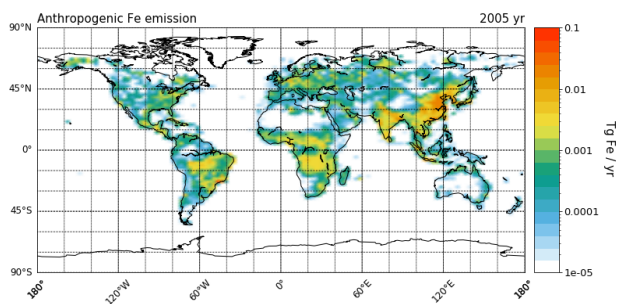


Fig. 2. Combustion Fe emissions  $[(TgFe)/yr]$  for the year 2005 for one of our simulations.

processing mechanism of iron is treated as a kinetic process accounting for 1) a proton-promoted and 2) an oxalate-promoted.

For the proton-promoted mechanism, the dissolution rate of minerals in aerosol and cloud water is calculated by applying an empirical parameterization [17], taking the saturation degree of the solution, the type of mineral as well as the reactivity of Fe species and the ambient temperature into account. The ligand-promoted dissolution scheme follows what has been experimentally proposed by other studies [18] and directly depends on oxalate concentration.

### C. Planned work and expected outcome

Time-slice experiments for the present-day and future climate scenarios will be carried out. Our goal is to understand and compare the responses and regional variations of soluble Fe deposition under different periods, with different anthropogenic emissions, and quantify the contribution of the different sources to the soluble Fe deposition.

Present-day simulations will be evaluated with available compilations of total Fe deposition and concentration measurements.

This study will allow to do fully coupled simulations that account for the effect of bioavailable Fe variations upon the carbon and nitrogen cycles (currently ESM base the determination of the soluble Fe deposited over ocean on climatological information). The implementation of a explicit mineralogy for dust in EC-Earth will also allow a more detailed exploration of other effects of dust on climate.

## II. ACKNOWLEDGMENT

We thankfully acknowledge the computer resources at Marenostrum4, and the technical support provided by the BSC and the Computational Earth Sciences team of the Earth Sciences Department. This work was funded by the Ministerio de Economía y Competitividad (MINECO) as part of the NUTRIENT project (CGL2017-88911-R). The research leading to this work has also received support by the ERC Consolidator Grant “FRAGMENT” (grant agreement No. 773051), and the AXA Chair on Sand and Dust Storms funded by the AXA research Fund. Finally, also acknowledge financial support for this work by the National Observatory of Athens research grant (n° 5065).

## REFERENCES

- [1] P. Ciais and C. Sabine, “Carbon and other biogeochemical cycles,” in *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, 2013.
- [2] M. Behrenfeld *et al.*, “Climate-driven trends in contemporary ocean productivity,” in *Nature*, vol. 444, 01 2007, pp. 752–5.
- [3] T. Jickells *et al.*, “Global iron connections between desert dust, ocean biogeochemistry, and climate,” in *Science (New York, N.Y.)*, vol. 308, 05 2005, pp. 67–71.
- [4] A. R. Baker and T. D. Jickells, “Mineral particle size as a control on aerosol iron solubility,” in *Geophysical Research Letters*, vol. 33, no. 17, 2006.
- [5] J. A. Rizzolo *et al.*, “Soluble iron nutrients in saharan dust over the central amazon rainforest,” in *Atmospheric Chemistry and Physics*, vol. 17, no. 4, 2017, pp. 2673–2687.
- [6] G. Zhuang *et al.*, “Link between iron and sulfur cycles suggested by fe(ii) in remote marine aerosol,” in *Nature*, vol. 355, 02 1992, pp. 537–539.
- [7] H. Matsui *et al.*, “Anthropogenic combustion iron as a complex climate forcer,” in *Nature Communications*, vol. 9, 04 2018, paper 1593.
- [8] W. Hazeleger *et al.*, “Ec-earth v2.2: Description and validation of a new seamless earth system prediction model,” in *Climate Dynamics*, vol. 39, 12 2011, pp. 1–19.
- [9] V. Huijnen *et al.*, “The global chemistry transport model tm5: description and evaluation of the tropospheric chemistry version 3.0,” *Geoscientific Model Development*, vol. 3, no. 2, pp. 445–473, 2010.
- [10] T. Claquin *et al.*, “Modeling the mineralogy of atmospheric dust sources,” in *Journal of Geophysical Research*, vol. 104256, 09 1999, pp. 243–22.
- [11] E. Journet *et al.*, “A new data set of soil mineralogy for dust-cycle modeling,” in *Atmospheric Chemistry and Physics*, vol. 14, 04 2014, pp. 3801–3816.
- [12] J. Kok, “A scaling theory for the size distribution of emitted dust aerosols suggests climate models underestimate the size of the global dust cycle,” in *Proceedings of the National Academy of Sciences of the United States of America*, vol. 108, 01 2011, pp. 1016–21.
- [13] C. P. García-Pando *et al.*, “Predicting the mineral composition of dust aerosols: Insights from elemental composition measured at the izaña observatory,” in *Geophys. Res. Lett.*, 2016.
- [14] A. Ito *et al.*, “Radiative forcing by light-absorbing aerosols of pyrogenetic iron oxides,” in *Scientific Reports*, vol. 8, no. 1, may 2018.
- [15] “Ipcc, 2014: Climate change 2014: Synthesis report. contribution of working groups i, ii and iii to the fifth assessment report of the intergovernmental panel on climate change [core writing team, r.k. pachauri and l.a. meyer (eds.)],” *IPCC*, p. 151.
- [16] S. Myriokefalitakis *et al.*, “Changes in dissolved iron deposition to the oceans driven by human activity: a 3-d global modelling study,” in *Biogeosciences*, vol. 12, 07 2015.
- [17] A. C. Lasaga *et al.*, “Chemical-weathering rate laws and global geochemical cycles,” in *Geochimica et Cosmochimica Acta*, vol. 58, 1994, pp. 2361–2386.
- [18] R. Paris *et al.*, “Variability of dust iron solubility in atmospheric waters: Investigation of the role of oxalate organic complexation,” *Atmospheric Environment*, vol. 45, pp. 6510–6517, 11 2011.



**Elisa Bergas-Massó** received her BSc degree in Physics with a mention in fundamental physics from Universitat de Barcelona (UB) in 2018. She completed her MSc degree in Meteorology in the same university in 2019. While doing the MSc degree, she did an internship in the Atmospheric Composition group of the Earth Sciences department of the Barcelona Supercomputing Center (BSC) working on mineral dust emission. Since September 2019, in the scope of a Ph.D., she has focused her work on the iron cycle and its implementation in climate models.

# Field Observations of the Variability of Dust Emission, its Size-Spectrum and Mineralogy

Cristina González-Flórez<sup>#1</sup>, Martina Klose<sup>#2</sup>, Carlos Pérez García-Pando<sup>#&3</sup>

<sup>#</sup> Earth Sciences Department, Barcelona Supercomputing Center (BSC), Barcelona, Spain

& ICREA, Pg. Lluís Companys 23, 08010 Barcelona, Spain

{<sup>1</sup>cristina.gonzalez, <sup>2</sup>martina.klose, <sup>3</sup>carlos.perez}@bsc.es

**Keywords**— Mineral dust, size distribution, mineral composition, field campaign, Morocco

## EXTENDED ABSTRACT

### A. Introduction

Atmospheric mineral dust consists of tiny mineral particles that are produced by the wind erosion of arid and semi-arid surfaces of the Earth, and it is one of the most important aerosols in terms of mass in the global atmosphere [1].

The physical and chemical properties of dust, that is, its particle size distribution (PSD), mineralogical composition, shape and mixing state determine its impact on the Earth's system. Dust mineralogy in particular has been identified by the Intergovernmental Panel on Climate Change (IPCC) as a key uncertainty in the overall contribution of aerosols to radiative forcing [2] and many studies over recent years have shown its potential importance [3,4].

Despite this, Earth System models typically assume dust aerosols to have a globally uniform composition, neglecting the known local and regional variations in the mineralogical composition of the sources [5,6] and therefore, preventing further understanding of the role of dust in the Earth system. However, this simplification is justified by the current incomplete understanding of the physical processes at emission, the lack of coincident measurements of individual mineral PSDs for emitted dust and the parent soil, the fundamental disagreements among existing dust emission schemes on multiple aspects, the limited global knowledge of soil mineral content and the insufficient knowledge of the mixing state of the minerals.

The ERC Consolidator Grant called FRAGMENT (FRontiers in dust minerAlOGical coMposition and its Effects upoN climaTe) aims to address these limitations and to better understand and constrain the global mineralogical composition of dust along with its effects upon climate. This ambitious and multidisciplinary project combines theory, field measurements, laboratory analyses, remote spectroscopy and modelling.

### B. Objectives

As part of FRAGMENT, my research aims to improve our fundamental and quantitative understanding of the emitted dust PSD and mineralogy along with their relationship to the parent soil properties by fulfilling the following specific objectives:

- Study the size-resolved dust fluxes under different meteorological and soil conditions and its relationship with the emission mechanism.
- Overcome our limited understanding of the size-resolved mineralogy of the emitted dust.

Given the incomplete understanding of the physical processes and the paucity and incompleteness of available

measurements, new experimental data are required in order to answer these key scientific questions, and to test, evaluate and extend currently available theories and models.

### C. Methodology

The FRAGMENT team is performing an unprecedented set of coordinated field campaigns in Morocco, California and Iceland. These regions meet several key criteria including accessibility, variety of soil types, textures and landforms, local/regional collaborators that help with the logistics of the campaigns. This combination of measurement locations and conditions will allow FRAGMENT to tackle the aforementioned and other key research objectives.

The first field campaign took place in September 2019 at “El Bour”, a dry lake located at the edge of the Sahara, approximately 15 km west of M'Hamid El Ghizlane in Morocco and 30 km north of the Moroccan-Algerian border. Fig. 1 shows some photos of our experimental set up, which includes a wide variety of instruments to perform a detailed characterization of the soil, dust emission and meteorology. These instruments include, e.g., low volume samplers, cascade impactors, passive aeolian sediment traps, saltation sensors, aerosol spectrometers, water content reflectometers, a fully equipped meteorological tower, a radiometer, an aethalometer and a polar nephelometer. Results presented here focus mainly on the time variation of the particle size distribution, in terms of mass concentrations, and its relationship with meteorological variables.



Fig. 1 Photos of the experimental setup at “El Bour” (Morocco) in September 2019 and small insert showing a map with the measurement site.

### D. Preliminary results of the field campaign in Morocco

A large number of dust events of varying intensity were recorded during this one-month measurement period (Fig. 2).

The average maximum air temperature at 2 m height during the field campaign was very close to 40°C and the minimum above 20°C. In addition, the average daily maximum 2 m

wind speed was around 12 m/s and peak wind gusts reached 21.68 m/s.

A few moist convective storms occurred during our period of measurements. After a moist-convective storm on 6<sup>th</sup> September 2019, results of which we present in the following, the surrounding of our site was flooded to large degree.

Fig. 2 shows an example of data collected on 6<sup>th</sup> September 2019 by an aerosol spectrometer (Fidas) and a wind sensor placed at 2 m height. The top panel depicts the temporal evolution of the two-minute average Particulate Matter (PM) levels along with their respective ratios (multiplied by a factor of 1000), while wind speed is represented at the bottom. As can be observed, from 8 to 12 UTC, average wind speeds varied around 7.5 m/s, and PM values exhibited several consecutive peaks, the highest one reaching the following values: 5517.09  $\mu\text{g}/\text{m}^3$  of PM total (Particulate Matter 27.38  $\mu\text{m}$  or less in diameter), 3493.72  $\mu\text{g}/\text{m}^3$  of PM<sub>10</sub>, 1639.57  $\mu\text{g}/\text{m}^3$  of PM<sub>4</sub>, 795.51  $\mu\text{g}/\text{m}^3$  of PM<sub>2.5</sub> and 118.22  $\mu\text{g}/\text{m}^3$  of PM<sub>1</sub>. If we zoom in this temporal interval, we can first identify very similar patterns between the peaks of wind and PM levels, and second, we observe that the ratios of PM values remain almost constant. The latter gives an idea of the evolution of particle size distribution, which in this case seems not to be much affected by dust emission. From 12UTC onward, both PM levels and wind speed gradually decrease until experiencing a sharp rise at around 16 UTC, which corresponds to one of the strongest dust events recorded during the field campaign. At the beginning of this event, considered as a haboob, we could clearly identify a wall of blowing dust formed from the outflow of a strong convective storm. During this haboob, which lasted around two hours, PM levels reached the following values: 30732.32  $\mu\text{g}/\text{m}^3$  of PM total, 20286.24  $\mu\text{g}/\text{m}^3$  of PM<sub>10</sub>, 9555.64  $\mu\text{g}/\text{m}^3$  of PM<sub>4</sub>, 4505.46  $\mu\text{g}/\text{m}^3$  of PM<sub>2.5</sub> and 597.99  $\mu\text{g}/\text{m}^3$  of PM<sub>1</sub>. Besides, PM<sub>tot</sub>/PM<sub>10</sub> and PM<sub>10</sub>/PM<sub>2.5</sub> ratios remained almost constant whereas the ratio PM<sub>2.5</sub>/PM<sub>1</sub> experienced a gradual rise. However, the most striking and steep increase in PM<sub>10</sub>/PM<sub>2.5</sub> and PM<sub>2.5</sub>/PM<sub>1</sub> ratios, which reflect a coarser particle size distribution, occurred around 20UTC when the two-minute average wind speed was below 5 m/s and there were only some gusts of wind close to 7 m/s.




Fig. 2 Data from 6<sup>th</sup> September 2019. Top (bottom): Two minute-average PM values and their ratios (wind speed) at 2 m height.

## E. Outlook

The next steps will focus on (a) analysing the size distribution of measured dust concentration, (b) calculating size-resolved dust emission fluxes, (c) combining the results on dust emission and size-distribution with results on mineralogy determined by other FRAGMENT team members, (d) evaluating existing dust emission theories against the measurement results and (e) evaluating exiting frameworks for mineralogy and expand as needed.

## F. ACKNOWLEDGEMENTS

FRAGMENT has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 773051). The entire FRAGMENT team has contributed to this research as part of the field campaign in Morocco. We thank Cristina Reche for processing part of the data presented in this abstract.

## References

- [1] Textor, C., Schulz, M., Guibert, S. et al., "Analysis and quantification of the diversities of aerosol life cycles within AeroCom". Atmospheric Chemistry and Physics, Vol. 6, Pp. 1777-1813, May 2006.
- [2] Boucher, O., Randall, D., Artaxo, P. et al., "Clouds and aerosols. Climate change 2013: The physical science basis. Contribution of working group I to the fifth assessment report of the intergovernmental panel on climate change". Cambridge University Press, Cambridge, 2013.
- [3] Atkinson, J. D., Murray, B. J., Woodhouse, M. T. et al. "The importance of feldspar for ice nucleation by mineral dust in mixed-phase clouds". Nature, Vol. 498, no. 7454, Pp. 355-358, Aug 2013.
- [4] Di Biagio, C., Formenti, P., Balkanski, Y. et al. "Complex refractive indices and single scattering albedo of global dust aerosols in the shortwave spectrum and relationship to iron content and size". Atmos. Chem. and Phys. Disc., (10.5194/acp-2019-145) (hal-02368521), 2019.
- [5] Claquin, T., Schulz, M., and Balkanski, Y. "Modeling the mineralogy of atmospheric dust sources". Journal of Geophysical Research: Atmospheres, Vol. 104, no. D18, Pp. 22243-22256, Sep.1999
- [6] Journet, E., Balkanski, Y., and Harrison, S. P. "A new data set of soil mineralogy for dust-cycle modeling". Atmos. Chem. Phys., Vol. 14, no. 8, Pp. 3801-3816, Apr. 2014.

## Author biography



**Cristina González Flórez** was born in Madrid, Spain, in 1995. She received her BSc degree in Physics (2017) and MSc degree in Meteorology and Geophysics (2018) from the Complutense University of Madrid, Spain. She is currently pursuing her PhD in Environmental Engineering at Universitat Politècnica de Catalunya (UPC) and performing her research at the facilities of the Barcelona Supercomputing Center (BSC), under the supervision of Dr. Carlos Pérez García-Pando and Dr. Martina Klose.

# Multi-model assessment of the late-winter extra-tropical response to El Niño and La Niña

Bianca Mezzina<sup>1</sup> (✉ bianca.mezzina@bsc.es), Javier García-Serrano<sup>1,2</sup>, Ileana Bladé<sup>2</sup>

<sup>1</sup>Barcelona Supercomputing Center (BSC), Barcelona, Spain

<sup>2</sup>Group of Meteorology, Universitat de Barcelona (UB), Barcelona, Spain

**Keywords**—*ENSO, teleconnections, climate variability, atmospheric models*

## I. INTRODUCTION

El Niño-Southern Oscillation (ENSO) is a natural phenomenon in the tropical Pacific and the dominant mode of climate variability on interannual timescales. The first term, El Niño, refers to a recurring warming of the tropical Pacific Ocean (every 2-7 years), while the opposite phase, an anomalous cooling, is called La Niña. These variations in sea surface temperature (SST) are accompanied by changes in the tropical atmospheric circulation (Southern Oscillation), thus making ENSO a coupled phenomenon involving ocean-atmosphere interactions. Furthermore, ENSO can affect climate in regions far from the tropical Pacific, producing a cascade of global impacts through so-called ‘teleconnections’. Understanding the extra-tropical impacts of ENSO is important to improve seasonal forecasts, for which it represents the most important source of predictability.

In the North Atlantic-European (NAE) sector, the ENSO teleconnection is still controversial in several aspects. A first cornerstone was set in a review by Brönnimann (2007) [1], who concluded that a robust, ‘canonical’ ENSO signal exists over the NAE region in late winter (January to March, JFM): a dipole in sea-level pressure (SLP) with centers over the mid-latitude and high-latitude North Atlantic. While Brönnimann described this pattern as “close to symmetric” for El Niño and La Niña, recent studies deliver contradictory results, with some reporting a symmetric signal (e.g. [2] [3] [4]) and others claiming asymmetry (e.g. [5] [6] [7]). The actual linearity of the ENSO-NAE teleconnection thus remains unresolved, and addressing this issue is the primary objective of this study. We will also investigate another key aspect of the ENSO-NAE teleconnection that is nothing but settled: the dynamical mechanism leading to the ‘canonical’ SLP dipole.

The underlying idea of this study is to use idealized experiments with atmospheric models forced by symmetric anomalous SST patterns representing El Niño and La Niña to diagnose symmetries and asymmetries in the extra-tropical response. A multi-model approach is used, as the experiments analyzed here are run with the same protocol using three state-of-the-art models. We aim not only at diagnosing asymmetries in the extra-tropical ENSO-related SLP signal, but also at understanding their cause by examining all the steps involved in the atmospheric response, starting from the tropical Pacific.

## II. METHODS

All experiments are atmosphere-only simulations. The multi-model ensemble consists of the atmospheric components of three state-of-the-art models:

- EC-EARTH3.2 (T255L91, 0.01 hPa) [8]
- CNRM-CM6-1 (T127L91, 0.01 hPa) [9]
- CMCC-SPS3 (~110 km, L46, 0.3 hPa) [10]

The set of experiments include a control simulation and two perturbed runs. The control simulation (CTL) is run with climatological SSTs from observations (HadISST2.2). The El Niño (EN) experiment is performed with SST anomalies that mimic a strong, canonical El Niño event (Fig. 1), while the La Niña experiment (LN) has identical prescribed pattern but with flipped-sign SST anomalies. The forced atmospheric response associated with El Niño (La Niña) is estimated by computing the difference between the ensemble mean of the 50 winters in EN (LN) and CTL. The target season is JFM.

## III. MAIN RESULTS

Asymmetries arise in the SLP response over both the North Pacific (Aleutian Low) and NAE sector (North Atlantic dipole): the response to cold (La Niña) SST anomalies tend to be weaker and shifted westward with respect to the one associated with warm (El Niño) anomalous forcing (Fig. 2). The same behaviour is present in the upper troposphere (200-hPa geopotential height) in the large-scale Rossby wave train that is ENSO’s dominant extra-tropical response.

The response of tropical convection to the SST forcing is underlying the extra-tropical asymmetries: warm (cold) SST anomalies superimposed to the mean state enlarge (restrict) the region suitable for the triggering of deep convection (SST above 27°C) and increase (decrease) the amount of available heating, while the longitude of maximum convection is found east (west) of the Date Line due to the different SST gradient. The convective response is followed by anomalous divergent




Fig. 1. JFM average of the anomalous SST pattern prescribed in EN.

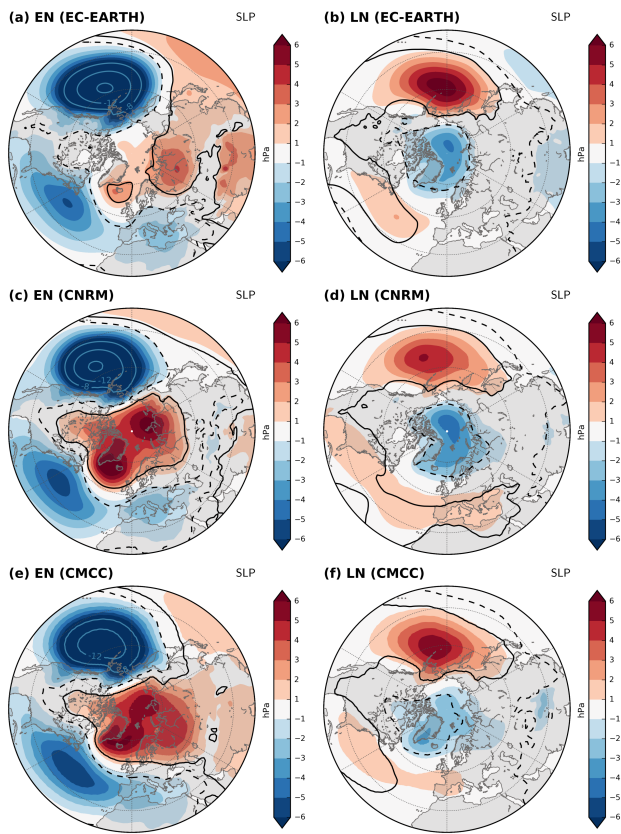


Fig. 2. Ensemble-mean SLP anomalies for (left) EN and (right) LN with respect to CTL in JFM: EC-EARTH (top), CNRM (middle), CMCC (bottom). Blue contours show values exceeding the color scale limit at -8, -12, -16 hPa. Black contours indicate statistically significant areas at the 95% confidence level.

wind with similar properties, but in order to explain the longitudinal shift of the extra-tropical SLP signal, the anomalous divergence needs to be considered in tandem with the mean flow, namely diagnosing the Rossby wave source.

The ENSO surface signal in the NAE sector is the ‘canonical’ dipole between mid and high latitudes, with asymmetries in terms of amplitude and longitude but not structure. These asymmetries are not indicative of different mechanisms driving the teleconnection for El Niño and La Niña; instead, in both cases the ENSO teleconnection to the North Atlantic is mainly associated with the large-scale tropospheric Rossby wave train and its westward tilt with height.

#### IV. ACKNOWLEDGMENTS

We are thankful to the other co-authors of this study: Froila M. Palmeiro (UB), Lauriane Batté (CNRM), Constantin Ardilouze (CNRM), Marianna Benassi (CMCC), and Silvio Gualdi (CMCC). This work was supported by the MEDSCOPE project, co-funded by the European Union (Grant 690462), and by the Spanish DANAE project (CGL2015-68342-R).

#### REFERENCES

- [1] S. Brönnimann, “Impact of El Niño-Southern Oscillation on European climate.” *Rev Geophys*, vol. 45, no. 3, p. RG3003, 2007.
- [2] C. Deser *et al.*, “The Northern Hemisphere Extratropical Atmospheric Circulation Response to ENSO: How Well Do We Know It and How Do We Evaluate Models Accordingly?” *J Clim*, vol. 30, no. 13, pp. 5059–5082, 2017.
- [3] B. Ayarzagüena *et al.*, “Intraseasonal Effects of El Niño-Southern Oscillation on North Atlantic Climate.” *J Clim*, vol. 31, no. 21, pp. 8861–8873, 2018.
- [4] I. Weinberger *et al.*, “The salience of nonlinearities in the boreal winter response to ENSO: Arctic stratosphere and Europe.” *Clim Dyn*, no. 53, pp. 4591–4610, 2019.
- [5] P. Trascasa-Castro *et al.*, “On the Linearity of the Stratospheric and Euro-Atlantic Sector Response to ENSO.” *J Clim*, vol. 32, no. 19, pp. 6607–6626, 2019.
- [6] S. C. Hardiman *et al.*, “The Impact of Strong El Niño and La Niña Events on the North Atlantic.” *Geophys Res Lett*, vol. 46, no. 5, pp. 2874–2883, 2019.
- [7] B. Jiménez-Estève and D. I. V. Domeisen, “Nonlinearity in the North Pacific Atmospheric Response to a Linear ENSO Forcing.” *Geophys Res Lett*, vol. 46, no. 4, pp. 2271–2281, 2019.
- [8] R. Haarsma *et al.*, “HighResMIP versions of EC-Earth: EC-Earth3P and EC-Earth3P-HR. Description, model performance, data handling and validation.” *Geosci Model Dev Disc*, vol. 2020, pp. 1–37, 2020.
- [9] A. Voldoire *et al.*, “Evaluation of CMIP6 DECK Experiments With CNRM-CM6-1.” *J Adv Model Earth Syst*, vol. 11, no. 7, pp. 2177–2213, 2019.
- [10] A. Sanna *et al.*, “CMCC-SPS3: The CMCC Seasonal Prediction System 3.” *CMCC Tech Rep RP0285*, 2019.



**Bianca Mezzina** holds a BSc in Physics and a MSc in Terrestrial and Environmental Physics from the University of Trieste, Italy. Following mysterious signs from the universe, in 2016 she started a PhD at the Barcelona Supercomputing Center, within the Earth Sciences Department. No regrets so far.



# An integrated model system tool to evaluate the impact of urban mobility policies on air pollution: Barcelona case study

Daniel Rodriguez-Rey\*, Marc Guevara\*, Josep Casanovas\*<sup>†</sup>

\*Barcelona Supercomputing Center, Barcelona, Spain

<sup>†</sup>Universitat Politècnica de Catalunya, Barcelona, Spain

E-mail: {daniel.rodriguez, marc.guevara}@bsc.es

**Abstract**—Air pollution remains as a key unresolved problem in many urban areas. Cities with such problem are gradually implementing Traffic Management Strategies (TMS) to reduce the total kilometers travelled by vehicles and subsequently decrease emissions. However, a prior evaluation of such TMS is needed if the target goals want to be achieved. In this sense, the combination of traffic simulation with emissions and air quality models can be of great use to assess the potential impacts of such policies. This study presents an integrated modelling system tool for Barcelona that allows to estimate the changes induced by the implementation of TMS on traffic activity, emissions and air quality levels at a very spatial (street level) and temporal (hourly level) resolution.

**Keywords**—Traffic emissions, Traffic modelling, traffic management strategies.

## I. EXTENDED ABSTRACT

### A. Introduction

In the city of Barcelona (Spain), chronic nitrogen dioxides (NO<sub>2</sub>) and fine particular matter (PM<sub>2.5</sub>) concentrations exceeding the 2008/50/EC EU Ambient Air Quality Directive (AQD) and the World Health Organization (WHO) air quality guidelines (AQGs) are being recorded in urban traffic stations [1]. Several studies have also highlighted the impacts of air pollution on public health [2]. Local authorities are focusing on mobility policies (e.g. implementation of Low Emission Zones, traffic calming) that try to reduce the number of circulating vehicles within the city. In this sense, the application of numerical models is highlighted in the AQD as a fundamental tool to better assess and manage air quality, encouraging their use in the evaluation of air quality plans. In order to simulate the effect of restrictions on the traffic activity across a city, an integrated and dynamical system that links a transportation model with an emissions model and an air quality model is needed.

### B. Related work and research contribution

During the last years, several works have performed the assessment of the effect of TMS in air pollution by coupling traffic models with air quality systems. Most of these studies have used mesoscale air quality systems, which are limited to spatial resolutions of 1km<sup>2</sup> and therefore cannot depict the strong urban pollutant concentration gradients. Works evaluating the impact of TMS at street-level have been limited to restricted domains (e.g. [3]). The system presented in this work allows to estimate street-level emission values in Barcelona

by a dynamic approach able to estimate the effect of different TMS. The complete workflow is detailed on section I-C.

### C. Methods

The integrated modeling system is based on: (i) a detailed multimodal transport model of Barcelona (VML) in VISUM [4] and (ii) the HERMESv3 emission model, which computes street-level and hourly road traffic emissions [5]. The simulation domain of the system comprises the Primary Crown of Barcelona, including Barcelona and other municipalities from the metropolitan area (95 km<sup>2</sup>). The dynamical information on traffic flow and travelling speed modelled at the road link level by VML is passed to HERMESv3, which combines the aforementioned information with the emission factors reported by COPERT V [6]. In the present work, a business as usual (BAS) and a traffic restriction scenario (TR) are simulated to illustrate the capabilities of the system. The traffic restriction scenario represents the superblocks policy which is gradually being applied all over Barcelona. This consists on the traffic calming of certain streets within an area comprised by several blocks. Mobility demand was supposed to be constant during the introduction of superblocks, which means that the same amount of vehicles is loaded into the network. Under the restrictions applied, new vehicle routes will be generated through the non-restricted streets.

### D. Results

Total NO<sub>x</sub> 2016 annual emissions for the BAS scenario are represented in figure 1a. The total annual NO<sub>x</sub> road transport emissions were compared against the latest available local emission inventory done by Barcelona Regional (BR) ([7]) corresponding to the year 2017. Estimated NO<sub>x</sub> emissions by the integrated model are a 9% higher than the estimated by BR. This difference is a consequence of the (i) different years simulated, (ii) the BR study uses a correction of COPERT IV, which differs from COPERT V emission factors, and (iii) HERMESv3 uses hourly speed profiles, while BR uses an average daily speed. Following that, the superblocks scenario was simulated for the 9 AM rush hour (figure 1b). The blue links represent a reduction on emissions while red ones indicate an increase. Although total NO<sub>x</sub> emissions at the marked superblock area are very similar between both scenarios, strong emission gradients appear between neighbouring streets. For example, Street 1 (marked on figure 1b) suffered a reduction of traffic flow of 92% which lead a reduction in NO<sub>x</sub> emissions

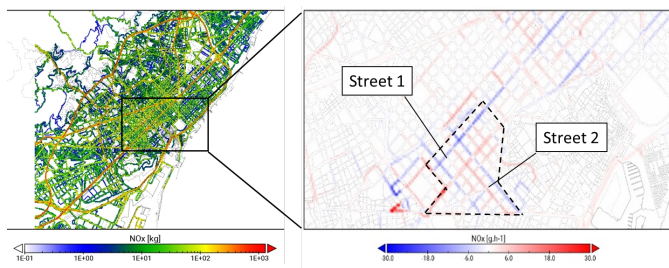


Fig. 1. a) Aggregated NOx emissions for the BAS scenario at 30 meters resolution; b) Difference in NOx emission between the BAS and the TR scenarios at 30 meters resolution, both at 9 AM local time. The dotted line represents the superblock area.



**Daniel Rodriguez-Rey** Holds a MSc in Air Pollution Management and Control by the University of Birmingham, and a BSc degree in Chemical Engineering by the UPC. Currently he is doing a Ph.D in Environmental Engineering with Marc Guevara (BSC-ES), Josep Casanovas (BSC-UPC) and M<sup>a</sup> Paz Linares (inLab) in the evaluation of the impact of mobility policies in Barcelona's air quality by the development of an integrated model.

of 93%. On the other hand, street 2, which crosses street 1, had an increase in NOx emissions of 95% as a consequence of traffic increase of 80%. Additionally, the dynamic system allows to observe emission variances in other areas of the city as a consequence of the simulated superblocks. On further steps of this work other TMS will be simulated together with air quality results.

## II. ACKNOWLEDGMENT

Daniel R. Rey acknowledges the Ministerio de Economía, Industria y Competitividad of Spain for the FPI research grant BES-2016-078116.

## REFERENCES

- [1] Agència de Salut Pública de Barcelona, "Informe de qualitat de l'aire de Barcelona," Tech. Rep., 2018. [Online]. Available: [https://www.aspb.cat/wp-content/uploads/2019/09/Informe\\_qualitat-aire-2018.pdf](https://www.aspb.cat/wp-content/uploads/2019/09/Informe_qualitat-aire-2018.pdf)
- [2] J. Sunyer, M. Esnaola, M. Alvarez-Pedrerol, J. Forn, I. Rivas, M. López-Vicente, E. Suades-González, M. Foraster, R. Garcia-Esteban, X. Basagaña, M. Viana, M. Cirach, T. Moreno, A. Alastuey, N. Sebastian-Galles, M. Nieuwenhuijsen, and X. Querol, "Association between Traffic-Related Air Pollution in Schools and Cognitive Development in Primary School Children: A Prospective Cohort Study," *PLoS Medicine*, vol. 12, no. 3, 2015.
- [3] R. Borge, B. Artñano, C. Yagüe, F. J. Gomez-Moreno, A. Saiz-Lopez, M. Sastre, A. Narros, D. García-Nieto, N. Benavent, G. Maqueda, M. Barreiro, J. M. de Andrés, and Á. Cristóbal, "Application of a short term air quality action plan in Madrid (Spain) under a high-pollution episode - Part I: Diagnostic and analysis from observations," *Science of the Total Environment*, 2018. [Online]. Available: <https://doi.org/10.1016/j.scitotenv.2018.03.149>
- [4] L. Montero, P. Linares, J. Salmerón, G. Recio, E. Lorente, and J. José Vázquez, "Barcelona Virtual Mobility Lab The multimodal transport simulation testbed for emerging mobility concepts evaluation," pp. 2–5, 2018.
- [5] M. Guevara, C. Tena, M. Porquet, O. Jorba, and C. Pérez García-Pando, "HERMESv3, a stand-alone multiscale atmospheric emission modelling framework - Part 1: global and regional module," *Geoscientific Model Development Discussions*, pp. 1–35, 2019.
- [6] EMEP/EEA, "Air pollutant emission inventory guidebook 2016 – Last Update June 2017," Tech. Rep. June, 2017.
- [7] Barcelona Regional, "Informe dels Resultats del Balanç d'Emissions i la Modelització de la Qualitat de l'Aire de la ZBE (Zona de Baixes Emissions) de Barcelona i Municipis Propers," Tech. Rep., 2019. [Online]. Available: <https://urbanaccessregulations.eu/>

# Real-Time Communication Support for Over-water Wireless Multi-hop Networks

Miguel Gutierrez Gaitan<sup>#1</sup>, Pedro M. Santos<sup>\*2</sup>, Luis Almeida<sup>#3</sup>

<sup>#</sup>CISTER/FEUP, Universidade do Porto, Porto, Portugal

<sup>1</sup>mgg@fe.up.pt, <sup>3</sup>lda@fe.up.pt

<sup>\*</sup>CISTER/ISEP, Instituto Politécnico do Porto, Porto, Portugal

<sup>2</sup>pss@isep.ipp.pt

**Keywords**— Real-Time Systems, QoS, Wireless Communication.

## EXTENDED ABSTRACT

**Motivation.** The prospect scenario for wireless communications and networking technologies in aquatic environments is nowadays promising. The growing interest around this subject in the last decades has recently been accelerated due to the more powerful capabilities of a number of sensing, control and communication devices. Moored, fixed, drifting, and vehicular nodes form now a rich ecosystem of autonomous embedded systems potentially connected in a multi-hop (and over-water) fashion, which demand innovative solutions to satisfy the ever-increasing requirements of reliability, bandwidth, latency and cost [1]. The efforts in this direction, mostly as a result of the push from the Internet-of-Thing (IoT) and related communication paradigms, are now at an early stage, and thus still pose significant, technical and research challenges, especially from the perspective of communication and networking for applications involving real-time and/or multimedia networking traffic.

**Research Challenges.** In effect, many issues arise when wireless communication occurs above water [2]. The different radio propagation behaviour, when compared to the over-land case, prevent a direct application of most general-purpose solutions. The radio signal reflections are much stronger, and this may lead to much more severe interference on the link quality. The natural water movements (such as tides and waves), among other distinctive phenomena, increase the prior difficulties and thus add extra considerations to the already challenging scenario. The mitigation of these factors by means of a careful design of link parameters (e.g., distance, height, polarization), as well as with diversity (and non-diversity) strategies is a well-studied topic in communications; but, mostly focused on the long-range. The case of short and medium-range distances with antennas at a few meters above surface still shows non-conclusive ideas, e.g., in terms of the most suitable propagation model to guide the network design labour. Moreover, the impact of tides (and waves) cycles in this more restricted settings show little or null effort in the literature. The relevant techniques that, e.g., have been proposed to mitigate the so-called tidal fading [3], although effective on the long-range, do not show clear applicability when water level fluctuations are in the order of the antenna height, and/or the link distances are relatively short. The general situation is further aggravated if extended to multi-hop networks, where the effect of water-level variations over-the-time can significantly degrade the connectivity of several links simultaneously; an issue that still deserves further studies.

**Related Work.** In the literature, different strategies at physical, link and network levels have been proposed to cope with problems of a similar kind. In maritime communications, e.g., dynamic/adaptive routing schemes helping to undertake the intermittent ship-to-ship connectivity issues due to varying sea-state levels have extensively been explored (e.g., [4]). At

the link level, protocol adaptations, e.g., to enable more reliable and/or predictable packet-level behaviour have also been considered. Likewise, physical layer techniques such as multiple antenna systems and beamforming have been explored too. These works, as well as several others in different (but related) fields, although do not show straightforward match with our primary research direction, can act as a fruitful source of inspiration, and thus represent an important starting point.

**Research Proposal.** In this research, we focus on the communication and networking aspects of over-water multi-hop networks aiming at support real-time and/or multimedia (audio/video) traffic using IEEE 802.11 (WiFi) commodity technologies. Special attention is devoted to the impact of cyclic water-level variations (such as tides and waves) on the overall network performance, and how an integrated approach to (i) network design, (ii) protocol adaptation and (iii) routing can contribute to mitigating such an issue.

Fig. 1 graphically summarizes our research proposal.




Fig. 1 An schematic illustration of the research approach taken.

## A. Network Design

This research line considers both, the geometrical basis of the two-ray propagation model [5] and the foreseeable nature of water movements as the key inputs to provide better strategies to link design and network planning. To this aim, our first step is to experimentally assess the over-water channel, particularly at near-shore/low-altitude (marine or freshwater) areas affected by noticeable over-time water-level variations. Along this line, our preliminary works in [6,7] show a considerable consistency between RSSI packet-based measurements (using WiFi COTS) and the theoretical predictions of the two-ray model. These works, strengthen our initial intuitions in [8,9] that suggest the two-ray as an effective method to estimate the large-scale fading in the presence of strong surface reflections. Thus, as convenient to guide the design of over-water network deployments. The next steps in this direction consider further experiments using both, a broader distance/height granularity and different antenna polarization. Once properly assessed the channel, we target to provide better network planning strategies based on discrete optimization techniques (closed-form and/or heuristic) that maximize/minimize a given QoS metric, while minimizing antenna height or another relevant parameter.

## B. Protocol Adaptation

This line of proposes the family of reconfigurable-and-adaptive TDMA overlay protocols, also known as RA-TDMA [10-12], as the primary source of related work to build upon. Here, we aim at extending them by mechanisms driven/aware of the water-level cycles, while oriented to improve the performance of real-time/multimedia applications. The RA-TDMA is a prior work in our group, which have shown promising results in fields such as teams of mobile robots [10], aerial multi-hop networks [11], and vehicular platoons [12]. We expect to leverage on this research and further explore the extension to mesh networks, a line of work still open. The performance evaluation of the propose mechanisms is expected to be done with OMNeT++/INET framework while incorporating the tides/waves features to the simulator.

## C. Routing

This research line conceives the development of novel methods for proactive/reactive routing, e.g., route-selection, forwarding, that adaptively respond to network topology dynamics. The end goal is to improve some overall QoS metric, regardless of the impact of water-level changes. We refer to [4] as a relevant work on this subject, providing some initial intuitions for the water-level dynamics mitigation in over-water multi-hop networks. Likewise, we refer to the work in [13], as with prior/promising insights for multi-hop routing on TDMA dynamic networks. The evaluation of these methods (as well as those in the previous direction), from the real-time network perspective, is expected to be done by means of analyses of the worst-case end-to-end delays and schedulability, using the approach described in [14,15].

## D. Conclusion and Future Work

This research envisions future reliable and real-time network scenarios in coastal areas, where ubiquitous and heterogeneous nodes communicate mostly above water. In this respect, this paper summarizes the main research lines, key insights and prior research efforts done so far.

## E. ACKNOWLEDGEMENTS

This work was partially supported by National Funds through FCT/MCTES (Portuguese Foundation for Science and Technology), within the CISTER Research Unit (UIDB/04234/2020) and within the AQUAMON project (PTDC/CCI-COM/30142/2017).

## References

- [1] A. Zolich, D. Palma, K. Kansanen, K. Fjortoft, J. Sousa, K. H. Johansson, Y. Jiang, H. Dong, and T. A. Johansen, "Survey on communication and networks for autonomous marine systems," *Journal of Intelligent & Robotic Systems*, vol. 95, no. 3-4, pp. 789–813, 2019
- [2] J. Wang, H. Zhou, Y. Li, Q. Sun, Y. Wu, S. Jin, T. QS Quek, and C. Xu. Wireless channel models for maritime communications. *IEEE Access*, 6: 68070–68088, 2018.
- [3] A. Macmillan, M. K. Marina, and J. T. Triana, "Slow frequency hopping for mitigating tidal fading on rural long distance over-water wireless links," in *IEEE Conf. on Computer Communications (INFOCOM) Workshop*, pp. 1– 5, IEEE, 2010.
- [4] C.W. Ang and S. Wen. "Signal strength sensitivity and its effects on routing in maritime wireless networks."

In 2008 33rd IEEE Conference on Local Computer Networks (LCN), pp. 192-199. IEEE, 2008.

- [5] T. Rappaport, *Wireless Communications: Principles and Practice*. USA: Prentice Hall PTR, 2nd ed., 2002.
- [6] M.G. Gaitan, P. M. Santos, L. Pinto and L. Almeida, "Wi-Fi-based network systems design over freshwater: Experimental evaluation using COTS devices," in 15th Doctoral Symposium in Informatics Eng. (DSIE), 2020.
- [7] M.G. Gaitan, P. M. Santos, L. Pinto and L. Almeida, "Experimental evaluation of the two-ray model for near-shore WiFi-based network systems design," in 91st IEEE Vehicular Technology Conf. (VTC), 2020 (to appear).
- [8] M.G. Gaitan, L. Pinto, P. M. Santos, and L. Almeida, "An Analysis of the Two-Ray Propagation Model to Support Near-Surface Overwater Wireless Sensor Networks Design," in 3rd Doctoral Congress in Eng. (DCE), 2019.
- [9] M.G. Gaitan, L. Pinto, P. M. Santos, and L. Almeida, "On the two-ray model analysis for overwater links with tidal variations," in 2019 11th National Symposium on Informatics (INFORUM), 2019.
- [10] F. Santos, L. Almeida, and L.S. Lopes. "Self-configuration of an Adaptive TDMA wireless communication protocol for teams of mobile robots." in *IEEE Int. Conference on Emerging Technologies and Factory Automation (ETFA)*. IEEE, 2008.
- [11] L. R. Pinto and L. Almeida. "A robust approach to TDMA synchronization in aerial networks." *Sensors* 18, no. 12 (2018): 4497.
- [12] A. Aslam, L. Almeida, and F. Santos. "A flexible tdma overlay protocol for vehicles platooning." *Int. Workshop on Communication Technologies for Vehicles*. Springer, Cham, 2018.
- [13] L. Oliveira, L. Almeida, and P. Lima. "Multi-hop routing within TDMA slots for teams of cooperating robots." 2015 *IEEE World Conference on Factory Communication Systems (WFCS)*. IEEE, 2015.
- [14] M.G. Gaitán and P.M. Yomsi. "FF-DBF-WIN: On the Forced-Forward Demand-Bound Function Analysis for Wireless Industrial Networks." *Proceedings of the Work-in-Progress Session (ECRTS 2018)* (2018): 13-15.
- [15] M.G. Gaitán, P.M. Yomsi, P.M. Santos, and L. Almeida. "Work-in-Progress: Assessing supply/demand-bound based schedulability tests for wireless sensor-actuator networks." *IEEE World Conf. on Factory Communication Systems (WFCS)*. IEEE, 2020 (to appear).

## Author biography



**Miguel Gutiérrez Gaitán** is a Ph.D. student in Electrical and Computer Engineering at the Faculty of Engineering of the University of Porto (FEUP), Porto, Portugal, and a research fellow at CISTER/FEUP (Research Centre in Embedded & Computing Systems at FEUP). He holds an Electronics Engineering degree from the Pontifical Catholic University Valparaíso (PUCV), Valparaíso, Chile, and an M.Sc. degree in Telecommunications Engineering from the Polytechnic University of Turin (POLITO), Turin, Italy. He is also an assistant professor at the Andrés Bello University (UNAB) in Santiago, Chile, where he serves as a lecturer since 2011. Among other positions, Miguel has served as an RF/MW engineer at Huawei Technologies in Santiago, Chile, and as a research trainee at Telecom Italia Labs in Turin, Italy.

# Game Theoretic Analysis of the Slurm Scheduler Model

Wilmer Uruchi Ticona\*<sup>†</sup>,

\*Barcelona Supercomputing Center, Barcelona, Spain

<sup>†</sup>Universitat Politècnica de Catalunya, Barcelona, Spain

E-mail: wilmer.uruchi@bsc.es

*Keywords*—*SLURM, Scheduling, High Performance Computing, Game Theory.*

## I. EXTENDED ABSTRACT

In the context of High Performance Computing, scheduling is a necessary tool to ensure that there exists acceptable quality of service for the many users of the processing power available. The scheduling process can vary from a simple *First Comes First Served* model to a wide variety of more complex implementations that tend to satisfy specific requirements from each group of users. Slurm is an open source, fault-tolerant, and highly scalable cluster management system for large and small Linux clusters [1]. MareNostrum 4, a High Performance Computer, implements it to manage the execution of jobs sent to it by a variety of users [2]. Previous work has been done from an algorithmic approach that attempts to directly reduce queuing times among other costs [3][4]. We consider that there is utility at looking at the problem also from a Game Theoretic perspective to define clearly the mechanics involved in the system, and also those that define the influx of tasks that the scheduler manages. We model the Slurm scheduling mechanism using Game Theoretic concepts, tools, and reasonable simplifications in an attempt to formally characterize and study it. We identify variables that play a significant role in the scheduling process and also experiment with changes in the model that could make users behave in a way that would improve overall quality of service. We recognize that the complexity of the models might derive in difficulty to theoretically analyze them, so we make use of usage data derived from real usage from BSC-CNS users to measure performance. The real usage data is extracted from Autosubmit [5], a workflow manager developed at the Earth Science Department at BSC-CNS. This is a convenient choice, given that we also attempt to measure the influence of an external agent (e.g. a workflow manager) could have in the overall quality of service if it imposes restrictions, and the nature of these restrictions.

### A. Slurm Overview

The Slurm scheduling mechanism has two main components: Priority and Scheduler. **Priority:** A value calculated based on data from the user, and the jobs that the user sends to the High Performance Computer (HPC). The calculation tries to give higher priority to those users that have less usage. **Scheduler:** Once the jobs have been received and their Priority calculated, there are certain rules in Slurm that determine when a job is sent for execution. As a result, the job with the highest

priority is not always executed first, but the order is altered so resource usage is optimized.

**Users Hierarchy:** The users are organized in a hierarchical tree structure, specifically in a Rooted Plane Tree [6], where on top of it we have a main **root** account. Then, the leafs are users and the internal nodes are the accounts to which they are associated. In the Slurm documentation we encounter many references to the term **account**, we consider it equivalent to the term **group**.

**Priority Calculation:** There are four factors involved in the calculation of this value: **Age**, a value from 0.0 to 1.0. The longer a job sits in the queue and is eligible to run, the bigger this value gets. **Size**, a value from 0.0 to 1.0 determined by the number of nodes a job requests, the more nodes a job requests the bigger this value gets. **Fair-share**, a value from 0.0 to 1.0 determined by [7]. **QoS**, a value from 0.0 to 1.0 determined by the priority given to different *QoS* defined in the Slurm implementation.

**Scheduling Mechanism:** The **Priority** value effectively produces an execution sequence; however, this ordering can result in sub optimal resource allocation. For example: *Consider a large (in the size of nodes required) job with high priority that is waiting to be scheduled, this job will take 25% of the nodes in the HPC and it has a planned (supplied by the user) running time of 10 hours; furthermore, it is in the front of the queue. After this job, we have a number of smaller jobs requiring a number of nodes from 1% to 2% of the total nodes in the HPC, and expected times lower than 1 hour.* Working under this standard configuration, the scheduler is going to wait for enough resources to be available and then schedule the large job for execution, and this is not optimal because resources will be idle. To avoid this scenario (to some extent), there exists the **backfill** mode. In this mode, the scheduler will start lower priority jobs if that does not delay the start of higher priority jobs.

### B. Data Overview

Slurm receives jobs, these jobs come from experiments on which the users are working on. A typical experiment can be modeled as a directed acyclic graph (DAG). It starts with jobs that retrieve or set information, or they might compile software that will be used in the later stages of the experiment. Then, there are some heavy computation tasks that usually consist on simulations that implement parallel processes and subsequently require many nodes and long running time. As we mentioned,

these experiments can be modeled as a **DAG**  $G = (V, E)$  where we have  $V = \{1, \dots, n\}$  tasks and  $V^i$  as the list of tasks sorted in topological order with sizes  $w_1, w_2, \dots, w_n$  measured in the number of **HPC** nodes they require, and  $t_1, t_2, \dots, t_n$  as the planned time in minutes they will need to complete. Typically, a number of vertex at the beginning and end of  $V^i$  will require less computation resources than the rest in average. We will avoid using the word "node" to name the vertex in a graph to avoid confusion with **HPC** nodes.

### C. Game Theoretic Perspective

We begin by assuming that the backfill scheduling mechanism subject of study is working as a **BF\_MOD** scheduler as defined in [3]. Under this assumption, resources are reserved when a job reaches the top of the queue and is about to be scheduled and the next jobs in the queue can be used for backfilling, but if in the next scheduling event a job with higher priority takes position at the top of the queue, the previous reservation is discarded and a new reservation is executed for the newly arrived top priority job. On the other hand we have **BF\_UNMOD** where the top job does not change even if a higher priority job arrives, this is also defined in [3]. We assume our scheduler under **BF\_MOD** configuration for further discussion.

There are  $M$  nodes in the HPC, there are  $U$  users that handle experiments represented as **DAGs**  $G_u = (V_u, E_u)$  that result in a topological order of jobs  $V_u^i = \{1, 2, \dots, n\}$  for each experiment  $G_u \in G$  where  $G$  is the set of all experiments from users  $U$ . Users will send job  $v_i \in V_u^i$  when  $\alpha_i = \{v_j | (v_j, v_i) \in E_u \wedge \text{status}(v_j) \neq 5\} = \emptyset$ , where  $\text{status}()$  is a function that returns 5 if the input job  $v$  has status **COMPLETED**, meaning that all preceding jobs of  $v_i$  must have been completed successfully. Then, we define the set of jobs sent to the scheduler by user  $u$  at a given iteration as  $\nu_u = \{v_i | v_i \in V_u^i \wedge \alpha_i = \emptyset\}$ .

$$N = \bigcup_{u \in U} \nu_u$$

Think of the scheduler as an agent  $\Lambda$  that receives  $N$  jobs, each job  $n \in N$  has attributes  $w_n$  for its weight or size,  $t_n$  as the planned time (supplied by the user) in minutes that the job will take to complete,  $a_n$  for the time it has been waiting for execution, and  $p_n$  for its **Priority** calculated using the previously mentioned attributes. Agent  $\Lambda$  uses the **Priority**  $p_n$  as the main ordering principle to define a list  $P_k : N \rightarrow \{1, \dots, n\}$  of execution order. We now proceed to give a glimpse of the analysis of games modeled using these definitions.

**Single Attempt Game:** We have  $N$  jobs ordered in a priority list  $P_k$  that defines the order in the queue of jobs to be scheduled for execution. A **strategy profile**  $s$  is an assignment of  $d$  jobs to  $|M|$  nodes by an agent  $\Lambda$ . In a single backfill scheduling attempt, agent  $\Lambda$  takes the first  $d$  jobs in the ordering  $P$  that minimize  $\varepsilon = |M| - \sum_{i=1}^d w_i$ . Then, the utility  $u$  of agent  $\Lambda$  is  $u(s) = |M| - \varepsilon$ . Then, we proceed to analyze the Nash Equilibrium of this setup as a foundation stone for the analysis of more complex games.

**Multiple Attempt Backfill Game:** Agent  $\Lambda$  actions under a sequence of backfill iterations. Clearly, it is under this scenario that the backfill mechanism comes into play.

**Submit to Scheduler Game:** Instead of analyzing a single agent  $\Lambda$ , we have players  $U$  submit jobs to an independent scheduler. We define a utility function for each player based on some of the variables that the user can modify when sending a job.

**Multiple Submit to Scheduler Game:** Players  $U$  submit jobs to an independent scheduler under a sequence of iterations.

### D. Experimental Environment

The workflow manager Autosubmit [5] saves the submit, start, and finish time of each job it manages. We will use this data to model synthetic data and test the behavior that, according to the Game Theoretic analysis, shows promising results. Slurm implements a node simulator that can be used for this purpose. We will also try to simulate different traffic situations in the arrival of jobs in order to test worst case scenarios, if possible.

### E. Conclusion

In this study, we try to use the Game Theoretic approach to analyze an existing system in order to detect opportunities for the implementation of configurations or tools that potentially result in an improvement of quality of service. However, this study might also reveal that there design choices or configurations that make any kind of collaboration or control not convenient.

## II. ACKNOWLEDGMENT

This work is being developed in close collaboration with the Computational Earth Science team at BSC-CNS.

## REFERENCES

- [1] "Slurm documentation." [Online]. Available: <https://slurm.schedmd.com/archive/slurm-17.11.7/overview.html>
- [2] "Marenostrum4 user's guide." [Online]. Available: <https://www.bsc.es/user-support/mn4.php>
- [3] Baraglia R., Capannini G., Pasquali M., Puppini D., Ricci L., Techouba A.D, "Backfilling strategies for scheduling streams of jobs on computational farms," *Making Grids Work*, 2008.
- [4] Sergei Leonenkov, Sergey Zhumatiy, "Introducing new backfill-based scheduler for slurm resource manager," *Procedia Computer Science*, vol. 66, pp. 661–669, 2015.
- [5] E. S. D. at BSC-CNS, "Autosubmit." [Online]. Available: <https://autosubmit.readthedocs.io/en/latest/>
- [6] S. G. W. Edward A. Bender, *Foundations of Combinatorics with Applications*, 2006.
- [7] "Fair tree fairshare algorithm." [Online]. Available: [https://slurm.schedmd.com/archive/slurm-17.11.7/fair\\_tree.html](https://slurm.schedmd.com/archive/slurm-17.11.7/fair_tree.html)



**Wilmer Uruchi Ticona** received his BSc degree in Systems Engineering from Private University of Tacna (UPT), Tacna in 2011. Then, he started working in a variety of large scale software development projects for his alma mater and other private companies. In 2018 he decided to pursue further education taking the Masters in Innovation and Research in Informatics (Advanced Computing specialization) program at Universitat Politècnica de Catalunya (UPC), Spain. Since 2019, he has been with the Computational Earth Science group of Barcelona Supercomputing Center (BSC), Spain.

# Memory-Coherence between host and devices in a runtime

Ruben Cano, Daniel Jiménez, Vicenç Veltran

\*Barcelona Supercomputing Center, Barcelona, Spain

†Universitat Politècnica de Catalunya, Barcelona, Spain

E-mail: [ruben.canodiaz, daniel.jimenez, vicenc.beltran]@bsc.es

Special thanks to Carlos Álvarez.

**Keywords**—Runtime, Nanos6, Main memory, Coherence, High-performance computing.

## I. EXTENDED ABSTRACT

As the end of the Moore’s law approaches, more specific devices such as GPUs, FPGAs or AI accelerators tend to steal the workload that was traditionally run on the CPU, allowing with this offload more specific solutions that improve the execution time of specific applications. One of the main problems that arise with this approach, is that now, the data is not centralized in one main memory, but distributed among the different accelerators which need a correct and coherent data to perform its operations. This can potentially limit the performance an accelerator can achieve, as well as delegates the programmer the task of enforcing the coherence between memories.

To relieve this model, in which the programmer has to take into account the memory of devices, models like NVIDIA Unified Memory[1] manage the hard work of maintaining the memory-coherence, potentially hurting performance but making the per-device memory management much easier.

In this work, the main objective is to develop an extension for a task-based runtime, which maintains the coherence between SMP and multiple devices in the system, using the dependency information of a task, acting as a Translation-Allocation Layer between the multiple memory spaces defined by the accelerators.

For our development, we use Nanos6 Runtime[2] as the target of our implementation. Nanos6 is a Runtime that implements the OmpSs-2 programming model, it is an SMP task-based runtime, with cluster support and is able to run CUDA tasks using unified memory. With this implementation, our objective is to extend Nanos6 capabilities to allow running CUDA with distributed memory, as well as other kind of accelerators such as FPGAs.

### A. Memory abstraction for devices

To manage the different memories, each memory should provide a interface that allows to: copy from and into it, allocate memory and free memory, in our system, we will manage the memories using this abstraction, which will allow to map any kind of esoteric memory into the same system that will be shared by all the devices.

### B. Task dependencies and symbols

A task dependency from a memory perspective, is a region of host memory that a task needs to coherently access, to write or read, in order to be executed, however, a symbol can be a superset of various dependency regions, and has the particularity that each symbol needs to be continuous in memory, this is due to the different already-compiled algorithms take into account the different offsets while accessing the data so we don’t only has to have the data valid on the device, but has to be in the way the already-existing software expects them to be. This means that the dependency can be multiple parts of a symbol, but the whole symbol has to be in memory.




Fig. 1. In this figure we can see how, for a Symbol, we have only two dependencies that are not contiguous, however, we have to maintain the offset address value between the two dependencies because the algorithm could expect the data to be in a certain offset.

### C. Non-blocking Software Memory Management Directory

The Directory is a range-based multipurpose software system that given multiples address spaces, manages the allocation, copies and translation from host memory to any other address space. So, for the runtime, it’s only a translation layer where you give an address and a range, and it returns a list of copy operations that are needed to satisfy the access, and an address to the data in the given address space. When a task finishes using that region, that information is passed to the directory, and the access is released. However, due to the possibility of an already-existing region on an address-space being part of a symbol that needs more contiguous allocation, we can ensure that when the directory returns an address, that address will remain valid and accessible until the task finishes using it, but consequent tasks that may access that region, may have different addresses. So, the coherence of a returned address is not enforced, and a returned address shouldn’t be used outside its task scope, and all accesses have to be checked by the Directory.

For explaining the functionality without going into too much detail, we will divide the Directory duties in three main steps:

1) *Allocation Step*: The directory checks if for a symbol region, there is an already-valid allocation on the address-space. If there is no allocations, it will try to allocate a paged-size chunk in the address-space. If there is a partial allocation, it will allocate a new paged-size sufficient to allocate all the symbol again, and generate the copies of the data to the new address, and, lastly, if there is no space for a new allocation, will return an error and the runtime will need to try again later when some data has been released.

2) *Data Copies Generation Step*: When an access has an already-valid allocation, we need to maintain the coherency between all the address-spaces. This means, for a range, we can have multiple entries on our directory, each one with its own state. We use a MESI-Like protocol in each entry, which allows us to have read-valid data in more than one device, and invalidate if necessary the data in any address space.

This means that if the data is invalid in the given address-space, we will check where the last version of the data is, and copy from there to the address space, updating the state or invalidating if necessary. However, this is not that simple, since we are not synchronously getting the data, but generating the copy that will be performed asynchronously. For this, we create a "promise" of validity, put the entry status in a transitory state, and only changing the status to valid when the data has finished copying. This has implications on following accesses to this region, that will, instead of creating a new copy information, create a directory lookup, to check if the data is already copied or not.

3) *Data free step*: When a taskwait arrives, the directory must ensure that all the data the tasks have modified, it's valid again on the host memory, invalidating and freeing allocations on the devices memories.

## II. CUDA COMPARISONS BETWEEN UNIFIED MEMORY - PREFETCH - DIRECTORY

To compare the performance achieved, we compared the performance in GFLOPS using a Blocked Matmul kernel with three memory models, our Directory model implementation, the Prefetch method and Unified memory model with Nanos6 Runtime. The algorithm has already a Tiled matrix, and will call our CUDA MATMUL BLOCK function, which is a Nanos6 task that performs the block multiplication on GPU, having as an in dependence  $[BS*BS]tileA[i][k]$  and  $[BS*BS]tileB[k][j]$  and as an out dependence  $[BS*BS]tileC[i][j]$ .

```
for (int i = 0; i < BS; i++)
  for (int j = 0; j < BS; j++)
    for (int k = 0; k < BS; k++)
      CU_MAT_BLOCK(tileC[i][j],
                  tileA[i][k],
                  tileB[k][j], BS);
}
```

Listing 1. Matmul code which shows how the different tasks are created for performing a tiled-matmul.

The Directory model, is the one that is mentioned in this document, the Prefetch model is using unified memory, but telling the CUDA Driver which regions of memory we are going to use (Basically, prefetching the symbols using the

CUDA API), and the Unified memory is calling CUDA without any modification nor notification, as is.

The experiment setup is using one node of Marenstrum 4 Power9 Cluster which contains 2 x IBM Power9 8335-GTH at max frequency of 3.8GHz, 512GB RAM and we used one of the four NVIDIA V100 (Volta).

The Nanos6 and mercurium versions are the master branch from 8 april 2020, and an experimental Nanos6 branch with support for prefetching CUDA memory.




Fig. 2. Graphic representation of the performance results, in BLUE, Directory approach, in RED, prefetcher approach and in YELLOW, unified memory approach, as we can see, the horizontal axis contains the block size, and the vertical axis the performance in GFLOPS, higher is better, where we can see the benefits of our method.

As we can see on the figure 2, for almost all the cases, our directory outperforms the unified memory with, or without prefetcher, with the prefetcher, unified memory seems to be more close to our approach, and with unified memory, without any kind of prefetching, is the clear loser in terms of performance

### A. Conclusion

There is a lot of work to be done in this subject, but seeing the preliminary performance results, we can see that it has a huge potential as a transparent way to the user to improve the performance of its applications and allowing multiple devices to communicate independently of its nature, in fact, this software solution is meant to work in heterogeneous systems, with any kind of device that can compute data.

Implementing this inside Nanos6 will allow to improve performance for already-existing applications built for the OmpSs-2[3] Programming model.

### REFERENCES

- [1] "NVIDIA training webpage." [Online]. Available: <https://devblogs.nvidia.com/unified-memory-cuda-beginners/>
- [2] "Nanos6 runtime github." [Online]. Available: <https://github.com/bsc-pm/nanos6>
- [3] "OmpSs-2 webpage." [Online]. Available: <https://pm.bsc.es/omps-2>



**Ruben Cano Díaz** received his BSc degree in Computer Engineering from Universitat Politècnica de Catalunya (UPC), Barcelona, Spain in 2018. He currently is coursing a Master in High Performance Computing at UPC, while working in programming models at the Barcelona Supercomputing Center, Spain.



# Design and implementation of an architecture-aware hardware runtime for heterogeneous systems

Juan Miguel de Haro Ruiz<sup>\*†</sup>, Jaume Bosch Pons<sup>\*†</sup>, Daniel Jiménez-González<sup>\*†</sup>

<sup>\*</sup>Barcelona Supercomputing Center, Barcelona, Spain

<sup>†</sup>Universitat Politècnica de Catalunya, Barcelona, Spain

E-mail: {juan.deharoruiz, jbosch, djimenez}@bsc.es

*Keywords—heterogeneous systems, task-dependence analysis, High-performance computing, FPGA, task-based programming models.*

## I. EXTENDED ABSTRACT

Parallel computing has become the norm to gain performance in multicore and heterogeneous systems. Many programming models allow to exploit this parallelism with easy to use tools. In this work we focus on task-based programming models. The parallelism is expressed with pieces of work called tasks that have data dependencies among them, and therefore have to be executed in a certain order. However, tasks that don't depend on any other running task can be executed in parallel.

### A. Picos

A software runtime that handles task dependencies and schedules them is able to exploit high levels of parallelism with moderate size tasks. Nevertheless, with fine-grained tasks it suffers from performance degradation due to the overhead and thread contention.

Picos [1] is a hardware runtime that aims to solve this problem. By moving the dependence management to a hardware designed and optimized for that, the overhead is reduced dramatically. It can be used in different scenarios, for instance, one of the prototypes was implemented on an FPGA with hardware accelerators attached to it. Picos would take tasks from a host, and distribute them among the accelerators, which are designed to do specific calculations. It was also implemented on a RISC-V multicore system embedded in an FPGA. In both cases Picos was integrated into the OmpSs programming model.

All the previous work with Picos required adapting the programming model and a lot of hand-crafting to integrate the accelerators with Picos and OmpSs. The next step involved doing the inverse, adapting Picos to the programming model OmpS@FPGA [2], an extension to OmpSs with support for automatic generation of hardware accelerators on FPGAs. This allows faster testing and easier deployment of the hardware.

However, Picos exposed some limitations regarding area and timing when being implemented along big accelerators or many of them. It uses a significant amount of resources, and its design has a long critical path. This influences negatively the routing algorithm when the FPGA has a high resource usage rate. The consequence is that designs without Picos

could scale in area much more than with it, thus achieving better performance even with higher runtime overheads.

### B. New design

To overcome these limitations we decided to build a new Picos based on two main principles: to be configurable and use minimum resources. The previous Picos used always a fixed amount of memory and had support for at most 15 dependencies per task. This is not ideal for some applications, which may need more memory, or on the other hand, they may work well with a small amount of memory and/or dependencies. Adding parameters to tune the resource usage of Picos gives the flexibility to adapt to every application. Also using less overall resources gives it a bigger margin for the user to add more computation resources.

For maintainability purposes, the new Picos (called Picos Daviu) does not include the scheduler, it relies on the scheduler already present in the OmpSs@FPGA hardware runtime, Smart OmpSs Manager (SOM). The latter includes similar features as the old version of Picos (called Picos++) without dependence management, which has to be performed on a CPU. However, it includes new features such as task creation on a hardware accelerator [3]. This complements with Picos Daviu, which is integrated in SOM, since it allows the hardware runtime to solve the dependencies of tasks created inside the FPGA. Therefore, the CPU does not have to take care of these tasks. Everything is managed internally, thus removing communication between FPGA and CPU. By reducing device communication, the new hardware runtime, called Picos OmpSs Manager (POM) improves dramatically the performance of systems with high latency interconnections such as cloud.

Figure 1 a) and b) show the internal design of POM and Picos Daviu compared with Figure 1 c) that shows the design of Picos++. Since POM provides more features, it contains more modules than Picos++, but in fact it consumes less resources when Picos Daviu is configured with the same parameters as Picos++. This configuration allows to store 256

TABLE I: Post-synthesis resource utilization of POM and Picos++ for Xilinx FPGAs & maximum frequency

Hardware runtime	LUT	FF	BRAM	LUTRAM	F (MHz)
POM	9492	10184	48.5	177	300
Picos++	17692	17365	142.5	518	200

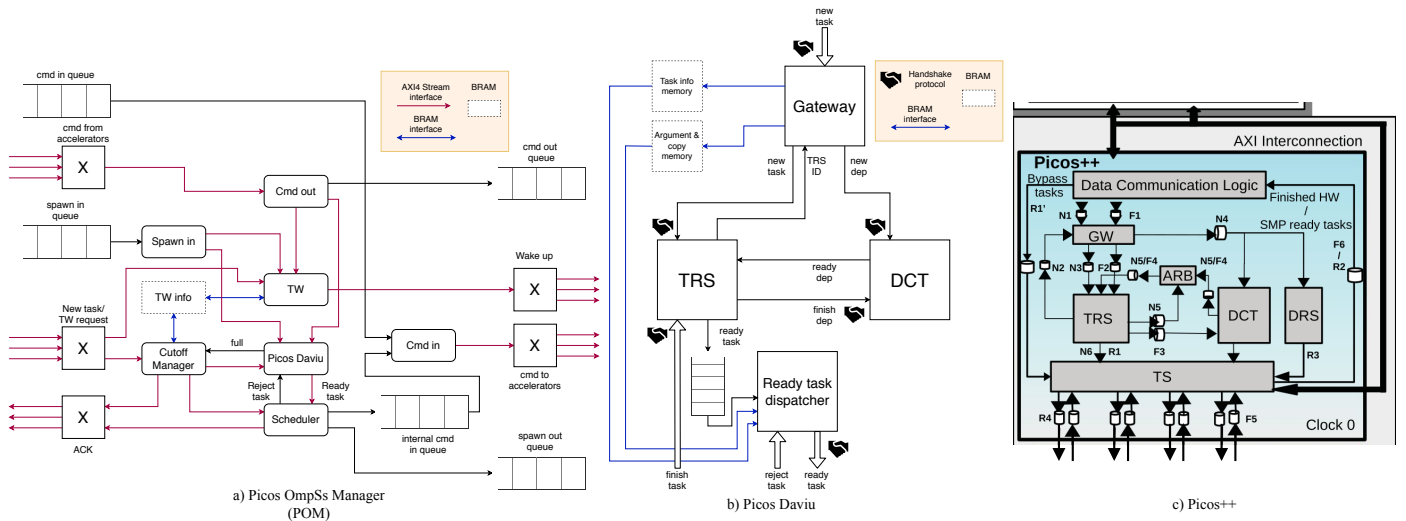


Fig. 1: POM, Picos Daviu and Picos++ internal designs

in-flight tasks, with a memory to store dependencies of 512 slots and another memory to store dependence relationship with the same number of slots. The maximum number of dependencies, arguments and copies (between main memory and the accelerator internal memory) per task is 15. The comparison of resources is shown in table I. The results have been obtained by hardware synthesis for Xilinx FPGAs. The usage of each resource is significantly decreased in POM, which shows a very high improvement in memory usage. It uses nearly 3 times less BRAMs and LUTRAMs than Picos++. The main difference is the removal of many queues used by Picos++ to communicate between its internal modules and the accelerators. In addition, the new design improves the maximum working frequency of Picos. While in the previous design a maximum of 200 MHz where achieved the new design has been successfully tested at 300 MHz.

### C. Results

Both Picos versions have been tested with an N-Body like algorithm (Picos++ and Picos Daviu inside POM). In this specific application, the performance of the algorithm is directly proportional to the number of parallel operators that can be deployed in the FPGA (as it is an embarrassingly parallel application). Our preliminary tests show that the old Picos code was able to support 3 accelerators with 32 operators each (so a total of 96 operators) or 4 accelerators with 14 operators each (56 operators). With more accelerators, even if the FPGA fabric was not fully utilized, the routing tool was not able to obtain a viable design due to communication constraints between the different elements.

In contrast, POM, using exactly the same accelerators code was able to support 3 accelerators with 48 operators (144 operators in total) or 4 accelerators with 36 operators (also 144 parallel operators). Correspondingly the execution time diminished from the old Picos to the new from 292s and 173s to 116s. This performance improvement is obtained thanks to the better suitability of the new design to be deployed in an FPGA fabric and its improved connectivity.

### D. Conclusion


In this work, we design a hardware runtime for task-based programming models which is configurable and uses low resources compared to its predecessor Picos. We compare both designs integrated in the OmpSs@FPGA programming model, and observe that the new design allows to attach more hardware accelerators, thus getting better performance.

## II. ACKNOWLEDGMENT

This work has been supervised by Carlos Álvarez Martínez (Barcelona Supercomputing Center, Universitat Politècnica de Catalunya, Barcelona, Spain, calvarez@ac.upc.edu). It received funding from Spanish Government (projects SEV-2015-0493 and TIN2015-65316-P), and from Generalitat de Catalunya (contracts 2017-SGR-1414 and 2017-SGR-1328).

## REFERENCES

- [1] X. Tan *et al.*, "A hardware runtime for task-based programming models," *IEEE Transactions on Parallel and Distributed Systems*, vol. 30, no. 9, pp. 1933–1946, 2019.
- [2] J. Bosch *et al.*, "Application acceleration on fpgas with ompss@fpga," in *International Conference on Field Programmable Technology*, ser. FPT, 2018, pp. 70–77.
- [3] —, "Breaking master-slave model between host and fpgas," in *PPoPP '20: 25th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, San Diego, California, USA, February 22-26, 2020*, R. Gupta and X. Shen, Eds. ACM, 2020, pp. 419–420. [Online]. Available: <https://doi.org/10.1145/3332466.3374545>



**Juan Miguel de Haro** received his BSc degree in Computer Engineering from Universitat Politècnica de Catalunya (UPC), Spain in 2018. He completed his bachelor thesis at the Department of Electronics and Information Systems in Ghent University, Belgium also in 2018. Later that year he started his MSc degree in High Performance Computing (HPC) at UPC, and started working with the OmpSs@FPGA group at the Barcelona Supercomputing Center (BSC).

# Accelerating Atmospheric Models using GPUs

Christian Guzman-Ruiz, Mario Acosta, Oriol Jorba

Barcelona Supercomputing Center, Barcelona, Spain

E-mail: {christian.guzman, mario.acosta, oriol.jorba}@bsc.es

**Keywords**—*Modeling and prediction, Chemistry, Parallelism and concurrency, Performance*

## I. EXTENDED ABSTRACT

Environmental models are simplified representations of an object or a process [1]. These models provide valuable information on the nature of real-world phenomena and systems [2], with many applications in science and engineering [3]. For example, environmental models play an increasingly important role in understanding the potential implications of climate change [4].

There are many types of models in the environmental sciences [5]. These models are often associated with large computational costs because of their complexity [6]. The model studied in this work, the *Multiscale Online Nonhydrostatic Atmosphere Chemistry* model (MONARCH), is an atmospheric model that currently runs in the MareNostrum supercomputer of the Barcelona Supercomputing Center (BSC), one of the Top-500 supercomputers in the world [7] [8]. MONARCH provides regional mineral dust forecasts to the World Meteorological Organization's (WMO) Barcelona Dust Forecast Center (BDFC) and the Sand and Dust Storm Warning Advisory and Assessment System (SDS-WAS). MONARCH also provides global aerosol forecasts to the International Cooperative for Aerosol Prediction (ICAP) initiative.

However, MONARCH does not currently make use of all the supercomputing resources available in the Marenostrum cluster. The MONARCH source code was originally prepared for clusters exclusively using CPUs. However, Marenostrum comprises two supercomputers with different system architectures: Marenostrum IV (MN4), a CPU-only cluster, and CTE-POWER, a GPU-accelerated cluster [9]. As MONARCH lacks GPU-specific code, as is often the case for environmental system models, it is currently only run on MN4.

For many years, the CPU-centric approach was accepted by the scientific community as individual GPUs have low computational abilities compared with CPUs, and the effective use of GPUs presents challenges to scientific software developers. However, with the advent of new, more powerful GPUs and the growth of Big Data [10], GPUs have begun assuming some of the computational responsibilities of CPUs on supercomputers for specific applications. Many studies have found improved efficiency applying GPU technology to traditional CPU-only software in supercomputers [11], and more-powerful GPU-accelerated supercomputers are appearing with increasing frequency [12].

In response to these recent developments, new approaches to GPU-based modeling have been developed. The molecular

modeling community was among the earliest adopters of GPU-based computing, providing in some cases a performance of two-orders-of-magnitude faster than CPU-only implementations [13]. Atmospheric models have primarily focused on the parallelization of dynamical process calculations [14], as does the Nonhydrostatic Icosahedral Atmospheric Model (NICAM), which has been parallelized for GPUs and reports a 7–8× performance speedup compared with the CPU version [15]. Chemical solvers, another important component of atmospheric models, have also been ported to GPUs, by either fully integrating GPU usage in the solver with a reported 20× speedup [16], or porting components of the solver to GPUs to achieve a 3–4× speedup [17]. Our work follows the latter approach, as we port only part of the chemical solving to GPUs. However, previous work [17] did not take into account the number of solver iterations in the performance assessment, and the grid-cell iterations are computed sequentially for both CPU and GPU versions. Here, we are able to reduce the number of grid-cell iterations for the GPU version to one for an idealized case. Moreover, our implementation is based on the innovative *Chemistry Across Multiple Phases* (CAMP) framework, which allows multiple chemical processes to be solved simultaneously as a single system [18].

In this work, we investigate the potential of a GPU-based chemical treatment by selectively adapting computationally costly portions of the MONARCH code to GPUs, to reduce the overall execution time. We achieve this objective by performing the following set of tasks:

- 1) Select the most time-consuming MONARCH component by analyzing the model from a computational point of view.
- 2) Study the efficiency of porting the selected component to GPUs, to evaluate the potential for improvement in the execution time using a GPU implementation.
- 3) Propose possible GPU implementations and general optimizations for both GPU and CPU implementations.
- 4) Apply these proposals and study their effect on execution time, in terms of the expected and actual improvement.
- 5) Evaluate the heterogeneous implementation, including transfers from CPU to GPU and GPU to CPU.

A computational profiling analysis showed that chemistry-related components of MONARCH are computationally expensive. His future chemistry solver (CAMP) will take more than 48.6% of the execution time, making it an interesting objective to optimize and our main focus. To optimize this module, this work focused around a GPU implementation in

order to exploit the available parallelization of the chemical reactions and the big amount of GPU resources available at Marenostrum at the moment.

As our main optimization, we adapted CAMP to allow us to solve chemical processes simultaneously for multiple grid cells in a given spatial region as a single system, as opposed to the common practice of looping through individual grid cells. This solution avoids repeating initialization steps, memory misses, solver iterations and allows further vectorization during solving. Our adaptation only requires a single solver call per chemistry time-step, allowing the solver to compute internally the chemistry for all grid cells and achieves a factor of 12 time-speedup for the configuration tested. Furthermore, we have started to port some of the most time-consuming functions to GPUs, splitting chemical reaction computation amongst individual threads. The selected function represents 30% of the chemistry computation time. Preliminary GPU results of such functions show a speedup of x12, ending in a total module speedup of 18x from the base version. With more possible optimizations pending, our work denotes the GPU potential and the importance of developing from a concurrent point of view.


## II. ACKNOWLEDGMENT

Special thanks to Matthew Dawson, Kim Serradell, and all the people involved from BSC.

## REFERENCES

- [1] S. Hartmann and R. Frigg, "Scientific models," Jan. 2005.
- [2] J. D. Sterman, "All models are wrong: reflections on becoming a systems scientist," *System Dynamics Review*, vol. 18, no. 4, pp. 501–531, 2002. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/sdr.261>
- [3] M. Poznic, "Models in Science and Engineering," Tech. Rep., Apr. 2017.
- [4] M. Carmen Lemos *et al.*, "Narrowing the Climate Information Usability Gap," *Nature Climate Change*, vol. 2, Oct. 2012.
- [5] E. Holzbecher, *Environmental Modeling*, May 2012.
- [6] N. D Bennett *et al.*, "Performance evaluation of environmental models," Jul. 2010.
- [7] "top500 Marenostrum." [Online]. Available: <https://www.top500.org/system/179067>
- [8] O. Jorba *et al.*, "Potential significance of photoexcited NO<sub>2</sub> on global air quality with the NMMB/BSC chemical transport model," *Journal of Geophysical Research D: Atmospheres*, vol. 117, p. D13301, Jul. 2012.

- [9] "Technical Information." [Online]. Available: <https://www.bsc.es/marenostrum/marenostrum/technical-information>
- [10] E. Mizell and R. Biery, "Introduction to GPUs for Data Analytics," p. 39.
- [11] Y. Jararweh *et al.*, "Gpu-based personal supercomputing," *2013 IEEE Jordan Conference on Applied Electrical Engineering and Computing Technologies (AEECT)*. [Online]. Available: [https://www.academia.edu/10680401/Gpu-based\\_personal\\_supercomputing](https://www.academia.edu/10680401/Gpu-based_personal_supercomputing)
- [12] "New GPU-Accelerated Supercomputers Change the Balance of Power on the TOP500 | TOP500 Supercomputer Sites." [Online]. Available: <https://www.top500.org/news/new-gpu-accelerated-supercomputers-change-the-balance-of-power-on-the-top500/>
- [13] "GPU-accelerated molecular modeling coming of age. - PubMed - NCBI." [Online]. Available: <https://www.ncbi.nlm.nih.gov/pubmed/20675161>
- [14] M. Govett *et al.*, "Parallelization and Performance of the NIM Weather Model on CPU, GPU, and MIC Processors," *Bull. Amer. Meteor. Soc.*, vol. 98, no. 10, pp. 2201–2213, Oct. 2017. [Online]. Available: <http://journals.ametsoc.org/doi/10.1175/BAMS-D-15-00278.1>
- [15] H. Sato *et al.*, "Extreme Big Data with TSUBAME2 and Beyond," p. 20.
- [16] M. Alvanos and T. Christoudias, "GPU-accelerated atmospheric chemical kinetics in the ECHAM/MESy (EMAC) Earth system model (version 2.52)," *Geoscientific Model Development*, vol. 10, no. 10, pp. 3679–3693, Oct. 2017. [Online]. Available: <https://www.geosci-model-dev.net/10/3679/2017/>
- [17] J. Sun *et al.*, "Computational Benefit of GPU Optimization for the Atmospheric Chemistry Modeling," *J. Adv. Model. Earth Syst.*, vol. 10, no. 8, pp. 1952–1969, Aug. 2018. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1029/2018MS001276>
- [18] M. D. Dawson, Matthew, "CAMP: A Scalable, Portable, Gas–Aerosol Chemistry Treatment for Atmospheric Models," AAAR, Portland, OR, USA, Oct. 2019.



**Christian Guzman** received his Bachelor's degree in Computer Engineering plus a bachelor's degree in Telecommunication Electronic Engineering, and a Master in Modelling for Science and Engineering by the Autonomous university of Barcelona (UAB). He specialized in techniques for High-Performance Computing (HPC) and is actually developing his pre-doctoral studies on the Barcelona Supercomputing Center (BSC), working on the development of Chemistry Across Multiple Phases (CAMP) module alongside the Multiscale Online Nonhydrostatic Atmosphere Chemistry Model (MONARCH), contributing to the most-computational and logic part from a performance point of view and integrating multitudinal ways of GPU computation in search of speeding-up the system.

# Non-linear MHD simulations of magnetically confined plasma using OpenFOAM

David Garrido González\*, Daniel Suárez Cambra†, Shimpei Futatani†

\*Department of Particle Physics, Universidade de Santiago de Compostela, Santiago de Compostela, Spain

†Department of Physics, Universitat Politècnica de Catalunya, Barcelona, Spain

E-mail:† shimpei.futatani@upc.edu

**Keywords**—MHD, Plasma, OpenFOAM, Magnetic confinement, mhdFoam.

## I. INTRODUCTION

Investigation of alternative energy resources for future is utmost important topic. One of the possibilities is nuclear fusion which powers the sun, being this very attractive since it is clean, safe and virtually unlimited energy. The achievement of controlled nuclear fusion will require a wide variety of fields such as plasma physics, materials physics, electrical engineering, heat transfer, etc. The computational fluid dynamics (CFD) contribute to understand the behaviour of confined plasma and in guiding experiments. This project aims to assess the feasibility of open source software, OpenFOAM to study the physics of magnetically confined plasma, and to expect realistic modelling of fusion plasma as a long-term project objective. The OpenFOAM magnetohydrodynamics (MHD) solver has been applied to solve plasma dynamics.

## II. OPENFOAM.

OpenFOAM [1] is a free and open source CFD software released and developed under the GPL license. Being free and open source implies that one can find a large experienced community and a very collaborative environment. This, together with an apparently user-friendly program, has been sufficient claim to verify the viability of this code as a potential high-performance solver for simulations close to fusion plasmas. OpenFOAM has been applied to the study of liquid metal dynamics of breeding blankets [2]. In this work, OpenFOAM is applied to magnetically confined plasma.

### A. The OpenFOAM MHD solver.

The simulations of plasma dynamics which is represented by MHD equations have been performed using mhdFoam, the OpenFOAM's MHD solver based on the PISO algorithm. This code provides numerical solutions for the single-fluid non-linear visco-resistive MHD equations for the pressure-constant density-constant (also magnetic diffusivity and kinematic viscosity are constant) incompressible approach [3]:

$$\frac{\partial \mathbf{v}}{\partial t} = \mathbf{j} \wedge \mathbf{b} + \nu \nabla^2 \mathbf{v} - (\mathbf{v} \cdot \nabla) \mathbf{v}, \quad (1)$$

$$\frac{\partial \mathbf{b}}{\partial t} = (\mathbf{b} \cdot \nabla) \mathbf{v} - (\mathbf{v} \cdot \nabla) \mathbf{b} + \lambda \nabla^2 \mathbf{b}, \quad (2)$$

$$\nabla \cdot \mathbf{b} = \nabla \cdot \mathbf{v} = 0; \quad (3)$$

where  $\mathbf{j} = \nabla \wedge \mathbf{b}$  is the charge current density. Here, all variables are normalized to Alfvén units. The dimensionless kinematic viscosity,  $\nu$ , and magnetic diffusivity,  $\lambda$ , are the

inverse of Reynolds and magnetic Reynolds number, respectively.


The mhdFoam solver has been extended to include an external magnetic field by presenting two part; a fixed part which is time-independent and a perturbation part which evolves by time,

$$\mathbf{b}(t, \mathbf{r}) = \mathbf{b}_{\text{Fixed}}(\mathbf{r}) + \mathbf{b}_{\text{Vble}}(t, \mathbf{r}). \quad (4)$$

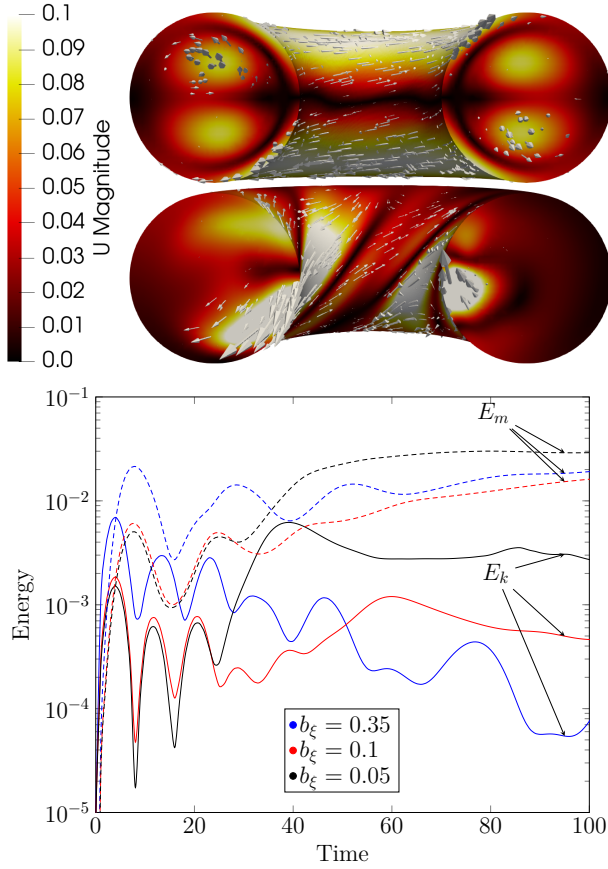
The equations on which resistiveMhdFoam (the modified mhdFoam) works are the result of replacing (4) in (1)-(3).

### B. Benchmark with Orszag-Tang test.

The Orszag-Tang test [4] has been performed to assess the robustness of the OpenFOAM code for magnetohydrodynamics, mhdFoam, to benchmark with [5]. This is a simple and classic test to check the validity of MHD codes based on simple, non-random and periodic initial conditions (that is, at  $t = 0$ ). The time evolution of the magnetic and the velocity fields are calculated in a cube (whose size is  $2\pi \times 2\pi \times 2\pi$ ) with periodic boundary conditions and without external sources. The MHD simulations have been studied for three cases;  $\nu = \lambda = 0.01$  with  $N^3 = 64^3$  cells (using 2 CPUs,  $\sim 2$  hours of wall-time),  $\nu = \lambda = 0.005$  with  $N^3 = 128^3$  cells (2 CPUs,  $\sim 2$  days of wall-time) and  $\nu = \lambda = 0.001$  with  $N^3 = 256^3$  cells (8 CPUs,  $\sim 1$  month of wall-time). First and second order discretization's were used for the temporal and spatial derivatives, respectively.



**Fig. 1:** Top panel, velocity and magnetic fields profiles for  $\nu = \lambda = 0.01$  in the cube domain at  $t = 2$ . Bottom panel, time evolution of the maximum charge current density (left) and energy dissipation (right) has been benchmarked with [5].



**Fig. 2:** Top panel, contour plot of velocity profiles for  $b_\theta = 0.35$ ,  $\nu = \lambda = 0.002$  and  $b_\xi = 0.35$  at  $t = 30$ ,  $b_\xi = 0.1$  at  $t = 70$ . Bottom panel, time evolution of kinetic (solid) and magnetic (dashed) energies  $b_\xi = 0.35$  (blue),  $b_\xi = 0.1$  (red),  $b_\xi = 0.05$  (black), keeping  $b_\theta = 0.35$  for all cases.

Figure 1 shows the contour plots of the magnetic end velocity fields profiles during the MHD decay process and the results for maximum charge current density,  $\max \|\mathbf{j}\| = \max \sqrt{j_x^2 + j_y^2 + j_z^2}$ , and the energy dissipation,  $\epsilon = \nu \langle \omega^2 \rangle + \lambda \langle \mathbf{j}^2 \rangle$ , where  $\omega = \nabla \wedge \mathbf{v}$  represents the vorticity and  $\langle \dots \rangle$ , the volumetric average. The magnetic and the velocity fields interact through non-linear MHD dynamics process, forming the filamentary structures correlating each other. The benchmark has been carried out with the time evolution of  $\max \|\mathbf{j}\|$ . The  $\max \|\mathbf{j}\|$  evolution in an early phase turns out to be independent from the  $\nu$  and  $\lambda$  chosen. The case with  $\nu = \lambda = 0.001$  shows a dependence of  $\sim t^3$  from  $t \approx 0.5$  to  $t \approx 1$ . A saturation of  $\max \|\mathbf{j}\|$  is observed beyond  $t \approx 1$ . It is inverse proportion against the transport parameters  $\nu$  and  $\lambda$ . The  $\max \|\mathbf{j}\|$  plot is consistent with the observations of [5]. The plot of the time evolution of  $\epsilon$  which corresponds to small vortex structures of the field is different from [5]. The possible reason could be due to the numerical methods used in the work.

### III. MHD SIMULATIONS IN TOROIDAL GEOMETRY.

The MHD simulations have been applied to a toroidal geometry ( $\sim 700k$  cells, 3 to 8 CPUs,  $\sim 1-3$  days of wall-time). The geometry consists of a toroid of circular cross section of a major radius  $R_0 = 0.55\pi$  and a minor radius  $r_0 = 0.3\pi$ . Two magnetic fields, poloidal and toroidal, are

imposed by an external source assuming the external coils. Since the interspacing between the coils is smaller inside than outside, the toroidal field decreases with the inverse of the distance as one moves away from the axis of symmetry,  $\mathbf{b}_\xi = b_\xi R_0 / R \hat{\xi}$ . The poloidal field,  $\mathbf{b}_\theta = b_\theta r / r_0 \hat{\theta}$ , is produced by the circulation of a toroidal current. Here,  $b_\theta = b_\theta|_{(r=r_0)}$  is the value of the poloidal field at the boundary and  $b_\xi = b_\xi|_{(R=R_0)}$  is the value of the toroidal magnetic field at centre of the cross section. The toroidal magnetic field has been scanned,  $b_\xi = 0.35, 0.1$  and  $0.05$ , keeping the poloidal magnetic field fixed to  $b_\theta = 0.35$ . The simulations were performed with  $\nu = \lambda = 0.002$ . An initial velocity field is given by random perturbation of  $E_k \approx 1.24 \times 10^{-5}$ . Figure 2 shows the contour plots of velocity field of  $b_\xi = 0.35$  at  $t = 30$  and  $b_\xi = 0.1$  at  $t = 70$  keeping  $b_\theta = 0.35$  for both cases. They show four counter rotating vortices in the case of  $b_\xi = 0.35$  and a helical structure formed in the case of  $b_\xi = 0.1$ . The ratio of  $b_\theta$  and  $b_\xi$  characterizes the MHD stability. The energy evolutions show oscillations before  $t = 100$ . In the case of low  $b_\xi$ , the MHD dynamo is induced forming the helical structures. These results are consistent with the observations of the work which has been carried out by other code [6].

### IV. CONCLUSIONS AND PERSPECTIVES

In this work, the viability of OpenFOAM as a potential code for high performance simulations oriented to fusion plasmas has been assessed. The results of Orszag-Tang test are comparable to the studies of [5]. In the toroidal domain, the energy plots turn out to be in good agreement with [6]. The non-linear MHD simulations via OpenFOAM gives plausible results although there is a room to improve, for example, the MHD modelling and physics parameters.

The generation of simulation mesh has been applied to other geometries. The capability of generating arbitrary mesh is potential interest for numerical studies of fusion research. The study of those geometries will be carried out in the future.

### REFERENCES

- [1] openfoam.org
- [2] E. Mas de les Valls Ortiz *et al.*, Fusion Eng. Des. **86** (2011) 2326.
- [3] S. Cappello *et al.* Phys. Rev. Lett. **85** (2000) 3838-41.
- [4] S. A. Orszag and C. M. Tang, J. Fluid Mechanics **90** (1979) 129-143
- [5] P. Mininni *et al.* Phys. Rev. Lett. **97** (2007) 244503; J. Morales *et al.* J. Comput. Phys. **274** (2014) 64-94.
- [6] S. Futatani *et al.* Phys. Plasmas **22** (2015) 052503.



**David Garrido González** is a Physics bachelor student at Universidade de Santiago de Compostela (USC). He is at his last year and is currently carrying out his bachelor thesis about non-linear MHD simulations of magnetically confined plasma using OpenFOAM in close collaboration with Universitat Politècnica de Catalunya (UPC), Spain.





**Barcelona  
Supercomputing  
Center**  
*Centro Nacional de Supercomputación*



**EXCELENCIA  
SEVERO  
OCHOA**

## **Barcelona Supercomputing Center**


Jordi Girona, 31 - Torre Girona  
08034 Barcelona (Spain)

education@bsc.es  
www.bsc.es

@BSC\_CNS 

/BSCCNS 

/BSC\_CNS 

/barcelona-supercomputing-center 

/BSCCNS 