

4th BSC Severo Ochoa Doctoral Symposium

2nd, 3rd and 4th, May 2017

Book of Abstracts

Book of Abstracts

4th BSC Severo Ochoa Doctoral Symposium

Editors

Nia Alexandrov

María José García Miraz

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

www.bsc.es/doctoral-symposium-2017

published by

Barcelona Supercomputing Center

supported by

The “Severo Ochoa Centres of Excellence” programme

4th Edition, May 2017

ACKNOWLEDGEMENTS

The BSC Education & Training team gratefully acknowledges all the PhD candidates, Postdoc researchers, experts and especially the Keynote Speaker Alfonso Valencia and the tutorial lecturer Darío Garcia, for contributing to this Book of Abstracts and participating in the 4th BSC Severo Ochoa Doctoral Symposium 2017. We also wish to thank expressly the volunteers that supported the organisation of the event: Silvia M. Giménez Santamarina, Bruno Cuevas and Marc Eixarch.

BSC Education & Training team
education@bsc.es

EDITORIAL COMMENT

We are proud to present the Book of Abstracts for the 4th BSC Severo Ochoa Doctoral Symposium.

During more than ten years, the Barcelona Supercomputing Center has been receiving undergraduate, master and PhD students, and supporting them to develop research and professional skills for a successful career through an extensive education and training program. Many of those students are now researchers and experts at BSC and in other internationally renowned research institutions.

With the years, the number of students and the span of their research areas have increased and the Center recognised that these highly qualified students and young researchers needed a forum to present their findings and fruitfully exchange ideas. As a result, in 2014, the first BSC Doctoral Symposium was born.

Last year, a total of 43 presentations were given, including 21 posters, and we had more than 90 attendees. The tutorials, given by researchers from Computer Science Department, reviewed the fundamental Algorithms and Techniques for Data and Computationally Intensive Problems (Prof. Vassil Alexandrov, ICREA-BSC) and introduced Scientific Visualisation of Data (Dr. Javier Espinosa) as well as provided some practice regarding these concepts. The keynote speaker, Prof. Francisco Doblas Reyes, gave the lecture: Big Data for the Study of Climate Change and Air Quality.

In this fourth edition of the BSC Severo Ochoa Doctoral Symposium, we have planned a keynote by Prof. Alfonso Valencia, the Director of Life Sciences Department in BSC, whose research is in the area of Bioinformatics. The tutorial will be on Deep Learning by Dr Darío Garcia Gasulla.

The talks will be held in five different sessions and will tackle the topics of: HPC & Novel Computer Architectures, Mathematics, Algorithms and Computational Science, Simulations and Modelling, Programming Models, Performance Analysis and Software Tools. The posters will be exhibited and presented during four poster sessions that will give the authors the opportunity to explain their research and results.

WELCOME ADDRESS

I am delighted to welcome all the PhD students, Postdoc researchers, advisors, experts and attendees participating in the 4th BSC International Doctoral Symposium.

The goal of the event, year after year continues to provide a networking opportunities to the participating junior researchers and allowing them to share their research results. The Symposium provides an interactive forum for PhD students, post Docs and MSc students and an opportunity to acquaint themselves with the research projects developed by their colleagues across areas as diverse as Novel Computer Architectures, HPC Simulations in Sciences and Engineering, applications of Computational Science and use of Supercomputing technology in Life Sciences, Environmental studies and Medicine.

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. Consequently, 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 students' participation in the Symposium. Moreover, I would like to thank specially the keynote speaker Alfonso Valencia and the invited lecturer of this year's tutorial Darío Garcia Gasulla, for their willingness to share with us their knowledge and expertise.

Finally, I would like to thank all PhD students and Postdoc researchers for their presentations 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

KEYNOTE SPEAKER

Alfonso Valencia

Head of Life Sciences Department, BSC

Personalised Medicine as a Computational Challenge

Personalized Medicine represents the adoption of Genomics and other –omics technologies to the diagnosis and treatment of diseases. PerMed is one of the more interesting and promising developments resulting from the genomic revolution.

The treatment and analysis of genomic information is tremendously challenging for a number of reasons that include the diversity and heterogeneity of the data, strong dependence of the associated metadata (i.e, experimental details), very fast evolution of the methods, and the size and confidentiality of the data sets. Furthermore, the lack of a sufficiently developed conceptual framework in Molecular Biology makes the interpretation of the data extremely difficult.

In parallel, the work in human diseases will require the combination of the –omics information with the description of diseases, symptoms, drugs and treatments: The use of this information, recorded in Electronic Medical Records and other associated documents, requires the development and application of Text Mining and Natural Language processing technologies.

These problems have to be seen as opportunities, particularly for students and Post Docs in Bioinformatics, Engineering and Computer Sciences. During this talk, I will introduce some of the areas in which the Life Sciences Department works, from Simulation to Genome Analysis, including text mining, pointing to what I see as the most promising future areas of development.

Alfonso Valencia is the founder and President of the International Society for Computational Biology and Co-Executive Director of the main journal in the field (Bioinformatics of Oxford University Press). During the last ten years, Valencia has been the Director of the National Institute of Bioinformatics (Salud Carlos III Institute platform (INB-ISCI)) and node of ELIXIR the European Infrastructure of Bioinformatics), worked in the National Oncology Research Center, CNIO, where he was also the Vice-Director of Basic Research and Director of the Structural Biology and BioComputing Program.

Alfonso Valencia's research is in the area of Bioinformatics and Computational Biology. Particularly, the application of Computational Methods for Genome Analysis to Precision Medicine. Alfonso Valencia is a Member of the Scientific Advisory Committee of the Swiss Institute of Bioinformatics, EBI chemical and protein domain databases, IRB, UPF_DCEX, Greek ELIXIR-Node, amongst others; Associate Editor of eLIFE, PeerJ, FEBS Letters, and co-leader of the new journal f1000 "Bioinformatics, Biomedical Informatics & Computational Biology". In addition to being a member of the European Molecular Biology Organisation (EMBO) and Founder of the BioCreative challenge in text mining.

In the opinion of the new head of the Life Sciences Department, "Facing the enormous biomedical challenges of the future will only be possible with the coordination of the incredible scientific and technical resources of the BSC, in the rich scientific environment of Barcelona, and in combination with both National (INB-ISCI) and European scientific infrastructures (ELIXIR)"

TUTORIAL

Darío Garcia Gasulla

Computer Science Department, BSC

Deep Learning

In this tutorial, we will review the basic concepts of Convolutional Neural Networks (CNN), we will define a few neural network models (including CNN), train them and evaluate their performance for image classification and some other related tasks. CNNs have become the state-of-the-art methods for any image-processing task. Their performance on image classification problems (reaching levels beyond human capacities) has motivated their application to many other challenges (image segmentation, image clustering, volumetric space processing, style transfer, synthetic image generation, etc.).

The agenda:

Day 1:

Theory 1.1: Introduction to artificial neurons, activation functions, fully connected networks, back propagation algorithm and other basic components.

Hands on 1.1: Train a network with 2 fully connected layers for solving MNIST

COFFE BREAK

Theory 1.2: Limitations of fully connected layers. Computational cost.

Hands on 1.2: Train a network with 2 fully connected layers for solving CIFAR100

Day 2:

Theory 2.1: Introduction to convolutional layers, pooling and typical architectures

Hands on 2.1: Train a network with 1 convolutional, 1 fully connected layer for solving CIFAR100

COFFE BREAK

Theory 2.2: Style transfer

Hands on 2.2: Use a pre-trained network to perform image style transfer

Dr. Darío Garcia-Gasulla is a Post Doc at BSC since 2015. Before that, he worked as an assistant researcher for four years at the KEMLg group (UPC), participating in several research projects related to Artificial Intelligence (particularly in Knowledge Representation and Reasoning, Machine Learning and Data Mining). His PhD thesis tackled the large-scale graph-mining problem, bringing together AI and HPC topics. As a Post Doc, he is now leading a research project for the integration of Deep Learning and graph mining technologies.

4th BSC Severo Ochoa Doctoral Symposium AGENDA

Day 1 (2nd of May)

Start	Activity	Speaker/s	Chair
8.30 h	Registration		
9.00h	Welcome and opening	Mateo Valero, BSC Director	Maria-Ribera Sancho
9.20h	Event Photo		
	<i>Coffee break</i> & First Poster Session 9:40 until 10:30		Osman Unsal
1	Supporting Real-Time Visual Analytics in Neuroscience	Enrique Javier Arriaga Varela	
2	Advanced Vector Architectures for Future Applications	Adrián Barredo Ferreira	
3	Leveraging FPGAs to Accelerate the Query Processing of SQL-Based DataBases	Behzad Salami	
4	Time-Predictable Parallel Programming Models	Maria A. Serrano	
5	Protein modelling for enzyme engineering	Rubén Cañadas	
	First Talk Session: HPC & Novel Computer Architectures 10:30 until 12:50		Petar Radojkovic
1	Performance Impact of a Slower Main Memory: A case study of STT-MRAM in HPC	Kazi Asifuzzaman	
2	Aggregating and Managing Memory Capacity Across Computing Nodes in Cloud Environments	Luis Angel Garrido	
3	Energy Optimizing Methodologies On Heterogeneous Data Centers	Rajiv Nishtala	
4	Machine Learning Performance Prediction Model for Heterogeneous Systems	Daniel Nemirovsky	
5	Performance Analysis on the Intel Knights Landing Architecture	Michael Wagner	
12.50h	<i>Lunch break</i>		
14.00h	Tutorial Deep Learning: The tutorial will review the basic concepts, will define a few neural network models (including CNN), train them and evaluate their performance for image classification and some other related tasks.		Dario Garcia Gasulla
16.00h	<i>Coffee break</i>		
	Tutorial part 1 continues		Dario Garcia Gasulla
18.00h	Adjourn		

Day 2 (3rd May)			
Start	Activity	Speaker/s	Chair
9.00h	Opening of the second day		
Second Talk Session: Mathematics, Algorithms & Computational Sci. 9:10 until 10:30			Vassil Alexandrov
6	Discovering Ship Navigation Patterns towards Environmental Impact Modeling	Alberto Gutiérrez Torre	
7	Identification and Characterization of Recurrent Deletions in the Human Genome Promoted by Expression of Transposase-Derived Gene Across Different Tumour Types	Elias Rodriguez-Fos	
8	Functional implications of the structural genomic rearrangements in cancer	Luisa Fernanda Delgado Serrano	
9	Exploring the use of mixed precision in NEMO	Oriol Tintó-Prims	
10	ORCHESTRA: An Asynchronous Non-Blocking Distributed GVT Algorithm	Tommaso Tocci	
Coffee break & Second Poster Session 10:30 till 11:20			Javier Espinosa
1	Prediction of binding energies upon mutation in 3D-structure-known complexes through PyDock scoring functions	Bruno Cuevas Zuviria	
2	Impact of Aerosol Microphysical Properties on Mass Scattering Cross Sections	Vincenzo Obiso	
3	On the suitability of Time-Randomized Processors for Secure and Reliable High-Performance Computing	David Trilla	
4	A case for code-representative microbenchmarks	Calvin Bulla	
Third Talk Session: Programming Models, Performance analysis & Software Tools 11:20 until 12:40			Xavier Martorell
11	Programming models for mobile environments	Francesc-Josep Lordan Gomis	
12	GUIDANCE: An Integrated Framework for Large-scale Genome and PhenomeWide Association Studies on Parallel Computing Platforms	Marta Guindo-Martínez	
13	Testing simple models for street wind conditions in Barcelona	Jaime Benavides	
14	Web-based tool for the annotation of pathological variants on proteins: PMut 2017 update	Víctor López Ferrando	
15	Python for HPC geophysical electromagnetic applications: experiences and perspectives	Octavio Castillo-Reyes	
12.40h Lunch break			
14.00h	Tutorial Deep Learning - part 2		Dario Garcia Gasulla
16.00h Coffee break			
	Tutorial part 2 continues		Dario Garcia Gasulla
18.00h	End of the Tutorials and Adjourn		

Day 3 (4th May)			
Start	Activity	Speaker/s	Chair
9.00h	Opening of the third day		
Fourth Talk Session: Simulations & Modeling 9:10 until 10:30			Eduard Ayguade
16	Data-Driven Crowd Simulation	Hugo Perez	
17	A new Reliability-Based Data-Driven approach to Simulation-Based Models	Jacobo Ayensa Jiménez	
18	Level of Detail for Complex Urban Scenes with Varied Animated Crowds, using XML	Leonel Antonio Toledo	
19	Effect of Terrain Relief on Dust Transport over Complex Terrains in West Asia	Lluís Vendrell	
20	pyDockDNA: A new approach for protein-DNA docking	Luis Ángel Rodríguez Lumbreras	
Coffee break & Third Poster Session 10:30 until 11:20			Sara Basart
1	Deep and Cognitive Learning applied to Precision Medicine: the initial experiments linking (epi)genome to phenotypes-disease characteristics.	Davide Cirillo	
2	Stabilization of Microturbulence by Fast Ions in ASDEX Upgrade	Felipe Nathan de Oliveira	
3	Main Memory in HPC: Do We Need More or Could We Live with Less?	Darko Zivanovic	
4	Excited state gradients within a polarizable QMMM formulation	Maximilian Menger	
5	Vortex Induced Vibration (VIV) of circular cylinders at high Reynolds numbers	Daniel Pastrana	
11:20h	Keynote Alfonso Valencia: Personalised Medicine as a Computational Challenge, The treatment and analysis of genomic information is tremendously challenging for a number of reasons that include the diversity and heterogeneity of the data, strong dependence of the associated metadata (i.e experimental details), very fast evolution of the methods, and the size and confidentiality of the data sets. These problems have to be seen as opportunities, particularly for students and Post Docs in Bioinformatics, Engineering and Computer Sciences. During this talk, I will point to what I see as the most promising future areas of development.		Maria-Ribera Sancho
13.00h Lunch break			
Fifth Talk Session: Mathematics, Algorithms & Computational Sci. 14:20 until 16:00			Mariano Vazquez
21	Protein modelling for enzyme engineering	Gerard Santiago	
22	Classification of retrotransposition during somatic variant calling in cancer genomes	Jordi Valls Margarit	
23	Relevant aspects of the seismic hazard in Colima, Mexico	Armando Aguilar-Meléndez	
24	SMuFin2: generating and implementing new and more efficient search engines integrated in particular hardware architectures	Mercè Planas-Fèlix	
25	Uncertainty in near-surface wind speed trends at seasonal time scales	Verónica Torralba	
26	Production Planning and Scheduling Optimization Model: A case of study for a Glass Container Company	Laura Hervert-Escobar	
Coffee break & Fourth Poster Session 16:00 until 16:50			Victor Guallar
1	Exploring the Relationship between Gene Expression and Topological Properties of Arabidopsis thaliana Interactome Network.	Silvia M. Gimenez	
2	Simulations of Alfvénic Modes in TJ-II Stellarator	Allah Rakha	
3	Composite materials calculation using HPC-based multiscale technique.	Guido Giuntoli	
4	Application of Reduced order Modelling in Geophysics	Prattya Datta	
5	Saiph, a Domain Specific Language for Computational Fluid Dynamics simulations	Sandra Macià Sorrosal	
16:50h	Conclusions		Maria-Ribera Sancho



Talk Abstracts

Relevant aspects of the seismic hazard in Colima, Mexico

Armando Aguilar-Meléndez^{#1,*2}, Josep De la Puente^{#1}, Héctor E. Rodríguez-Lozoya^{**3}, Alejandro Córdova-Ceballos^{*4}, Alejandro García-Elías^{*5}, Sergio N. González-Rocha^{*6}, Amelia Campos-Ríos^{***7}

[#]CASE Department, Barcelona Supercomputing Center, Jordi Girona 29, Barcelona, Spain

¹armaguilar@uv.mx, ²josep.delapuente@bsc.es,

^{*}Faculty of Civil Engineering-Faculty of Chemical Sciences, University of Veracruz (UV),

Av. V. Carranza S/N, Poza Rica, Veracruz, México

^{#4}acordova@uv.mx, ^{#5}alejagarcia@uv.mx, ^{#6}n.gonzalez@uv.mx

^{**}Faculty of Engineering, Autonomous University of Sinaloa, Mexico

³rolohel@yahoo.com.mx

^{***7}Services in Engineering, private consulting, Azueta 10, Tuxpan, Ver., Mexico

⁷amelia.campos.r@gmail.com

Keywords— Seismic hazard, Colima, local sites, CRISIS2015

ABSTRACT

Perform probabilistic assessment of seismic hazard (PSHA) is a fundamental activity that offer valuable information to the seismic risk management of cities. In the present work some relevant aspects about a recent analysis about the seismic hazard of Colima, Mexico are mentioned.

A. Introduction

The state of Colima is located aside of the Pacific Ocean and about of one third of the total perimeter of this state of Mexico is delimited by the sea. In this region the interaction of tectonic plates that are in a subduction process below the Continental Plate, produces significant earthquakes. In 2003 an important earthquake generates significant damage in the region [1]. For this reason, in the present document some relevant aspects of the preliminary results of a PSHA of Colima in development are included.

B. Seismicity in Colima

The seismicity related to the subduction process that generate the major part of the earthquakes that affect to the Colima state can be considered between moderate and high. As a reference of the kind of earthquakes that affect to the Colima state, it is possible to mention some recent earthquakes: a) Colima earthquake of 31 October 2007; b) Tecomán, Mexico earthquake of 22 of January 2003 with magnitude $M_w=7.4$ (2003) [1]. The earthquake of October had a magnitude equal to 5.1, but it trigger a Peak Ground Acceleration (PGA) of 1368 gals in a station located to 37 km from the epicenter (Figure 1). The Tecomán earthquake generated important damage in the towns of Tecomán and Armería, and it left 21 persons dead.

Figure 2 shows some of the different results that can be obtained after a seismic record of the 2007 earthquake was processed: a) the accelerogram of the original seismic record, b) the Fourier Spectrum of the original seismic record, c) the seismic record corrected of accelerations, d) the Fourier Spectra of the seismic record corrected, e) the diagram of velocities versus time and the diagram of displacement versus time. At the same time it is possible to observe in Figure 3 the different response spectra related to the same seismic record corrected.

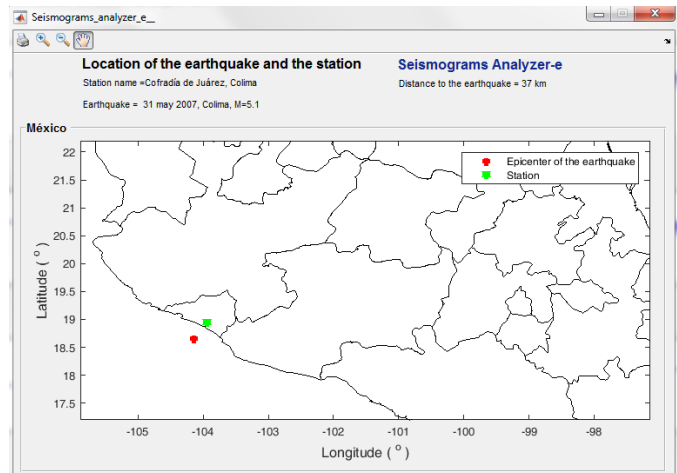


Figure 1. Location of the station Cofradía de Juárez, Colima, and the epicenter of the earthquake of magnitude 5.1 that occurred on 31 may 2007 [2].

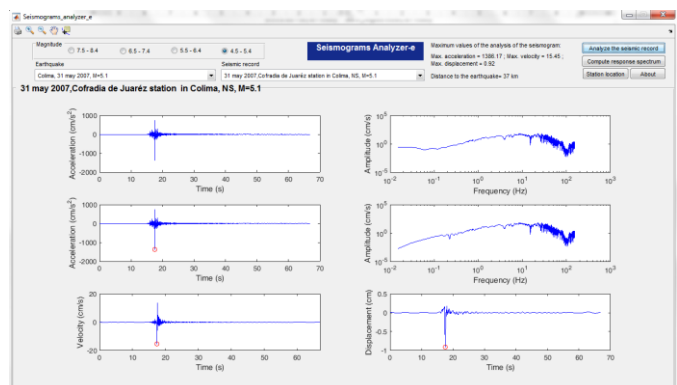


Figure 2. Graph of the original seismic record and graphs related to the corrected seismic record [2].

In the present work a PSHA has been performed in order to determine the seismic hazard of Colima. A particular contribution of the present research is the fact that the effect of local sites has been considered in the assessment of the seismic hazard.

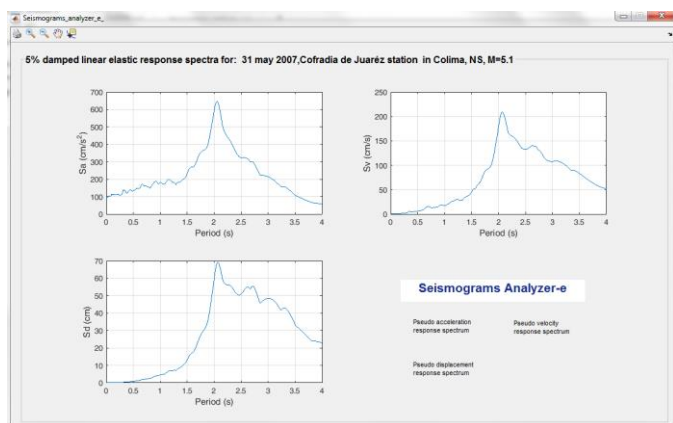


Figure 3. Response spectra related to the seismic record of the component North-South obtained in the Cofradía de Juárez Colima station, during the earthquake of magnitude 5.1 occurred on 31 May 2007 [2].

C. Probabilistic seismic hazard assessment in Colima

In order to perform the PSHA for Colima the computer code CRISIS2015 was applied [3, 4]. Originally, 43 seismic sources were considered to analyze the seismic hazard in Colima. According to the results only 14 seismic sources have a contribution to the seismic hazard of the city of Colima (Fig. 4).

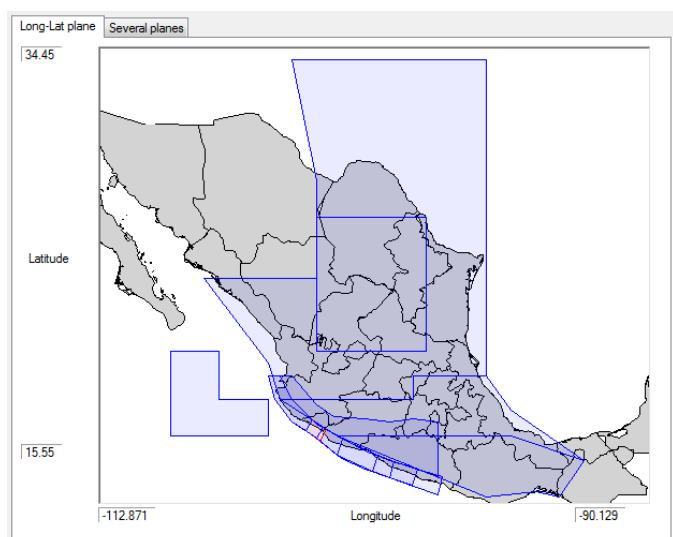


Fig. 4. Fourteen seismic sources that contribute to the seismic hazard of Colima (image of a screen of CRISIS2015 where the seismic hazard of Colima has been computed).

Figure 5 shows the seismic hazard curve of the city of Colima for a rock site. According to this curve, in a rock site of the city of Colima the intensity of 292 gals has a return period of 475 years.

On the other hand, exist evidence that confirm the importance of consider the site response effects as a part of a seismic hazard assessment. In the case of Colima recent studies have been oriented to determine the importance of site response effects in significant cities of Colima [5]. Therefore, with the purpose of consider local site effect in the assessment of the seismic hazard of the city of Colima amplification factors were considered in the present work. According to Gutierrez and colleagues [6], it is possible to expect amplification factors until 6 in some regions of the city of Colima. Therefore, this value was considered in the

assessment of the seismic hazard of the city of Colima applying CRISIS2015. Fig. 5 shows the seismic hazard curve for a soil site of the city of Colima. According to this curve, in a soil site in the city of Colima with important levels of amplification the intensity that has a return period of 475 years can reach a value of 385 gals.

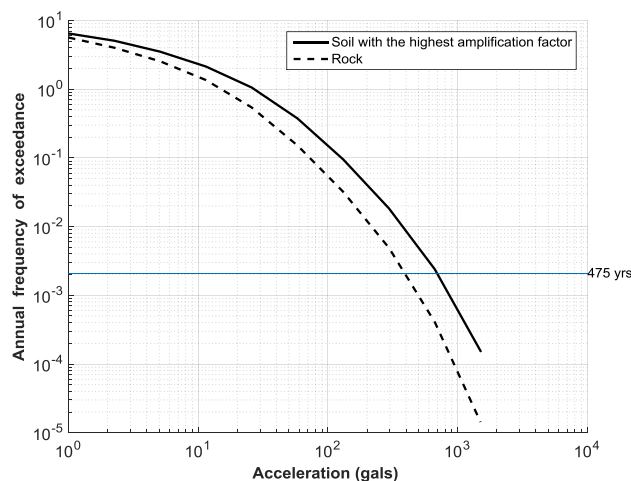


Fig. 5. Seismic hazard curves of the city of Colima, for a rock site and for a soil with important local site effects, computed by CRISIS2015.

D. Conclusion and Future Enhancement

As a preliminary conclusion it is possible to highlight that according to the preliminary results, it is very important to take into account the local site effects in the assessment of the seismic hazard of the city of Colima. At the same time has been programmed to continue with the analysis of the seismic hazard of the city Colima, considering additional data and different computation to the considered until now.

E. ACKNOWLEDGEMENT

We thanks to the CONACYT and PRODEP for the partial support for the present research. We also thanks the collaboration of Alicia Abril Zaleta Alarcón e Ingrid Aylin Cruz Martínez.

References

- [1] Singh, S. K., Pacheco, J. F., Alcántara, L., Reyes, G., Ordaz, M., Iglesias, A., ... & Reyes, C. (2003). A preliminary report on the Tecomán, Mexico earthquake of 22 January 2003 (Mw 7.4) and its effects. *Seismological Research Letters*, 74(3), 279-289.
- [2] Aguilar-Meléndez, A., Pujades, L., De la Puente Josep, et al. "Seismograms Analyzer-e", Software to analyze seismic records. 2017.
- [3] Ordaz, M., F. Martinelli, F., Aguilar, A., Arboleda, J., Meletti, C., D'Amico, V. (2015). CRISIS2015. Program for computing seismic hazard. Last accessed: 2017/02/31 <https://sites.google.com/site/codecrisis2015/>
- [4] Development and validation of software CRISIS to perform probabilistic seismic hazard assessment, with emphasis in the recent CRISIS2015. *Computación y Sistemas*, Vol. 21, No. 1, Pp. 67-90. doi: 10.13053/CyS-21-1-2578
- [5] Dominguez, R. T., Rodríguez-Lozoya, H. E., Sandoval, M. C., Sanchez, E. S., Meléndez, A. A., Rodríguez-Leyva, H. E., & Campos, R. A. (2017). Site response in a

representative region of Manzanillo, Colima, Mexico, and a comparison between spectra from real records and spectra from normative. *Soil Dynamics and Earthquake Engineering*, 93, 113-120.

Faculty of Civil Engineering at Universidad Veracruzana (UV).

- [6] Gutiérrez, C., Masaki, K., Lermo, J., Cuenca, J. (1996). *Relative Amplification and Dominant Periods Map for Seismic Motion in Colima City, México*. Eleventh World Conference on Earthquake Engineering, Paper No. 1650. ISBN: 0 08 042822 3.

Author biography

Armando Aguilar-Meléndez received his PhD in earthquake engineering and structural dynamic at Universitat Politècnica de Catalunya (UPC); His Master of Engineering (Structures) from UNAM, Mexico. Currently, he is in a postdoctoral research stay at CASE Department in BSC in Barcelona, Spain. This postdoctoral research stay is supported by CONACyT.

Josep De la Puente received his PhD in the Ludwig Maximilian University in 2007. His main topic research is about computational seismology. Currently, he is the manager of the geosciences application group of the Barcelona Supercomputing Center.

Héctor E. Rodríguez-Lozoya received PhD degree in Seismology from the CICESE, México. Currently he is a Full time Professor in the Faculty of Engineering at the Autonomous University of Sinaloa. He is a member of the National Researcher System (SNI-Mexico), Level 1.

Alejandro Córdova-Ceballos is a civil engineer graduated from the Universidad Veracruzana (UV-Veracruz, México). He obtained both the specialist degree on construction in 1999 and the master degree in construction in 2007 from the University of Veracruz. He is part of the board of peer evaluators of the committee of accreditation of engineering teaching (CACEI), in the civil engineering area.

Alejandro García-Elías obtained both the specialist degree on construction and the master degree in construction from the Universidad Veracruzana. He is member of the research group UV-CA-215 Structures of the University of Veracruz. He has been honored as Perfil Deseable PRODEP (2015-2018). He is a full time professor of the Faculty of Civil Engineering of the University of Veracruz.

Sergio Natan González-Rocha received his PhD in Environmental Management for development from Universidad Popular Autónoma de Veracruz in Mexico (UPAV); His M.Sc. in Environmental Science and his M.Sc. in Computer Science from Universidad Veracruzana (UV), Mexico. Currently, he is in a postdoctoral research stay in Earth Sciences Department in BSC in Barcelona Spain. This postdoctoral research stay is supported by CONACyT.

Amelia Campos-Rios. She is a civil engineer graduated from National Autonomous University of Mexico (UNAM). She studied the Master Degree Program on Construction Management. She has been professor of the Faculty of Engineering at UNAM and she also has been professor at the

Performance Impact of a Slower Main Memory: A case study of STT-MRAM in HPC

Kazi Asifuzzaman^{*†}, Milan Pavlovic^{*†}, Milan Radulovic^{*†}, David Zaragoza^{*†},
Ohseong Kwon[‡], Kyung-Chang Ryoo[‡] and Petar Radojković^{*}

^{*}Barcelona Supercomputing Center, Barcelona, Spain

[†]Universitat Politècnica de Catalunya, Barcelona, Spain

[‡]Samsung Electronics Co., Ltd., Memory Planning Group, Seoul, South Korea

Keywords—*STT-MRAM, Main memory, High-performance computing.*

I. EXTENDED ABSTRACT

Memory systems are major contributors to the deployment and operational costs of large-scale HPC clusters [1][2][3], as well as one of the most important design parameters that significantly affect system performance. In addition, scaling of the DRAM technology and expanding the main memory capacity increases the probability of DRAM errors that have already become a common source of system failures in the field. It is questionable whether mature DRAM technology will meet the needs of next-generation main memory systems. So, significant effort is invested in research and development of novel memory technologies. A potential candidate for replacing DRAM is Spin Transfer Torque Magnetic Random Access Memory (STT-MRAM).

In this paper, we explore whether STT-MRAM is a good candidate for HPC main memory systems. To that end, we simulate and analyze performance of production HPC applications running on large-scale clusters with STT-MRAM main memory, and compare the results with DRAM. Our results show that, despite being 20% slower than DRAM at the device level, STT-MRAM main memory delivers performance comparable to DRAM — for most of the applications under study, STT-MRAM introduces a slowdown below 1%.

A. STT-MRAM

The storage and programmability of STT-MRAM revolve around a Magnetic Tunneling Junction (MTJ). An MTJ is constituted by a thin tunneling dielectric being sandwiched between two ferro-magnetic layers. One of the layers has a fixed magnetization while the other layer's magnetization can be flipped. As Figure 1(a) and (b) depict, if both of the magnetic layers have the same polarity, the MTJ exerts low resistance therefore representing a logical “0”; in case of opposite polarity of the magnetic layers, the MTJ has a high resistance and represents a logical “1”. In order to read a value stored in an MTJ, a low current is applied to it. The current senses the MTJ's resistance state in order to determine the data stored in it. Likewise, a new value can be written to the MTJ through flipping the polarity of its free magnetic layer by passing a large amount of current through it [4].

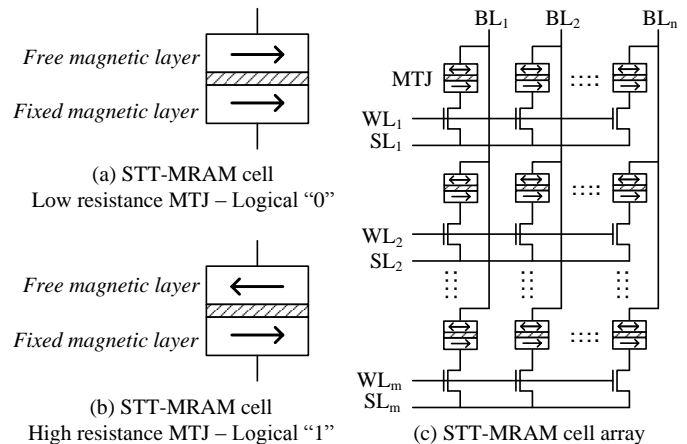


Fig. 1. STT-MRAM cell and cell-array

B. Experimental Environment

We evaluated STT-MRAM main memory on HPC applications included in the Unified European Application Benchmark Suite (UEABS) [8]. UEABS is the latest benchmark suite distributed by Partnership for Advanced Computing in Europe (PRACE) and it represents a good coverage of production HPC applications running on European Tier-0 and Tier-1 HPC systems. All UEABS applications are parallelized using Message Passing Interface (MPI) and they are regularly executed on hundreds or thousands of processing cores. UEABS also includes input data-sets that characterize production use of the applications. In our experiments, we executed UEABS applications with Test Case A, input data-set that is designed to run on Tier-1 sized systems, up to thousands of x86 cores. Table I summarizes applications used in the study.

We collected traces of UEABS applications running on MareNostrum supercomputer [10]. MareNostrum contains 3056 compute nodes (servers) connected with the Infiniband network. Each node contains two Intel Sandy Bridge-EP E5-2670 sockets that comprise eight cores operating at 2.6 GHz. Although Sandy Bridge processors support hyper-threading at core level, this feature is disabled, as in most of the HPC systems. Sandy Bridge processors are connected to main memory through four channels and each channel is connected to a single 4GB DDR3-1600 DIMM.

TABLE I. UEABS APPLICATIONS USED IN THE STUDY

Application	Scientific area	Selected number of cores
ALYA	Computational mechanics	1024
BQCD	Particle physics	1024
CP2K	Computational chemistry	1024
GADGET	Astronomy and cosmology	1024
GENE	Plasma physics	1024
GROMACS	Computational chemistry	1024
NEMO	Ocean modeling	1024
Quantum Espresso	Computational chemistry	256

We analyze main memory system in which the DRAM modules are simply replaced with STT-MRAM modules with same capacity and organization (4×4 GB DIMMs per socket), without requiring any modification in the rest of the system architecture [7]. Memory controller and channel latencies were set to 30 ns while the average DRAM device latency was simulated with 15 ns. These parameters correspond to average latencies of HPC applications running on real HPC systems [10]. Memory planning group of Samsung Electronics Co., Ltd. estimates that the high-density STT-MRAM main memory devices will be approximately 20% slower than conventionally used DRAM, so the average STT-MRAM access time was simulated with 18 ns¹. Like Suresh *et al.* [11] we used symmetrical read/write for STT-MRAM operation, which is in compliance with several scientific studies and products released recently [12][13][14].

C. Results

The performance comparison between STT-MRAM and DRAM main memory is presented in Figure 2. For each application, different bars correspond to different simulated CPI of 0.5, 1 and 2. The solid bars represent the average STT-MRAM slowdown over DRAM, and the error bars show the standard deviation for various application processes and main-loop iterations. For ALYA and GROMACS, we detect almost no performance difference between STT-MRAM and DRAM main memory systems. Four out of the remaining six applications, CP2K, GADGET, QE and BQCD, experience less than 1% slowdown. Finally, GENE slowdown ranges between 1.5% and 2%, while the slowdown of NEMO is around 2.5%. Overall, the impact of higher STT-MRAM latency on the HPC application performance is very low — for six out of eight

¹Micro-architecture and detailed timings of Samsung high-density STT-MRAM main memory devices can not be disclosed due to confidentiality issues. Samsung memory planning group estimates that capacity of high-density STT-MRAM devices will be comparable with DRAM modules.




Fig. 2. STT-MRAM slowdown with respect to DRAM main memory

applications the slowdown is below 1% and it is only 2.6% in the worst case.

D. Conclusion

We simulate and analyze performance of production HPC applications running on large-scale clusters with STT-MRAM main memory and compare the results with DRAM. Our results reveal that, although being 20% slower than DRAM at the device level, STT-MRAM main memory induces a performance degradation below 1% for most of the HPC applications under experiment.

REFERENCES

- [1] P. Kogge *et al.*, “ExaScale Computing Study: Technology Challenges in Achieving Exascale Systems,” DARPA, Sep. 2008.
- [2] A. Sodani, “Race to Exascale: Opportunities and Challenges,” Keynote Presentation at the 44th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO), Dec. 2011.
- [3] R. Stevens *et al.*, “A Decadal DOE Plan for Providing Exascale Applications and Technologies for DOE Mission Needs,” Presentation at Advanced Simulation and Computing Principal Investigators Meeting, Mar. 2010.
- [4] Y. Xie, “Modeling, Architecture, and Applications for Emerging Memory Technologies,” *IEEE Design Test of Computers*, 2011.
- [5] C. Kim *et al.*, “Magnetic Random Access Memory,” 2013.
- [6] H. Kim *et al.*, “Magneto-resistive memory device including source line voltage generator,” 2013.
- [7] H. Oh, “Resistive Memory Device, System Including the Same and Method of Reading Data in the Same,” 2014.
- [8] *Unified European Applications Benchmark Suite*, Partnership for Advanced Computing in Europe (PRACE), 2013.
- [9] A. Rico *et al.*, “On the Simulation of Large-scale Architectures Using Multiple Application Abstraction Levels,” *ACM Trans. Archit. Code Optim.*, 2012.
- [10] Barcelona Supercomputing Center, “MareNostrum III System Architecture,” <http://www.bsc.es/marenostrum-support-services/mn3>, 2013.
- [11] A. Suresh *et al.*, “Evaluation of Emerging Memory Technologies for HPC, Data Intensive Applications,” in *IEEE International Conference on Cluster Computing (CLUSTER)*, 2014.
- [12] H. Noguchi *et al.*, “A 250-MHz 256b-I/O 1-Mb STT-MRAM with Advanced Perpendicular MTJ Based Dual cell for Nonvolatile Magnetic Caches to Reduce Active Power of Processors,” in *Symposium on VLSI Technology (VLSIT)*, 2013.
- [13] R. Nebashi *et al.*, “A 90nm 12ns 32Mb 2T1MTJ MRAM,” in *IEEE International Solid-State Circuits Conference*, 2009.
- [14] Everspin Technologies, Inc., “Everspin Enhances RIM Smart Meters with Instantly Non-Volatile, Low-Energy MRAM Memory,” <http://www.everspin.com/everspin-embedded-mram>.



Kazi Asifuzzaman was born in Dhaka, Bangladesh. He received his Bachelor of Science (BSc) degree in Computer Engineering from North South University (NSU), Bangladesh in 2008. The following year, he worked at the IT department of Shimizu Denetsu Kogyo Co. Ltd (SEAVAC) in Japan. He completed his Master of Science (MSc) degree in Electronic Design from Lund University, Sweden in 2013. Since 2014, he has been with the Memory Systems group of Barcelona Supercomputing Center (BSC) as well as a PhD student at Universitat Politècnica de Catalunya (UPC), Spain. His research primarily focuses on analyzing suitability and feasibility of using emerging non-volatile memories as the main memory of high performance computing (HPC) systems.

A new Reliability-Based Data-Driven approach to Simulation-Based Models

Jacobo Ayensa Jiménez*, Mohamed Hamdy Doweidar* and Manuel Doblare Castellano*

Mechanical Engineering Department, University of Zaragoza, Spain;

Aragón Institute of Engineering Research (I3A), University of Zaragoza, Spain;

Centro de Investigación Biomédica en Red en Bioingeniería, Biomateriales y Nanomedicina (CIBER-BBN), Spain.

Campus Río Ebro, Edificio I+D, Mariano Esquillor S. N.

Email:jacoboaj@unizar.es; mohamed@unizar.es; mdoblare@unizar.es

Abstract—Data Science has burst into simulation-based engineering sciences with an impressive impulse. However, data are never uncertainty-free and a suitable approach is needed to face data measurement errors and their intrinsic randomness in problems with well-established physical constraints. We face this problem by hybridizing a standard mathematical modeling approach with a new data-driven solver accounting for the phenomenological part of the problem and able to handle the uncertainty of input data in an intelligent way. The reliability-based data-driven procedure performance is evaluated in a simple but illustrative unidimensional problem showing, in contrast with other data-driven solvers, better convergence, higher accuracy, clearer interpretation and major flexibility.

I. INTRODUCTION

Despite the wide application of Data Science in areas such as marketing and e-commerce [1], social sciences [2] or healthcare [3], there are other fields where very little has been done. An example are the disciplines where physical models and the corresponding mathematical and numerical simulation tools are well established like Computational Physics, Computational Chemistry or Computational Engineering (simulation-based engineering and sciences -SBES-). A straightforward application of these techniques is dynamic data-driven systems (DDDAS) [4], in which the idea is providing both predictive and learning capabilities to the control system from data acquired from a sufficient set of sensors. This paradigm was settled down by Kalman [5] in the sixties with his groundbreaking filter and is still nowadays a hot topic of research opening up a huge range of possibilities [6].

Unlike these approaches, based on the direct treatment of available data, SBES incorporates, in addition to data, some *a priori* characteristic physical knowledge of the analyzed system. At this point, it is crucial to distinguish between two kinds of knowledge. On one hand, physical general principles, such as conservation and thermodynamic laws that are universally accepted as able to describe the underlying universe structure. On the other hand, phenomenological models, such as macroscopic material constitutive relations. Even if it would be possible to derive the real mechanistic constitutive relations also from first physical principles, the overwhelming number of degrees of freedom involved in the relevant spatial-temporal scales needed for real applications makes this possibility intractable [7]. *Data Analytics* techniques would be very useful

in SBES to extrapolate the phenomenological submodel, but now constrained with the mathematical expression of first principles.

One idea in this direction was introduced by T. Sussman [8] for hyperelastic isotropic materials using splines interpolation in what is now known as "What You Prescribe Is What You Get" (WYPWYG) philosophy. Recently, F. Chinesta and coworkers defined a strategy for data-driven Computational Mechanics [9], combining Manifold Learning techniques and a (possibly optimized) directional search strategy inspired in the LaTin method [10]. M. Ortiz and his group presented a material model-free method based on the minimization of the distance between the searched solution and a set of experimental data, using a proper energy norm, while remaining in the equilibrium manifold, or equivalently, a well-posed penalty approach [11]. None of these works take into consideration the inherent inaccuracy of the data. Here, a new family of methods, called reliability-based data-driven solvers are defined. Data-driven solver methodology naturally allows incorporating reliability along the statement of the modeling. With this, simulations become sensitive to measurements precision and incorporate uncertainty considerations. An easy but illustrative one-dimensional problem is used to compare results and to show improvements using this methodology, emphasizing the added value with respect to existing methods.

II. METHODOLOGY

Data-driven solvers may be seen as iterative solvers searching for the intersection of an empirical (data based) manifold and a physical manifold. The first one is in many practical applications experimentally based and has, therefore, a discrete nature. The second is usually established in terms of sound laws particular to the problem in hands, but otherwise, derived from first principles universally accepted as the basis of Physics. For the sake of simplicity, let us consider the elastic three-dimensional problem. In that case, the physical manifold is the set of states that verify global and local equilibrium (i.e. conservation of linear and angular momentum), that in the static case (negligible inertial effects) is written in differential form as $\nabla \cdot \sigma = \mathbf{0}$, where σ is the stress tensor. This equation is usually approximated and solved in a discrete form using numerical methods like Finite Elements (FEM). After a

convenient discretization, we can state $\mathbf{B}\mathbf{y} = \mathbf{0}$, where \mathbf{y} is a finite dimensional vector containing the full stress tensor field information related to a given discretization (for FEM, this vector contains the components of the stress tensor for all the integration points) and \mathbf{B} is a matrix encoding the geometry and connectivity of the domain.

The empirical manifold is defined via a set, $\mathcal{E} = \{(\mathbf{x}^j; \mathbf{y}^j)\}_{j=1, \dots, m}$, of data points, resulting from experimental measurements. The set \mathcal{E} may be seen as a representation of the underlying material behavior in the following asymptotic sense: (i) if \mathcal{E} approximates a mathematical manifold and (ii) the uncertainty of each point approximates to zero.

A. Problem formulation

We postulate that a model-free engineering problem may be defined in terms of state variables (X, Y) that are related through a latent and unknown relationship $F(X, Y) = 0$. Returning to the discretized elastic problem, \mathbf{x}_i is the vector containing all strain components, ε_{kl} , at the point i and \mathbf{y} the vector containing all stress components, σ_{kl} , at the point i . It is now necessary to define a metric distance in the state space for \mathbf{x}_i and \mathbf{y}_i . As we are considering engineering problems, we have physical constraints. For the sake of simplicity, but without any conceptual limitation, we shall consider linear constraints only, so they can be written as $\mathbf{A}\mathbf{x} = \mathbf{a}$ and $\mathbf{C}\mathbf{y} = \mathbf{c}$.

At each point i , we have a trial set \mathcal{E}_i that may be thought as the result of experimental tests. We then define a local penalty function for each point i as:

$$F_i(\mathbf{x}_i, \mathbf{y}_i) = \min_{(\mathbf{x}', \mathbf{y}') \in \mathcal{E}_i} \{d_{x,i}(\mathbf{x}_i, \mathbf{x}') + d_{y,i}(\mathbf{y}_i, \mathbf{y}')\} \quad (1)$$

Therefore, a data-driven simulation-based engineering problem is defined by a constrained optimization problem where two steps are required:

- Local search of a minimum of the penalty function F_i for each element i using the nearest neighbor algorithm. This search looks for the most representative datum in the empirical discrete manifold.
- Global resolution of a linear system constraining the searched points to remain on the physical manifold.

B. Reliability-based data-driven solver

Now, each of the pairs $U^j = (X^j, Y^j)$ is considered to have random nature. Returning to the discrete case, $U = (X, Y) = (\mathbf{X}_i, \mathbf{Y}_i)_{i=1, \dots, N}$, where \mathbf{X}_i , and \mathbf{Y}_i are now random vectors whose dimension, n , is the size of the state vector and, as before, $\mathcal{N} = N \times n$ is the number of scalar state variables. We may, therefore, define a stochastic analogous problem to the deterministic one.

With this formulation, the solution candidate $\mathbf{u} = (\mathbf{x}, \mathbf{y})^T$ is not random, while $F(\mathbf{x}, \mathbf{y}|\mathcal{E}) = F(\mathbf{u}, \mathcal{E})$ is a random variable due to the random nature of \mathcal{E} , which can be characterized by means of expected value $\boldsymbol{\mu}$ and the variance-covariance matrix $\boldsymbol{\Sigma}$. A crucial point is to select a suitable norm for

this stochastic approach of the problem. A very recommended one is Mahalanobis distance [12], equivalent to choose $\mathbf{M} = 2(\boldsymbol{\Sigma})^{-1}$ as metric matrix.

Under normality conditions, $D_{\mathcal{E}}^2$ follows a non-central chi-squared distribution with $n = 2\mathcal{N}$ degrees of freedom and non-centrality parameter $\lambda = (\mathbf{u} - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1} (\mathbf{u} - \boldsymbol{\mu})$. In particular, expected value writes $\mu(D_{\mathcal{E}}^2) = 2\mathcal{N} + d_{\mathcal{E}}^2$ and variance $\sigma^2(D_{\mathcal{E}}^2) = 2(2\mathcal{N} + 2d_{\mathcal{E}}^2)$.

III. RESULTS

We evaluate the performance of different data-driven solvers, including the reliability-based one proposed herein. As it could be predicted, the main problem of the linearization approach appears when dealing with irregular (non-smooth) empirical manifolds. This is typical in Physics when working with models that have discontinuities, like in many mechanical problems such as plasticity, damage, fracture and contact problems. A very basic unidimensional trivial problem exemplifies well their main pathologies.

For a simple uniaxial loaded rod, with $F = 100$ kN, $A = 200$ cm² and $L = 10$ m, this problem may be easily solved through traditional model-based techniques. The solution is based on the combination of three equations. Equilibrium equation, $\sigma A = F$, compatibility equation, $\varepsilon = \frac{u}{L}$. For this problem to be mathematically closed, we need a mathematical relation, i.e. a model, relating the internal (state) variable stress, σ , and the measurable variable strain, ε , what is known as a constitutive relation of the material $\sigma = f(\varepsilon)$.

Let us consider that the constitutive relation is not known and the material behavior could be linear, smoothly nonlinear or non-smoothly nonlinear. In any case, what we have to describe the material behavior is a considerable amount of experimental pair values (ε, σ) , $\mathcal{E} = \{(\varepsilon^j; \sigma^j)\}_{j=1, \dots, m}$. For testing data-driven solvers based on linearization, let us compare the computed results when considering a non-smoothly nonlinear behavior and using the well-known iterative tangent Newton-Raphson method, with the analytical results obtained through the exact linear model.

Figure 1a shows the considered empirical set, the equilibrium manifold and the constitutive manifold built for some fitting techniques (linear interpolation, natural spline interpolation and 5 degree polynomial regression). The vertical dashed line shows initial points considered for the Newton-Raphson solver. Figure 1b shows the empirical set, the equilibrium manifold and final point for each solver. Both reliability-based data-driven (RBDD) and data-driven (DD) solvers converge to the same point. Even if polynomial regression is not sensitive to noise, convergence is not achieved for classic solvers because of the untrue local convexity of the built manifold. Besides, due to noise, natural splines suffer spurious oscillations provoking bad convergence. This can be avoided using linear interpolation, but in this case, non-smoothness of the broken line is incompatible with a tangent-based solver, which in turns results in non-convergence.

Figure 2 shows solution points for the DD and RBDD solvers for a material where the uncertainty associated with

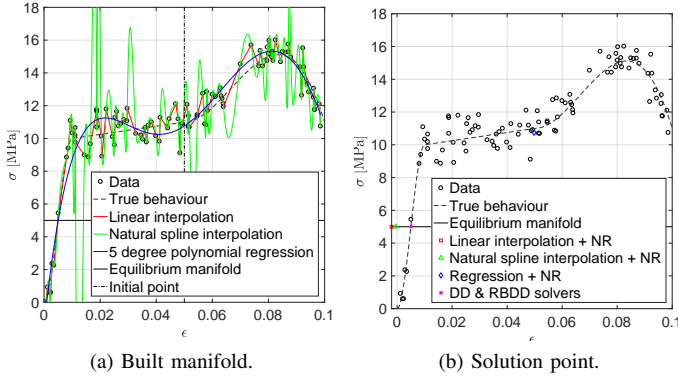


Fig. 1. Performance of different tested solvers.

the actual material behavior is not homogeneous: in the elastic zone, where the material is very well characterized, uncertainty is low, but it increases when strains are higher. RBDD solver is sensitive to this variation, while DD solver is not. For complete information, Figure 2 is complemented by the statistical properties summarized in Table I.

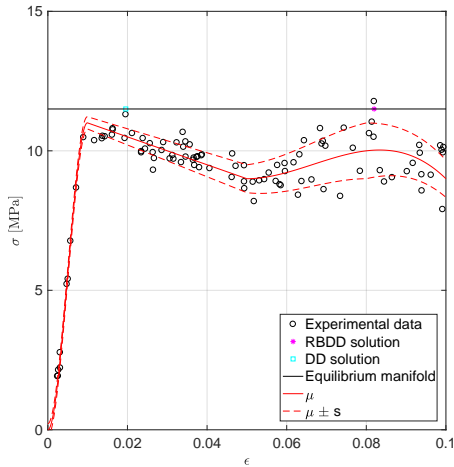


Fig. 2. Performance of DD and RBDD solvers for heterogeneous uncertainty.

TABLE I
STATISTICAL PROPERTIES FOR BOTH SOLVERS.

DD	Squared distance	$3.90 \cdot 10^2$
RBDD	Squared distance	0.08
	Expected value	3.33
	Variance	11.32
	95%-Confident bound	10.06

IV. CONCLUSION

A new reliability-based data-driven (RBDD) solver has been formulated for data-driven simulation-based engineering problems, allowing uncertainty considerations in the input data that are, therefore, not considered as uncertainty-free, but of random nature. The data-driven simulation problem is defined as a constrained stochastic optimization problem.

It has been shown that selecting a proper uncertainty dependent distance, the Mahalanobis distance, results in very good statistical properties as well as easily interpretable optimal distances. RBDD solvers present a meeting point between theoretical sciences, through epistemological constraints, and experimental sciences, through uncertain real world data. The elegance of the mathematical formulation enables many analysis and theoretical considerations for the whole spectrum of Continuum Physics. The ease of combining the presented concepts with all trendy Data Science and Deep Learning tools opens up huge possibilities for facing the most challenging problems.

ACKNOWLEDGMENT

The authors gratefully acknowledge the financial support from the Spanish Ministry of Economy and Competitiveness (MINECO MAT2016-76039-C4-4-R, AEI/FEDER, UE), the Government of Aragon (DGA) and the Biomedical Research Networking Center in Bioengineering, Biomaterials and Nanomedicine (CIBER-BBN). CIBER-BBN is financed by the Instituto de Salud Carlos III with assistance from the European Regional Development Fund.

This work is part of the paper *A new reliability-based data-driven approach for noisy experimental data with physical constraints*, in publication process.

REFERENCES

- [1] S. Hill, F. Provost, and C. Volinsky, "Network-based marketing: Identifying likely adopters via consumer networks," *Statistical Science*, pp. 256–276, 2006.
- [2] C. S. Aneshensel, *Theory-based data analysis for the social sciences*. Sage, 2013.
- [3] W. Raghupathi and V. Raghupathi, "Big data analytics in healthcare: promise and potential," *Health Information Science and Systems*, vol. 2, no. 1, p. 1, 2014.
- [4] F. Darema, "Dynamic data driven applications systems: A new paradigm for application simulations and measurements," in *International Conference on Computational Science*. Springer, 2004, pp. 662–669.
- [5] R. E. Kalman, "A new approach to linear filtering and prediction problems," *Journal of basic Engineering*, vol. 82, no. 1, pp. 35–45, 1960.
- [6] B. Peherstorfer and K. Willcox, "Dynamic data-driven reduced-order models," *Computer Methods in Applied Mechanics and Engineering*, vol. 291, pp. 21–41, 2015.
- [7] S. Xiao and T. Belytschko, "A bridging domain method for coupling continua with molecular dynamics," *Computer methods in applied mechanics and engineering*, vol. 193, no. 17, pp. 1645–1669, 2004.
- [8] T. Sussman and K.-J. Bathe, "A model of incompressible isotropic hyperelastic material behavior using spline interpolations of tension-compression test data," *Communications in numerical methods in engineering*, vol. 25, no. 1, pp. 53–63, 2009.
- [9] R. Ibanez, E. Abisset-Chavanne, J. V. Aguado, D. Gonzalez, E. Cueto, and F. Chinesta, "A manifold learning approach to data-driven computational elasticity and inelasticity," *Archives of Computational Methods in Engineering*, pp. 1–11, 2016.
- [10] P. Ladevèze, "The large time increment method for the analysis of structures with non-linear behavior described by internal variables," *COMPTES RENDUS DE L'ACADEMIE DES SCIENCES SERIE II*, vol. 309, no. 11, pp. 1095–1099, 1989.
- [11] T. Kirchdoerfer and M. Ortiz, "Data-driven computational mechanics," *Computer Methods in Applied Mechanics and Engineering*, vol. 304, pp. 81–101, 2016.
- [12] R. De Maesschalck, D. Jouan-Rimbaud, and D. L. Massart, "The mahalanobis distance," *Chemometrics and intelligent laboratory systems*, vol. 50, no. 1, pp. 1–18, 2000.



Jacobo Ayensa Jiménez is a Ph.D. student at the Mechanical Engineering Department, University of Zaragoza, Spain. He received his B.Sc. in mathematics and M.Sc. in Civil Engineering at the CFIS center, in the Polytechnic University of Catalonia, Spain, and his M.Sc. in Advanced Design in Mechanical Engineering in the University of Seville, Spain. He began his professional career in the research department of Abengoa working on the uncertainty incorporation in mechanical systems design, such as CSPP troughs and heliostats. Then he moved to the academic world for starting his Ph.D at the Aragón Institut of Engineering Research, where he is focused in mathematical modeling of tumour microenvironments.

Testing simple models for street wind conditions in Barcelona

Jaime Benavides, Oriol Jorba, Albert Soret, Marc Guevara
 Barcelona Supercomputing Center (BSC-ES)
 jaime.benavides@bsc.es

Abstract—Street wind speed and direction drive models to estimate air quality levels at street scale. In this study, simple models are combined with a mesoscale meteorological model to provide wind conditions at street level. Then, wind speed and direction are evaluated using observations collected during an experimental campaign in April 2013 at street level in Barcelona, Spain. Overall, models considering street geometry give better estimates for both wind speed and direction than those assuming homogeneous terrain. For light winds, models tend to produce a large amount of error estimating wind direction.

I. INTRODUCTION

The focus of this PhD work is to develop a coupled air quality modelling system in which a regional photochemical model, a mesoscale meteorological model, an emissions model and a street scale dispersion model are coupled to give air quality estimates of street level air pollutants. This coupled modelling system is driven by surface meteorological parameters. Typically in urban air quality modelling, these parameters are generated by adapting airport measurements to local winds using empirical equations. This approach results in less precise air pollutant level estimates compared to using local wind conditions as meteorological input [1].

II. OBJECTIVE

We aim to use specific wind conditions for each street segment to drive a dispersion model. In this abstract, a comparison of simple models to obtain street-specific wind conditions combined with WRF, a mesoscale meteorological model, is presented.

III. DATA

Wind speed and direction estimates are evaluated using meteorological observations from an experimental campaign conducted by CSIC in April 2013 in Barcelona [2]. During the campaign, mobile laboratories placed at the parking lane of several street segments measured air quality and meteorological parameters at 3 meters (m) height. For this study, data gathered every 30 minutes at Industria Road No. 213, Industria Road No. 309 and Valencia Road No. 445 were used. Street geometries in these sites are similar: street width at Industria Road and Valencia Road is 20 m and average building height is approx. 20 m. As input for road configuration, we use HERMES emission model road links [3]. For building geometry, we use district geometry from Barcelona City Council [4].

IV. METHODS

Adapting regional winds to street level using simple models can be divided in two sub-problems: adapting mesoscale wind to surface level and projecting wind flow into the street segment. In this abstract, two approaches to adapt WRF outputs to surface level using a logarithmic wind profile and taking into account atmospheric stability are evaluated: WRF surface layer parametrization (originally implemented to diagnose winds at 10 m) was adapted to estimate




Figure 1. Tested models for street wind conditions. a) Channel, approach purposed by Soulhac et al. [5]; b) OSPM, developed by Hertel and Berkowicz [6]; c) H-H, designed by Hotchkiss and Harlow [7] for transverse flow and extended by Yamartino et al. [8] for along canyon flow.

winds at 3 m height; and RLINE parametrization to estimate surface level winds [9]. These methods assume that terrain is homogeneous. In order to project wind into each street segment, three approaches are evaluated (sketches shown in Figure 1): a) channel, which projects wind speed into the street segment using θ , the angle between over roof wind direction and street axis. It estimates wind speed as $u_{street}(\theta) = u_{street}(\theta = 0^\circ) \cos(\theta)$ following Soulhac et al. [5] method. Wind direction is set to either up street or down street depending on θ angle cosine value; b) OSPM, which adapts roof wind speed to surface level using a logarithmic profile and does not consider atmospheric stability [6]. For the analysed street segments, OSPM assumes that street wind direction is mirrored inside the street compared to over roof wind direction (see Sub-figure b in Figure 1); and c) H-H model, which incorporates a model for transverse wind flow in street canyons [7] and a logarithmic wind profile for along canyon flow following Canyon Plume-Box Model methodology [8].

V. EXPERIMENT SETUP

Two WRF simulations were executed over the Catalanian domain at 1 km x 1 km horizontal resolution, for the entire month of April 2013. The first simulation used the current CALIOPE Air Quality Forecast System configuration in urban areas with surface roughness (z_0) equal 0.5. In the second simulation, z_0 was set to 1.0, 1/20 of the average building height (approx. 20 m) in Eixample District. Then,

Table I

MEAN GROSS ERROR (MGE) COMPARISON OF WIND SPEED (M/S) AND DIRECTION (DEG.) FOR EACH TIME PERIOD. MODELS WITH LOWER MGE FOR EACH PERIOD ARE HIGHLIGHTED. * DATA ENDS 24TH OF APRIL. $MGE = \frac{1}{n} \sum_{i=1}^n |Model_i - Obs_i|$

Site	Period	WRF-Open		RLINE-Open		WRF-Chan		RLINE-Chan		OSPM		H-H	
		Speed	Dir.	Speed	Dir.	Speed	Dir.	Speed	Dir.	Speed	Dir.	Speed	Dir.
Industria No. 213	3 – 11	0.57	91	0.63	91	0.51	51	0.53	51	0.73	62	0.77	62
	12 – 18	0.40	92	0.52	92	0.39	85	0.41	85	0.53	104	0.59	104
	19 – 24*	0.77	114	0.86	114	0.73	64	0.78	64	0.98	65	1.04	54
Industria No. 309	3 – 11	0.50	81	0.54	81	0.51	56	0.51	56	0.63	57	0.66	51
	12 – 18	0.37	118	0.49	118	0.39	121	0.38	121	0.48	82	0.52	88
	19 – 30	0.67	91	0.88	91	0.56	64	0.69	64	0.86	49	1.00	47
Valencia No. 445	3 – 11	0.55	76	0.52	76	0.65	53	0.62	53	0.66	58	0.72	56
	12 – 18	0.61	93	0.48	93	0.81	99	0.75	99	0.72	89	0.82	104
	19 – 30	0.68	87	0.75	87	0.75	61	0.81	61	0.88	66	1.05	61

wind conditions at 10 m using WRF and RLINE approaches were evaluated using observations from the airport station. At urban sites, WRF simulations were combined with four street wind methods: open terrain, using either WRF surface level wind (WRF-open) or RLINE surface wind parametrization (RLINE-open); channel, combined with either WRF or RLINE surface wind (WRF-Chan and RLINE-Chan); OSPM and H-H. Both OSPM and H-H, were driven by WRF first layer wind as over roof wind (middle of the first layer at 20.23 m). Hourly wind conditions were estimated for the entire period of April 2013. Additionally, this period was subdivided into three sub-periods to evaluate model outputs under low wind speeds and high wind speeds. From the 12th to the 18th of April 2013, the slow wind speed period (e.g. airport average wind speed was 3.3 m/s), there was a high pressure system over the Iberian Peninsula that brought atmospheric stability and calm winds to Barcelona [10]. In contrast, during periods 1st to 11th of April and 19th to 30th higher wind speeds were measured (e.g. average wind speed at airport 4.3 m/s and 4.8 m/s). In this abstract, Mean Gross Error (MGE) is used for model comparison. MGE is: $MGE = \frac{1}{n} \sum_{i=1}^n |Model_i - Obs_i|$.

VI. RESULTS

At the airport, hourly wind speed and direction outputs from WRF $z_0 = 1$ (WRF1) simulation were compared to WRF $z_0 = 0.5$ (WRF.5) for the entire April 2013. WRF.5 produced better wind speed estimates (MGE 2.09 m/s) than WRF1 (MGE 2.38 m/s) at the airport while wind direction results were similar (both simulations approx. MGE = 105 degrees). Then, WRF1 and WRF.5 ability to simulate wind speed combined with simple models was analysed. WRF1 produced more accurate estimates. Thus, WRF1 simulation was used as input for the comparison of simple models for street wind conditions. Regarding urban sites, a preliminary data exploration showed that Valencia Road No. 445 wind speed measurements were higher (1st period average wind speed was 1.2 m/s; 2nd 1.1 m/s; 3rd 1.06 m/s) than Industria Road No. 213 (1st period average wind speed was 0.71 m/s; 2nd 0.36 m/s; 3rd 0.52 m/s) and Industria Road No. 309 (1st period average wind speed was 0.90 m/s; 2nd 0.40 m/s; 3rd 0.54 m/s). This fact may be caused by the lower height of Valencia Road No. 445 surrounding buildings compared to Industria Road sites. Model results at street segments are summarized in Table I, where lower MGE (higher accuracy) for each period is highlighted. In the table, we see that wind speed at Valencia Road calculated by WRF-Open and RLINE-Open were more precise. In contrast, for most of the sub-periods in Industria Road sites WRF-Chan and RLINE-Chan were more accurate than other approaches. Regarding wind direction, models considering street geometry produced more precise estimates in every site and period. Additionally, under low

wind speed (12th to 18th) wind direction estimates were less precise than under high wind speed conditions.

ACKNOWLEDGEMENTS

This work is funded with a grant from the FPI Programme by the Spanish Ministry of the Economy and Competitiveness, call 2014.

REFERENCES

- [1] S. Vardoulakis, B. Fisher, K. Pericleous, and N. Gonzalez-Flesca, "Modelling air quality in street canyons: A review," *Atmospheric Environment*, vol. 37, no. 2, pp. 155–182, 2003.
- [2] F. Amato, A. Karanasiou, P. Cordoba, A. Alastuey, T. Moreno, F. Lucarelli, S. Nava, G. Calzolari, and X. Querol, "Effects of Road Dust Suppressants on PM Levels in a Mediterranean Urban Area," *Environmental Science and Technology*, vol. 48, pp. 8069–8077, 2014.
- [3] M. Guevara, F. Martínez, G. Arévalo, S. Gassó, and J. Baldasano, "An improved system for modelling Spanish emissions: HERMESv2.0," *Atmospheric Environment*, vol. 81, pp. 209–221, 2013.
- [4] Barcelona City Council, "Cartobcn," 2011. [Online]. Available: <http://w20.bcn.cat/cartobcn/>
- [5] L. Soulhac, R. J. Perkins, and P. Salizzoni, "Flow in a Street Canyon for any External Wind Direction," *Boundary-Layer Meteorology*, vol. 126, pp. 365–388, 2008.
- [6] O. Hertel and R. Berkowicz, "Modelling Pollution from Traffic in a Street Canyon. Evaluation of Data and Model Development. Technical Report 129," *National Environmental Research Institute. DMU LUFT-A.*, 1989.
- [7] R. Hotchkiss and F. Harlow, "Air pollution transport in street canyons," vol. EPA-R4-73-029, 1973.
- [8] R. Yamartino and G. Wiegand, "Development and evaluation of simple models for flow, turbulence and pollutant concentration fields within an urban street canyon," *Atmospheric Environment*, vol. 20, pp. 2137–2156, 1986.
- [9] M. Snyder, A. Venkatram, D. Heist, S. Perry, W. Petersen, and V. Isakov, "RLINE: a line source dispersion model for near-surface releases," *Atmospheric Environment*, vol. 77, pp. 748–756, 2013.
- [10] M. T. Pay, F. Martínez, M. Guevara, and J. M. Baldasano, "Air quality forecasts on a kilometer-scale grid over complex Spanish terrains," *Geoscientific Model Development*, vol. 7, no. 5, pp. 1979–1999, 2014.



Jaime Benavides has a Msc. degree in Soft Computing and Intelligent Systems with emphasis on Data Mining and a BSc. degree in Civil Engineering both from University of Granada. Currently, he is doing a PhD under the supervision of Dr. Oriol Jorba, Dr. Albert Soret and Dr. Marc Guevara dealing with the development and evaluation of a street scale air quality modelling system over Barcelona within the Earth Sciences Department at BSC.

Python for HPC geophysical electromagnetic applications: experiences and perspectives

Octavio Castillo-Reyes, Josep de la Puente and José María Cela
Geosciences Applications Group

Computer Applications in Science & Engineering

Barcelona Supercomputing Center-Centro Nacional de Supercomputación (BSC-CNS)

Abstract—Nowadays, the electromagnetic modelling are a fundamental tool in geophysics due to their wide field of application: hydrocarbon and mineral exploration, reservoir monitoring, CO₂ storage characterization, geothermal reservoir imaging and many others. In particular, the 3D CSEM forward modelling (FM) is an established tool in the oil & gas industry because of the hope that the application of such methods would eventually lead to the direct detection of hydrocarbons through their insulating properties. Although 3D CSEM FM is nowadays a well-known geophysical prospecting tool and his fundamental mathematical theory is well-established, the state-of-art shows a relative scarcity of robust, flexible, modular and open-source codes to simulate these problems on HPC platforms, which is crucial in the future goal of solving inverse problems. In this talk we describe our experience and perspectives in the development of an HPC python code for the 3D CSEM FM, namely, PETGEM. We focus on three points: 1) 3D CSEM FM theory from a practical point of view, 2) PETGEM features and Python potential for HPC applications, and 3) Modelling results of real-life 3D CSEM FM cases. These points depict that PETGEM could be an attractive and competitive HPC tool to simulate real-scale of 3D CSEM FM in geophysics.

Index Terms—Python, HPC, electromagnetic modelling, geophysics.

I. 3D CSEM FM THEORY

Regardless of the scope or scale of the surveys, exploration geophysics methods are based on studying the propagation of the different physical fields within the Earth. In the context of geophysical exploration, the main target of these methods is to build a constrained model of geology, lithology and fluid properties based upon which commercial decisions about reservoir exploration, development and management can be made [7]. Nowadays, the three main technologies in applied geophysics are: seismic methods, potential field methods (magnetic and gravity approaches) and electromagnetic methods. Each of these methods processes a set of data within an integrated framework, so that the resultant Earth model is coherent with all data used in its construction.

In the oil & gas sector, the seismic methods have become a standard technique for obtaining high-resolution images of structure and stratigraphy of the Earth. However, seismic data have extremely poor sensitivity to changes in the type of fluids, such as water, brine, oil & gas. It is the main reason why in some scenarios it is difficult, if not impossible, to determine fluid properties from seismic data. The drawback of the seismic method of determining the presence of oil in a formation, encouraged the development of new methods

aimed to strengthen the models and reduce uncertainty. In this sense, the electromagnetic methods (EM) have received special attention from industry and academia. On top of that, the last decade has been a period of rapid growth for EM in geophysics, mostly because of their industrial adoption. In particular, the three-dimensional marine controlled-source electromagnetic forward modelling (3D CSEM FM) method has become an important technique for reducing ambiguities in data interpretation in hydrocarbon exploration.

In 3D CSEM FM a deep-towed electric dipole transmitter is used to produce a low frequency electromagnetic signal (primary field) which interacts with the electrically conductive Earth and induces eddy currents that become sources of a new electromagnetic signal (secondary field). The two fields, the primary and the secondary one, add up to a resultant field, which is measured by remote receivers placed on the seabed. Since the electromagnetic field at low frequencies, for which displacement currents are negligible, depends mainly on the electric conductivity distribution of the ground, it is possible to detect thin resistive layers beneath the seabed by studying the received signal [7]. Operating frequencies of transmitters in CSEM may range between 0.1 and 10 Hz, and the choice depends on the dimensions of a model. In most studies, typical frequencies vary from 0.25 to 1 Hz, which means that for source-receiver offsets of 1012 km, the penetration depth of the method can extend to several kilometres below the seabed [5], [7]. The main disadvantage of 3D CSEM FM is its relatively low resolution compared to seismic imaging. Therefore, marine CSEM is often used in conjunction with seismic surveying as the latter helps to constrain the resistivity model.

II. HPC PYTHON TOOL FOR 3D CSEM FM

In order to be able to predict the electromagnetic signature of a given geological structure, modelling tools provide us with synthetic results which we can then compare to real data. Additionally, in the multi-core and many-core era, parallelization is a crucial issue. Edge Finite Element Method (EFEM) offer good scalability potential. Its low DOF number after primary/secondary field decomposition make them potentially fast, which is crucial in the future goal of solving inverse problems which might involve over 100,000 realizations. As consequence, in past 2 decades the modelling tools have become one of the pillars for the simulation

of numerical methods which main goal is elucidating the fundamental mechanisms behind simplified abstractions of complex phenomena in different areas. The 3D CSEM FM in geophysics is no exception and the scientific community has developed important contributions in this field. In this regard some examples of modelling tools for geophysical prospecting are [1], [6], [7], [8], [9], [10], [11]. However, the tools that full fit needs for the solution of real models are commercial and often are inaccessible to the wider scientific community. Furthermore, details of their implemented methods are generally hidden behind a black box, which could lead to a situation in which the formulation could be unknown. Due to the discretization method employed, not all codes that are affordable to community are capable of dealing with complex geometries such as models including bathymetries. Additionally there are few parallel codes that are flexible, modular, efficient, scalable and can deliver good performance. On top of that, we have developed and documented a new open-source modelling code for 3D CSEM FM in geophysics using Python and its parallel and vectorized techniques on HPC platforms, namely, the Parallel Edge-based Tool for Geophysical Electromagnetic Modelling (PETGEM). Its based on tetrahedral meshes, as these are the easiest to scale-up to very large domains or arbitrary shape, and is written mostly in Python with heavy use of mpi4py and petsc4py packages for parallel computations. Other scientific Python packages used include: H5py for binary data format support, Numpy for efficient array manipulation and Scipy algorithms. PETGEM allow users to specify edge-based variational forms of $H(\text{curl})$ for the simulation of electromagnetic fields in real 3D CSEM FM on shared-memory/distributed-memory HPC platforms. Many features have gradually been included, such as modules for EFEM data structures and a set of Python wrappers for the use of efficient solvers and preconditioners suitable for the resulting matrix system. PETGEM is now a complete package particularly suited for the 3D CSEM FM aiming to foster our understanding about EM in geophysics and its coupling with HPC technologies. Since it was intended tackle realistic problems, its data structure was designed to cope simultaneously three key requirements: accuracy, flexibility and efficiency. In addition the adopted algorithms has the possibility to easily add or remove components without having to rewrite large parts of the code. This approach leads to optimal performance in terms of development and computation time, in other words, PETGEM uses Python as a glue language for interconnecting different modules of codes written in compiled languages. By exploiting this methodology, complex scientific modelling codes can take advantage of the best attributes of both worlds: the efficient high-level data structures and a simple but effective approach to object-oriented programming of Python, and the well-know efficiency of compiled languages for numerically intensive computations.

III. RESULTS AND CODE AVAILABILITY

We have verified the PETGEM solutions by comparison to well-established CSEM models. In [2] we described the simulation results for a canonical model of an off-shore hydrocarbon reservoir. Additionally, in [3] we presented a 3D CSEM

FM with bathymetry. This model is especially interesting because the EFEM allow precise and efficient representations of arbitrarily complex geological structures such as seafloor bathymetry without critically increasing problem sizes. In both cases, the numerical solutions obtained with PETGEM were found in good agreement with reference models. Finally, in [4] we depicted some performance improvements for the parallel assembly of EFEM matrices. The code is available at <http://petgem.bsc.es> or by requesting the author (octavio.castillo@bsc.es). The code is supplied in a manner to ease the immediate execution under Linux platforms.

IV. CONCLUSIONS

The 3D CSEM modelling is well established and widely used in industry and academy to define and characterize bodies by its electric resistivity, which help us to conduct exploration campaigns with a significant reduction of costs and risks. On the other hand, simulation and modelling tools help us to formalize and simplify the complexity we observe in nature. This simplification together to HPC advances allow us to render natural phenomena treatable and testable. We developed and documented the PETGEM as a new open-source tool to promote the use and understanding of 3D CSEM FM in geophysics using HPC architectures and Python. Future work will focus on the study of code performance in other modeling scenarios. This effort, we hope, will foster our understanding about the 3D CSEM modelling and its coupling with HPC technologies.

ACKNOWLEDGEMENT

This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 644202. The research leading to these results has received funding from the European Union's Horizon 2020 Programme (2014-2020) and from Brazilian Ministry of Science, Technology and Innovation through Rede Nacional de Pesquisa (RNP) under the HPC4E Project (www.hpc4e.eu), grant agreement No. 689772. In addition, authors gratefully acknowledge the support from the Mexican National Council for Science and Technology (CONACyT).

REFERENCES

- [1] **Alumbaugh, D. L. and Newman, G. A. and Prevost, L. and Shadid, J. N. (1996)**. Three-dimensional wideband electromagnetic modeling on massively parallel computers. Radio Science. Wiley Online Library.
- [2] **Castillo-Reyes, O., de la Puente, J., Puzyrev, V., Cela, J. M. (2016)**. Edge-based parallel framework for the simulation of 3D CSEM surveys. ICE, Society of Exploration Geophysicists.
- [3] **Castillo-Reyes, O., de la Puente, J., Cela, J. M. (2017-submitted)**. PETGEM V1.0: Parallel code for 3D CSEM forward modelling in geophysics using edge finite elements Geoscientific Model Development Journal.
- [4] **Castillo-Reyes, O., de la Puente, J., Cela, J. M. (2016)**. Improving edge finite element assembly for geophysical electromagnetic modelling on shared-memory architectures. IEEE.
- [5] **Hanif, N. H. and Hussain, N. and Yahya, N. and Daud, H. and Yahya, N. and Noh, M. (2011)**. 1D Modeling of Controlled-Source Electromagnetic (CSEM) Data using Finite Element Method for Hydrocarbon Detection in Shallow Water. Proceedings of the International MultiConference of Engineers and Computer Scientists.

- [6] **Key, K. and Weiss, C. (2006)**. Adaptive finite-element modeling using unstructured grids: The 2D magnetotelluric example. *Geophysics*.
- [7] **Koldan, J. (2013)**. Numerical solution of 3-D electromagnetic problems in exploration geophysics and its implementation on massively parallel computers. PhD thesis, Polytechnic University of Catalonia.
- [8] **Li, Y. and Dai, S. (2011)**. Finite element modelling of marine controlled-source electromagnetic responses in two-dimensional dipping anisotropic conductivity structures. *Geophysical Journal International*.
- [9] **Mackie, R. L. and Smith, J. T. and Madden, T. R. (1994)**. Three-dimensional electromagnetic modeling using finite difference equations: The magnetotelluric example. *Radio Science*. Wiley Online Library.
- [10] **Newman, G. A. and Alumbaugh, D. L. (2002)**. Three-dimensional induction logging problems, Part 2: A finite-difference solution. *Geophysics*.
- [11] **Xiong, Z. and Raiche, A. and Sugeng, F. (2000)**. Efficient solution of full domain 3D electromagnetic modelling problems. *Exploration Geophysics*.



Octavio Castillo-Reyes has his bachelor's in Computer Systems engineering from Xalapa Institute of Technology, Mexico and M.Sc. in Networks and Telecommunications from Atenas Veracruzana University, Mexico. He has previously worked as lecturer at the University of Veracruz, particularly in the Master in Telematic Engineering and Bachelor in Administrative Computer Systems. His scientific interests range in the broad fields of computational methods, finite element method, multiprocessor architectures, parallel programming, python, performance and workload characterization. Octavio Castillo Reyes is currently associate PhD student at Barcelona Supercomputing Center under the supervision of PhD. José María Cela Espn and PhD Josep de la Puente. His doctoral thesis is about Edge-based finite element method for the solution of electromagnetic problems in geophysics and its coupling on HPC architectures. Octavio Castillo is supported by a studentship from Mexican National Council for Science and Technology (CONACYT).

Aggregating and Managing Memory Capacity Across Computing Nodes in Cloud Environments

Luis A. Garrido^{#1}, Paul Carpenter^{#2}

[#]Computer Science Department, Barcelona Supercomputing Center

¹luis.garrido@bsc.es, ²paul.carpenter@bsc.es

Keywords— memory management, virtualization, tmem

EXTENDED ABSTRACT

Managing memory capacity in cloud environments is a challenging issue, mainly due to the temporal variability in virtual machine (VM) memory demand. The Virtual Machine Manager or the hypervisor allocates a portion of the physical memory to the VMs, and it can change their allocation dynamically, depending on their needs. In a cloud-coomputing infrastructure, every computing node has an instance of a hypervisor. In many cases, the VMs demand for memory creates too much pressure on the memory resources of a node, prompting the need to make more memory resources available to the computing node. In our research, we have addressed and provided solutions for the two following problems: 1) how to efficiently manage the memory capacity available to a hypervisor? 2) how to aggregate memory capacity across multiple nodes?

A. Memory Management in the Hypervisor

Many solutions exist for dynamic memory management in a single node executing many VMs. Some of these are memory ballooning and memory hotplug, but these are slow to respond and do not provide suitable interfaces for memory aggregation across multiple nodes. Transcendent memory (tmem) was introduced to improve responsiveness in memory provisioning to the VMs by pooling the node's idle and fallow memory in the hypervisor. These pages are given to the VMs on demand through a key-value store.

Still, tmem presents some limitations of its own. State-of-the-art hypervisors do not implement any efficient way to manage the tmem capacity, letting VMs compete for it. Thus, it is possible for VMs to take up a disproportionate amount of tmem capacity, creating an unfair imbalance in the memory allocation, which reduces overall performance.

With our research, we demonstrated the need to implement high-level tmem capacity management, and for this we have designed a mechanism called SmarTmem. This mechanism integrates coarse-grained user-space memory management with fine-grain allocation and enforcement at the virtualization layer.

SmarTmem consists of the following components: 1) hypervisor support, 2) a tmem kernel module (TKM), and 3) user-space process called the Memory Manager (MM). The hypervisor support consists a mechanism to enforce the allocation of tmem pages as dictated by the MM and assign them to the VMs. Additionally, it includes gathering information capabilities regarding the memory utilization of the VMs. The hypervisor sends this information to the MM through the TKM via a virtual interrupt request issued every second, approximately. The Memory Manager keeps track of the status of the nodes over time and reallocates memory dynamically based on its high-level memory management policies. So far, we have implemented the following policies besides the default unmanaged (greedy) way: 1) static allocation, 2) reconfigurable static allocation, 3) and a smart allocation (SM) policy that seeks to reduce the swapping rate

proportionally to the sampled demand of the VMs. SM increases the allocation to a VM in increments of P% of the available tmem capacity, and deallocates them by a similar proportion. We tested SM for different values of P.

B. Experimental Platform and Results for SmarTmem

We evaluated SmarTmem using a VirtualBox image with Xen 4.5 and Ubuntu 14.04 with kernel 3.19 in every domain. The VirtualBox environment had two processor cores enabled, 6GB of RAM, 2GB of swap disk and 32GB of disk storage. The physical machine running VirtualBox had an Intel Core i7 processor at 2.1GHz, 8GB of RAM, 4GB of swap disk and 320GB of disk storage. The MM was implemented in C.

We use Cloudsuite benchmarks as our test applications, configuring one VM for each. We run multiple DomUs with different benchmarks, sometimes using the same benchmark multiple times. Table 1 summarizes the scenarios we used. Scenario 1 runs in-memory-analytics twice in every DomU. Scenario 2 runs graph-analytics once in every DomU and Scenario 3 runs graph-analytics in two DomUs and in-memory-analytics in the third one.

TABLE I
EVALUATION SCENARIOS FOR SMARTMEM

Scenarios	VM parameters
Scenario 1	VM1: 1GB RAM, 1 CPU, VM2: 1GB RAM, 1 CPU, VM3: 1GB RAM, 1 CPU
Scenario 2	VM1: 512MB RAM, 1 CPU, VM2: 512MB RAM, 1 CPU, VM3: 512MB RAM 1 CPU
Scenario 3	VM1: 512MB RAM, 1 CPU, VM2: 512MB RAM, 1 CPU, VM3: 1GB RAM, 1 CPU

In Figure 1 we present the running times only for Scenario 3, in which we are able to obtain a peak of 35% improvement over the case with no memory management using SM, and 40% improvement when using static allocation.

C. Aggregating Memory Resources Across Multiple Nodes

In many cases, the memory resources in a node become under pressure due to increase demand of VMs. In light of this, new computer architectures have introduced hardware support for a shared global address space with fast interconnects. These two features are exploited to enable computing nodes to share their resources, resulting in the memory capacity becoming a global rather than a local resource. This helps relieve the pressure on the resources of the node.

In order to aggregate memory capacity across multiple computing nodes, we introduced a mechanism based loosely on our initial approach for memory management within a single node. This mechanism for remote memory aggregation and management we call it AR-Tmem, since it exploits the tmem interface to aggregate memory capacity, an interface very much suitable for these purposes.

AR-Tmem consists of the following three components: 1) hypervisor support for remote memory accesses, page ownership and page transfers, 2) a tmem kernel module for communication between the user-space processes and the

hypervisor, 3) user-space process called the Remote Memory Manager (RMM). The RMMs in every node perform most of the work of AR-Tmem by cooperating to: 1) distribute memory owned by each node among its VMs, 2) distribute memory capacity among nodes, 3) implement the flow of page ownership among nodes, 4) enable nodes to join and leave, and handle failures.




Fig. 1 Average running times for Scenario 3 with different memory management policies.

In our current implementation, we have a centralized approach in which one node has an RMM master (RMM-M) which processes the memory requests of other nodes and forwards them to other nodes when necessary. The RMM-M tracks the status of all nodes by periodically receiving status information from the other RMMs

The hypervisor support includes organization of the memory it owns using a zoned Buddy allocator, with a separate zone for each node in the system (including itself) from which it has ownership of at least one page.

In order to pool the memory resources across multiple nodes, AR-Tmem requires the underlying hardware to have the following: 1) a fast interconnect, providing a synchronous interface across the NUMA architecture, 2) Direct Memory Access from the hypervisors to all the memory available, 3) disable remote access to a node's page on hardware boot, 4) ability to identify the nodes from the physical address of the pages. AR-Tmem implements one high-level policy for memory management aside from the default unmanaged un-aggregated way (greedy-local), which we refer to as Two-Tier Memory Management (TTM). TTM allocates pages by a percentage %P (similar to SM) of the pages owned by the node (local or remote) to a VM depending on its swapping rate, which it attempts to minimize. We refer to greedy-remote to the case when we aggregate memory, but disable any high-level memory management policies.

D. Experimental Platform and Evaluation for AR-Tmem

We tested our memory aggregation mechanism in an experimental platform consisting of three computing nodes (Nodes 1, 2 and 3). All nodes and VMs run Ubuntu 14.04 with Linux Kernel 3.19+ and Xen 4.5. The RMMs communicate through TCP/IP sockets. We make use of Cloudsuite but with different scenarios. Nodes 1 and 3 execute the scenario under evaluation, while Node 2 executes only the RMM-M. Table II summarizes the scenarios. Scenario 1 runs in-memory-analytics in three DomUs and Scenario 2 executes graph-analytics in three DomUs.

In Figure 2, we show the average running times only for Scenario 1 in Node 1 and Node 3. Greedy-remote obtains a 23.5% maximum improvement over greedy-local, and TTM is able to obtain a 6% maximum improvement over greedy-remote.

TABLE II
EVALUATION SCENARIOS FOR AR-TMEM

Scenarios	VM parameters
Scenario 1	VM1: 768MB RAM, 1 CPU, VM2: 768MB RAM, 1 CPU, VM3: 1GB RAM, 1 CPU
Scenario 2	VM1: 512MB RAM, 1 CPU, VM2: 512MB RAM, 1 CPU




Fig. 2 Average running times for Scenario 1 with different memory management policies. Ttm-2 is the case where P=2.


E. Current Research Efforts

We are developing sophisticated memory management policies that would work well within each node and globally. Resiliency and reliability are one of our main concerns. In addition, we are also studying how our mechanism for memory management and aggregation relates to other dynamic memory management mechanisms.

References

- [1] D. Magenheimer, C. Mason, D. McCracken and K. Hackel, "Transcendent Memory and Linux", in Proceedings of the Linux Symposium, pp. 191-200, Citeseer, 2009
- [2] Y. Durand, P.M. Carpenter, S. Adami, A. Bilas, D. Dutoit, A. Farcy, G. Gaydadjiev, J. Goodacre, M. Katevenis, M. Marazakis, E. Matus, I. Mavroidis and J. Thomson, "Euroserver: Energy Efficient node for European Microservers," in Digital System Design (DSD), 2014 17th Euromicro Conference on, pp.206-2013, IEEE, 2014.

Author biography



Luis Angel Garrido was born in Panama City, Panama. He received the B.E. degree in Electronics and Telecommunications engineering from the Technological University of Panama, Panama, in 2011, and the Master of Science degree in Electrical Engineering and Computer Science from National Chiao Tung University (NCTU), Taiwan, in 2013.

Since July 2014, he has been part of the research team at the Barcelona Supercomputing Center as PhD student. His current research interests include computer architecture, memory management, virtualization technology and cloud computing.

Discovering Ship Navigation Patterns towards Environmental Impact Modeling

Alberto Gutiérrez, David Buchaca, Josep Lluís Berral
 firstname.lastname@bsc.es

Barcelona Supercomputing Center, Data Centric Computing department

Index Terms—Pattern Mining, Machine Learning, Neural Networks, Ships, Environmental modeling.

In this work a data pipe-line to manage and extract patterns from time-series is described. The patterns found with a combination of Conditional Restricted Boltzmann Machine (CRBM) and k-Means algorithms are then validated using a visualization tool. The motivation of finding these patterns is to leverage future emission model.

I. INTRODUCTION

According to the European Community Shipowners Associations (ECSA) in 2015, the maritime traffic has become a key component for European economy [1], being sea transportation more fuel-efficient than other modes of transport (e.g. trucks and trains). Also, according to a recent report by the International Maritime Organization (IMO), it is expected that this form of transport will continue increasing in the future due to globalization and the increase of global-scale trade [2], but at the same time, it is considered an important contributor to primary atmospheric emissions in coastal areas [3] and subsequently to European coastal air quality degradation [4], especially in the North Sea and the Mediterranean basin. Maritime traffic is also responsible for about 2.5% of global greenhouse gas (GHG) emissions, which are predicted to increase between 50% and 250% by 2050 [2].

Automatic Identification System (AIS) is the Global Positioning System (GPS) based tracking system used for collision avoidance in maritime transport, as a supplement to marine radars. AIS provides information such as a unique identification for each transport (MMSI identifier), the position as latitude and longitude (GPS positioning), the course and speed (from the on-board gyrocompass). Such information is used by maritime authorities to track and monitor vessel movements and transmitted through standardized VHF transceivers, mainly to prevent collisions amongst ships.

Discovering which patterns ships perform according to the data provided will help to give explanation to: 1) air pollutant concentrations in coastal zones and cities; and 2) degradation of sea life, by detecting unusual or even criminal activities from fishing fleets working in special sea-life protection zones.

AIS data can be considered a *time-series*, as each input updates the vessel status in time. There are several approaches for mining patterns for time series, from stream mining methods for learning on time-changing data [5], to series-aware neural network methods like *Recurrent NNs* and Hidden Markov

Models [6]. Here we are focusing on a simplistic pipeline consisting in CRBMs to deal with time dimensions [7], and a classical clustering method like *k-Means* [8]. The reasons for choosing CRBMs is because our analytics goal passes to determine patterns through dimensionality reduction attempting to simplify clustering and pattern mining processes, and CRBMs have the ability to encode multidimensional input data and its history into a dimension and time aware *k*-length code, easier for feeding simplistic clustering techniques.

A. Dataset

The currently used dataset has been provided by the Spanish Ports Authority (*Puertos del Estado*), from their vessel monitoring database collecting the AIS signals from all registered ships navigating national waters. The dataset used for current experiments is a slice of data concerning the coastal area of Barcelona, including a week of maritime traffic. It is composed by more than 1.5 million entries and indicating 19 features, including the vessel identification, the position in longitude and latitude degrees, speed over ground, facing position, and other vessel properties like vessel category.

The data is cleaned and processed in order to ensure that the data is of enough quality for analyzing it. Several features are then generated from the original for using them in the analysis, e.g. rotation attribute generated from the GPS traces.

B. Methodology

The methodology here presented implements a data pipeline consisting in the preparation of data, then passing the data through a CRBM for data encoding and reduction of dimensions, then clustering / classifying it through a classical method like *k-Means*. Following subsections explains each step of the chain, also depicted on Figure 1.

The CRBM is trained with sample series of data, structured as explained before. The CRBM is not aware of time by itself, but is our *history* input data what provides such notion. This allows training it through data batches and without forcing data order, as the notion of order is already present on each new instance. Best practices in modeling and prediction require to split training data with validation and testing data, to prevent the auto-verification of the model, so for this reason we performed this training process with a subset of the available time-series.

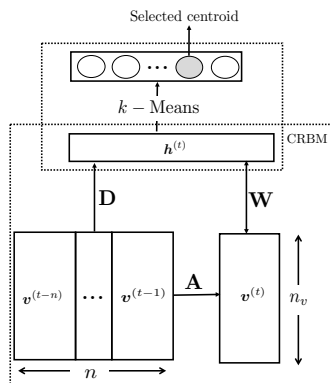


Fig. 1. Schema of the data pipeline, given a sequence at a particular time t (and its n predecessor values) the hidden activations of the CRBM are computed. Then the hidden activations are fed to a k -Means in order to cluster the sequence at time t .




Fig. 2. Example of the visualizing tool for traces and categorization. In this case patterns are distinguished using colors over the traces of the ship.

Once CRBM and k -Means models are trained, new data can be fed to the pipe-line, encoding and classifying each input into a status. At this time, similitude between patterns are checked visually using a tool for ship trace visualization, created on behalf this and future analyses in our center, as can be seen in Figure 2.

The visualization tool also allows the superposition of traces for different ships, allowing us to detect geographical regions where clustering labels *cluster*, indicating where behaviors are caused by geographical causes.

II. CONCLUSION

Detecting and discovering patterns in maritime traffic is an important topic for modeling air quality in coastal urban zones and sea-life. Maritime emissions, combined with urban emissions (industry and road traffic), are responsible of pollution in such areas.

In this paper we presented a methodology for characterizing maritime traffic, understanding it as time series, and using an ensemble of CRBMs and clustering techniques like k -Means to reduce dimensionality of data while considering time, then clustering it into common patterns of traffic. Such methodology implies pre-processing data, knowing that AIS provides error-prone data. Such datasets can be cleaned using standard techniques, also aggregated features can be derived from the most reliable ones, i.e. GPS traces.

CRBMs have proven to be useful for reducing such dimensionality, as most time series contained more than 3000 observations, even after pre-processing and reducing the time scale from seconds to minutes. When tuning the CRBM hyper-parameters, we observed that it is not required to introduce a large history window (< 20 minutes) or a high number of hidden units (~ 20) to achieve good results. Also k -Means appeared as a simple but effective approach to cluster the reduced outputs of CRBMs, comparable to real ship statuses.

By using the presented methodology, we observed identifiable patterns for real use cases, like vessel discrimination and operation modes. Such patterns can be used to complement or correct missing or erroneous data from AIS, trace ship behaviors and recognize their activity, and define geographical regions with common operation modes and behaviors. We also provided a tool for data and patterns visualization, available to the general public.

ACKNOWLEDGMENT

We would like to thanks to Marc Guevara and Albert Soret from Earth Sciences department for providing us with the use case, related information and their help on this project and Spanish Ports Authority (Puertos del Estado) for providing the data for this study. This project has received funding from the European Research Council (ERC) under the European Union Horizon 2020 research and innovation programme (grant agreement No 639595). This work is partially supported by the Ministry of Economy of Spain under contracts TIN2012-34557, 2014SGR1051, and Severo Ochoa Center of Excellence SEV-2015-0493-16-5.

This paper has been submitted to a Data Mining conference and is under review process.

REFERENCES

- [1] E. C. S. A. (ECSA), "The economic value of the eu shipping industry. update," February 2015.
- [2] T. S. et al., "Third imo greenhouse gas study," 2014.
- [3] F. D. Natale and C. Carotenuto, "Particulate matter in marine diesel engines exhausts: Emissions and control strategies," *Transportation Research Part D: Transport and Environment*, vol. 40, pp. 166 – 191, 2015. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1361920915001169>
- [4] M. Viana, P. Hammingh, A. Colette, X. Querol, B. Degraeuwe, I. de Vlieger, and J. van Aardenne, "Impact of maritime transport emissions on coastal air quality in europe," *Atmospheric Environment*, vol. 90, pp. 96 – 105, 2014. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1352231014002313>
- [5] A. Bifet and R. Gavaldá, "Learning from time-changing data with adaptive windowing," in *In SIAM International Conference on Data Mining*, 2007.
- [6] Z. Ghahramani, "Hidden markov models." River Edge, NJ, USA: World Scientific Publishing Co., Inc., 2002, ch. An Introduction to Hidden Markov Models and Bayesian Networks, pp. 9–42. [Online]. Available: <http://dl.acm.org/citation.cfm?id=505741.505743>
- [7] V. Mnih, H. Larochelle, and G. E. Hinton, "Conditional restricted boltzmann machines for structured output prediction." in *UAI*, F. G. Cozman and A. Pfeffer, Eds. AUAI Press, 2011, pp. 514–522. [Online]. Available: <http://dblp.uni-trier.de/db/conf/uai/uai2011.html#MnihLH11>
- [8] P. K. Agarwal and C. M. Procopiuc, "Exact and approximation algorithms for clustering," in *Proceedings of the Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, ser. SODA '98. Philadelphia, PA, USA: Society for Industrial and Applied Mathematics, 1998, pp. 658–667. [Online]. Available: <http://dl.acm.org/citation.cfm?id=314613.315040>

Production Planning and Scheduling Optimization Model: A case of study for a Glass Container Company

Laura Hervert-Escobar
and Oscar A. Esquivel-Flores

Instituto Tecnológico de Estudios Superiores de Monterrey
Monterrey, Nuevo Leon, Mexico

Abstract—Based on a case study, this paper deals with the production planning and scheduling problem of the glass container industry. This is a facility production system that has a set of furnaces where the glass is produced in order to meet the demand, being afterwards distributed to a set of parallel molding machines. Due to huge setup times involved in a color changeover, manufacturers adopt their own mix of furnaces and machines to meet the needs of their customers as flexibly and efficiently as possible. In this paper we proposed an optimization model that maximizes the fulfillment of the demand considering typical constraints from the planning production formulation as well as real case production constraints such as the limited product changeovers and the minimum run length in a machine. The complexity of the proposed model is assessed by means of an industrial real life problem.

I. INTRODUCTION

Currently, the activities of the planning and control of companies are becoming increasingly complex. The managers of this area are constantly pressured to reduce operating costs, maintain inventories at adequate levels, to fulfill the demand of customers, and to respond effectively to changes. There are plenty tools that help to meet these goals, however, they do not consider real factors that have and impact in the result, consequently, additional decisions should be make to complete the process.

II. PROBLEM DEFINITION

Consider a manufacturing system that produces glass containers. The process begins with the mixtures of raw materials which determines the color of the glass, this mixture includes the recycling glass. The mixture is transported into a furnace where it is melted. The glass paste is fed to a set of parallels glass moulding machines m that shape the product. The formed containers are then passed through a strict quality inspection. Containers found to be defective are discarded and melted down in the furnace as 'cullet'. Once they have been quality approved, the containers are packed on pallets at the end of the product lines. If it is required, the product pass through a decoration process, otherwise the product is stored in the warehouse or shipped for sale.

III. MATHEMATICAL SECTION

Mathematical formulation of the model is defined in the following equations. Objective function given in Equation 1 maximizes the benefits of the fulfillment of demand v_j^t by penalizing the unsatisfied demand f_j^t . Thus, the cost h of not being able to satisfied the demand is the same for all shapes while the benefit is differentiated first by SKU \hat{g}_s and consequently by shape $g_j = \sum_{j(s)} \hat{g}_s$. Additionally, a normalized cost $a_{j,m}$ of production is considered. This cost is associated with the performance of the combination of product-machine $x_{j,m}^t$. Shapes and SKU's are interrelated, where a subset of SKU's can be obtained from one type of shape. In this way from each j there are n products, we will refer the number of products generated by some shape j as $j(s)$.

$$\text{Maximize}(z) = \sum_J \sum_T g_j v_j^t - \sum_J \sum_M \sum_t a_{j,m} x_{j,m}^t - \sum_J \sum_T h f_j^t \quad (1)$$

Constraints for this model are as follows:

$$x_{j,m}^t \leq 1 \quad \forall j \in J, m \in M, t \in T \quad (2)$$

A machine m can produce only one type of shape j during period t this condition is checked in Equation 2.

$$x_{j,m}^t \leq b_{j,m}^t \quad \forall j \in J, m \in M, t \in T \quad (3)$$

Therefore, the availability of a machine for a certain product in period t according with a given timetable $b_{j,m}^t$ is outlined in Equation 3. This constraint allows to prove different campaigns of color in the furnaces. Also, it is helpful to define a maintenance program for the planning horizon. During this case study, the minimum run length per machine allowed is three days. This condition is given in Equations 4 to 6.

$$x_{j,m}^{t-1} + x_{j,m}^{t+1} \geq x_{j,m}^t \quad \forall j \in J, m \in M, t = 2, \dots, |T| - 1 \quad (4)$$

$$x_{j,m}^{t-2} + x_{j,m}^{t+1} \geq x_{j,m}^{t-1} \quad \forall j \in J, m \in M, t = 3, \dots, |T| - 1 \quad (5)$$

$$x_{j,m}^{t-2} + x_{j,m}^t \geq x_{j,m}^{t-1} \quad \forall j \in J, m \in M, t = 3, \dots, |T| \quad (6)$$

Staff for changeovers is typically a team of highly skill workers. Therefore, this resource is limited in the plant. Also, due to the time required to change the product and stabilize the process, the changeover $c_{j,m}^t$ are limited to one during weekdays.

$$x_{j,m}^t - x_{j,m}^{t-1} \leq c_{j,m}^t \quad \forall j \in J, m \in M, t = 2, \dots, |T| \quad (7)$$

$$x_{j,m}^t - x_{j,m}^{t-1} \leq -1 + 2 * c_{j,m}^t \quad \forall j \in J, m \in M, t = 2, \dots, |T| \quad (8)$$

$$\sum_J \sum_M c_{j,m}^t \leq 1 \quad \forall t \in W \quad (9)$$

$$\sum_J \sum_M c_{j,m}^t = 0 \quad \forall t \in \hat{W} \quad (10)$$

Equations 7 and 8 allow to identify the day of a changeover in a machine. While, Equations 9 and 10 limits to one changeover product in the whole plant during the weekdays, and prevent changes during the weekend. The production $p_{j,m}^t$ is given by the rate $r_{j,m}$ of the selected machine, however, when a change of product happens, the rate decreases 25%, and this condition is given in Equation 11.

$$r_{j,m} * x_{j,m}^t - 0.25 * r_{j,m} * c_{j,m}^t = p_{j,m}^t \quad \forall j \in J, m \in M, t \in T \quad (11)$$

Inventory balance equation is given in Equation 12. The balance of satisfied (v_j) and unsatisfied demand (f_j) is given in Equation 13.

$$\sum_M p_{j,m}^t + i_j^{t-1} = i_j^t + v_j^t \quad \forall j \in J, t \in T, i_j^0 = v_j^0 \quad (12)$$

$$v_j^t + f_j^t = d_j^t \quad \forall j \in J, t \in T \quad (13)$$

Inventory i_j^t is bounded by two situations, first, by the inventory policy of the SKU, given in Equation 14, and second by the capacity of the warehouse, given in Equation 15.

$$i_j^t \leq u_j \quad \forall j \in J, t \in T \quad (14)$$

$$\sum_J i_j^t \leq q \quad \forall t \in T \quad (15)$$

IV. SOLUTION METHODS

Glass container sales are highly seasonable and variable. The products are classified into three categories: A, B, C regarding the variability and the volume of the demand. Furnaces handle only one color at the time. The color and availability of the furnaces and machines is given in a predefined calendar. The first furnace (FA) has a capacity of 260 ton and has two parallel molding machines (m_1, m_2) to shape the product. Second furnace (FB) has a capacity of 360 ton and has three parallel molding machines (m_3, m_4, m_5) to shape the product. Finally, the third furnace (FC) also has a capacity of 360 ton but has 4 molding parallel machines to shape the product (m_6, m_7, m_8, m_9). The molding machines are able to produce 55 types of shapes or sub-products that are later turned into 135 types of final products (SKU's). This transformation takes place in the decoration process, which has unlimited capacity.

ID	% Demand		ID			WH	CO	RL (days)		ET
	Sat-is-fied	Un-sat-is-fied	A	B	C			MIN	MAX	
E0	98.99%	1.01%	14	46	66	88%	186	3	151	0.67
E1	99.15%	0.85%	13	49	60	86%	204	3	202	2.7
E2	99.00%	1.00%	14	52	53	89%	186	3	275	1.14
E3	99.16%	0.84%	16	52	105	126%	211	3	180	0.83
E4	99.57%	0.43%	14	48	65	77%	198	3	175	4.12
E5	99.39%	0.61%	10	42	70	67%	183	3	312	2
E6	99.12%	0.88%	13	52	57	87%	197	3	182	16.2
E7	98.68%	1.32%	26	62	83	89%	186	3	268	12

TABLE I: Business metrics results

Several scenarios were designed according to the needs of the company. The scenarios were designed under different inventory policies for products A and B which represent 95% of demand. Products under classification C are Make-To-Order and do not have inventory policy. The horizon planning is for 550 days with a bucket of one day. The scenario E0 use the history data of the company. The rest of the scenarios use a forecast demand for the following 550 days. Then, scenario E1 helps to find a solution in case the forecast demand is actually 5% higher. Scenario E2 tests the forecast demand with a change in the upper bound of the mix of products. Scenario E3 considers an unlimited capacity in the warehouse, in the case of study, the company has a limited space for store, this scenario analyzes the result of the metrics of interest when such limitation does not exist. Another situation of interest in this case of study is when the current capacity of the warehouse is reduced by external situations. Then, scenario E4 tests a reduction of 12.5% of the capacity and E5 considers a 25% of reduction. Scenario E6 tests the reduction of the upper bound for the products make-to-stock, or with classification A and B. Finally, scenario E7 tests the stop of operations at the end of the year.

The proposed model was implemented using AMPL to call the optimizer GUROBI 5.6. The results of the proposed methodology are given in terms of business metrics and the execution time.

V. RESULTS

The results for business metrics are showed in table 1. Where rows are used to identify the scenario and columns denote the business metric of interest. The table provides the result for each scenario and business metric. The business metrics of interest are the service level given by the satisfied and unsatisfied demand, the average inventory days ID for each type of product, the final inventory level, which is given as a percentage of utilization of the warehouse, the product changeovers in a year CO , and the minimum and maximum run-length RL of a product, and the execution time ET of the model.

As table 1 shows, most of the scenarios improved the service level in comparison with the current result. Due to the data confidentiality, it is not possible to present the real numbers of production and benefits. Although the improvement in terms of service level seems small, the results in economic benefits and the utilization of the warehouse fully justify the implementation of the tool. The solving time increases significantly as the capacity of the production plant and/or the warehouse is reduced. The proposed model generates approximately 82,000 continuous variables and 110,000 integer variables. But even in those cases, the model is able to achieve an optimal solution within a reasonable time for the user of the tool. It is expected that the company uses the tool at most once per month and at least once per semester in order to readjust strategies.

VI. CONCLUSIONS

This paper presents an approach to solve a real life production planning and scheduling problem. The model represents as possible the operation of a glass container industry. In this research, we provide a tool that provides a strategy for the planning production as well as detailed information (scheduling plan) to apply the proposed strategy. The fact that we are deriving long-term plans might reduce the importance of the computation times. The tool is flexible enough to test different scenarios under a reasonable time according to the availability of the final user. Next step in this research is to include the scheduling of the color campaign in the furnaces and the load of product machines simultaneously. This will lead to ability to analyze different approaches for the solution since the size and complexity of the problem will increase significantly.

VII. ACKNOWLEDGEMENTS

This research has been submitted for presentation at CLAIO 2016. The authors are grateful to Sintec for financial and technical support during the development of this research. Sintec is the leading business consulting firm for Supply Chain, Customer and Operations Strategies with a consultative model in Developing Organizational Skills that enable their customers to generate unique capabilities based on processes, organization and IT. Also, we appreciate the financial support of CONACYT-SNI program in order to promote quality research.

REFERENCES

- [1] J. F.; Caravilla M. A. Almada-Lobo, B.; Oliveira. Production planning and scheduling in the glass container industry: A vns approach. *International Journal of Production Economics*, 114(1):363–375, 2008.
- [2] Michael O. Ball. Heuristics based on mathematical programming. *Surveys in Operations Research and Management Science*, 16(1):21 – 38, 2011.
- [3] Marcio; Yukio Bressan Hossomi Marcelo; Almada-Lobo Bernardo. Fabiano Motta Toledo, Claudio; da Silva Arantes. Mathematical programming-based approaches for multi-facility glass container production planning. *Computers & Operations Research*, 74:92–107, 2016.
- [4] Claudio Fabiano Motta Toledo, Márcio Da Silva Arantes, Renato Resende Ribeiro De Oliveira, and Bernardo Almada-Lobo. Glass container production scheduling through hybrid multi-population based evolutionary algorithm. *Appl. Soft Comput.*, 13(3):1352–1364, March 2013.

- [5] Bernardo Almada-Lobo; Diego Klabjan; Maria Antónia Carravilla; José F. Oliveira. Multiple machine continuous setup lotsizing with sequence-dependent setups. *Computational Optimization and Applications*, 47(3):529–552, 2010.

BIO

Name:

LAURA

Family Name:

HERVERT-ESCOBAR

Institution

INTITUTO TECNOLÓGICO DE ESTUDIOS SUPERIORES DE MONTERREY

Department:

Research

Group:

Industrial Engineering and numerical methods



Short Bio:

(max. 150 words)

Postdoctoral researcher at Instituto Tecnológico de Estudios Superiores de Monterrey, her research include, but are not limited to the development, analysis and implementation of metaheuristic techniques for solving complex real-life combinatorial problems with several objectives. Currently, the principal emphasis will be the study and development of advanced mathematical methods and algorithms both stochastic and hybrid (stochastic/deterministic) ones for Linear Algebra (Matrix Inversion, Solving Large Systems of Linear Equations, etc.) and multi-objective multi-constrained optimization. Also the modeling of complex systems using stochastic and hybrid approaches as well as Network Science techniques, mainly parallel algorithms and parallel computing. Work experience includes the development of different projects including new product launch, design and opening of workshops, business model re-engineering, etc.

Talk/Poster's title:

Production Planning and Scheduling Optimization Model: A case of study for a Glass Container Company

Web-based tool for the annotation of pathological variants on proteins: PMut 2017 update.

Víctor López-Ferrando¹, Xavier de la Cruz², Modesto Orozco^{1,3}, Josep Ll. Gelpí^{1,3}

¹Barcelona Supercomputing Center (BSC). Joint Program BSC-CRG-IRB research Program for Computational Biology. Barcelona. Spain.

¹victor.lopez.ferrando@bsc.es

²Vall d'Hebron Research Institute, Universitat Autònoma de Barcelona, Barcelona, Spain.

²xavier.delacruz@vhir.org

³Dept. of Biochemistry and Molecular Biomedicine, University of Barcelona. Spain.

³modesto.orozco@irbbarcelona.org, gelpi@ubedu

Keywords— Pathology prediction of SAV, Machine learning, Single amino acid variants

EXTENDED ABSTRACT

Assessing the impact of amino acid mutations in human health is an important challenge in biomedical research. As sequencing technologies are more available, and more individual genomes become accessible, the number of identified variants has dramatically increased. PMut, released back in 2005 [1], has been one of the popular predictors in this field. PMut was a neural-network-based classifier using sequence data to provide a pathology score for point mutations in proteins.

We now release a new, revised, and much more powerful version of PMut. It features PyMut prediction engine, a Python module that includes numerous machine learning capabilities aimed at the analysis of protein variant pathology annotation. We also release PMut2017 predictor, a full update of the PMut predictor based on the SwissVar [2] variation database. It achieves an accuracy of 82% and a Matthews Correlation Coefficient (MCC) of 0.62, and matches the most popular predictors' performance. The engine is implemented in Python using MongoDB engine for data management. It has been adapted to run at the HPC level to cover large scale annotation projects.

A. PMut prediction engine

PMut predictor engine (PyMut), is a Python 3 module, based on the popular scientific computing libraries NumPy (www.numpy.org), Scipy (www.scipy.org), Pandas (pandas.pydata.org), Scikit-learn (scikit-learn.org), Matplotlib (matplotlib.org), and Seaborn (seaborn.pydata.org).

PyMut performs all the operations involved in the machine learning process, such as: features computation and distribution analysis, most informative features selection, classifier training and evaluation using different metrics and Receiver Operating Characteristic (ROC) curves, the training and use of predictors.

The source code of PyMut is publicly available at <https://github.com/vlopezferrando/pymut>, and can be installed locally from the official Python package repository (<https://pypi.python.org/pypi/pymut>).

B. PMut2017 predictor

PMut2017 is a predictor trained using the SwissVar [2] variation database (October 2016 release), containing 27,203 disease and 38,078 neutral mutations on 12,141 proteins. 215 features were computed for all these mutations, describing protein interactome information, physical property differences between the wild type and mutated amino acids, and sequence conservation. Conservation features are derived from PSI-Blast [3] searches over UniRef100 and UniRef90 [4] cluster databases and Multiple Sequence Alignments performed by Kalign2 [5]. 12 of these features were selected by an iterative algorithm and used to train a Random Forest classifier. The Random Forest score (between 0 and 1) was analyzed and translated to a statistically meaningful reliability score.

PMut2017 was evaluated using two different approaches. First, we run a classic 10-fold cross-validation with 50% sequence identity exclusive folds (Fig. 1), getting an MCC of 0.62, which increases to 0.69 and 0.77 when keeping the most reliable predictions.




Fig. 1 Receiver Operating Characteristic (ROC) curve of a 10-fold cross-validation with 50% sequence identity exclusion between folds.

Secondly, we performed a blind validation using the SwissVar entries added during 2016. For this purpose, we trained a predictor using the SwissVar data of December 2015, and evaluated the predictions of the 1,656 pathological and 1,510

neutral variants on 762 proteins added during the year 2016. We compared the predictions to those of other widely used predictors (Table 1).

TABLE I

COMPARATIVE PERFORMANCE OF PMut PREDICTOR

Method	% Cov.	Acc.	Spec.	Sens.	AUC	MCC
SIFT	89.6	0.61	0.33	0.88	0.60	0.25
Polyphen2	92.1	0.64	0.35	0.91	0.63	0.32
FATHMM	90.5	0.55	0.45	0.64	0.55	0.09
CADD	95.0	0.65	0.33	0.94	0.64	0.35
M-CAP	91.5	0.60	0.19	0.95	0.57	0.22
Condel	91.0	0.63	0.40	0.84	0.62	0.26
PON-P2	42.4	0.72	0.52	0.9	0.71	0.45
PROVEAN	91.5	0.64	0.41	0.87	0.64	0.31
LRT	95.1	0.73	0.58	0.87	0.73	0.47
MutationTas.	95.1	0.65	0.31	0.96	0.64	0.36
MutationAss	95.1	0.63	0.46	0.78	0.62	0.26
MetaSVM	95.1	0.63	0.51	0.74	0.62	0.26
MetaLR	95.1	0.6	0.46	0.73	0.60	0.20
PMut	100.0	0.71	0.65	0.76	0.71	0.42
PMut (85%)	81.0	0.76	0.76	0.77	0.76	0.53
PMut (90%)	51.2	0.81	0.78	0.84	0.81	0.62

Coverage, Accuracy, Specificity, Sensitivity, Area Under the ROC Curve and Matthews Correlation Coefficient of the predictions of different methods of the new variants added to SwissVar during 2016 (3,166 mutations).

C. PMut web portal

The PMut web portal provides access to all the PMut functionalities. The portal is divided in 3 sections: 1) a data repository, with a set of precalculated features and predictions, 2) single-protein and batch analysis requests and 3) a frontend to the PyMut engine allowing the user to train their own custom predictors. Figures 2 and 3 show screenshots of the interface.

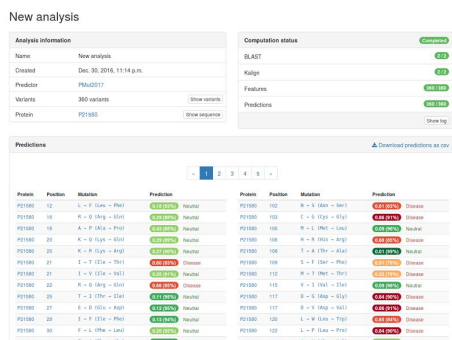


Fig. 2 Sample prediction analysis.

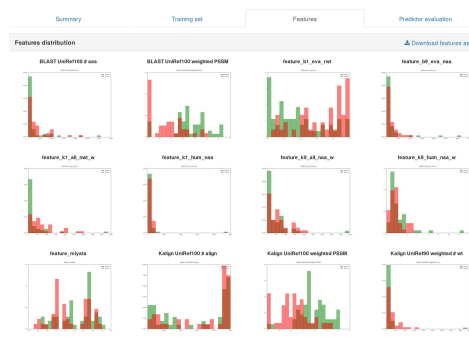


Fig. 3 Distributions of features computed to train a custom predictor.

D. Conclusions

The 2017 release of PMut includes an up to date predictor engine which matches the performance of state-of-the-art predictors, allows an easy access to all its capabilities via an intuitive web interface and offers a full set of tools bundled in the PyMut module to apply machine learning methods to the prediction of protein mutation pathology.

E. ACKNOWLEDGEMENT

We thank Dr. Patrick Aloy for providing protein interactome data. We thank Jose A. Alcantara, and Dr. Adam Hospital, for technical assistance.

References

[1] Ferrer-Costa, et al. «PMUT: A Web-Based Tool for the Annotation of Pathological Mutations on Proteins». *Bioinformatics* 21, num. 14 (15 July 2005): 3176-78.
 [2] Mottaz, et al. «Easy Retrieval of Single Amino-Acid Polymorphisms and Phenotype Information Using SwissVar». *Bioinformatics* 26, num. 6 (15 March 2010): 851-52.
 [4] Altschul, et al. «Gapped BLAST and PSI-BLAST: A New Generation of Protein Database Search Programs». *Nucleic Acids Research* 25, num. 17 (9 January 1997): 3389-3402.
 [3] Suzek, et al. «UniRef: Comprehensive and Non-Redundant UniProt Reference Clusters». *Bioinformatics* 23, num. 10 (15 May 2007): 1282-88.
 [5] Timo Lassmann, Oliver Frings, and Erik L. L. Sonnhammer, «Kalign2: High-Performance Multiple Alignment of Protein and Nucleotide Sequences Allowing External Features», *Nucleic Acids Research* 37, no. 3 (January 2, 2009): 858-65, doi:10.1093/nar/gkn1006.

Author biography



Victor López Ferrando was born in Castelló de la Plana in 1990. He studied Mathematics and Computer Engineering in Universitat Politècnica de Catalunya, having an special interest in algorithmics. His degree thesis was titled “Topology and Time Synchronization algorithms in wireless sensor networks”. After finishing his studies, he worked as a software engineer for two years at Talaia Networks, a spin-off of the Department of Computer Architecture of UPC. In September 2014 he got a La Caixa Severo Ochoa Fellowship and started his PhD in Bioinformatics in the BSC.

Programming models for mobile environments

Francesc Lordan^{#1}, Rosa M. Badia^{#2}

[#]Computer Science Department, Barcelona Supercomputing Center (BSC-CNS), Barcelona, Spain

¹flordan@bsc.es, ²rosa.m.badia@bsc.es

¹Departament d'Arquitectura de Computadors (DAC), Universitat Politècnica de Catalunya (UPC), Barcelona, Spain

²Artificial Intelligence Research Institute (IIIA), Spanish National Research Council (CSIC), Barcelona, Spain

Keywords— Mobile Cloud Computing, distributed computing, programming model, Android, Offloading

EXTENDED ABSTRACT

In the recent years, we have assisted to a revolution in IT technologies. The traditional centralized paradigm, where the whole application is hosted in local resources, has evolved into a distributed model where users have a simple device to interact with the application, but the heavy-weighted computation is performed remotely.

On one end, smartphones and tablets are devices with little computing capability that stand out for their high mobility and the wide range of possibilities to interact with the user: multi-touch screens, cameras or a large set of sensors such as GPS, light, movement,... People permanently bring a mobile device that connects them to the Internet and provides them with immediate access to all kind of services that support them in their work or daily life. For instance, a doctor who is visiting interned patients in their rooms can check on a tablet the results of previous tests and the patient evolution within the last hours to decide the most suitable treatment.

On the other end, the Cloud has emerged as the response to the growing need of computing power. Cloud technologies allow any person or organization to use an infinity of computing resources. These services have reduced the costs of owning a large computing infrastructure by converting the expenses of purchase, maintenance and operation into a pay-as-you-go bill.

Mobile Cloud Computing (MCC) brings together the benefits of both: it gathers the immediacy of access of mobile devices with the infinite computing capacity of the Cloud. Thus, mobile users can increase the computing capacity of their devices and solve more complex computational problems. Instead of consulting the evolution of the patients, doctors could simulate the impact of several treatments on them and pick the most suitable one.

Developing applications that fully exploit MCC is not straight-forward. To achieve a high performance on complex applications, developers must face all the concerns of parallelizing the application and the distribution of its components. Besides, the developer has to deal with the rapid variability of the network conditions induced by the high mobility of the mobile device. Applications should adapt their execution according to the current conditions; thus avoids harming the energy-efficiency and performance of the application. Facing these issues taking into account all the variables requires a high level of expertise. For experts in distributed computing, dealing with them means to increase the development time of the application. For developers without the expertise, they are an impassable wall.

In this article, we present COMPSs-Mobile: a framework that aims to ease the development of MCC applications by freeing the developer of all these concerns. At packaging time,

applications, written following its own programming model, are modified to invoke a runtime toolkit included in the application bundle as a library. This toolkit manages the parallelization and distribution of the application execution aiming to reduce the execution time and the energy consumption of the mobile device. To the best of our knowledge, COMPSs-Mobile is the first framework that applies the automatic parallelization and distribution of applications on mobile environments.

The main contributions can be summarized in 9 points:

- Transparent instrumentation of Android applications
- The re-designed architecture of the COMPSs runtime targeting MCC environments
- Execution of tasks within the computing devices embedded on the mobile device (CPU, GPU and accelerators)
- Mechanism to offload task execution on cloud resource
- Modeling the temporal, economic and energetic cost to decide whether to offload a task execution
- Scheduling algorithm to optimize the performance, energy consumption and price of the execution
- The data-sharing mechanism via a distributed data directory
- Fault tolerance mechanisms to handle network disruptions and a checkpointing mechanism to avoid the re-execution of the whole application
- Secured communications (secrecy with federated identities)

A. ACKNOWLEDGEMENT

This work has been supported by the Spanish Government (contracts TIN2012-34557, TIN2015-65316-P and grants BES-2013-067167, EEBB-I-15-09808, EEBB-I-16-11272 of the Research Training Program and SEV-2011-00067 of Severo Ochoa Program), by Generalitat de Catalunya (contract 2014-SGR-1051) and by the European Commission (ASCETiC project, FP7-ICT-2013.1.2 contract 610874).

B. List of publications

Lordan, F., Tejedor, E., Ejarque, J., Rafanell, R., Álvarez, J., Marozzo, F., Lezzi, D., Sirvent, R., Talia, D., Badia, R.: «Services: An interoperable programming framework for the cloud» in *Journal of Grid Computing*, 12(1), 67{91 (2014). DOI 10.1007/s10723-013-9272-5. URL: <http://dx.doi.org/10.1007/s10723-013-9272-5>

Lordan, F., Badia, R.: «COMPSs-Mobile: Parallel Programming for Mobile-Cloud Computing», in *Cluster, Cloud and Grid Computing (CCGrid)*, 2016 16th IEEE/ACM International Symposium on, 497-500, 2016, IEEE

Lordan, F., Badia, R.: «COMPSs-Mobile: Parallel Programming for Mobile-Cloud Computing», in Journal of Grid Computing (Under peer-review)

Lordan, F., Jensen, J., Badia, R.: « COMPSs-Mobile: Distributed Computing with Single Sign On from Mobile», in Journal of Grid Computing (Under peer-review)

Djemame, K., Kavanagh, R., Armstrong, D., Sirvent, R., Ejarque, J., Lordan, F., Macias, M., Guitart J., and Badia R.: «Energy Efficiency Support through Intra-Layer Cloud Stack Adaptation» in Proceedings of the 13th International Conference on Economics of Grids, Clouds, Systems and Services (GECON'2016), Athens, Greece, September 2016. (To appear)

Lordan, F., Ejarque, J., Sirvent, R., Badia, R.: «Energy-Aware Programming Model for Distributed Infrastructures» in 24th Euromicro International Conference on Parallel, Distributed, and Network-Based Processing, 2016.

Author biography



Francesc Lordan was born in Barcelona, Spain, in 1987. He received the B.E. degree in informatics engineering from the Universitat Politècnica de Catalunya, Barcelona, Spain, in 2010, and the M.Tech. Degree in computer architecture, network and systems from the Universitat Politècnica de Catalunya (UPC), Barcelona, Spain, in 2013.

Since February 2010, he has been part of the Workflows and Distributed Computing group of the Barcelona Supercomputing Center (BSC) and developing extensions of the COMPSs programming model to support the execution on Cloud environments. Currently, he is working on his PhD thesis, advised by Rosa M. Badia, studying the problems of distributed programming for mobile cloud computing.

Machine Learning Performance Prediction Model for Heterogeneous Systems

Daniel Nemirovsky*, Tugberk Arkose*, Osman Unsal*, Adrian Cristal* and Mateo Valero*

*Barcelona Supercomputing Center

{daniel.nemirovsky, tugberk.arkose, osman.unsal, adrian.cristal, mateo.valero}@bsc.es

Recent advances in machine learning techniques and specialized hardware has enabled a resurgence in the interest and applicability of powerful artificial neural network based prediction systems. However, as of yet no significant leaps have been taken towards applying machine learning in heterogeneous scheduling in order to maximize system throughput. As heterogeneous systems become more ubiquitous, computer architects will need to develop new CPU scheduling approaches applying novel techniques capable of exploiting the diversity of computational resources. However, non-heterogeneous aware schedulers like the current Linux Completely Fair Scheduler (CFS) [3] cannot take advantage of diverse system resources. Heterogeneous scheduling approaches have been previously proposed by V. Craeynest et al. [6] and Markovic et al. [2] which intend to provide full and fair utilization of the different hardware resources for all threads. These approaches yielded significant performance benefits compared to the CFS on heterogeneous architectures. In this extended abstract, we describe a novel performance prediction model which is the first of its kind to utilize machine learning performance predictors at the granularity of scheduling quanta. We then highlight how a heterogeneous system scheduler may be improved by the addition of this model.

The prediction model is composed of the following:

- Statistical information about each thread's state and which core type it has been executing on.
- A next quantum behavior predictor (NQP) that predicts what will be a thread's behavior during the next scheduling quantum.
- Machine learning based performance predictors which use a thread's behavior statistics to estimate its performance for a given core type.

Recognizing and exploiting the variations in program behavior is instrumental for effective schedulers in achieving optimal mapping schemes to maximize system performance. While not all programs exhibit the same behavior, studies [1], [4] have shown that the behavioral periodicity in different applications is typically consistent. In fact, the behavioral periodicity has been shown to be roughly on the order of several millions of instructions and is present in various different and even non correlated metrics stemming from looping structures inside of applications.

In addition, in order to provide contextual awareness to a CPU scheduler, certain thread and hardware statistics should

be periodically collected. These may include values indicating a thread's state (e.g. running, ready, or stalled), execution time, number of instruction executed, types of instructions executed, number of memory operations, cache accesses as well as hit or misses, and available cores and their types. The amount of statistics needed to be collected depends upon the complexity and optimization scheme of the scheduler.

Several novel approaches such as [5] have been proposed which predict program behavior based upon various statically or dynamically collected program statistics. However, for our model, we propose using a next quantum thread behavior predictor (NQP) that will always predict the next behavior to be equal to the previous quantum behavior. It takes as input a selection of thread statistics and by utilizing the "previous-value" prediction, it then passes these same thread statistics as the input parameters to the machine learning based performance predictors.

The machine learning architectures that the predictors are based upon can include a diverse selection. Artificial neural networks, decision trees, and clustering schemes such as k-means or nearest neighbors are examples of machine learning (ML) models that can be utilized alone or in conjunction as an ensemble to improve the prediction accuracy and training generalization. A separate ML based predictor (whether individual or ensemble based) should be implemented and trained for each different core type present in the target system. Therefore, the ML predictors take as input the thread statistics (which describe its expected behavior for the next quantum) and predict an estimated IPC value. This IPC value is the performance expected of executing that thread on the core type for which the predictor has been trained. The central features to research in detail are: the hyper-parameters of the ML predictors, the target applications of the system (or what benchmarks provide enough generalization for most programs), and which thread statistics to collect (parameter engineering). The hyper-parameters include items such as the number of hidden units and hidden layers of a neural network, the training algorithms, and which regularization techniques to use. These may be fine tuned by hand or through a machine learning approach as well and should be further evaluated by the amount of overheads that they would add. For instance, a leaner but less precise neural network may lead to a more optimal all around predictor when considering overheads compared to a very deep or complex ensemble of predictors.

Identifying target applications is also critical in the sense that

the data collected for training, validation, and testing should be composed of execution results from benchmarks which reflect these target applications. These data should be gathered by running the predetermined benchmarks on the different core types and sampling the IPC result and other thread statistics periodically after each scheduling quantum. The model could also be trained online to improve its ability to predict for newly executed applications dynamically.

When a scheduling quantum is completed, the statistics of each thread passes through the NQP and then the predictors of all the different core types (with the exception of the core type the thread last executed on since the NQP prediction will suffice in this case). In this manner, the system throughput of all the different thread mappings on the heterogeneous architecture can be compared. Assuming that the prediction error of the model is not onerously high, then these comparisons will offer a heterogeneous system scheduler ample knowledge in order to constantly maximize the system throughput. Issues to be improved upon and included in the future consist of specific ML model implementations, parameter tuning, including thread interference effects, and possibly targeting other scheduling goals such as reducing energy consumption.

I. BIO

Daniel Nemirovsky is a PhD candidate at the Universitat Politècnica de Catalunya (UPC) and a researcher in the Computer Architecture for Parallel Paradigms Group at the Barcelona Supercomputing Center. His research includes improving heterogeneous systems using machine/deep learning techniques, CPU schedulers alternative cache configurations, and advanced heterogeneity software/hardware abstractions. Daniel holds an MSc in computer architecture, networks, and systems from the

UPC and dual bachelors in electrical engineering and political science from the University of Michigan.



REFERENCES

- [1] E. Duesterwald, C. Cascaval, and S. Dwarkadas, "Characterizing and predicting program behavior and its variability," in *Parallel Architectures and Compilation Techniques, 2003. PACT 2003. Proceedings. 12th International Conference on*. IEEE, 2003, pp. 220–231.
- [2] N. Markovic, D. Nemirovsky, V. Milutinovic, O. Unsal, M. Valero, and A. Cristal, "Hardware round-robin scheduler for single-isa asymmetric multi-core," in *European Conference on Parallel Processing*. Springer, 2015, pp. 122–134.
- [3] C. S. Pabla, "Completely fair scheduler," *Linux Journal*, vol. 2009, no. 184, p. 4, 2009.
- [4] T. Sherwood, E. Perelman, G. Hamerly, S. Sair, and B. Calder, "Discovering and exploiting program phases," *IEEE micro*, vol. 23, no. 6, pp. 84–93, 2003.
- [5] O. S. Unsal, I. Koren, C. Khrihna, and C. A. Moritz, "Cool-fetch: A compiler-enabled ipc estimation based framework for energy reduction," in *Interaction between Compilers and Computer Architectures, 2004. INTERACT-8 2004. Eighth Workshop on*. IEEE, 2004, pp. 43–52.
- [6] K. Van Craeynest, S. Akram, W. Heirman, A. Jaleel, and L. Eeckhout, "Fairness-aware scheduling on single-isa heterogeneous multi-cores," in *Proceedings of the 22nd international conference on Parallel architectures and compilation techniques*. IEEE Press, 2013, pp. 177–188.

Energy Optimizing Methodologies On Heterogeneous Data Centers

Rajiv Nishtala^{*†}, Paul Carpenter^{*}, Vinicius Petrucci[‡], Xavier Martorell^{*†}

^{*} Barcelona Supercomputing Center, Barcelona, Spain
 {rajiv.nishtala, paul.carpenter, xavier.martorell}@bsc.es

[†] Universitat Politècnica de Catalunya, Barcelona, Spain

[‡] Federal University of Bahia, Salvador, Brazil
 {vpetrucci}@ufba.br

Abstract—In 2013, U.S. data centers accounted for 2.2% of the country’s total electricity consumption, a figure that is projected to increase rapidly over the next decade. Many important workloads are interactive, and they demand strict levels of quality-of-service (QoS) to meet user expectations, making it challenging to reduce power consumption due to increasing performance demands.

I. INTRODUCTION

In 2013, U.S data centers consumed 91 billion kilowatt-hour, which is enough to power every single house in New York City twice [1]. This is approximately 50 billion kg of greenhouse gas emissions, equivalent to the total volume of gas emitted by the United Kingdom, the 5th largest economy in the world [2]. A significant proportion of power consumed within a data center is attributed to the servers, and a large percentage of that is wasted as workloads compete for shared resources. Many data centers host interactive workloads (e.g., web search or e-commerce), for which it is critical to meet user expectations and user experience, called Quality of Service (QoS). There is also a wish to run both interactive and batch workloads on the same infrastructure to increase cluster utilization and reduce operational costs and total energy consumption. Although much work has focused on the impacts of shared resource contention, it still remains a major problem to maintain QoS for both interactive and batch workloads. The goal of this thesis is twofold. First, to investigate how, and to what extent, resource contention has an effect on throughput and power of batch workloads via modeling. Second, we introduce a scheduling approach to determine on-the-fly the optimal configuration to satisfy the QoS for a latency-critical job on any architecture.

To achieve the above goals, we first propose a modeling technique to estimate server performance and power at run-time called Runtime Estimation of Performance and Power (REPP). REPPs goal is to allow administrators control on power and performance of processors. REPP achieves this goal by estimating performance and power at multiple hardware settings (dynamic frequency and voltage states (DVFS), core consolidation and idle states) and dynamically sets these settings based on user defined constraints. The performance

monitoring counters (PMCs) required to build the models are available across architectures, making it architecture agnostic.

We also argue that traditional modeling and scheduling strategies are ineffective for interactive workloads. To manage such workloads, we propose Hipster that combines both a heuristic, and a reinforcement learning algorithm to manage interactive workloads. Hipsters goal is to improve resource efficiency while respecting the QoS of interactive workloads. Hipster achieves its goal by exploring the multicore system and DVFS. To improve utilization and make the best usage of the available resources, Hipster can dynamically assign remaining cores to batch workloads without violating the QoS constraints for the interactive workloads.

We implemented REPP and Hipster in real-life platforms, namely 64-bit commercial (Intel SandyBridge and AMD Phenom II X4 B97) and experimental hardware (ARM big.LITTLE Juno R1). After extensive experimental results, we have shown that REPP successfully estimates power and performance of several single-threaded and multiprogrammed workloads. The average errors on ARM, AMD and Intel architectures are, respectively, 7.1%, 9.0%, 7.1% when predicting performance, and 6.0%, 6.5%, 8.1% when predicting power. Similarly, we show that when compared to prior work, Hipster improves the QoS guarantee for Web-Search from 80% to 96%, and for Memcached from 92% to 99%, while reducing the energy consumption by up to 18% on the ARM architecture.

In this paper, we summarize the results for Hipster.

II. HIPSTER

In the last BSC PhD symposium [3], we presented a proof-of-concept version of this work. In the current version, Hipster is a hybrid reinforcement learning approach with a short learning phase captured by a heuristic algorithm and the exploitation phase of a reinforcement learning algorithm. Hipster aims to deliver the best balance between QoS guarantee and energy reduction compared to heuristic policies, like Octopus-Man [4]. In practice, to best optimize for energy efficiency, and to improve QoS, Hipster needs a short learning phase. Figure 1 shows the QoS guarantee and energy distribution




Fig. 1. QoS Guarantees of HipsterIn and Octopus-Man. Each data point represents the QoS guarantees over 100 s intervals.

over 100 s intervals for Web-Search, for both Hipster and Octopus-Man. Each data point in the graph refers to a 100-second interval. The learning phase is set to 200 s for Hipster. As can be seen, Hipster quickly learns during the heuristic phase, which improves QoS guarantees. On the other hand, for Octopus-Man, the QoS guarantees are consistently around the 80% mark, since it does not use past decisions and their associated effects to improve the future decisions.

III. RELATED WORK

Octopus-Man [4] was designed for big.LITTLE architectures to map workloads on big and small cores at highest DVFS using a feedback controller in response to changes in measured latency. Heracles [5] uses a feedback controller that exploits collocation of latency-critical and batch workloads while increasing the resource efficiency of CPU, memory and network as long as QoS target is met. However, this work is limited to modern Intel architectures due to its extensive use of cache allocation technology (CAT) and DRAM bandwidth monitor, which are available from Broadwell processors released after 2015. Pegasus [6] achieves high CPU energy proportionality for low latency workloads using fine-grained DVFS techniques. TimeTrader [7] and Rubik [8] exploit request queuing latency variation and apply any available slack from queuing delay to throughput-oriented workloads to improve energy efficiency. Quasar [9] uses runtime classification to predict interference and collocate workloads to minimize interference.

IV. CONCLUSION

One of the paradigms we discussed in the thesis and this paper is one which allocates a latency-critical workload to cores such that the real-time performance guarantees are met while improving resource utilization.

In this context, we proposed Hipster: a hybrid approach to solve this problem. Our system takes as input the performance guarantee, current latency, load, for a latency-critical job, and the static ordering of configurations of the system in order of power efficiency. Typically, a latency-critical job is allocated resources in response to changes in measured latency, and any remaining resources are allocated to batch jobs to maximize resource efficiency. We make the observation that a different latency-critical jobs have different orderings of power efficient configurations that maximize the energy efficiency while meeting real-time performance guarantees. The question becomes:

how do we detect on-the-fly the optimal configurations for a latency-critical job on any architecture? We developed a hybrid scheme that combines heuristics and reinforcement learning to manage heterogeneous cores with DVFS control for improved energy efficiency. In this abstract, we have shown that Hipster performs well across workloads and interactively adapts the system by learning from the QoS/power/performance history to best map workloads to the heterogeneous cores and adjust their DVFS settings. When only latency-critical workloads are running in the system, Hipster reduces energy consumption by 13% in comparison to prior work. In addition, to improve resource efficiency in shared data centers by running both latency-critical and batch workloads on the same system, Hipster improves batch workload throughput by 2.3 compared to a static and conservative policy, while meeting the QoS targets for the latency-critical workloads.

V. FIRST AUTHOR PAPERS

- A Hipster approach for Improving Cloud System Efficiency (Under review in ACM TOCS 2017)
- Hipster: Hybrid Task Manager for Latency-Critical Cloud Workloads. (HPCA 2017)
- REPP-H: Runtime Estimation of Performance-Power on Heterogeneous Data Centers. (SBAC-PAD 2016)
- REPP-C: Runtime Estimation of Performance-Power with Workload Consolidation in CMPs. (IGSC 2016)
- A Methodology to Build Models and Predict Performance-Power in CMPs. (ICPPW 2015)

VI. AUTHOR BIOGRAPHY

Rajiv Nishtala is a final year PhD student at Barcelona Supercomputing Center and Universitat Politècnica de Catalunya. His research interests include energy efficient (green) computing and thread scheduling.



REFERENCES

- [1] P. Delforge and J. Whitney, "Data center efficiency assessment," *Natural Resources Defense Council (NRDC)*, May 2014.
- [2] NRDC, "The power of efficiency to cut data center energy waste," June 2016.
- [3] R. Nishtala, X. Martorell, and P. Carpenter, "Runtime estimation of performance-power in CMPs under QoS constraints," May 2015, BSC PhD symposium 2015.
- [4] V. Petrucci, M. Laurenzano, J. Doherty, Y. Zhang, D. Mosse, J. Mars, and L. Tang, "Octopus-Man: QoS-driven task management for heterogeneous multicores in warehouse-scale computers," in *HPCA-2015*.
- [5] D. Lo, L. Cheng, R. Govindaraju, P. Ranganathan, and C. Kozyrakis, "Improving Resource Efficiency at Scale with Heracles," *ACM Trans. Comput. Syst.*, vol. 34, no. 2, pp. 6:1–6:33, May 2016.
- [6] L. David, L. Cheng, R. Govindaraju, L. Barroso, and C. Kozyrakis, "Towards energy proportionality for large-scale latency-critical workloads," *ACM SIGARCH Computer Architecture News*, 2014.
- [7] V. Balajee, S. Hamza, H. Jahangir, and T. N. Vijaykumar, "TimeTrader: exploiting latency tail to save datacenter energy for online search," in *MICRO-2015*.
- [8] H. Kasture, D. Bartolini, N. N. Beckmann, and D. Sanchez, "Rubik: fast analytical power management for latency-critical systems," in *MICRO-2015*.
- [9] C. Delimitrou and C. Kozyrakis, "Quasar: resource-efficient and QoS-aware cluster management," in *ASPLOS-2014*.

Data-Driven Crowd Simulation

Hugo Perez, Jorge Ramirez, Isaac Rudomin

Barcelona Supercomputing Center, Barcelona, Spain

hugo.perez@bsc.es, jorge.ramirez@bsc.es, isaac.rudomin@bsc.es

Keywords— HPC, CGI, Big Data, Machine Learning

INTRODUCTION

Our objective is to simulate entire cities with the most realistic possible scenario. This kind of systems require a lot of processing power, therefore we use hybrid computer clusters with graphics cards (GPUs). GPUs allow us to accelerate calculations and visualization.

In order to fully exploit the equipment we overlap the visualization with the calculation, in addition we apply different techniques of parallelism up to four levels, these techniques allow us to execute the simulation in several nodes of the cluster making use of all available resources (CPUs and GPUs), in addition the simulation can scale to as many nodes as necessary and available in the cluster without modifying the source code. We combine different parallel programming models like MPI, OmPss and CUDA for communication and synchronization between and within the cluster nodes [1].

Some of the applications of this kind of systems are: training in emergency situations like earthquakes, hurricanes and epidemics. Planning and urban logistics in special events such as concerts, protest march, parades or daily events in large cities such as vehicular traffic. In the area of entertainment are used in video games and movies, among many others.

DEVELOPMENT

The basis of our project is the High Performance Computing HPC, on this basis we have developed algorithms of machine learning that allow us to learn behaviors based on real data (Big Data), for example GeoLife GPS trajectory dataset released by Microsoft [2], with data of 182 users, 17,621 trajectories of ca. 1.2 million Km. and 48,000+ hours, or the history of trajectories of New York taxis during 6 years or the data that will be released by Uber soon about the movement of its passengers in different cities.

To simulate realistic urban environments we use sources such as OpenStreetMap OSM and Natural Hearth [3], [4] from them we generate 3D geometry of buildings and get data of streets, parks, points of interest, etc. To complement the data of the environment, we use Open Data, these data are published by the public administration of each city, describing aspects such as the city's infrastructure, pollution, population, among others, however, they are not standardized, therefore, each city publishes different data with its own format and structure.

The combination of different data sources such as the ones just mentioned involves a process of cleaning and extracting data before they can be useful. For this task we use Hadoop, Pig, Spark, among others.

The system can be accessed from a web browser using different types of devices such as desktop computers, tablets, mobiles, etc. For this purpose we use Websockets to transfer the final image that combines the scene with the simulation (character's motion). we receive from the client the camera position, since our system is interactive, allowing to the user manage the camera. Figure 1 shows the system architecture.

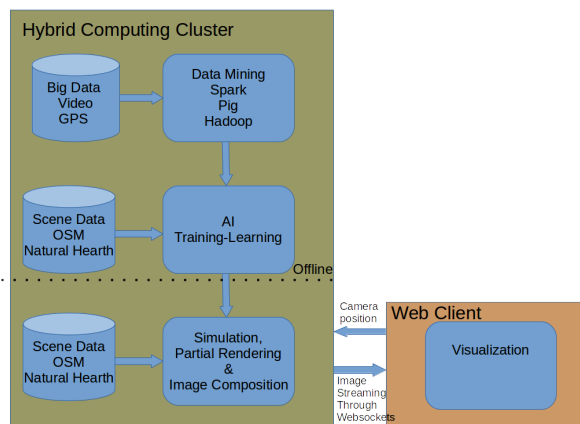


Fig. 1 System Architecture

Every new generation of embedded systems like Jetson TX2 of NVidia or Myriad 2 of Movidius-Intel has more computing resources, make them capable of compute and render part of the simulation in the client side, making more efficient procedures and communication, this concept is called Fog Computing [5]. If the client device has enough computing resources we can render the composition on the client side. For this task we use Tangram which is based on the use of OpenGL / WebGL that queries the OSM data, in the past we have also used Cesium with good results, however, to render 3D buildings, Tangram has more functionality, at the moment we were working on this task, the Cesium module in charge of rendering 3D buildings "3D Tiles" was not yet released, however, in the future will be reviewed again as it is an excellent tool.

Conclusion and Future Enhancement

The research areas covered by the project have made significant progress in recent years: HPC, CGI, Machine

Learning and Big Data.

In the area of high-performance computing and visualization, technology is constantly evolving, this implies the adaptation of the system to new computing architectures and the corresponding programming models that are used to manage it, in order to exploit resources to the maximum. As future work we plan to explore the use of OmpSs Cluster that would allow us to perform the communication within the cluster in a simpler and more efficient way, Vulkan that is the successor of OpenGL handles multithreaded rendering, so it may have a better interaction with OmpSs that will allow us to launch several threads with visualization tasks within a GPU. We also will try WebGL2 to perform visualization tasks on the client device.

Machine learning can be applied to many different tasks, for example: developing varied characters combining features of a dataset, animating characters realistically with movements learned from video data, among others. We will explore different possibilities.

To improve the interface we are developing a version in virtual reality that will allow the user to have a total immersion in the simulation and interact in a natural way using devices such as Oculus Rift and Leapmotion.

We will test different embedded systems that improve the simulation and allow greater functionality.

A. ACKNOWLEDGEMENT

This research was partially supported by: CONACyT doctoral fellowship 285730, BSC-CNS Severo Ochoa program (SEV-2011-00067), CUDA Center of Excellence at BSC, the Spanish Ministry of Economy and Competitiveness under contract TIN2012-34557, and the SGR programme (2014-SGR-1051) of the Catalan Government.

References

- [1] Hugo Pérez, Benjamín Hernández, Isaac Rudomin, Eduard Ayguade. Scaling Crowd Simulations in a GPU Accelerated Cluster. High Performance Computer Applications: 6th International Conference, Mexico City, Mexico, Revised Selected Papers. Springer 2016.
- [2] Zheng, Y., Xie, X., Ma, W.-Y.: GeoLife: A Collaborative Social Networking Service among User, location and trajectory. IEEE Data Engineering Bulletin. 33, (2010).
- [3] Menzel, J. Robert and Middelberg, Sven and Trettner, Philip and Jonas, Bastian and Kobbelt, Leif. City Reconstruction and Visualization from Public Data Sources. Eurographics Workshop on Urban Data Modelling and Visualisation 2016.
- [4] Ari Seff, Jianxiong Xiao. Learning from Maps: Visual Common Sense for Autonomous Driving. Computer Vision and Pattern Recognition 2016.
- [5] Shanhe Yi, A Cheng Li, A Qun Li. "A Survey of Fog


Computing: Concepts, Applications and Issues". Proceedings of the 2015 Workshop on Mobile Big Data.

Author biography



Hugo Perez got his B.S. degree in Electronic Engineering from National University in Mexico (UNAM). He got his M.Sc. degree in Computers Architecture, Networks and Systems from Universitat Politècnica de

Catalunya BarcelonaTech. Currently he is working in the eXtreme Computing Group at the Barcelona Supercomputing Center as PhD student. His research project entitled "Crowd Simulation and Visualization" which aims to represent the most realistic possible scenarios in a city, these kind of systems allow: urban planning, simulating disasters, simulate epidemics, among other applications. The project combines different areas research such as: High Performance Computing, Parallel Programming Models, Computer Graphics, Big Data, Machine Learning, Virtual Reality, between others.



Isaac Rudomin is a senior researcher at the Barcelona Supercomputer Center, which he joined in 2012. His focus is on crowd rendering and simulation including generating, simulating, animating, and rendering large and varied crowds using GPUs in consumer-level machines and in HPC heterogeneous clusters with GPUs. Previously, Isaac was on the faculty at Tecnológico de Monterrey Campus Estado de Mexico (from 1990 to 2012). He finished his Ph.D. at the University of Pennsylvania under Norman Badler on the topic of cloth modeling.

pyDockDNA: A new approach for protein-DNA docking

¹Luis Ángel Rodríguez-Lumbreras, ²Brian Jiménez-García, ³Juan Fernández-Recio

Life Sciences - Protein Interactions and Docking, Carrer de Jordi Girona, 29-31, 08034 Barcelona (Spain)

¹luis.rodriquez@bsc.es, ²brian.jimenez@bsc.es, ³juanf@bsc.es

Keywords— *ab-initio* docking, protein-DNA interaction, docking scoring function.

ABSTRACT

Structural prediction of protein-DNA interactions can contribute to the understanding of essential cell processes at molecular level, especially those related to gene expression and regulation. While some of the existing protein-protein docking methods, such as HADDOCK [1], can handle nucleic acids, very few algorithms specifically developed for protein-DNA docking have been reported so far [2].

In this context, our pyDock docking scheme, which has been successfully applied to model a significant number of protein-protein cases [3], can be also applied to model protein-RNA complexes [4]. However, the modeling of protein-DNA interactions with pyDock had not been fully explored so far.

Here we present pyDockDNA, which is based on the pyDock program, with a new module for reading and parsing DNA molecules. The protocol is composed of two major steps: sampling and scoring. The first sampling step consists in the generation of 10,000 protein-DNA docking models by FTDock [5]. This program takes a protein and a nucleic acid coordinate file, discretizes the molecules into corresponding 3D grids, and computes their geometric correlation by using Fast Fourier Transform algorithms to speed up the translations between the two molecules.

For the second scoring step, we have used a modification of the default scoring function implemented in pyDock [6] in order to reevaluate the protein-DNA docking models generated during the sampling step. In the new scoring function within pyDockDNA, the binding energy of each docking model is based on van der Waals (VDW), desolvation and electrostatics energy, using the charges for the DNA molecule as defined in amber 94 [7] force-field.

$$\begin{aligned}
 U_{pyDock} &= \Delta U_{electrostatic} + \Delta U_{desolvation} + \Delta U_{vanderWaals} \\
 &= \frac{1}{4\pi\epsilon} \sum_i^{N_{rec}} \sum_j^{N_{lig}} \frac{q_i \cdot q_j}{r_{ij}^2} + \\
 &+ \sum_i^{N_{rec}} \Delta ASA_i \cdot \sigma_i + \sum_j^{N_{lig}} \Delta ASA_j \cdot \sigma_j + \\
 &+ \sum_i^{N_{rec}} \sum_j^{N_{lig}} \epsilon(i, j) \left[\left(\frac{r_m(i, j)}{r_{ij}} \right)^{12} - 2 \left(\frac{r_m(i, j)}{r_{ij}} \right)^6 \right]
 \end{aligned}$$

Fig. 1: pyDock scoring function, with details of each energetic term.

We have thoroughly tested this protocol on an available protein-DNA docking benchmark [8], comparing several versions of the scoring function, and different combinations of the electrostatics, VDW, solvation energy and the SCScore (FTDock geometric score). The preliminary results show that a combination of electrostatics and VDW might be the best option to re-score the different protein-DNA decoys.

References

- [1] Dominguez, C., et al., *HADDOCK: A Protein-Protein Docking Approach Based on Biochemical or Biophysical Information*. J. Am. Chem. Soc., 2003. **125**(7): p. 1731-1737.
- [2] Banitt, I. and H.J. Wolfson, *ParaDock: a flexible non-specific DNA-rigid protein docking algorithm*. Nucleic Acids Res, 2011. **39**(20): p. e135.
- [3] Pallara, C., et al., *Expanding the frontiers of protein-protein modeling: from docking and scoring to binding affinity predictions and other challenges*. Proteins, 2013. **81**(12): p. 2192-200.
- [4] Pérez-Cano, L. and J. Fernández-Recio, *Optimal protein-RNA area, OPRA: A propensity-based method to identify RNA-binding sites on proteins*. Proteins: Structure, Function and Bioinformatics, 2010. **78**(1): p. 25-35.
- [5] Gabb, H.A., R.M. Jackson, and M.J. Sternberg, *Modelling protein docking using shape complementarity, electrostatics and biochemical information*. J. Mol. Biol., 1997. **272**(1): p. 106-120.
- [6] Cheng, T.M.-K., T.L. Blundell, and J. Fernandez-Recio, *pyDock: electrostatics and desolvation for effective scoring of rigid-body protein-protein docking*. Proteins, 2007. **68**(2): p. 503-515.
- [7] Wendy D. Cornell, P.C., Christopher I. Bayly, Ian R. Gould, Kenneth M. Merz, David M. Ferguson, David C. Spellmeyer, Thomas Fox, James W. Caldwell, Peter A. Kollman, *A Second Generation Force Field for the Simulation of Proteins, Nucleic Acids, and Organic Molecules*. J. Am. Chem. Soc., 1995. **119**(19): p. 5179-5197.
- [8] van Dijk, M. and A.M.J.J. Bonvin, *A protein-DNA docking benchmark*. Nucleic Acids Research, 2008. **36**(14): p. e88-e88.

Presenting author biography



Luis Ángel was born in Cadiz, Spain, in 1985. He received a BSc degree in biochemistry from the Universidad de Sevilla, Spain, in 2013, and a Master in Bioinformatics for Health Sciences from the Universitat Pompeu Fabra, Barcelona, Spain, in 2015.

In October 2014, he joined the Protein Interactions and Docking group at the Life Sciences department of the Barcelona Supercomputing Center, first as a Visitor Student, and a year later as a PhD candidate. His current research includes software development and optimization for protein-protein and protein-DNA interactions.

Protein modelling for enzyme engineering

Gerard Santiago^{#1}, Victor Guallar^{*#2}[#]Joint BSC-CRG-IRB Research Program in Computational Biology, Barcelona Supercomputing Center, 08034 Barcelona, Spain¹ gerard.santiago@bsc.es,^{*}Institució Catalana de Recerca i Estudis Avançats (ICREA), 08010 Barcelona, Spain² victor.guallar@bsc.es**Keywords**— laccase, aniline, protein engineering, arylamines, esterases, promiscuity, accessible volume

EXTENDED ABSTRACT

ACTIVITY PREDICTION

Laccases (EC 1.10.3.2) are a polyfacetic enzyme type with multiple applications in various fields: food industry, manufacture of anticancer drugs or ingredients in cosmetics, bioremediation, etc.. Their ability to oxidize a great variety of molecules and the versatility of laccase-mediator systems open interesting perspectives to this biocatalyst. Moreover, laccases require only oxygen to function and produce water as the reaction byproduct, which makes them a perfect green catalyst for sustainable industrial applications. Furthermore, soft buffers, like phosphate, can be used which are less aggressive for the industrial material and safer for the personnel. Arylamines, such as aniline, are a group of compounds with a large industrial applications, for example aniline polymers (PANI) have a broad range of uses. Actually PANI production requires hard chemicals as ammonium peroxydisulfate and strongly acid conditions. Some laccases are able to perform aniline oxidation for PANI production but with a low efficiency and requires protein engineering for improve activity above aniline.

By experimental directed evolution it was possible to produce an improved laccase variant for aniline oxidation and PANI production. Using computational modeling, we aim at further improve this initial design for a higher PANI production.. In particular, two positions were selected, 207 and 263, and mutated to serine and aspartic acid respectively. The aim of both mutations is to modify the binding cavity environment, improving the substrate oxidation potential. Computational predictions indicated an activity increase in the double mutant in comparison with the improved laccase variant from directed evolution. Experimental validation of double mutation confirmed a 2-fold Kcat increase for aniline oxidation and a maintenance of protein stability parameters ¹.





Fig. 1 Representative binding mode for aniline interaction with parent laccase



(A) and double mutant (B).

Fig. 2 On the left the enzymatic polymerization of 15 mM aniline along time (with 5 mM SDBS as template) as shown by the increase of absorbance at 800 nm, and on the right green PANI produced in (C) by parental and DM laccases.

PROMISCUITY PREDICTION

Enzyme promiscuity characterizes the range of substrates that react with a biocatalyst. The growing industrial needs in sustainable chemistry drives the search for new biocatalysts capable to perform a large variety of reactions. Deciphering an enzyme promiscuity, however, requires a huge amount of experimental tests and it's really difficult to obtain in an accurately manner. In this part of the work, by combining structural parameters, including the active site volume and solvent accessible surface area (SASA), we show that it is possible to predict substrate promiscuity. For the analysis we used esterases (EC 3), a well-known and widely used (including industry) hydrolase enzyme that splits esters into an acid and an alcohol, by means of a chemical reaction with water called hydrolysis.

Our previous studies revealed that the SASA and the cavity volume had key roles in substrate diffusion and enzyme-substrate interaction in different esterases with opposite promiscuity profiles. Based on this knowledge the analysis was expanded to a list of 98 esterases tested with 96 substrates each. Figure 3 presents the Accessible (or Effective) Volume, the active site volume divided by the catalytic triad SASA, for all 98 esterases versus the promiscuity, the total number of compounds hydrolyzed. The data reveals that for an Accessible Volume higher than 62,5 Å³ the esterase will be capable to hydrolyze more than 20 substrates. The base idea behind is the cavity volume as a major component of enzyme promiscuity, the bigger is the cavity the more substrate can accommodate. At some point, however, the cavity is too large, becoming too solvent exposed and unable to accommodate (retain) substrates. SASA (adimensional value) corrects the volume and allow us to obtain the Accessible Volume. That test has a precision of 95% becoming a useful tool to predict esterases promiscuity

and opening the possibility to extent the analysis to other enzymatic types.




Fig. 3 The figure shows the relationship between the reaction cavity volume and catalytic triad SASA (accessible volume) with enzyme promiscuity (number of substrates hydrolyzed). When the esterase is capable to hydrolyze more than 20 substrates has an accessible volume (Catalytic Volume/Triad SASA) higher than 62.5.

ACKNOWLEDGEMENT

This study was supported by the INDOX (KBBE-2013-7-613549) EU-project and the OxiDesign (CTQ2013-48287-R) Spanish project. G.S. acknowledges an FPI grant from the Spanish Ministry of Competitiveness.

References

- [1] Santiago G, de Salas F, Lucas MF, Monza E, Acebes S, Martínez A T, Camarero S, Guallar V Computer-aided laccase engineering: toward biological oxidation of arylamines. *ACS Catal* 2016, 6, 5415–5423.

Author biography



Gerard Santiago obtained his Biochemistry degree by University of Barcelona (UB) in 2013, and a Bioengineering Master by Institut Químic de Sarrià (IQS) in 2015. He is currently a PhD candidate in Barcelona Supercomputing Center (BSC-CNS). His

scientific career revolves around different levels of Green Biotechnology, such as: new enzyme discovery for industrial applications in extreme environments, production of high-value compounds on eukaryotic strains grown in waste materials and enzyme design (focused on oxidoreductases) for eco-friendly industry.

Exploring the use of mixed precision in NEMO

Oriol Tintó Prims^{#1}, Miguel Castrillo^{#2}

[#]*Earth Sciences Department, Barcelona Supercomputing Center, Barcelona (Spain)*

¹oriol.tinto@bsc.es, ²miguel.castrillo@bsc.es

Keywords— Mixed precision, Ocean model, Climate Sciences

EXTENDED ABSTRACT

It has been a widely extended practice in scientific computing to use 64-bit to represent data without even considering which level of precision is really needed. In many applications, 32-bit should provide enough accuracy, and in other cases 64-bit is not enough. In climate science, the inherent difficulties collecting data imply a considerable level of uncertainty, which suggest that the general use of 64-bit to represent the data may be a waste of resources, while on the other hand, some specific algorithms could benefit from an increment of the precision used. These factors suggest that in the future more attention has to be paid to the precision used in scientific software, to use the resources wisely and also avoid losing accuracy. In this work we question whether the precision used in the oceanic model NEMO is necessary and sufficient, and the potential benefits of adjusting this precision.

A. The model

The Nucleus for European Modelling of the Ocean (NEMO) is a state-of-the-art modelling framework for oceanographic research, operational oceanography, seasonal forecast and climate studies, which is used by a large community: about 100 projects and 1000 registered users around the world. It is controlled and maintained by an European consortium, made up by CNRS and Mercator-Ocean from France, NERC and Met Office from the United Kingdom and CMCC and INGV from Italy.

Several millions of computer hours are invested each year in simulations involving NEMO, resources valued in tenths, or even hundreds, of thousands of euros.

The core of NEMO is the OPA module that solves the Navier-Stokes equations from regional to global scales, using Euler first-order discretization methods on a three-dimensional (3D) grid. The model was parallelized and is able to be executed in both shared and distributed memory environments, using MPI to overcome data dependencies.

B. Is the use of 64-bit data justified in NEMO?

While it can be affirmed that high levels of precision could be used only on sensitive calculations, the generalized lack of background in numerical analysis provokes that it often has been much easier and cheaper to employ high-precision arithmetic than to take care of the necessary precision [1]. However, having the cost of scientific computations grown in several orders of magnitude, the possibility to employ mixed precision algorithms has won a lot of interest. These

algorithms can provide performance benefits while maintaining the accuracy [2][3].

Moreover, while it is clear that reducing the overall precision of a model will unequivocally increase the arithmetic and rounding errors, whether this errors are important or not has to be evaluated. In a computational model, arithmetic and rounding operations are not the only sources of error, float truncation and data uncertainty can also harm the accuracy of the results. If the errors coming from the other sources are important, spending resources in increasing the precision may not be worth.

This is the case in climate science, where the inherent difficulties to collect the data and the chaotic nature of the simulated systems suggest that minimizing arithmetic errors may not benefit the accuracy of the results, and therefore reducing the precision from the default 64-bit to 32-bit may not be as harmful as to renounce to the performance benefits that it can provide.

C. Precision and computational performance

There is a direct relation between the precision used and the cost of a computation. Most modern processors have vectorial instructions for 32-bit and 64-bit data, performing the first the double of operations per cycle than the second. Moreover, the gap between processor and memory speeds implies that the time to bring the data to the CPU can be as long as the time to compute the operations itself in the CPU and in some cases becomes the computational performance bottleneck. So finally in both CPU bound and memory bandwidth bound computations, the expected difference between using 32-bit and 64-bit representation is about a factor of 2.

D. Preliminary approach for NEMO

To explore how the use of mixed precision can benefit the computational performance of this model the first approach is to move the whole model from 64 to 32-bit wherever it is possible.

Although at the first stages of the model development the precision was intended to be easily selected by only changing one single parameter, later developments did not respect that approach and several changes had to be done to change the precision. As a consequence, some issues had to be solved before obtaining a compiling version of the code.

Nevertheless, this preliminary version of the model in 32-bit crashed in several places, but luckily enough, when the changes cause the model to crash, it was easy to track where the floating point exceptions are occurring. With this trial-and-error approach, we determined that the module corresponding to sea-ice was especially problematic and that

could be convenient to keep this part of the model in 64-bit due to the higher sensitivity to the precision used. After solving the issues derived from having variables of different precisions in the same code, it was possible to obtain a running version able to produce results.

E. Preliminary results

With this early-stage mixed precision version of the code, it is possible to evaluate the potential impact of the changes in the computational performance. To do so we compared the model throughput of the 64-bit and mixed-precision versions in three different use cases: a low-weight low resolution configuration in a work-station and a high resolution in a supercomputer. The results are exposed in Table 1, Table 2 and Figure 1.

TABLE I
MODEL THROUGHPUT IN A WORKSTATION

Precision	Throughput (Steps per Second)
Double	1.31
Mixed	2.06

Table 1 Model throughput of a low resolution ORCA2 simulation in an i7-4790 CPU @ 3.60GHz workstation using 4 of 8 cores.

TABLE II
MODEL THROUGHPUT IN MARENOSTURM III

Number of cores	Throughput (Steps per Second)	
	Double	Mixed
256	0.75	1.09
512	1.43	1.92
1024	2.39	3.15
2048	3.14	3.87

Table 2 Model throughput of a High Resolution ORCA025 configuration in Marenosturm III supercomputer at Barcelona Supercomputing Center..

It can be seen that in the workstation case (Table 1), the difference in throughput is 57% faster in mixed precision. In the high resolution case, with 256 cores the mixed precision version is 45% faster, while in the 2048 core case it is only 23%. This is explained because in low core-counts the computation has much more weight and in large core-counts the communication becomes more important.




Fig. 1 Model throughput for different number of cores with the double and the mixed precision codes with the high resolution ORCA025 configuration in the Marenosturm III supercomputer..

F. Issues

Although the model can run and produce results, these present significant differences compared to the double precision outputs. It is out of the scope of this work to evaluate these differences, but for sure it will be necessary for further studies.

G. Conclusions

The cost of the resources invested in simulations including NEMO can be estimated to be of the order of hundreds of thousands of euros. Adjusting the models to avoid the overuse of double precision representation generally has been proven to be a solution to reduce considerably the cost of the simulations. This exploratory work proved that using single instead of double precision speeded up the model a 59% in a workstation as also achieved an important improvement in high resolution simulations in Marenosturm III supercomputer, and therefore suggests that this must be studied further.

H. Future work

This study is only the necessary starting point to develop a mixed precision version of the code that uses the necessary precision in each part of the model, allowing an important reduction of the computational cost of the simulations and increasing the quality of the results. To do so it is necessary to develop methods for analyzing the algorithms and their response to the used precision. This kind of work has been done in other models and can be orientative.

References

- [1] Bailey, D.h. "High-Precision Floating-Point Arithmetic in Scientific Computation." *Computing in Science and Engineering* 7.3 (2005): 54-61.
- [2] Baboulin, Marc, Alfredo Buttari, Jack Dongarra, Jakub Kurzak, Julie Langou, Julien Langou, Piotr Luszczek, and Stanimire Tomov. "Accelerating scientific computations with mixed precision algorithms." *Computer Physics Communications* 180.12 (2009): 2526-533.
- [3] Leyffer, Sven , Stefan M. Wild, Mike Fagan, Marc Snir, Krishna Palem, Kazutomo Yoshii, Hal Finkel. "Doing Moore with Less -- Leapfrogging Moore's Law with Inexactness for Supercomputing", 2016

Author biography



Oriol Tintó Prims received his degree in physics from the Universitat Autònoma de Barcelona, his Masters degree in Modelling for Science and Engineering from the Universitat Autònoma de Barcelona and is currently doing his PhD in the Earth Sciences department from the Barcelona Supercomputing Center in collaboration

with the Computer Architecture and
Operating Systems department from the
Universitat Autònoma de Barcelona.

ORCHESTRA: An Asynchronous Non-Blocking Distributed GVT Algorithm

Tommaso Tocci

Barcelona Supercomputing Center (BSC), Sapienza University of Rome

tommaso.tocci@bsc.es

Abstract—Taking advantage of high computing capabilities of modern distributed architectures is fundamental to run large-scale simulation models based on the Parallel Discrete Event Simulation (PDES) paradigm. In particular, by exploiting clusters of modern multi-core architectures it is possible to efficiently overcome both the power and the memory wall. This is more the case when relying on the speculative Time Warp simulation protocol. Nevertheless, to ensure the correctness of the simulation, a form of coordination such as the GVT is fundamental. To increase the scalability of this mandatory synchronization, we present in this paper a coordination algorithm for clusters of share-everything multi-core simulation platforms which is both wait-free and asynchronous. The nature of this protocol allows any computing node to carry on simulation activities while the global agreement is reached.



1 INTRODUCTION

A classical technique to achieve high-performance simulation runs is to rely on Parallel-DES (PDES) [1]. PDES grounds on partitioning the simulation model into several distinct objects, known as Logical Processes (LPs), which concurrently execute simulation events, possibly on a distributed environment.

To deliver high-performance simulation runs, a core aspect of PDES systems is synchronization, which ensures causally-consistent (i.e. timestamp-ordered) execution of simulation events at each LP. Several synchronization protocols have been proposed, among which the optimism-oriented ones, such as the Time Warp protocol [2], are a viable solution to tackle simulation performance aspects. In Time Warp, events are processed speculatively, thus significantly exploiting parallelism, while causal consistency is guaranteed through rollback/recovery techniques, which restore the simulation model to a correct state upon the a-posteriori detection of consistency violations. These are originated when LP_a schedules a new event destined to LP_b having a timestamp lower than the one of some event already speculatively processed by LP_b . In case this occurs, the rollback of LP_b might also require undoing the send operation of events that were produced by LP_b during the rolled back portion of the computation. This is done via anti-messages (carrying anti-events), which annihilate the originally-sent events, thus possibly causing cascading rollbacks across chains of LPs.

This high level of independence among different LPs is the key to high-performance simulation. In fact, this execution model tries to capture time independence, without any manual intervention from the simulation model developer. Nevertheless, Time Warp has the need for a global notion of time. In fact, a core abstraction is the Global Virtual Time (GVT), which is defined as the smallest timestamp among events (or anti-events) that are still unprocessed, or that are currently being processed. The GVT allows to identify the

commitment horizon of the speculative simulation run—no LP can ever rollback to simulation time preceding the GVT value [2]. Its value is used both to execute actions that cannot be subject to rollback, such as displaying/inspecting intermediate simulation results [3], [4], and to reclaim memory [5] (the *fossil collection* operation).

To determine the GVT value, some sort of coordination among the computing nodes is required. This coordination can significantly affect the simulation performance, in a way that is directly affected by the organization of the computing environment. This is an aspect that must be explicitly taken into account, since computing architectures have recently hit some physical limits: the *power wall* [6] and the *memory wall* [7] have posed strict limitations on what can be done with out-the-shelf computers.

2 RELATED WORKS

Several GVT algorithm have been proposed in the last 30 years, and usually they are designed for a specific underlying architecture: either focusing on shared-memory (e.g. [8], [9]) or distributed computers. In 2008 Chen and Szymanski [10] presented an in-dept comparison and tried to classify them. As outlined in their study, some of these algorithms rely on assumptions or make use of some properties that limit their scalability. For instance, the usage of message acknowledgments (e.g. [11], [12], [13]) in order to cope with non observability of transient messages, is an obstacle to scalability, because it introduces big overhead in the communication. Mattern proposed a *two-cuts* [14] approach in which by using colors, it is possible to keep track of the transient messages without requiring explicit acknowledgments. Only recently, with the goal of reaching very large-scale simulations, some hybrid algorithms appeared. Lin and Yao [15] proposed the first GVT algorithm targeting a multi-thread PDES platform, that exploits the shared-memory capability of each machine while reaching a global virtual time agreement among them.

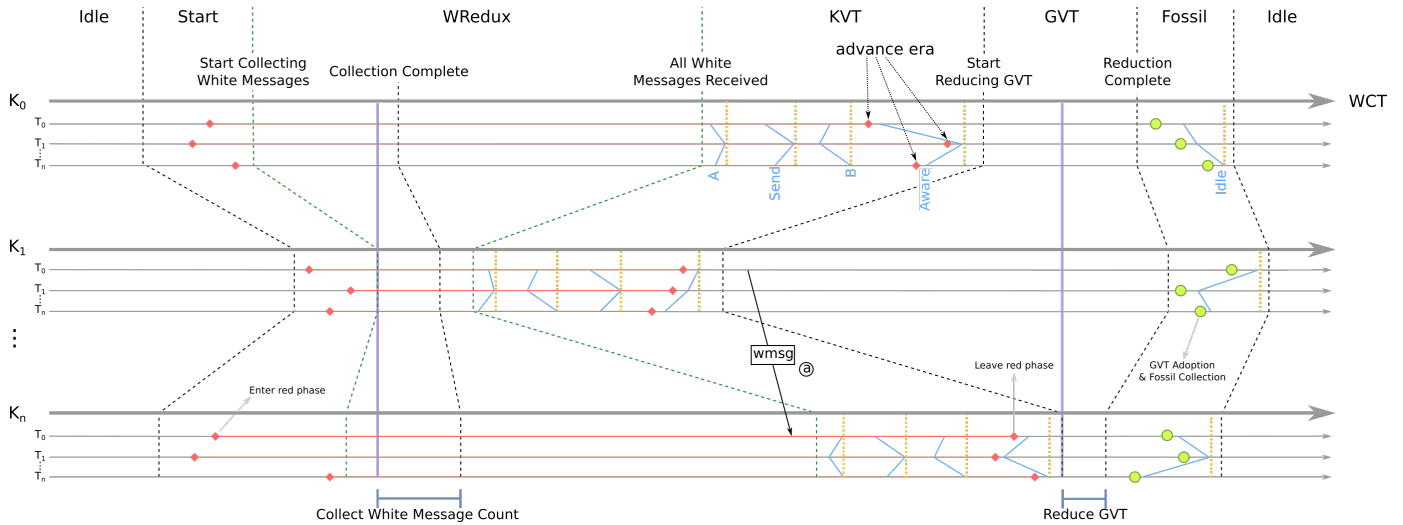


Fig. 1. Representation of phases of the distributed GVT algorithm

3 REFERENCE SYSTEM MODEL

Communication latency and cluster nodes topology can really impact on the simulation performance. To enforce highly-scalable simulations and give them the flexibility to adapt to the underlying network interconnection, it is important to design a truly distributed software architecture and to avoid any use of centralized algorithms.

Each computing node (i.e., a core in the distributed system) is associated with a thread. The LPs are distributed among all the simulation threads in such a way that each of the thread will carry on simulation activities of a subgroup of LPs bound to it. The threads running on the same machine are grouped together into the same simulation kernel and they share a portion of the main memory. Communication among threads running on different kernels need to pass through the interconnection network, while threads living within the same kernel will take advantage of shared-memory for local communication activities.

Unlike the organization used in [15], all the threads in our architecture have exactly the same role, meaning that no special thread is used to carry on communication tasks. This homogeneity between all the worker threads simplify the scheduling process and allows to better distribute the simulation workload among the nodes of the cluster.

4 ASYNCHRONOUS NON-BLOCKING GVT

The algorithm to compute the GVT that we present here works at two different levels. It uses the wait-free algorithm proposed in [9] to reach consensus between threads of a single kernel while implementing a variant of the two cuts Mattern's idea [14] to make all the kernels agree on a common global virtual time. The 2-color scheme [14] is used to solve the well known *transient message problem*. A thread becomes red colored while it is participating on a GVT round, while normally it is white. Every outgoing message takes the color of the sender thread. This coloring scheme allows to include all the messages that are in transit in the current GVT agreement round. In fact, while red messages are always accounted by the sender, the white

ones will be accounted by the receiver. As shown in Figure 1 each thread participating in a GVT round, pass through the phases of the algorithm that are actually determined by a combination of the state of the thread itself and the shared state of the kernel to which it belongs. At the beginning of each GVT round, during the *Start* phase, every thread switch its color to red. After that, every kernel collects the number of white messages that have been sent to it. Once the kernel has verified that no white message destined to it is still in transit, it triggers the start of the local *Kernel Virtual Time (KVT)* algorithm. This local algorithm, executed at each kernel, is the one presented in [9] where, by exploiting shared memory, all the threads reach an agreement upon the KVT. In addition to the original version, during this phase, every thread needs to leave from the red state and enter again into the white one. This allow to collect the minimum timestamps among all the red messages sent and to account for it in the current KVT calculation. Once all the KVTs are locally computed, all the kernels participate to a global reduction in order to elect the minimum among these values as the new GVT.

ACKNOWLEDGMENTS

This study has been carried on in close cooperation with Alessandro Pellegrini, Postdoctoral Researcher at the Department of Computer Science and System Engineering at "Sapienza" University of Rome. All the work has been supervised by Josep Casanovas from the Barcelona Supercomputing Center and supported by Toyotaro Suzumura from T.J. Watson Research Center and visiting professor at BSC. This research has been funded by JST, CREST Strategic Basic Research Programs in collaboration with Barcelona Supercomputing Center (BSC).

REFERENCES

- [1] R. M. Fujimoto, "Performance of Time Warp under synthetic workloads," 1990.
- [2] D. R. Jefferson, "Virtual Time," *ACM Transactions on Programming Languages and System*, vol. 7, no. 3, pp. 404–425, 1985.

- [3] F. Antonacci, A. Pellegrini, and F. Quaglia, "Consistent and efficient output-stream management in optimistic simulation platforms," in *Proceedings of the 2013 ACM SIGSIM Conference on Principles of Advanced Discrete Simulation*, ser. PADS. ACM, 2013, pp. 315–326. [Online]. Available: <http://dl.acm.org/citation.cfm?doid=2486092.2486133>
- [4] D. Cucuzzo, S. D'Alessio, F. Quaglia, and P. Romano, "A Lightweight Heuristic-based Mechanism for Collecting Committed Consistent Global States in Optimistic Simulation," in *DS-RT*, 2007, pp. 227–234.
- [5] S. R. Das and R. M. Fujimoto, "Adaptive Memory Management and Optimism Control in Time Warp," *ACM Transactions on Modeling and Computer Simulation*, vol. 7, no. 2, pp. 239–271, 1997.
- [6] H. Sutter, "The Free Lunch Is Over: A Fundamental Turn Toward Concurrency in Software," *Dr. Dobbs' Journal*, vol. 30, no. 3, pp. 202–210, 2005. [Online]. Available: <http://www.gotw.ca/publications/concurrency-ddj.htm>
- [7] S. a. McKee, "Reflections on the memory wall," *Proceedings of the first conference on computing frontiers on Computing frontiers - CF'04*, p. 162, 2004. [Online]. Available: <http://portal.acm.org/citation.cfm?doid=977091.977115>
- [8] R. M. Fujimoto and M. Hybinette, "Computing Global Virtual Time in Shared-Memory Multiprocessors," *ACM Transactions on Modeling and Computer Simulation*, vol. 7, no. 4, pp. 425–446, 1997.
- [9] A. Pellegrini and F. Quaglia, "Wait-Free Global Virtual Time Computation in Shared Memory TimeWarp Systems," in *2014 IEEE 26th International Symposium on Computer Architecture and High Performance Computing*. IEEE, oct 2014, pp. 9–16. [Online]. Available: <http://ieeexplore.ieee.org/document/6970641/>
- [10] G. G. Chen and B. K. Szymanski, "Time Quantum GVT: A Scalable Computation of the Global Virtual Time in Parallel Discrete Event Simulations," vol. 8, no. 4, pp. 423–435, 2008.
- [11] B. Samadi, "Distributed Simulation Algorithms and Performance Analysis," Ph.D. dissertation, Computer Science Department, University of California, Los Angeles, 1985.
- [12] S. Bellenot, "Global Virtual Time algorithms," in *Proceedings of the SCS Multiconference on Distributed Simulation*, 1990, pp. 122–127.
- [13] R. Baldwin, M. J. Chung, and Y. Chung, "Overlapping window algorithm for computing GVT in Time Warp.pdf," pp. 534–541, 1991.
- [14] F. Mattern, "Efficient Algorithms for Distributed Snapshots and Global Virtual Time Approximation." *Journal of Parallel Distributed Computing*, vol. 18, no. 4, pp. 423–434, 1993.
- [15] Z. Lin and Y. Yiping, "An asynchronous GVT computing algorithm in neuron time warp-multi thread," no. Gillespie 1977, pp. 1115–1126, 2015.

Tommaso Tocci chasing his passion for the informatics, he graduated in computer science at the Sapienza University of Rome. He is concluding the master of Science in Engineering in Computer Science at the same university, focusing on distributed systems. At the moment he is employed at Barcelona Supercomputing Center (BSC), working on a parallel discrete event simulation platform and developing distributed synchronization algorithms.

Level of Detail for Complex Urban Scenes with Varied Animated Crowds, using XML

Leonel Antonio Toledo Díaz^{#1}, Isaac Rudomín^{#1}

[#]Barcelona Supercomputer Center, Barcelona, Spain.

leonel.toledo.diaz@gmail.com, isaac.rudomin@bsc.com

Keywords— Crowd Simulation, GPU, Level of Detail.

EXTENDED ABSTRACT

Crowd simulation has gained attention recently in the movie and video game industry, still there are broader applications in which crowd simulation is associated, transportation and urbanism are some examples.

Even though hardware capabilities are surpassed constantly by each new generation of hardware, the number of polygons, memory and computational resources are not enough to properly simulate dense crowds composed by millions of characters. This cannot be achieved without the aid of rendering techniques such as Level of Detail (LOD). Interactive virtual crowds require high-performance simulation, animation and rendering techniques to handle numerous characters in real-time. These characters must be believable in their actions and behaviors. Real-time crowd simulation brings different challenges compared with systems that involve a small number of interacting characters (i.e. most contemporary computer games) and non-real-time (i.e. crowds seen in movies or visualizations of crowd evacuations after offline model computations). In comparison with single agent simulations, the main conceptual difference is the need for efficient variety management at every level, whether is visualization, motion-control, animation or sound rendering [1].

A. Hybrid Hierarchical LOD System for large crowds using configuration files

We present a system capable of handling several thousands of varied animated characters within a crowd. These characters are designed to have geometric, color animation and behaviour variety, nevertheless when a crowd becomes bigger, more memory is needed and is often difficult to achieve this objective [2]. To solve this problem, we implemented two complementary data structures. The first structure is a skeleton with associated octrees for each limb that is used for applying, transferring animations and generating variety for characters at any level of detail, as shown in figure 1. The second structure is a scene tiling with an associated quadtree that represents the environment. This structure is used for rendering, LOD selection and for combining characters in areas far away from the viewer to further reduce resource consumption, allowing us to handle dense crowds.

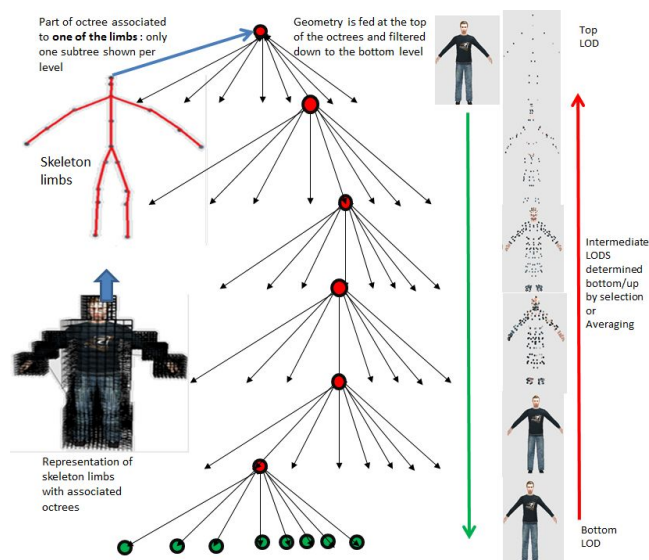


Figure 1: Creation of the skeleton and how LOD is built using our octree system. Each limb has an associated octree and data is stored inside the structure.

Extending this hybrid LOD method, and adapting it to use instructions from newer versions of OpenGL, we are implementing a complete authoring system based on XML configuration files with which we can create virtual environments that include buildings, trees, vehicles and people. We can declare the size of the crowd and create different groups, each of which will have different base geometry, rules for variety, distribution and behavior, and can be based on data that the user can provide to the system.

B. Results

Our LOD method outperforms traditional impostors and the most common variations in terms of memory requirements and/or computation. The work by Rudomin et al. [3] states that 12 megabytes of memory are used for each impostor, whereas the presented method only needs 75,532 kilobytes of memory to simulate a crowd composed of 65,536 characters. This is achieved because variety, animation and LOD are computed directly on the GPU, making possible to generate diverse crowds in real time. In comparison, traditional impostors would need $12 * 65536$ Megabytes of memory to display the same crowd. Our representation needs less memory or computation than any of the traditional impostor methods (at run-time, since there is an initial cost of

generating and filling hierarchical structures when loading the scene). Figure 2 shows the achieved results using our method.



Figure 2: A crowd composed by 65,536 varied and animated characters.

The system is not limited to crowds composed only by humans, it is possible for us to create urban environments and incorporate props into the simulation. Figure 3 shows an example of a simulation created using this method.



Figure 3: A scene created with the system presented in this paper. It is possible to combine different kinds of geometry to create complex urban environments

C. Conclusion and Future Enhancement

Some of the main challenges faced in this work are to remove the least perceptible details for the simulation to preserve the global vision aspects without compromising visual quality and meanwhile, significantly improve computation times. This work is mainly focused on discrete level of detail, but some variations are interesting to study as well, for example adaptive LOD which gives more importance to certain areas of the geometry or the environment, enhancing

visual quality without compromising the quality. Nevertheless LOD applications are broader than geometry simplification, in fact it is possible to make simulations more complex by implementing LOD variations that consider behavior, animation and collision avoidance just to mention a few.

References

- [1] Thalmann D.: Crowd Simulation. John Wiley Sons, Inc., 2007. URL: <http://dx.doi.org/10.1002/9780470050118.ecse676>, doi:10.1002/9780470050118.ecse676.
- [2] Toledo L., De Gyves O., Rudomin I.: Hierarchical level of detail for varied animated crowds. *The Visual Computer* 30, 6-8 (2014), 949-961. URL: <http://dx.doi.org/10.1007/s00371-014-0975-9>, doi:10.1007/s00371-014-0975-9.
- [3] RM06] Rudomin I., Millan E.: A comparison between impostors and pointbased models for interactive rendering of animated crowds.

Author biography



Leonel Toledo received his Ph.D from Instituto Tecnológico de Estudios Superiores de Monterrey Campus Estado de México in 2014, where he was a full-time professor from 2012 to 2014. He was an assistant professor and researcher and has devoted most of his

research work to crowd simulation and visualization optimization. He has worked at the Barcelona Supercomputing Center using general purpose graphics processors for high performance graphics. His thesis work was in Level of detail used to create varied animated crowds. Currently he is a researcher at Barcelona Supercomputer Center.



Isaac Rudomin is a senior researcher at the Barcelona Supercomputer Center, which he joined in 2012. His focus is on crowd rendering and simulation including generating, simulating, animating, and rendering large and varied crowds using GPUs in consumer-level machines and in HPC heterogeneous clusters with GPUs. Previously, Isaac was on the faculty at

Tecnologico de Monterrey Campus Estado de Mexico (from 1990 to 2012). He finished his Ph.D. at the University of Pennsylvania under Norman Badler on the topic of cloth modeling.

Uncertainty in near-surface wind speed trends at seasonal time scales

Verónica Torralba^{#1}, Francisco J. Doblas-Reyes^{*2} and Nube Gonzalez-Reviriego^{#3}

[#]Earth Sciences Department, Barcelona Supercomputing Center (BSC), Barcelona, Spain

^{*}Institució Catalana de Recerca i Estudis Avançats (ICREA), Barcelona, Spain

¹veronica.torralba@bsc.es, ²francisco.doblas-reyes@bsc.es, ³nube.gonzalez@bsc.es

Keywords— wind speed variability, trends, reanalyses, seasonal

EXTENDED ABSTRACT

Observational studies have identified wind speed trends in the last decades [1,2] attributed to several factors such as changes in the land use, aerosol emissions or atmospheric circulation. However, in spite of the potential impact of this long-term variability in wind energy activities, this type of variability has not been fully characterized yet. As a consequence such information is not currently incorporated in wind power decision-making processes related to planning and management.

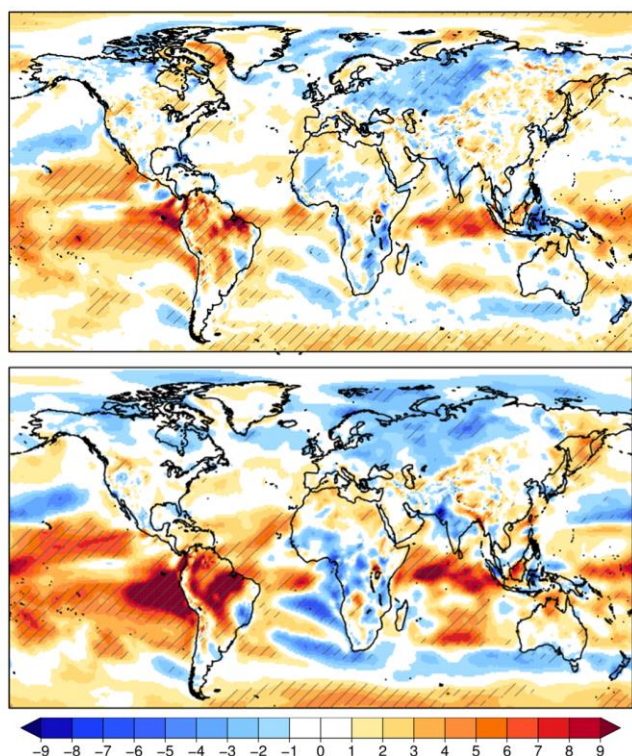


Fig. 1 Normalized linear trend (% per decade) calculated as the linear trend of ERA-I 10-m wind speed divided by the seasonal climatology of 10-m wind speed (top) and 850 hPa wind speed divided by the seasonal climatology of 850 hPa wind speed (bottom) over the period 1980-2015 in the December-January-February season. Hatched regions indicate where the trends are significant at the 95% confidence level.

The main limitation for the assessment of wind speed trends is the unavailability of long enough, homogeneous time series of historical data from observational measurements. These data products can be affected by discontinuities associated with changes in the measuring equipment or in the observing practices [3] which impacts on the data quality. To overcome these limitations, wind energy users have recently incorporated reanalysis products for the evaluation of the

long-term wind speed variability. For some of these users it is still difficult to identify the most suitable dataset for their specific needs, because a comparison of the quality of the wind speed data from different reanalyses at global scale is not readily available. For this reason, the present study investigates the wind speed long-term trends at global scale in the last decades (1981-2015) using three state-of-the-art reanalyses: ERA-Interim (ERA-I), the Japanese 55-year Reanalysis (JRA-55) and Modern Era Retrospective-Analysis for Research and Applications (MERRA-2).

Strong seasonal and regional variability of the wind speed trends have been identified. In the boreal winter (Fig. 1, top) an increase of the wind speed over the oceans, particularly in the tropical regions, and a decline in some continental areas such as Europe or India are shown. To investigate if these trends can be due to changes in the atmospheric circulation or to other forcings, the seasonal wind speed trends have been also computed at the level of 850 hPa (Fig. 1, bottom). The agreement among the trends at the two analysed levels illustrates the link between the trends in both near-surface and the atmospheric circulation, which seems to be one of the main drivers of the near-surface wind speed trends.




Fig. 2 Time series of the 90th (pink line), 50th (orange line) and 10th (green line) percentiles of the 10-m wind speed from ERA-I for December-January-February over the period 1980-2015 for a location in Poland [51.7° N and 18.9° E].

The characterization of the near-surface extreme wind speeds can provide extra information about the long-term changes in the frequency of unusual events in a particular season. The analysis of the seasonal extreme events is based on two indices, the 10th and 90th percentiles. These indices are based on 6-hourly near-surface wind speed data. Results for the DJF season in a particular point in Poland [51.7°N- 18.9°E] (Fig. 2) illustrates that the trend of the 90th percentile (-0.2173 m/s per decade) of the wind speed is twice higher than the 10th percentile (-0.112 m/s per decade). This indicates that the decline of the higher wind speeds is faster than the reduction in the lower values of wind speed.

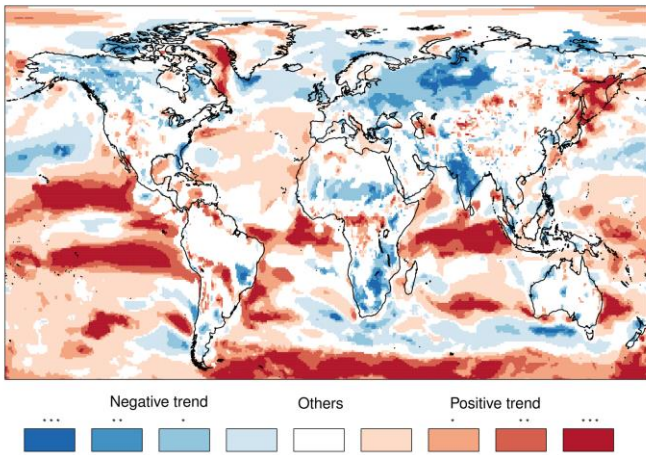


Fig. 3 Comparison of the 10-m wind speed trends obtained from ERA-I, JRA-55 and MERRA-2. Blue (reds) colours indicate agreement between the three reanalyses about the negative (positive) trends of the wind speeds for December-January-February in the period 1980-2015. Asterisks indicate that the trends in one (*), two (**), or the three (***) reanalyses are significant at the 95% confidence level.

The discrepancies and similarities of the wind speed trends from each reanalysis in DJF are summarized in Fig. 3. The inter-comparison of the near-surface wind speed trends shows agreement in many regions. The three reanalyses show an increase of the wind speed values in Tropical areas and a wind speed decline is found in Eurasia. The agreement in the sign of the trends in those regions among the three datasets indicate that these trends are robust and they should be caused by changes in the atmospheric circulation since other possible factors, such as the time changes in land use or aerosols, are considered differently in these reanalyses.

Despite the existence of many regions where the trends are consistent in the three reanalyses, important differences have also been identified. The most important discrepancy has been found for the JRA-55 reanalysis that show intense negative wind speed trends over land. This problem is caused by a negative bias that appears in the interpolation of the near-surface wind speed from the atmospheric lowermost level. The impact of the discrepancies of the wind speed trends in different reanalyses can lead to inconsistencies in the evaluation of long-term wind power estimations that use reanalyses information for planning and management of wind farms. For this reason these results should be taken into account by users that employ these datasets in their decision-making processes.

Finally we have explored the near-surface wind speed trends in the European Centre for Medium-Range Weather operational seasonal prediction system, System 4, to investigate if this system is able to reproduce the trends that have been observed in the reanalyses. The seasonal predictions of the near-surface wind speed show weakened trends, and in some regions the sign of the trends is different to the wind speed trends in the reanalyses. This is an important aspect from the forecast verification point of view because different trends in the predictions and in the reference dataset could lead to an important skill reduction.

References

- [1] Vautard, Robert, et al. "Northern Hemisphere atmospheric stilling partly attributed to an increase in surface roughness" *Nature Geoscience* 3.11 (2010): 756-761.
- [2] Bichet, A., et al. "Causes for decadal variations of wind speed over land: Sensitivity studies with a global climate model" *Geophysical Research Letters* 39.11 (2012).
- [3] Kaiser-Weiss, A. K., et al. "Comparison of regional and global reanalysis near-surface winds with station observations over Germany" *Advances in Science and Research* 12.1 (2015): 187-198.

Author biography



Verónica Torralba studied her Degree in Physics and a Master in Meteorology and Geophysics at Complutense University of Madrid (UCM) in Spain, where she acquired a background in climate sciences. She joined the Climate Forecasting Unit (CFU) at the Catalan Institute of Climate (IC3) in October 2013 where

she worked in the adaptation of the existing tools in the CFU unit to create the necessary material for the development of climate services for the wind energy sector. She is currently working in the Earth Sciences Department at the BSC where she is a PhD student involved in different European Projects (EUPORIAS, SPECS, NEWA) and national initiatives (RESILIENCE). The overall aim of her PhD is the quality assessment of wind speed seasonal predictions and the application of different statistical post-processing methods. This information will be tailored to specific needs from the wind energy industry.

Effect of Terrain Relief on Dust Transport over Complex Terrains in West Asia

Lluís Vendrell^{*1}, Sara Basart^{*2} and José María Baldasano^{##3}

^{*}Barcelona Supercomputing Center-Centro Nacional de Supercomputación (BSC-CNS), Barcelona, Spain

^{##}Environmental Modelling Laboratory, Technical University of Catalonia (UPC), Barcelona, Spain

¹lluis.vendrell@bsc.es, ²sara.basart@bsc.es, ³jose.baldasano@bsc.es

Keywords— Mineral dust, dust transport, NMMB/BSC-Dust, dust modelling, horizontal resolution and topography

EXTENDED ABSTRACT

This work investigates the impact of orography on dust transport using the multi-scale NMMB/BSC-Dust model. For this purpose, two model simulations at horizontal resolutions of $0.03^\circ \times 0.03^\circ$ (Low-resolution; LR) and $0.3^\circ \times 0.3^\circ$ (High-resolution; HR) are performed and analysed covering two intense dust storms that occurred in West Asia in March 2012. Differences between both simulations emerge when the dust storms reach the south and west Arabian Peninsula where its complex topography affected meteorology and dust fields in many ways. The HR simulation is better than the LR simulation at reproducing the topography and its topographic effects on meteorology, such as developing orographic clouds, wind speed bias reduction under the dust flows (larger than 5 m/s) and more accurate wind directions, as well as on dust fields, such as a more realistic representation of dust channeling/blocking. Consequently, it improves dust forecasts in the vicinity of complex terrains.

I. Introduction

Mineral dust is one of the most abundant atmospheric aerosols. Dust particles uplifted by strong winds from arid or semiarid regions can be injected into the atmosphere and, under favourable conditions, transported over thousands of kilometres away. Mineral dust plays an important role in the Earth system due to its impacts on radiation, clouds, precipitation, atmospheric chemistry, ecosystems, biogeochemical cycles and human health (e.g. respiratory and cardiovascular diseases). Dust storms can reduce visibility to a few meters and negatively impact goods and human activities by causing serious hazards in road and air transportation, reducing commercial solar energy production, and damaging crops and livestock [1].

In this framework, dust models have many applications on the Earth system because they are a powerful tool for predicting and simulating the dust cycle and its interaction in the climate-weather system, estimating the global or regional dust budgets, as well as complementing observations and improving our knowledge about dust processes.

The present work investigates the NMMB/BSC-Dust model's ability to reproduce dust transport in the vicinity of complex terrains. In order to do that, we perform two NMMB/BSC-Dust simulations at different horizontal resolutions covering two intense dust storms that occurred on 17-20 March 2012 and spanned over thousands of kilometres in West Asia where its topography affected dust propagation in many ways. Based on the literature, this dust event is considered among the most powerful ones [2]. The model results will be compared against ground-based (AERONET and weather) observations and satellite aerosol products (Aqua/MODIS and MSG/SEVIRI).

II. The NMMB/BSC-Dust model and model setup

The NMMB/BSC-Dust model [3]–[5] is one of the dust models in the Earth Sciences Department of the Barcelona Supercomputing Center (ES-BSC). This dust model is composed by the online coupling of the non-hydrostatic multiscale atmospheric NMMB model and the BSC-Dust dust module, which provides a unique framework to simulate and/or predict dust and meteorological fields at a wide range of spatio-temporal scales. The model provides daily operational dust forecasts over North Africa, the Middle East and Europe (NAMEE) at the Barcelona Dust Forecast Center (<http://dust.aemet.es/>), the first regional specialized meteorological center with activity specialization on atmospheric sand and dust forecast.

In this study, two NMMB/BSC-Dust configurations as LR and HR are performed; their main differences are shown in Table 1 and their main common features are explained below.

Table 1 The NMMB/BSC-Dust model configurations and their main features used in the present study

Features	Model configurations	
	LR (Low-resolution)	HR (High-resolution)
Grid Spacing	$0.33^\circ \times 0.33^\circ$ ($\sim 33\text{km}^2$ in the equator)	$0.03^\circ \times 0.03^\circ$ ($\sim 3\text{km}^2$ in the equator)
Grid Points	211x307	1001x1001
Box Domain (Lat, Lon)	$0^\circ - 70^\circ\text{N}$ $31^\circ\text{W} - 71^\circ\text{E}$	$10^\circ\text{N} - 40^\circ\text{N}$ $35^\circ\text{E} - 65^\circ\text{E}$
Timestep (s)	40	25
NetCDF (Gb/day of simulation)	~ 1.5	~ 15

The simulation period is 10 – 21 March 2012, and consists of daily forecasts (initialized at 0 UTC) with model outputs saved every 3 h. The initial meteorological state is supplied by the NCEP/Final Analyses (FNL; at $1^\circ \times 1^\circ$ horizontal resolution) and boundary conditions are updated every 6 h. The vertical resolution for both simulations is 40 σ -hybrid layers with the top of the atmosphere at 50 hPa. Dust concentration at 0 UTC is defined by the value at hour 24 of the previous day's dust forecast, except at 0 UTC 10 March when dust concentration is set to zero (cold start). The first six days are discarded and are only used as a warm up of the simulation.

III. Results and discussion

On 17 March, a dust storm originated in Iraq as a result of strong north-westerly Shamal winds and rapidly extended towards the Arabian Gulf, uplifting and transporting a dust cloud that dropped visibility below 500 m in its path. The dust storm then moved to the mountain ranges in the south and west Arabian Peninsula where it was initially blocked on 18 March, and between 18 and 19 March, satellite images observed dust plumes channeling through valleys between the




Fig. 1 The simulated dust load (g/m^2) for both NMMB/BSC-Dust configurations as LR (left panel) and HR (middle panel) at 0 UTC on 19 March 2012 and satellite MODIS Terra True Color snapshot (right panel) over the western Arabian peninsula and the Red Sea.

near-coast mountain ranges towards the Red Sea (Fig. 1) and the Arabian Sea. Both configurations reproduce the emission and transport of dust from Iraq to the mountain ranges in the south and west Arabian Peninsula, although overestimations in dust fields are found in comparison with observations (i.e. AERONET sun-photometers and Aqua/MODIS AOD) in the eastern Arabian Peninsula associated with a wind speed overestimation. Furthermore, both present a delay in the arrival of the dust front to those mountain ranges of about 8 h due to a wind speed underestimation in the inner Arabian Peninsula. Both presented wind speed biases are partly caused by the NCEP/FNL initial conditions. When the dust front reached the west and south Arabian Peninsula, both configurations represent the initial dust blocking by the topography, although the HR, in comparison with the LR, more realistically reproduces the steep mountains along the region (e.g. height and valleys). Compared to the satellite, the HR allows for a better representation of orographic clouds, dust blocking and dust channeling trough valleys (Fig. 1), and when compared to weather sites, the HR shows wind speed bias reduction under the dust flows (larger than 5 m/s) and more accurate wind directions.

IV. Conclusion

This work demonstrates how increasing horizontal resolution in dust simulations improves dust forecasts in the vicinity of mountains. This is an important result because most of the current operational dust models work at lower horizontal resolutions than the HR, due to its lower computational cost, and it consequently leads those models to regionally decrease their performance over complex topographies.

ACKNOWLEDGEMENT

The authors of this study acknowledge the contributions of the “Supercomputación and e-ciencia” project (CSD2007-0050) from the Consolider-Ingenio 2010 program, the CICYT project (CGL2013-46736) and Severo Ochoa (SEV-2011-00067) program of the Spanish Government. We would like to thank K. Serradell and F. Benincasa for their technical support, H. Becker for their comments on this work. The model simulations were performed on the MareNostrum supercomputer hosted by BSC.

References

[1] P. Knippertz and J.-B. W. Stuut, *Mineral Dust*.

- Dordrecht: Springer Netherlands, 2014.
- [2] P. Jish Prakash, G. Stenchikov, S. Kalenderski, S. Osipov, and H. Bangalath, “The impact of dust storms on the Arabian Peninsula and the Red Sea,” *Atmos. Chem. Phys.*, vol. 15, no. 1, pp. 199–222, 2015.
- [3] C. Pérez, K. Hausteijn, Z. Janjic, O. Jorba, N. Huneus, J. M. Baldasano, T. Black, S. Basart, S. Nickovic, R. L. Miller, J. P. Perlwitz, M. Schulz, and M. Thomson, “Atmospheric dust modeling from meso to global scales with the online NMMB/BSC-Dust model - Part 1: Model description, annual simulations and evaluation,” *Atmos. Chem. Phys.*, vol. 11, no. 24, pp. 13001–13027, 2011.
- [4] K. Hausteijn, C. Pérez, J. M. Baldasano, O. Jorba, S. Basart, R. L. Miller, Z. Janjic, T. Black, S. Nickovic, M. C. Todd, R. Washington, D. Müller, M. Tesche, B. Weinzierl, M. Esselborn, and A. Schladitz, “Atmospheric dust modeling from meso to global scales with the online NMMB/BSC-Dust model - Part 2: Experimental campaigns in Northern Africa,” *Atmos. Chem. Phys.*, vol. 12, no. 6, pp. 2933–2958, 2012.
- [5] S. Basart, L. Vendrell, and J. M. Baldasano, “High-resolution dust modelling over complex terrains in West Asia,” *Aeolian Res.*, vol. 23, pp. 37–50, 2016.

Author biography



Lluís Vendrell (LV’85) was born in Barcelona, Spain, in 1985. He received a Bachelor’s Degree in Physics (2010) and a Master’s Degree in Meteorology (2011) both at the University of Barcelona. He received the Obra Social Fundación la Caixa – Severo Ochoa fellowship (2013) to do a PhD at the ES-BSC. Since then, he has been working in

modelling the dust cycle, with a special focus on high-resolution models. His current research interests include dust modelling, high-resolution modelling, and meteorological processes that cause dust emissions. He was awarded as the best poster presentation in the 8th International Workshop on Sand/Duststorms and Associated Dustfall (2016), Lisbon, Portugal.

Performance Analysis on the Intel Knights Landing Architecture

Michael Wagner, Judit Gimenez

Index Terms—HPC, performance analysis, performance optimization, tools, Intel Knights Landing

High performance computing (HPC) systems provide enormous computational resources. But the increasing performance introduces more and more complexity, as well. Current leading edge HPC systems consist of millions of heterogeneous processing elements. They require consideration of parallel execution, network, system topology, and hardware accelerators as well as a variety of different parallel programming models such as message passing (MPI), threading and tasking (OpenMP), one-sided communication (PGAS), and architecture specific models to incorporate hardware accelerators such as GPUs. As a result, appropriate support tools have become inevitable in the development process.

Performance analysis tools assist developers not only in identifying performance issues within their applications but also in understanding their complex parallel behavior. Thereby, the development of such tools is facing their own challenges. On the one hand, there are challenges specific to the tools such as techniques to gather appropriate information from the running application while introducing minimal bias, the post-processing of the collected data to correctly correlate data from a plenitude of different data source, and finally methods to present the data or derived results in an appropriate manner to allow an accessible and meaningful analysis. On the other hand, the tools encounter challenges of similar nature as the application the support, for instance, scaling in all aspects (number of cores, file systems, parallel analysis operations) to current HPC system sizes, utilize and capture a variety of parallel programming models and concepts, and operate on and support new arising architectures.

One of the emerging architectures in HPC systems is Intel's Knights Landing (KNL) many core chip, which will also be part of BSC's next HPC installation MareNostrum 4. KNL is the code name of the second generation of Intel XEON Phi, a many integrated core architecture (MIC) with up to 72 cores with four-time hyper-threading. It includes up to 384 GB of DDR4 RAM and 8-16 GB of stacked MCDRAM, a version of high bandwidth memory. In addition, each core will have two 512-bit vector units and will support AVX-512 SIMD instructions.

To support HPC software developers at BSC and beyond in utilizing the capabilities of this new architecture, the BSC performance analysis tools (<https://tools.bsc.es>) are ready to record and analyze parallel applications on the new Knights Landing architecture. This presentation will demonstrate the functionality and capabilities of the tools on the KNL architecture. Based on a case study of an FFT kernel from

an electronic-structure calculations and materials modeling simulation framework, the presentation will highlight the application recording as well as analysis process starting from identifying the most severe performance issues using a parallel performance model, analyzing and understanding the uncovered performance issues with increasing levels of detail, and finishing with the advanced techniques to cluster and track different application phases with increasing scale. In addition, specific suggestions for potential code optimizations and a quick review of a proof-of-concept implementation of these suggestions conclude the presentation.

Within POP, the European Performance Optimisation and Productivity Centre of Excellence in Computing Applications (<https://pop-coe.eu/>) the tools and expertise are used to help customers to improve performance in various aspects of HPC: Code developers are assisted with an assessment of detailed actual behavior and suggestions of most productive directions to code re-factoring. Application users are supported with an assessment of achieved performance in specific production conditions, possible improvements modifying environment setup and evidence to interact with code providers. Finally, infrastructure operators benefit from an assessment of achieved performance in production conditions, possible improvements from modifying environment setup, information for scheduling and process allocation, as well as training of support staff.



Michael Wagner received the M.Ed. degree in mathematics and computer science and the Ph.D. degree in computer science from the Technische Universität Dresden, Dresden, Germany, in 2008 and 2015, respectively. He was a research scientist in performance analysis at the Center for Information Services and HPC (ZIH), Dresden, Germany and is currently a HPC application analyst at the Barcelona Supercomputing Center (BSC), Barcelona, Spain. His research interests are in high performance computing with emphasis on performance analysis.

SMuFin2: generating and implementing new and more efficient search engines integrated in particular hardware architectures


Mercè Planas-Fèlix, Jordà Polo, Alejandro Barberá-Mourelle, Nicola Cadenelli, David Carrera and David Torrents



Mercè Planas Fèlix: I was born in Barcelona in 1988. I received the Bachelors degree in Biology from the “Univeristat Auònoma de Barcelona - UAB”, Barcelona 2012. Then I performed a master in Bioinformatics for Genomics and drug design in UAB,Barcelona 2013. At the same time I was earning my degree, I did a bachelor’s internship and received a grant holder in the Evolutionary group at Genetics and Microbiology department, also I did a Bachelors internship at “Catalan Oncology Institute – ICO” on a Bioinformatics framework. I accomplished the master practices at the Structural Genomics groups led by Marc Marti-Renom at “Centre Nacional d’anàlisi genomic - CNAG”. Since September 2013 I am doing my PhD in Biomedicine (Bioinformatics) at Computational genomics group led by David Torrents at the BSC with a La Caixa fellowship. My thesis is involved in the improvement of SMUFIN generating and implementing new and more efficient search engines integrated in particular hardware architectures.

GUIDANCE: An Integrated Framework for Large-scale Genome and Phenome-Wide Association Studies on Parallel Computing Platforms


Marta Guindo-Martínez, Friman Sánchez, Silvia Bonàs-Guarch, Montserrat Puiggròs, Jorge Ejarque², Carlos Díaz², Enric Tejedor², Rosa M. Badia^{2,3}, Josep M. Mercader, David Torrents



Marta Guindo-Martínez: I was born in Barcelona, Spain, in 1988. I received the degrees in Biology and Biochemistry from the Universidad de Navarra (UNAV), Pamplona, Spain, in 2012, and the M.Sc. degree in Biomedical Research from the Universitat Pompeu Fabra (UPF), Barcelona, Spain, in 2013. Since February 2014, I am a PhD Student in the Life Sciences department at Barcelona Supercomputing Center (BSC). My research is focused on the genetic behind complex diseases and I am analyzing publicly available GWAS data for thousands of individuals in order to better understand why some people develop particular complex diseases and others do no.

Identification and Characterization of Recurrent Deletions in the Human Genome Promoted by Expression of Transposase-Derived Gene Across Different Tumor Types

Elias Rodriguez-Fos, Santi González, Montserrat Puiggròs, Anton G. Henssen, Alex Kentsis and David Torrents



Elias Rodriguez-Fos was born in Barcelona, Spain in 1986. I received the Bachelor’s degree in Biology with the speciality in Genetics and Cell Biology at the Universitat Autònoma de Barcelona (UAB) in 2012. I earned a Postgraduate degree in Bioinformatics at the Universitat Oberta de Catalunya (UOC) in 2014 and a Master’s degree in Genetics and Genomics at the Universitat de Barcelona (UB) in 2015. In 2014-2015, I did my internship in Dr. David Torrents group developing a new method

to perform the analysis of genomic data from Genome-wide association studies for complex diseases. Since 2016, I started my PhD in Biomedicine (Bioinformatics) in the Computational genomics group led by Dr. David Torrents at the BSC. My thesis is involved in the Identification and characterization of complex chromosomal rearrangements driven by transposase-derived genes across different tumor types.

Functional implications of the structural genomic rearrangements in cancer

Luisa Delgado-Serrano, David Torrents
luisa.delgado@bsc.es, david.torrents@bsc.es

Luisa Delgado-Serrano was born in Bogotá, Colombia in 1985. She received a BSc degree in Biology and Microbiology from Los Andes University Bogotá, Colombia in 2009, and a MSc degree in Biological Sciences from the same University in 2014. She was involved in different projects in the area of microbial ecology as bioinformatics specialist from 2009 to 2015. In July 2016, Luisa joined the Computational Genomics group of the BSC as PhD student. Her current research is in the context of the Pancancer Analysis of Whole Genomes study, and is focused on finding patterns in the structural variants associated with cancer diseases. She is particularly interested in studying chromosomal rearrangements in tumour genomes in order to determine their effect on genes, and how these rearrangements contribute to the development and progression of these human diseases.

Classification of retrotransposition during somatic variant calling in cancer genomes

Valls-Margarit J., Torrents D.
jordi.valls2@bsc.es, david.torrents@bsc.es



Jordi Valls Margarit was born in Capellades, Catalonia, Spain in 1988. He received the degree in Biology from the University of Barcelona, Catalonia, Spain in 2011 and the Master of Biomedicine at same university in 2014. Since June 2016, he is doing his PhD thesis at the Computational Genomics Group, at the Barcelona Supercomputing Center, where he studies the somatic structural variance in cancer genomes, in order to understand the different mechanisms that can be related to the process of oncogenesis.



Poster Abstracts

Supporting Real-Time Visual Analytics in Neuroscience

Enrique Arriaga-Varela¹, Javier A. Espinosa-Oviedo²³, Genoveva Vargas-Solar²³⁴, Rogelio Dávila Pérez¹

¹Universidad Autónoma de Guadalajara (UAG), Mexico

²Barcelona Supercomputing Centre (BSC-CNS), Spain

³French-Mexican Laboratory of Informatics and Automatic Control (LAFMIA), France

⁴French National Center for Scientific Research (CNRS), France

{enrique.arriaga, rogelio.davila}@edu.uag.mx, {gvargas, espinosa}@imag.fr

Keywords— Stream Processing, Data Visualization, Visual Databases

EXTENDED ABSTRACT

Neuroscience, as other experimental sciences, supports, refutes or validates hypotheses by conducting experiments on *living organisms*. For instance, by connecting electrical sensors to a cat's spinal cord, and monitoring its neurons activities, neuroscientists can determine whether capsaicin¹ (*chili pepper active component*) has the same effect as anesthesia in the presence of pain [1].

In a typical neuroscience experiment, a neuroscientist is responsible of: (i) preparing the subject; (ii) connecting and calibrating sensors; (iii) collecting and storing the experiment data (e.g. file, database); (iv) applying algorithms and statistics for discovering meaningful patterns. Because of the volume of the resulting data, and the complexity of the algorithms used for finding patterns, the data analysis is usually done *post-mortem*. Yet, given the complexity of setting an experiment (e.g., special and expensive equipment, juridical protocols concerning experiments using animals, gathering together field experts), its duration (e.g., 8 hours), and the fact that each experiment is unique (e.g., every subject has its own characteristics), neuroscientists require novel tools for processing and exploring data in real-time so that they can better control the progress of an experiment.

Although there has been a lot of progress in the domain of automatic knowledge discovery (e.g., deep learning), we believe humans play a central role in the data analysis task. Therefore, we propose to build a *visual stream processing system* for supporting the analysis and exploration of *data streams* in real-time, by exploiting human's natural ability for discovering patterns. Our work combines stream processing and data storage techniques [2, 3], with data visualization theory. We study strategies for visualizing different types of data considering constraints related to real time and data volume. This paper presents our stream processing and visualization system adapted to the requirements of the neuroscience domain.

A. State of the Art

Current solutions for visualizing data depend on traditional DBMSs for storing and retrieving raw data, and the use of custom visualization tool to process and render it [4]. For instance, ScalaR is a 3-layer based visualization system (GUI, web server, database) that dynamically performs resolution reduction when the expected result of a DBMS query is too large to be effectively rendered on a screen [5]. Instead of running the original query, ScalaR inserts aggregation, sampling or filtering operations to reduce the size of the result before plotting it. A similar example is ForeCache [6], a general-purpose tool for exploratory browsing of large datasets based on a lightweight browser interface, and a DBMS running

on a back-end server. For improving response times, ForeCache introduces the use of a cache system, for pre-fetching data as the user explores a dataset. Finally, Tableau² is an interactive interface to general OLAP queries.

As stated in [7], the decouple database-visualization tool has three main drawbacks: (i) the database is unaware of related queries and may recomputes the same results (e.g., slightly panning a map will issue a query to recompute the entire map, though most results are unchanged); (ii) visualization tools duplicate basic database operations, such as filtering and aggregation; (iii) visualization tools assume that all raw data and metadata fit entirely in memory, which is not the case for large datasets. There are some initial results for building interactive real-time visualizations over data streams [8, 9]. Yet, these works focus mainly on visualizing time-series. We study the full spectrum of data types (temporal, tabular, geo-spatial) and data visualization techniques, and propose a general-purpose visualization processing system adapted to the requirements of the neuroscience domain.

B. Visual Stream Processing System

Figure 1 gives an overview of our approach. In the figure, data are collected during a neuroscience experiment and continuously transmitted to our system for being processed in real-time. Then, depending on the type of analysis that a neuroscientist wants to conduct, he/she (i) defines queries using a set of operators and (ii) chooses the kind of visual representation that he/she requires. For instance, in our approach a neuroscientist can group the data into temporal windows of 1h. Then, for each window, he/she can choose different types of visualizations (e.g., point chart, histogram, start plot) for analyzing the correlation among the collected data.

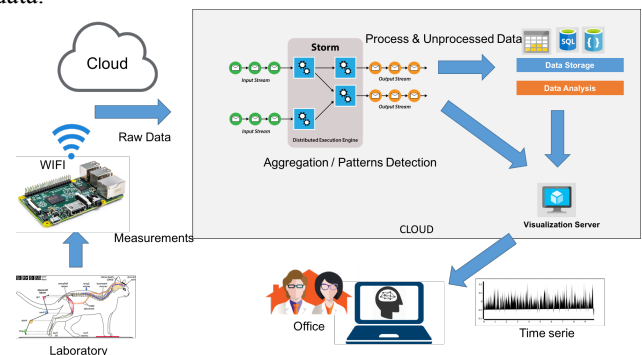


Fig. 1 Visual stream processing system for analyzing and visualizing data streams.

Our system is based on the notion of stream operators (e.g., fetch, sliding window, average, etc.). Figure 2 shows the general architecture of a stream operator. As shown in the figure, an operator communicates asynchronously with other

¹ <https://en.wikipedia.org/wiki/Capsaicin>

² <http://www.tableausoftware.com>

operators using a message oriented middleware. As data is produced, the operator fetches and copies the data to an internal buffer. Then, depending of the operator' logic, it applies an algorithm and sends the data to the next operator. In our approach, there is one operator for collecting the data from an experiment and one for plotting the data. The idea is that a neuroscientist defines the sequence of operators for processing and then plotting the data.




Fig. 2 Architecture of a visual stream processing system for analyzing and visualizing data streams.


C. Implementation and current work

We have implemented a first version of data stream operators and conducted an experimental validation using data from studies regarding pain. We are currently evaluating the capacity of the system for addressing data volume (with respect to online memory consumption) and data processing performance, while visualizing time-series continuously.

References

- [1] M. Martin et al., “A machine learning methodology for the selection and classification of spontaneous spinal cord dorsum potentials allows disclosure of structured (non-random) changes in neuronal connectivity induced by nociceptive stimulation,” *Frontiers in Neuroinformatics*, vol. 9, Aug. 2015.
- [2] T. Pelkonen, S. Franklin, J. Teller, P. Cavallaro, Q. Huang, J. Meza, and K. Veeraraghavan, “Gorilla: a fast, scalable, in-memory time series database,” *Proc. of the VLDB Endowment*, vol. 8, no. 12, pp. 1816–1827, Aug. 2015.
- [3] M. Zaharia, T. Das, H. Li, T. Hunter, S. Shenker, and I. Stoica, “Discretized streams: fault-tolerant streaming computation at scale,” in *Proc. of the 24th ACM Symposium on Operating Systems Principles (SOSP '13)*, 2013, pp. 423–438.
- [4] S. Idreos, O. Papaemmanouil, and S. Chaudhuri, “Overview of Data Exploration Techniques,” in *Proc. of the 2015 ACM SIGMOD International Conference*, 2015, pp. 277–281.
- [5] L. Battle, M. Stonebraker, and R. Chang, “Dynamic reduction of query result sets for interactive visualizaton,” in *Proc. of the 2013 IEEE International Conference on Big Data*, 2013.
- [6] L. Battle, R. Chang, and M. Stonebraker, “Dynamic Prefetching of Data Tiles for Interactive Visualization,” in *Proc. of the 2016 Int. Conference on Management of Data (SIGMOD '16)*, 2016, pp. 1363–1375.
- [7] E. Wu, L. Battle, and S. R. Madden, “The Case for Data Visualization Management Systems,” *Proceedings of the VLDB Endowment*, vol. 7, no. 10, pp. 903–906, Jun. 2014.
- [8] E. Wu, F. Psallidas, Z. Miao, H. Zhang, and U. C. Berkeley, “Combining Design and Performance in a Data Visualization Management System,” in *Proc. 8th Biennial Conference on Innovative Data Systems Research (CIDR '17)*, 2017.
- [9] J. Traub, N. Steenbergen, P. Grulich, T. Rabl, and V. Markl, “I²: Interactive Real-Time Visualization for Streaming Data,” in *Proc. of the 20th Int. Conference on Extending Database Technology (EDBT '17)*, 2017, pp. 526–529.

Author biography




Enrique Arriaga-Varela is master student at Universidad Autónoma de Guadalajara, in the Computer Science program. He obtained his bachelor in 2009 from Universidad del Valle de Mexico. In 2014 he realized a research internship at the laboratory of computational intelligence of

Ritsumeikan University in Shiga, Japan. He is currently research intern in the Barcelona Supercomputing Centre partially funded by the Mexican Council of Scientific Research (CONACyT). He worked previously as a software developer at Oracle Mexico.



Javier Espinosa is a postdoctoral research fellow at Barcelona Supercomputing Center (BSC) and member of the French-Mexican Laboratory of Informatics and Automatic Control (LAFMIA). He holds a PhD in Computer Science from University of Grenoble, France. His

research concerns databases and distributed systems. He is interested on Internet Technologies (e.g., Service-Oriented Architectures, Cloud Computing, Data Services) and NoSQL solutions for modern data management.



Genoveva Vargas-Solar received her first PhD on Computer Science from University Joseph Fourier and her second PhD from University Stendhal. She obtained her *Habilitation à Diriger des Recherches* (HDR - tenure) from University of Grenoble. Her research

interests in Computer Science concern distributed and heterogeneous databases, reflexive systems and service based database systems. She contributes to the construction of service based database management systems.



Rogelio Dávila Pérez is full professor at Universidad Autonoma de Guadalajara. He received his PhD and Master on Computer Science from University of Essex in 1996 and 1987 respectively. His research

interests concern Natural Language Processing (Spanish), Logic Programming, Semantic Web and Artificial Intelligence. He is founding member of the Mexican Laboratory of Advance Informatics (LANIA).

Advanced Vector Architectures for Future Applications

Adrián Barredo-Ferreira, Juan M. Cebrian, Miquel Moretó, Marc Casas, Mateo Valero
E-mail: {adrian.barredo, juan.cebrian, miquel.moreto, marc.casas, mateo.valero}@bsc.es
Barcelona Supercomputing Center

Abstract—Moore's Law predicted that the number of transistors on a chip would double approximately every two years. However, this trend is about to arrive at an impasse. Optimizing the usage of the available transistors within the thermal dissipation capabilities of the packaging is a pending topic. Different solutions have been proposed to deal with this problem, including multi-processors. Nevertheless, managing coarse-grain parallelism is not always an efficient solution, specially in applications with inherent instruction and/or data parallelism. Vectorization allows developers to exploit data-level parallelism. Current vector architectures have several problems, such as binary incompatibilities, waste of resources in case of long vector registers, or lack of instructions to successfully exploit certain applications. In this abstract we present an agnostic ISA, based on a generic vector length with new instructions to face emerging applications' problems.

Index Terms—Vector processor, SIMD, Multimedia Applications, Performance, Power Wall, Memory Wall.

I. INTRODUCTION

VECTOR processors are known to be very energy efficient and yield high performance whenever there is enough data-level parallelism (DLP). They typically operate with vector registers that hold multiple values instead of single-value registers as in superscalar processors.

Vector processors have a long and successful history in supercomputers where they are used for large scientific and engineering applications.

Large registers have several drawbacks. For example, if an application can not make full use of every register, then a hardware resource is being wasted. For this reason, fabricating chips with long vector registers has been considered senseless for a long time. Moreover, as authors in [1] explained, many applications have small data sets or iterate over an iteration space which is smaller than the vector register's length.

A. Current status and future trend

Nowadays, multimedia applications have seen an increasing usage which is predicted to keep proliferating next years. Furthermore, the amount of data to be computed is believed to grow. As a consequence, these applications will require an architecture that leads to high performance with low power consumption.

In this thesis we will study several representative vectorized applications and we will propose a modified architecture to improve their energy efficiency and performance. This way, we will evaluate the effect of different register's lengths in the program's performance as well as offer new techniques to

exploit this hardware. We will also analyze the impact to the memory hierarchy and try to mitigate the memory wall.

II. PRIOR WORK

Vector architectures, when utilized properly, offer many advantages over scalar ones. Commodity hardware manufacturers have been promoting the use of such architectures since the late 90s (Intel's MMX) and have improved them throughout new generations. Currently, the latest Intel's SIMD extension (AVX-512) processes up to 512-bit registers.

Recently, ARM presented the SVE [2] technology (Scalable Vector Extension) which was adopted by Fujitsu in a high performance processor [3].

ARM's SVE approach looks pretty similar to the idea we are developing in BSC. It provides the client with an agnostic vector length, allowing the usage of the whole vector register, no matter how long they are and without re-compiling the applications. It is also possible to deactivate unneeded vector register's elements using per-lane predication. However, we take a step forward and dynamically set the vector register's length and avoid the usage of the remaining elements saving energy in the process.

To the best of our knowledge, ForwardCom [4] is the only project which follows the same philosophy as ours. They intended to develop a new open ISA for high performance microprocessors. However, as they started from scratch, the software stack has not yet been programmed and there are no real implementations in hardware.

Another different approach to take advantage of longer vector registers was the one proposed by [5]. Their initial motivation had to do with extracting as much performance as possible out of supercomputers. This situation becomes a challenge when every architecture (e.g., Motorola's AltiVec vs. Intel's SSE2) provides different operations or restricted memory accesses. Their proposal was to implement a virtual vector instruction set architecture (Vector LLVA) which separated the native code generation and the processor-specific optimizations using what they called translators. Nevertheless, this abstraction layer needs to be programmed for every vector processor generation and it is limited to existing hardware, which hampers a short-term research.

III. CURRENT WORK

This section shows the current status of our research in the topic, explaining several details of the developed infrastructure and some results we obtained out of them.

Our evaluation framework is based on the gem5-x86, a full system simulator which models accurately each of the processor's components at every clock cycle. Our initial application set is based on the ParVec [6] suite, along with other kernels such as DGEMM, Vector reduction, etc. We modified Intel's SSE SIMD extension (128-bit registers) and implemented the AVX one (256-bit registers), using its instruction format to generate new ones for our own purpose. This way, just by changing AVX's instruction prefixes with a new one and gem5's decoder we were able to work with a generic vector instruction set. One of the main issues we had to deal with was the lack of compiler support, which prevented us from obtaining optimized code. In order to dynamically generate these new instructions, we developed a framework to make the programmer unaware of the specific underneath instruction features, such as the opcode. As a result, for every generic vector functionality a set of bytes is inserted into the binary at assembly level. We decided to reuse SSE and AVX register file, which was implemented with a 'vector register' type defined by its length. Functionally speaking, whenever a vector microinstruction is executed (e.g. an addition), all the register elements will be iterated and processed one by one. The number of elements is determined by the register size (in bits) and the data type. As this size can be changed at run time, just by setting it to the desired one and by increasing every vector register length to a maximum will force the simulator to deal with generic vector architectures.

Every benchmark contains what is commonly known as a Region of Interest (ROI), a set of functions the programmer is interested in measuring. In our case, that part of the ParVec suite contains the vector operations. For this reason, before the Region of Interest and, using new x86 instructions, we ask the architecture for the maximum vector length available and we set it in what we called 'MaxVectorLength' register. This situation will force the simulator to work with that vector length whenever the generic operations are executed.

Last but not least, it was also needed to tune several architectural components, such as the Load/Store Queue (LSQ), and different kernel operations, like context switches, to handle the new registers.

At this moment, we have managed to successfully run, in an out-of-order x86 processor, a generic version of selected benchmarks.

IV. RESULTS AND FUTURE WORK

In this section we show some of the results we obtained using the previously described infrastructure. Due to space limitations we only show the numbers for the blackscholes benchmark.

Figure 1 depicts the total instruction type count for scalar, SSE, AVX code as well as with the generic one with 128, 256, 512 and 1024-bit vector registers. It shows how SSE and AVX execute less instructions, due to the compiler support. However, if we focus on the generic results we appreciate a linear reduction on the instruction count as we increase register size.




Fig. 1. Absolute instruction type statistic




Fig. 2. Execution cycles per pipeline stage (logarithmic scale)

Figure 2 shows the total number of execution cycles in every pipeline stage for the same configurations. Wider registers reduce the pressure on the processor, as the number of instructions required to process the same amount of data elements is cut down.

Having these results in mind, there are several paths we could take and that we are still considering. One of the simplest proposals would be to reduce processor's frequency on those applications that are memory-bounded, that is, whenever the memory subsystem provides data at a lower rate it is consumed by the processor.

REFERENCES

- [1] L. Villa, R. Espasa, and M. Valero, "Effective Usage of Vector Registers in Advanced Vector Architectures," 1997.
- [2] ARM, "The Scalable Vector Extension SVE." [Online]. Available: <https://www.arm.com/products/processors/technologies/neon.php>
- [3] Fujitsu, "ARMv8, A Scalable Vector Extension for Post - K," 2016.
- [4] A. Fog, "ForwardCom Project." [Online]. Available: <http://forwardcom.info/>
- [5] R. L. B. Jr. and V. S. Adve, "Vector LLVA: A Virtual Vector Instruction Set for Media Processing," pp. 46–56, 2006.
- [6] "ParVec: Vectorized PARSEC Benchmarks." [Online]. Available: <http://www.ntnu.edu/ime/eecs/parvec>
- [7] M. Casas, M. Moretó, L. Alvarez, E. Castillo, D. Chasapis, T. Hayes, L. Jaulmes, O. Palomar, O. S. Unsal, A. Cristal, E. Ayguade, J. Labarta, and M. Valero, "Runtime-Aware Architectures," *Euro-Par 2015, Vienna, Austria*, pp. 16–27.
- [8] M. Valero, M. Moreto, M. Casas, E. Ayguade, and J. Labarta, "Runtime-Aware Architectures: A First Approach," *International Journal on Supercomputing Frontiers and Innovations (SuperFRI)*, vol. 1, no. 1, pp. 16–27, June 2014.



Adrián Barredo Ferreira was born in Santander, Cantabria (Spain) in 1993. He received his BSc degree in Telecommunications Engineering from Universidad de Cantabria (UC) in 2015. In August 2015, he joined the RoMoL team from BSC, and in September he started the Master's degree in Innovation and Research in Informatics from Universitat Politècnica de Catalunya (UPC).

His current on-going research is related to generic vector processors which involves adding support for them on an out-of-order x86 processor in the gem5 simulator, as well as a programming interface to deal with generic vector's ISAs.

A case for code-representative microbenchmarks

Calvin Bulla, Miquel Moreto Barcelona Supercomputing Center
 {calvin.bulla,miquel.moreto}@bsc.es

Abstract—Microbenchmarks are fundamental in the design of a microarchitecture. They allow rapid evaluation of the system, while incurring little exploration overhead. One key design aspect is the thermal design point (TDP), the maximum sustained power that a system will experience in typical conditions. Designers tend to use hand-coded microbenchmarks to provide an estimation for TDP. In this work we make the case for a systematic methodology to automatically generate code-representative microbenchmarks that can be used to drive the TDP estimation.

I. INTRODUCTION

The thermal design point (TDP) is a key aspect in the design of a microarchitecture. It indicates the maximum sustained power dissipation that a system will experience during typical runtime conditions and constrains the design of the cooling system. It is critical to have a good TDP estimation. If it is too low, power-hungry workloads will get throttled, which negatively impacts the overall system performance. If it is too high, we incur an increase in terms of packaging and cooling cost. As a consequence, a TDP estimation that is representative of real workloads is a main goal during the design process.

In a typical design environment it is infeasible to execute a wide range of applications to come up with a realistic TDP. Simulating cycle-intensive applications is a serious time investment, which prevents rapid design space exploration in the early stages of development. As a consequence, the industry relies on microbenchmarks to single out and evaluate the crucial parts of real applications. These microbenchmarks are often hand-crafted, L1-contained kernels with a low branch missprediction to maximize power dissipation and provide a conservative TDP estimation. Creating these microbenchmarks requires a significant engineering effort, and knowledge of the underlying architecture as well as the common workloads.

Previous work has successfully highlighted the benefit of microbenchmark generators for a wide range of purposes: cloning program behavior [5], [6], power profiling [4], [7], or to generate synthetic tests [1]. In this work we make the case for a framework to automatically extract code-representative microbenchmarks from real applications, which can then be used to effectively drive the TDP.

In this context, we have set the following goals:

- The generated microbenchmarks should provide a safe upper-bound for power. Therefore, we will force them to be L1-contained, with predictable branches.
- The generated instruction sequences should be code-representative, i.e. occur like this in the original benchmark.
- The methodology should be easily extendable, and applicable to a variety of architectures, without being tied to a specific simulation infrastructure.

II. FRAMEWORK

Fig. 1 gives an overview of the workflow for our framework. It can be divided into four stages: 1) Execution Tracing, 2) CFG Annotation, 3) Snippet Selection, and 4) Microbenchmark Synthesis.




Fig. 1: Overview of our proposed workflow.

A. Execution Tracing

We first trace the execution of the selected benchmark. To this end, we have developed a sampling based profiler using the `perf_event` interface. At regular intervals, we gather the current program counter (PC), as well as a set of configurable performance monitoring counters (PMCs), such as the branch and cache miss/hit ratio.

B. CFG Annotation

In this step we use the obtained runtime profile to annotate the control flow graph (CFG) of the application. A CFG is a graph representation of the program’s executional flow. Each node represents an uninterrupted sequence of instructions, or basic block (BB), each vertex the corresponding branch or fall-through at the end of a BB.

With the sampled program counter as key, we associate each sample to a specific BB. We then weigh each BB, based on the number of samples it generates. The measured PMCs can be used as an additional heuristic to increase the weight of impactful BBs.

C. Snippet Selection

Using the CFG annotation we select a subset of the heaviest BBs. In our initial approach, we ignore all BB, which sample count is below an α threshold (0.1 – 1.0% in our experiments).

We can now generate code snippets for each selected BB (seed). A code snippet is a short sequence of BBs between a loop (or function) entry and exit point. We form these snippets by finding all acyclic paths between the entry and exit BB that go through the seed. For each BB we only select the snippet with the highest accumulated weight. The snippets are validated against the instruction trace of the original application.

D. Microbenchmark Synthesis

Finally, we pass the snippets to a MicroProbe-based [2] backend that generates standalone executables. To this end we take the snippet kernel and embed it into a new loop. All branches are modified so that they directly jump to the next BB in the snippet chain. Where necessary, memory references are changed so that they access a previously allocated, L1-contained, memory region. The generated assembly code now can be compiled for the target architecture.

III. EVALUATION

We have tested our first approach on an IBM BladeCenter PS701 system, featuring an eight-core IBM POWER7 processor running at 3.8 GHz. The benchmarks have been selected from the SPEC CPU 2006 benchmark suite. We execute each benchmark in isolation, and use our toolchain to generate the corresponding microbenchmarks with a threshold $\alpha = 0.1$ and 1.0. We then execute the created microbenchmarks, and compare their power dissipation with the original application. We use AMESTER to access the on system power meters each millisecond. The shown values are normalized with respect to the maximum observed system power.

Fig. 2 shows the power histogram for two of the selected benchmarks and the corresponding microbenchmarks with different thresholds. We plot the percentage of samples/snippets (y -axis) observed for each power range (x -axis). We can observe that the original benchmarks tend to have two to three dominant power ranges, corresponding to the main kernels of the application. In the case of perlbench (a), our approach is able to replicate both spikes. As we decrease the threshold, our framework will consider more snippets and better approximate the original application. For zeusmp (b), a lower threshold produces an even more notable change in proportion of sampled power regions. We plan to investigate the most appropriate threshold in the future. It is important to note, that in all cases, we are able to provide a safe upper bound for power.

IV. CONCLUSION

In this work we have highlighted the need for microbenchmarks in estimating the TDP of a microarchitecture design. We have also presented our initial approach to automatically

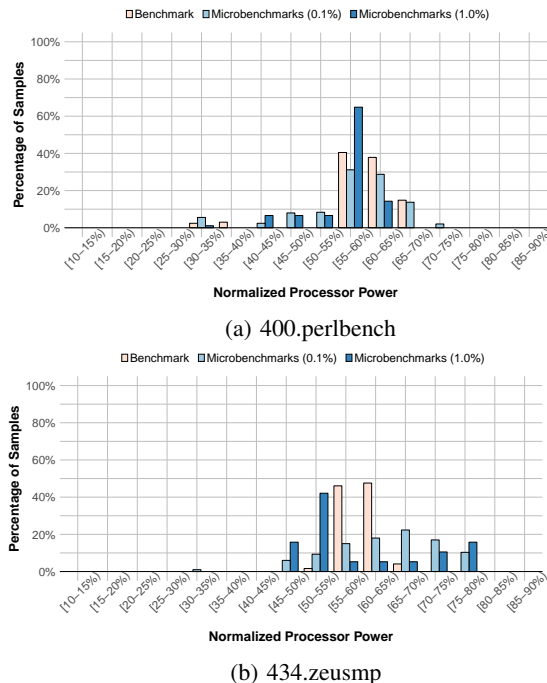


Fig. 2: Power histograms for benchmarks and their corresponding microbenchmarks

extract code-representative microbenchmarks from real applications using runtime profiling and CFG analysis. Finally, we present preliminary results on a POWER7 system.

For future work we plan to further test our methodology and apply it to parallel workloads and in the context of runtime-aware architectures [3], [8]. We also consider expanding the snippet selection phase with more sophisticated heuristics. Lastly, we plan to port our framework to different architectures and release it as an open-source toolchain.

REFERENCES

- [1] R. H. Bell and L. K. John. Improved automatic testcase synthesis for performance model validation. In *Proceedings of the 19th annual International Conference on Supercomputing*, ICS '05, pages 111–120, 2005.
- [2] R. Bertran, A. Buyuktosunoglu, M. S. Gupta, M. Gonzalez, and P. Bose. Systematic energy characterization of CMP/SMT processor systems via automated micro-benchmarks. In *Proceedings of the ACM/IEEE International Symposium on Microarchitecture*, 2012.
- [3] M. Casas, M. Moretó, L. Alvarez, E. Castillo, D. Chasapis, T. Hayes, L. Jaulmes, O. Palomar, O. Unsal, A. Cristal, et al. Runtime-aware architectures. In *European Conference on Parallel Processing*, pages 16–27. Springer, 2015.
- [4] C.-T. Hsieh and M. Pedram. Microprocessor power estimation using profile-driven program synthesis. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 17(11):1080–1089, Nov 1998.
- [5] A. Joshi, L. Eeckhout, R. Bell, and L. John. Performance cloning: A technique for disseminating proprietary applications as benchmarks. In *Proceedings of the IEEE International Symposium on Workload Characterization*, pages 105–115, Oct. 2006.
- [6] A. Joshi, L. Eeckhout, R. H. Bell, and L. K. John. Distilling the essence of proprietary workloads into miniature benchmarks. *ACM Transaction of Architecture and Code Optimization*, 5:10:1–10:33, September 2008.
- [7] L. Mukhanov, D. S. Nikolopoulos, and B. R. de Supinski. Alea: fine-grain energy profiling with basic block sampling. In *Parallel Architecture and Compilation (PACT), 2015 International Conference on*, pages 87–98. IEEE, 2015.
- [8] M. Valero, M. Moreto, M. Casas, E. Ayguade, and J. Labarta. Runtime-aware architectures: A first approach. *Supercomputing frontiers and innovations*, 1(1):29–44, 2014.



Calvin Bulla received his BSc degree in Computer Science from Universidad Las Palmas de Gran Canaria in 2015. Now, he is enrolled in the Master's degree of Innovation and Research in Informatics at Universitat Politècnica de Catalunya (UPC) with a mention in High Performance Computing. He joined the RoMoL team at the Barcelona Supercomputing Center in September 2015.

His research focuses on runtime-aware architectures, performance analysis, and system profiling. Currently he is developing a framework to trace applications and extract standalone microbenchmarks. These can then be used to guide the design characteristics of future microarchitectures, such as the thermal design point.

Esterases Computational Study

Ruben Cañadas^{#1}, Gerard Santiago^{#2}, Victor Guallar^{#*3}

[#]Joint BSC-CRG-IRB Research Program in Computational Biology, Barcelona Supercomputing Center, 08034 Barcelona, Spain

¹ ruben.canadas@bsc.es, ² gerard.santiago@bsc.es

^{*}Institució Catalana de Recerca i Estudis Avançats (ICREA), 08010 Barcelona, Spain

³ victor.guallar@bsc.es

Keywords— esterases, PELE, QM/MM, machine learning

EXTENDED ABSTRACT

Esterases are a group of alpha/beta hydrolases enzymes that split an ester into an alcohol and a carboxylic acid. These types of enzymes have special interest in industry, being used in food processing, perfume industry, degradation of synthetic materials..

Their catalytic mechanism is well-known in the literature¹. They have a catalytic triad typically formed by three amino acids: Serine, histidine and aspartic acid. In some cases we can find a glutamic acid instead of an aspartic acid. Many other hydrolases such as proteases use this mechanism, being on of the most conserved catalytic mechanism.

The histidine deprotonates the hydroxyl group of the serine. Then, this serine carries out a nucleophilic attack towards the ester. After that, a tetrahedral intermediate is formed followed by an electronic rearrangement, resulting in the formation of an alcohol. The next step is the entrance of a molecule of water. Thanks to this molecule, the second product of the reaction, a carboxylic acid, is formed.




Figure 1. Canonical esterase mechanism for hydrolysis of

methyl acetate, a typical ester

The Asp-His hydrogen bond is very important to maintain the catalytic triad in an active conformation. Using molecular dynamics, one can see that if this hydrogen bond is broken, this triad can adopt an inactive conformation². Another important aspect of the catalytic mechanism is the oxyanion hole. This is formed by two backbone hydrogens which stabilize the negative charge of the oxygen when the tetrahedral intermediate is formed.




Figure 2. Typical esterase catalytic triad.

In our project, we are focusing on two different esterases called LAE6 and LAE5. LAE6 is very promiscuous since it presents activity towards a wide range of different substrates, whereas LAE5 presents activity towards few substrates. This promiscuity can be due to the location of the binding cavity: in LAE6 the binding cavity is located buried inside the protein, so the substrates can be retained longer, correlating with a higher reactive probability. LAE5 binding cavity is situated on the surface of the protein, in direct contact with the solvent, which, in our working hypothesis, is the main reason for its much less promiscuity.

We are using several computational techniques to extract important descriptors in order to construct a mathematical model to predict the activity of these two enzymes towards different substrates. We normally start performing a Monte-Carlo-Metropolis calculation using an in-house software named PELE (Protein Energy Landscape Calculation). This method uses a pseudorandom number

generator to perturb the ligand randomly. Then the protein structure is adapted (relaxed) to such perturbation using protein structure prediction techniques. Once the system has been relaxed, a Metropolis criterion is used to decide whether this structure will be accepted or not. The overall technique, along with the atomic force field and implicit solvent is often biased towards low energy conformations, so we can get stuck in some local minimum, preventing us from getting over high energy barriers and consequently visiting other local minima. Thus, we introduced a machine learning adaptive technique to improve sampling.

After running PELE simulations, we can extract the best poses according to catalytic distances and binding energy profiles. QM/MM optimizations might also be done to refine structures since PELE uses a minimization algorithm based on molecular mechanics. Once the structure is minimized, we are ready to extract several descriptors. For example catalytic distances, deprotonation angles, charge distribution, molecular weight... Other methods are being used, such as molecular dynamics or metadynamics to deduce more accurate descriptors.

All these descriptors will be used in the future, to create a model, using machine-learning techniques. Ideally, this model will be able to predict the activity of a wide range of substrates. Thus, a correlation between experimental activity and predicted activity needs to be found.

References

- [1] A. Rauwerdink and R. J. Kazlauskas, "How the same core catalytic machinery catalyzes 17 different reactions: the serine-histidine-aspartate catalytic triad of alpha/beta hydrolase fold enzymes," *ACS catal.* 2015, 5, 6153-6176
- [2] E. Y. Lau and T. C. Bruice, "Consequences of breaking the Asp-His hydrogen bond of the catalytic triad: effects on the structure and dynamics of the serine esterase cutinase," *Biophysical Journal*, Vol. 77, Pp. 85-98, July 1999.

Author biography



Ruben Cañadas is a bachelor's degree student at his fourth year who is currently carrying out his degree final project at the atomic and electronic protein modeling group at the BSC (Barcelona Supercomputing Center). He is focused on the study of two esterases with relevant importance for the

industry.

Prediction of binding energies upon mutation in 3D-structure-known complexes through PyDock scoring functions

Bruno Cuevas, Miguel Romero, and Juan Fernandez Recio

Barcelona Supercomputing Centre, Nexus II Building c/Jordi Girona, 29 08034 Barcelona (Spain)

Abstract—The combination of emerging technologies in biology and the exponential increase of available biological data is leading to the new paradigm of personalized medicine where each patient will be diagnosed and treated by integrating all known personal features such as its genetic background.

In this context, one of the main challenges is the development of tools that can help to characterize pathological mutations. A pathological mutation can act in different ways. Some mutations can modify the overall structure of the protein, leading to incorrectly unfolded proteins that cannot play their functions, or that are targeted by the mechanisms of protein turnover. Other groups of mutations can play crucial roles in the interaction between proteins. In those cases, proteins could remain structurally unchanged, yet the formation of protein-protein complexes is disturbed, inhibiting different paths in which protein-protein interactions (PPIs) are important.

Different strategies have been employed to deal with this challenge, most of them using structural information about protein interactions. In this regard, it is possible to distinguish between two groups of methods: those that make use of machine learning methods, such as random forests or neural networks, applied directly at structure level, like PoPMuSiC [1], ELASPIC [2], BindProf ([3]), ZEMu [4], mCSM [5]; and those that make use of energy scoring functions [6].

The main objective of this work is to build a simple but robust predictor of binding energy changes upon mutation, once structural information is provided. Three different tools are being employed: First, given the wild type structure of a complex, mutations are modelled with Modeller, a powerful program extensively used in protein homology modeling. Using the workframe provided by the tool, 50 different models are created for each mutation. This model diversity helps to take into account protein flexibility and explore, more efficiently, the conformational space of interactions. Then, models are evaluated using Modeller DOPE assessment tool and pyDock scoring function. DOPE [7] is statistical-potential-based tool that evaluates the quality of a model. The pyDock scoring function [8] is formed by different energy terms (electrostatic, desolvation and van der Waals), and was originally designed to deal with protein-protein docking problems. However, due to its energetic basis, we wanted to test their ability to evaluate changes in binding energies.

The correlation between pyDock energy and experimental $\Delta\Delta G$ values is being tested on SKEMPI ([9]), a dataset that provides a large set of mutations in protein complexes for which there are both structural information and experimental measures of the changes in binding energy ($\Delta\Delta G$) upon mutation. A careful curation of the experimental data, coming from different experimental techniques, environmental conditions (e.g., temperature, pH) has been performed. Optimization of parameters and selection of unbiased experimental data are expected to lead

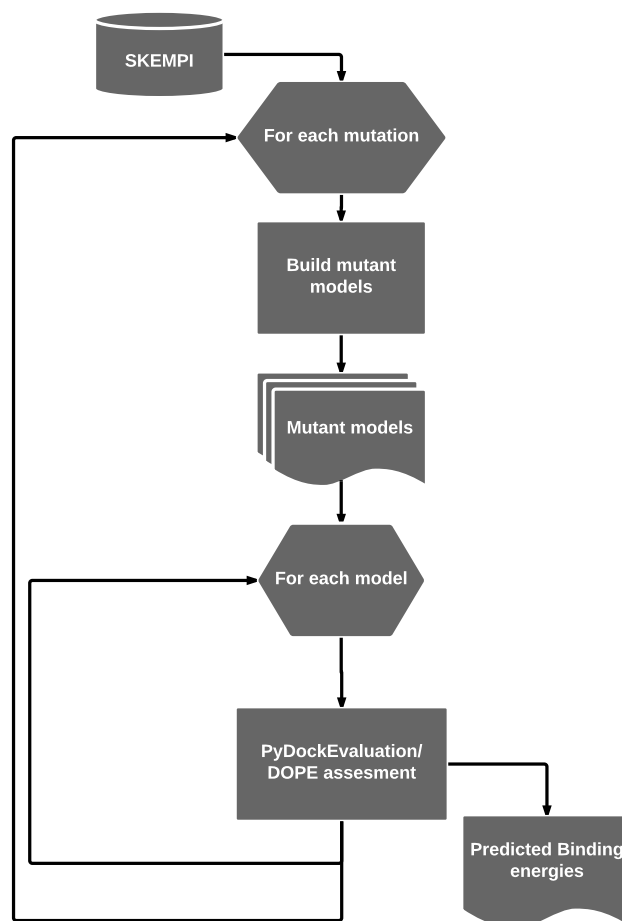
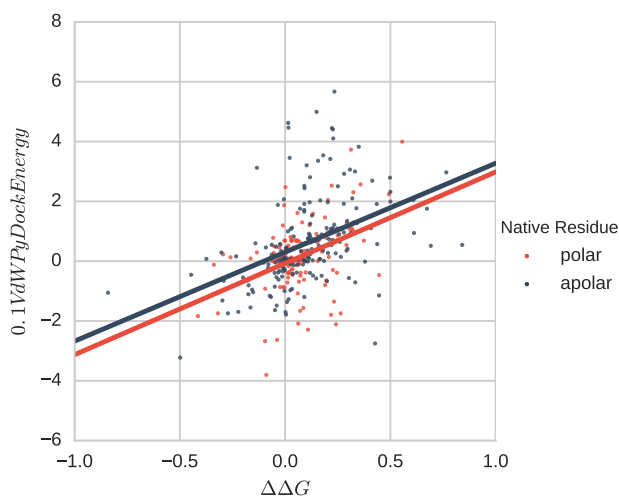


Fig. 1. Binding energy prediction workflow

to a high correlation coefficient between theoretical values and experimental ones.

Compared with other state-of-the-art methods based on machine-learning approaches, the method we here proposed, offers a much simpler workframe, keeping a scientific-knowledge base.

Index Terms—Docking, Personalized Medicine, Structural Biology, PPI



$$r = 0.43664777387110393$$

Fig. 2. Real experimental values against predicted ones using PyDock scoring functions when mutations and original residues are polar or apolar



Bruno Cuevas Zuviria Graduated in Biotechnology by Universidad Politécnica de Madrid in 2016, and current student of master in Modelling For Science and Engineering in Universidad Autonoma de Barcelona. He works at the protein-protein interaction and docking group in BSC, under direction of Juan Fernandez Recio. He got specialized in Computational Biology during the degree, and in Complex Systems during the master. His main interests are the systems biology, the structural biology and the epidemiology, field in which he used to work before.

REFERENCES

- [1] Y. Dehouck, A. Grosfils, B. Folch, D. Gilis, P. Bogaerts, and M. Rooman, "Fast and accurate predictions of protein stability changes upon mutations using statistical potentials and neural networks: PoPMuSiC-2.0," *Bioinformatics*, vol. 25, no. 19, pp. 2537–2543, 2009.
- [2] N. Berliner, J. Teyra, R. Çolak, S. G. Lopez, and P. M. Kim, "Combining structural modeling with ensemble machine learning to accurately predict protein fold stability and binding affinity effects upon mutation," *PLoS ONE*, vol. 9, no. 9, 2014.
- [3] J. R. Brender and Y. Zhang, "Predicting the Effect of Mutations on Protein-Protein Binding Interactions through Structure-Based Interface Profiles," *PLoS Computational Biology*, vol. 11, no. 10, pp. 1–25, 2015.
- [4] D. F. A. R. Dourado and S. C. Flores, "A multiscale approach to predicting affinity changes in proteinprotein interfaces," no. June, pp. 2681–2690, 2014.
- [5] D. E. V. Pires, D. B. Ascher, and T. L. Blundell, "MCSM: Predicting the effects of mutations in proteins using graph-based signatures," *Bioinformatics*, vol. 30, no. 3, pp. 335–342, 2014.
- [6] C. Geng, A. Vangone, and A. M. J. J. Bonvin, "Exploring the interplay between experimental methods and the performance of predictors of binding affinity change upon mutations in protein complexes," vol. 29, no. 8, pp. 291–299, 2016.
- [7] A. Sali and A. Sali, "Statistical potential for assessment and prediction of protein structures," *Protein Science*, pp. 2507–2524, 2006.
- [8] T. M.-K. Cheng, T. L. Blundell, and J. Fernandez-Recio, "pyDock: Electrostatics and desolvation for effective scoring of rigid-body protein-protein docking," *Proteins: Structure, Function, and Bioinformatics*, vol. 68, no. 2, pp. 503–515, apr 2007. [Online]. Available: <http://www.ncbi.nlm.nih.gov/pubmed/17444519> <http://doi.wiley.com/10.1002/prot.21419>
- [9] I. H. Moal and J. Fernández-Recio, "SKEMPI: A Structural Kinetic and Energetic database of Mutant Protein Interactions and its use in empirical models," *Bioinformatics*, vol. 28, no. 20, pp. 2600–2607, 2012.

Deep and cognitive learning applied to Precision Medicine: the initial experiments linking (epi)genome to phenotypes-disease characteristics.

D. Cirillo^{#1}, A. Valencia^{#2}

[#]Barcelona Supercomputing Center (BSC), c/Jordi Girona, 29, 08034 Barcelona (Spain)

Life Sciences Department, Computational Biology Group

¹davide.cirillo@bsc.es,

²alfonso.valencia@bsc.es

Keywords— Deep learning, Network analysis, Personalized medicine

EXTENDED ABSTRACT

Introduction

Present-day era of Big Data provides the unique opportunity to develop innovative approaches for data analysis to find new insights into specialized fields of biomedical research such as **Precision Medicine** [1]. Precision Medicine is defined as the integration of molecular research with clinical data in order to deliver better diagnoses and treatments tailored to the individual characteristics of each patient. Advanced analysis of health related data that is specific to a given individual must focus on both clinical information (e.g. clinical reports, medical images, patient histories) and biological data (e.g. gene and protein sequences, functions and pathways). This wealth of information has the potential to inspire systematic ways of making sense from the massive and heterogeneous stream of data and providing a unified view. In the regards, **Deep Learning** (DL) [2] and **Cognitive Computing** (CC) [3] are two branches of Artificial Intelligence (AI) representing convenient choices to tackle the problem of Big Data integration for Precision Medicine. DL comprises several machine learning techniques modeling multiple representations of data through many layers of nonlinear processing units. CC is a cross-disciplinary technology for adaptive and contextual knowledge representation and reasoning through sophisticated analytics aiming to mimic human learning mechanisms.

Objectives

1 - Ontologies of molecular and clinical annotations

A large fraction of biological knowledge is organized in the form of **ontologies**, i.e. sets of domain-specific categories (terms) with relations operating among them. The Gene Ontology (GO) [3] covers three domains: cellular components, biological processes and molecular functions. The Human Phenotype Ontology (HPO) [4] covers the domain of clinical signs or phenotypic anomalies in human diseases. Inspired by automatic language translation, specific Deep Learning approaches like “encoder-decoder” Recurrent Neural Networks (RNNs) with Long Short-Term Memory

(LSTM) architecture and Attention Mechanism (AM) (Fig. 1) can be used to model the “translation” of one ontology to the other and build a common function-disease reference framework. In the context of Personalized Medicine, this approach can be exerted to understand the relationships between the genetic variants found in a given patient and her/his specific set of disease-associated phenotypes and altered gene functions and pathways.

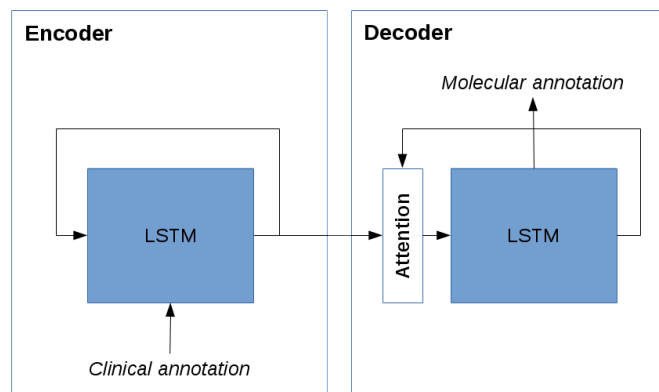


Fig. 1: Schematic view of Encoder-Decoder RNN that takes a Clinical annotation as an input and produces the corresponding Molecular annotation as an output.

2 – Genotype-phenotype relationships

A compelling application of RNNs techniques is to uncover of the determinants of **time-dependent biological processes** such as cellular differentiation. Recently, Carrillo de Santa Pau et al. [5] proposed a model to link epigenetic signatures to cell fate during hematopoiesis, i.e. the process of formation of blood cellular components starting from stem cells in the red bone marrow (Fig. 2). Along with changes in the epigenome, additional ‘omics’ data (e.g. chromatin conformations, expression levels, protein abundances) can be taken into account to model complex time series by means of deep RNNs. In the context of Personalized Medicine, this approach can be applied to the detection of patient specific disease related (epi)genomics modifications affecting cellular

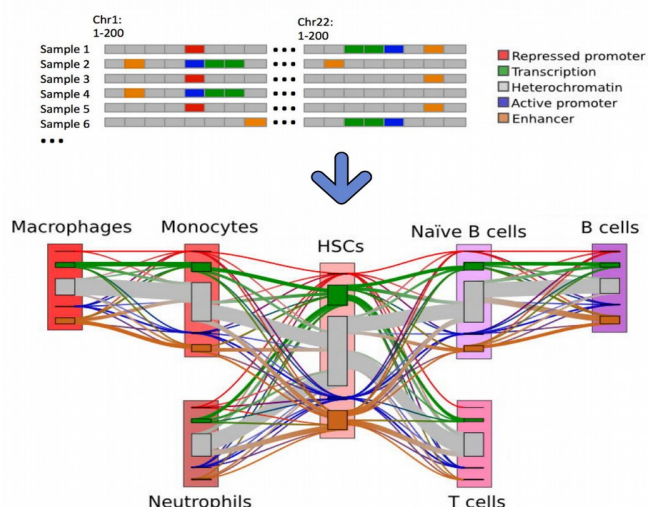


Fig. 2: Genome segmentation in 5 chromatin states across samples (upper panel). Changes of chromatin states between differentiation stages (lower panel; line thickness is proportional to the number of regions) (adapted from Carillo de Santa Pau 2016 [5]).

differentiation processes and will be potentially applicable to tumor evolution processes.

3 – Cognitive computing with IBM Watson

IBM Watson (<http://www.research.ibm.com/cognitive-computing/>) [6] is one of the most advanced AI-platform for CC. This technology differs from current computing applications in that it moves beyond pre-configured rules as it reasons based on broad objectives. IBM Watson system is adaptive, interactive, stateful and contextual, meaning that it learns as information changes by asking questions and finding additional inputs based on user's needs at that point in time and from different sources of information (syntax, domain regulations, etc.). As a result, IBM Watson is able to address complex situations that are characterized by ambiguity and uncertainty such as question answering (QA) tasks based on large unstructured collections of natural language documents. In collaboration with IBM we are applying Cognitive Computing approaches to the above mentioned scenario relating (epi)genome-phenome/disease information sources.

Discussion

In recent years, many AI techniques emerged that can effectively learn from for very large and complex data sets achieving human-level performances in image and speech recognition as well as natural language processing. Incorporating biomedical knowledge into such a new generation of learning algorithms is the current paradigm in biomedical research. Indeed, the development of algorithms for a comprehensive analysis of health-related Big Data represents one of the main challenges for the Computational Biology community and also a rational aid to experimental scientists and physicians. In particular, DL such as deep RNNs, and CC such as IBM Watson system, offer the possibility to address Personalized Medicine problems prompting functional hypotheses and guiding better diagnoses and treatments. Multiple data sources, namely large-scale clinical and molecular data, combined into integrative models designed to reveal unexpected relationships among biological entities will help to unveil the broader context in which physiological and pathological events are occurring.

References

- [1] Ashley EA. *Towards precision medicine*. Nat Rev Genet. 2016 Aug 16;17(9):507-22.
- [2] LeCun et al. *Deep learning*. Nature. 2015 May 28;521(7553):436-44.
- [3] Ashburner et al. *Gene ontology: tool for the unification of biology* (2000) Nat Genet 25(1):25-9.
- [4] Köhler et al. *The Human Phenotype Ontology in 2017* (2017) Nucleic Acids Res 45 (D1): D865-D876.
- [5] Carillo de Santa Pau et al. *Searching for the chromatin determinants of human hematopoiesis*. 206. bioRxiv doi: 10.1101/082917.
- [6] Chen et al. *IBM Watson: How Cognitive Computing Can Be Applied to Big Data Challenges in Life Sciences Research*. (2015) Clin Ther. 2016 Apr;38(4):688-701.

Author biography

Davide Cirillo was born in Rome, Italy, in 1985. He received the M.D. degree in Pharmaceutical Biotechnology from University of Rome 'La Sapienza', Italy, in 2011, and the Ph.D. degree in Biomedicine from Universitat Pompeu Fabra (UPF) and Center for Genomic Regulation (CRG), of Barcelona, Spain, in 2016. Since March 2017, he is a Recognized Researcher at computational Biology Group within the Life Sciences Department of Barcelona Supercomputing Center (BSC), Spain. His current research interests include Deep Learning, Network Analysis, Artificial Intelligence, Precision Medicine.

Exploring the Relationship between Gene Expression and Topological Properties of *Arabidopsis thaliana* Interactome Network.

Silvia M. Gimenez¹, Guillermo Rodrigo^{#2}, Brian Jiménez⁺³, Juan Fernandez-Recio⁺⁴.

[#]*Instituto de Biología Molecular y Celular de Plantas, Consejo Superior de Investigaciones Científicas – Universidad Politécnica de Valencia, 46022, Valencia, Spain*

⁺*Barcelona Supercomputing Center (BSC). Joint BSC-CRG-IRB Research Program in Computational Biology, Barcelona, Spain*

¹silvia.santam@gmail.com, ²guirotda@ibmcp.upv.es, ³brian.jimenez@bsc.es, ⁴juanf@bsc.es

Systems Biology, Protein-protein Interaction Network, Statistical correlations.

EXTENDED ABSTRACT

A. Background

Regulatory mechanisms as dynamical responses by a living cell, that is, how it reacts and adapts to a different environmental stimuli, is highly unknown yet at molecular level.

In the field of Systems biology research, models of interaction networks are a powerful tool to approach living complex systems. Presented as a method that facilitates the understanding the structure and dynamics of complex intercellular web of interactions that contribute to the structure and function of a living cell [1].

The protein-protein interaction (PPI) network of *Arabidopsis thaliana* L., previously validated through experimental procedures and published, has the characteristic properties of hierarchical networks [2, 3]. It is shown that PPI networks are dynamically organized into functional modules.

The aim of this study is to integrate and link up transcriptomic data with biological networks approaches. The main objective was to determinate the correlation of transcriptomic profiles with PPI topology, seeking to demonstrate relational or structural patterns within the network internal organization.

B. Results

Correlation coefficients between gene expression profiles derived from microarray data evaluated at steady-state RNA expression levels from several growth conditions, developmental stages biotic and abiotic stresses, and a variety of mutant genotypes [4], and topological properties, such as connectivity degree or clustering coefficient, in PPI context have been evaluated (Fig 1, section b). Genes with high and low expression profiles were located within the global PPI network (Fig 1, section a).

Based on microarray data analyzed, a matrix of coexpression between expressed genes was calculated. Twelve coexpressed nodes were identified and highlighted into the network (Fig 1, section c). Selecting its direct interactions and nodes, a model of functional modules was proposed. Table 1 contains a functional analysis of nodes involved in the proposed modules. Connector nodes between modules were detected and analyzed.

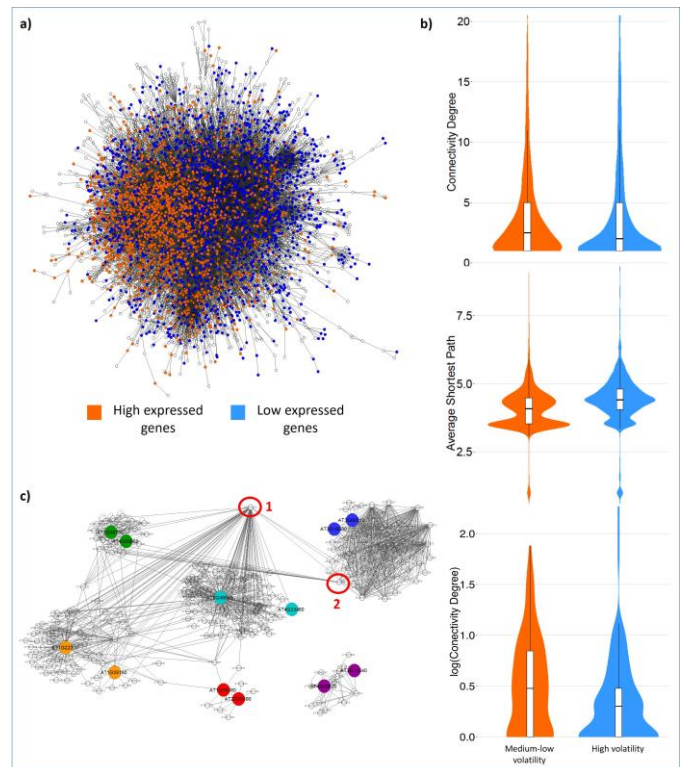


Fig 1 a) Highlight of nodes that represent genes with high and low gene expression level, included in correlation tests, located at its global *A. thaliana* interactome context. b) Violin plots combined with boxplots as representation of significant U-test evaluating correlations. $pV = 1.88e-6$; $pV < 2.2e-16$; $pV = 0.00053$. c) Highlight of coexpressed genes and its functional modules proposed. Circled potential modules regulatory targets, UBQ3 (1) and FBW2 (2) genes.

TABLE I
FUNCTIONAL MODULES NODES DESCRIPTION

Coexpressed nodes	Gene ID	Biological activity	N° genes
AT1G78380 AT2G29490	GST8 GST19	Glutation transferase activity	9
AT3G16580 AT3G60010	SK13 --	F-Box proteins. Unknown function	77
AT1G22300 AT1G35160	GF14 GF14 PHI	ATP binding	46
AT3G58780 AT4G09960	AGL1/SHP1 AGL11/STK	Transcriptional factors	22
AT3G13540 AT4G09820	MYB5 ATTT8/TT8	Transcriptional factors.	12
AT5G46630 AT4G23460	AP2M --	Protein intercellular transport	25

Author biography

C. Conclusion

Correlations between gene expression level and topological properties were determined, being positive for connectivity degree (C) and negative in relation with shortest path length (l). Furthermore, a negative correlation was determined for the relation between volatility coefficient of gene expression levels and connectivity degree (C). This conclusion implies that most connected nodes are the genes with higher expression level. At the same time they present the capability to transmit information through the network faster, connecting any other arbitrary node passing by less number of nodes. In addition, related to volatility, the most connected nodes present less fluctuating gene expression levels.


We identified 12 important genes showing relation of coexpression and direct interaction in the network; their functional modules were predicted and characterized. Determinant nodes, understood as connector elements between functional modules were detected. UBQ and FBW2 genes, detected as regulatory network nodes, have important roles in suppression of gene silencing and protein ubiquitination. Those genes are potential targets involved in the transmission procedures of biological information through the network.

D. Future develop/ways

The present study evidences the requirement of a time parameter integration for dynamic network future studies, to well understand how molecular regulatory mechanisms regulate the dynamic responses to environmental stimuli and intracellular signals.

References

- [1] H. Kitano. "Systems Biology: A Brief Overview". Science, Vol. 295, Pp. 1662-1664, 2002.
- [2] A.L. Barabási, Z.N. Oltvar. "Network biology: understanding the cell's functional organization". Nat. Rev. Genet, Vol. 5, No. 2, Pp. 101-113, 2004.
- [3] T. Ideker, N.J. Krogan. "Differential network biology". Mol. Syst. Biol., Vol. 8, Pp. 565, 2012.
- [4] J. Carrera, G. Rodrigo, A. Jaramillo, S. Elena. "Reverse-engineering the *Arabidopsis thaliana* transcriptional network under changing environmental conditions". Genome biology, Vol. 10, No 9, Pp. R96, 2009.



Silvia M. Giménez Santamarina was born in Valencia, Spain, in 1992. He received the B.E. degree in Biotechnology from the Catholic University of Valencia, Spain, in 2015. Nowadays she is finishing the MS. degree in Bioinformatics and Biostatistics from the Universitat Oberta de Catalunya (UOC) Barcelona, Spain, currently in the last semester.

Since 2014, she has been focused in to learn and growth in the field of Bioinformatics. Her first work experience in this field was in Brazil during a Bachelor's traineeship. After that, her Bachelor Final Thesis was focused on the study of the mode of action of a virus, integrating network biology approaches, and the design and implementation of a software to map proteomic data from non-model organism to a model organism. In 2015 she did a second Bachelor's traineeship in Uppsala, Sweden, at SLU University. She worked analyzing RNA-Seq data and differential gene expression. The presented study is her MS Final Thesis, which obtained the maximum punctuation. Currently she is working at the BSC, Life Science department, with the Protein Interactions and Docking team, lead by Juan Fernández-Recio.

Composite materials calculation using HPC-based multiscale technique.

Guido Giuntoli
Barcelona Supercomputing Center
CASE Department
Barcelona, Spain 08034
Email: guido.giuntoli@bsc.es

Mariano Vázquez
Barcelona Supercomputing Center
CASE Department
Barcelona, Spain 08034
Email: mariano.vazquez@bsc.es

Sergio Oller
Polytechnic University of Catalonia
Civil Engineering School
Barcelona, Spain 08034
Email: sergio.oller@upc.edu

Abstract—Calculating the behavior of large composite material structures still demanding large computational effort. The solution proposed to tackle the problem is to combine multi-scale homogenization methods with the high performance computational techniques available. This PhD work aims to implement inside the Alya HPC-code a multi-scale algorithm capable of solving this kind of problems in an efficient and accurate way.

I. INTRODUCTION

Simulation of large structures made of composite materials continues to be a challenge even with the huge development of computational technology (see [1]). This PhD work is focused on implementing a technique based on multi-scale calculation taking advantage on the high performance techniques in order to obtain accurate solutions of large non-linear problems.

In the present work we give a short explanation of the multi-scale method and we discuss the basics on material nonlinearities and the work flow that is been taken.

II. MULTI-SCALE CALCULATIONS

Generally the problems in mechanical analysis are expressed in a virtual work form and then are solved by the finite element method. In this case, for simplicity, taking a steady state problem and assuming the small deformation approach, the equations without considering boundary conditions can be state (see [2]) as:

$$\int_V \bar{\boldsymbol{\tau}} \cdot \delta \bar{\boldsymbol{\epsilon}} dV = \mathcal{R} \quad (1)$$

where $\bar{\boldsymbol{\tau}}$ is the Cauchy stress tensor, $\delta \bar{\boldsymbol{\epsilon}}$ is the virtual strain and \mathcal{R} is the external virtual work done.

The problem above consist on solving for the stress field $\bar{\boldsymbol{\tau}}$ that satisfies the virtual work Eq. (1) for every virtual strain $\delta \bar{\boldsymbol{\epsilon}}$. Commonly stress $\bar{\boldsymbol{\tau}}$ and strains $\bar{\boldsymbol{\epsilon}}$ are expressed as functions of the displacement $\bar{\boldsymbol{u}}$, i.e.:

$$\bar{\boldsymbol{\tau}} = \bar{\boldsymbol{\tau}}(\bar{\boldsymbol{u}}), \quad \bar{\boldsymbol{\epsilon}} = \bar{\boldsymbol{\epsilon}}(\bar{\boldsymbol{u}}), \quad (2)$$

and a constitutive law equation like

$$\bar{\boldsymbol{\tau}} = \mathbb{C} : \bar{\boldsymbol{\epsilon}}, \quad (3)$$

is added. Here \mathbb{C} is the four-order constitutive tensor and depends on the material. By this way the problem is to find the displacements $\bar{\boldsymbol{u}}$ that satisfy (1) and (3) at the same time.




Fig. 1. Multi-scale calculation representation. A general structure is discretized in a coarse mesh and during the assembly process at each integration point a problem over a representative volume element (RVE) that contains the micro-structure information of the composite is solved.

At this point it is possible to notice that during the assembly process for obtaining a detailed calculation of the solid that is being treated, the constitutive \mathbb{C} tensor should be defined on every point of the finite element mesh. Moreover, \mathbb{C} depends on $\bar{\boldsymbol{\epsilon}}$ and on the history making that the problem becomes non-linear and more difficult to solve.

Having this in mind, calculating the behavior of composite materials applying the finite element technique directly, specially on large domains such as an aircraft panel, can derive in an extremely large system of equations due to the huge number of elements used to represent such heterogeneous domain. On the other hand memory requirements become large when non linearities such as damage and plasticity should be modeled and a huge number of variables should be stored for every element in the mesh.

This situation brought scientist and engineers to developed the multi-scale calculation method capable to afford this problem reducing the computational cost and, at the same time, remaining the accuracy.

Multi-scale calculation method consist of reducing an heterogeneous problem to an homogeneous one, such process is outlined at Fig. (1). During the assembling stage of the global matrix that corresponds to the macro-scale problem, in each integration point a PDE over a representative volume element (RVE) is solved giving average values of \mathbb{C} and $\bar{\boldsymbol{\tau}}$.

It is important to remark that the RVE problem considers the periodical micro-structure of the composite material and

the boundary conditions that are given by the macro-problem.

A considerable part of this work is to acquire the capacity of modeling the most important phenomenons that occurs in composite materials which are going to be solved over each RVE during the multi-scale calculations. In the present work a damage model it is explained.

III. DAMAGE MODEL

Continuum damage models are widely used to represent the complex behavior of materials when they are subjected to large loads and they are out of the elastic and linear range.

During a multi-scale simulation on every RVE that it is subjected to large loads a damage model is going to be activated. By this way, an isotropic damage model based on [3] has been implemented in order to understand the basics of this phenomenon.

The general isotropic damage model (see [4]) can be formulated as:

$$\bar{\tau} = (1 - d)\mathbb{C}_0 : \bar{\epsilon}, \quad (4)$$

where d is the damage variable ($d \in [0, 1]$) and \mathbb{C}_0 is the undamaged constitutive tensor. Also a damage criterion, which determines when the material starts damaging, should be defined. For this case it adopts an inequality form like:

$$F(\tau, r) = G(\tau) - G(r) \leq 0, \quad (5)$$

where G is a function that depends on the material properties, τ is a scalar variable that determines the way of how the stress tensor $\bar{\tau}$ produce damage. When that inequality stops holding then the material starts damaging increasing the values of d .

In this work a 1D model was implemented using the finite element approach in order to test the objective response of the model, this means that the solution should not depend on the element size of the mesh. In Fig. 2 two solutions for a 1D case subjected to displacements boundary conditions are shown for two meshes that differ on the node number. As it can be seen, objectivity for this case has been achieved.




Fig. 2. Objectivity in a 1D finite element model implemented for damage calculations.

In Fig. 3 a 2D damage result is shown. At this point, some difficulties have been observed on the convergence of the Newton-Raphson method used for solving the non-linear equations. This is produced because the residue of the non-linear equations reaches a zero derivative point which is a situation where Newton-Raphson schemes fail. To solve this an Arc-Length method (see [5]) is been implemented and tested in order afford more complex problems.




Fig. 3. Damage calculation in a 2D traction test device.

IV. CONCLUSIONS

Multi-scale calculation concept has been presented here as the main line of work. The technique will be applied inside the Alya HPC-code after composite materials calculations have been done in smaller scale codes and tested. In this stage non-linear behavior that results from damage, plasticity and large-deformations should be dominated in order to be applied on multi-scale.

Up to now small scale codes have been developed in order to calculate the non-linear behavior of materials and to measure the performance of the numerical methods involved. The most important code takes advantage of the PETSc library to perform a domain decomposition finite element algorithm and to solve large systems of equations and problems that require big amounts of memory.

REFERENCES

- [1] S. Oller, *Simulación Numérica del Comportamiento Mecánico de los Materiales Compuestos*, 1era edición. Centro Internacional de Métodos Numéricos en Ingeniería, Barcelona, España, (CIMNE), 2003.
- [2] K.J. Bathe, *Finite Element Procedures*, 2nd edition. Prentice Hall, Pearson Education, 2016.
- [3] J. Oliver, M. Cervera, S. Oller and Lubliner, *Isotropic Damage Models and Smeared Crack Analysis of Concrete*, Second International Conference On Computer Aided Analysis and Design of Concrete Structures, 1990.
- [4] S. Oller, *Fractura Mecánica, un enfoque global*, 1era edición. Centro Internacional de Métodos Numéricos en Ingeniería, Barcelona, España, (CIMNE), 2001.
- [5] S. Oller, *Dinámica No-Lineal*, 1era edición. Centro Internacional de Métodos Numéricos en Ingeniería, Barcelona, España, (CIMNE), 2002.

Saiph, a Domain Specific Language for Computational Fluid Dynamics simulations

Sandra Macià, Vicenç Beltran, Daniel Mira and Sergi Mateo
Barcelona Supercomputing Center (BSC-CNS)

sandra.macia@bsc.es, vicenc.beltran@bsc.es, daniel.mira@bsc.es and sergi.mateo@bsc.es

Abstract—Nowadays, High-Performance Computing (HPC) is assuming an increasingly central role in scientific research while computer architectures are becoming more and more heterogeneous and using different parallel programming models and techniques. Under this scenario, the only way to successfully exploit an HPC system requires that computer and domain scientists work closely towards producing applications to solve domain problems, ensuring productivity and performance at the same time. Facing such purpose, Saiph is a Domain Specific Language designed to ease the task of solving couple and uncouple Partial Differential Equations (PDE's), with a primary focusing on Computational Fluid Dynamics (CFD) applications. Saiph allows to model complex physical phenomena featured by PDE's, easing the use of numerical methods and optimizations on different computer architectures to the users.

I. INTRODUCTION

THIS project aims to ease the development of scientific applications by allowing domain experts to transcribe their equations into the code Saiph and then generating HPC-ready code that efficiently exploits the computational resources of modern heterogeneous supercomputers while dealing with all the specific aspects of solving systems of PDE's. To achieve that, Saiph provides a high-level syntax that directly maps with concepts of the domain, hiding from the user all the complexities related to numerical methods and HPC systems. Users only have to translate their equations to Saiph language and specify some physical and numerical parameters; initial and boundary conditions and post-processing strategy. Later, domain optimizations can be internally applied providing extra efficiency and correctness boosting the workflow productivity. The final specialized system of equations is then solved in parallel using MPI and applying intra-node parallelization techniques (using OpenMP/OmpSs) to achieve high computational performance.

II. SAIPH OVERVIEW

This section introduces the design and underlying technology used for the development of Saiph and the resulting high-level language. We briefly describe the Saiph project and its state of development.

A. Saiph design and underlying technologies

Saiph, as a DSL, has been designed to be simple, efficient, largely applicable in Computational Mechanics problems and in particular on CFD applications. It has been implemented




Fig. 1. Underlying design and technologies of the compilation process

as an embedded compiler in Scala[1] using the Lightweight Modular Staging (LMS)[2] as a DSL development platform and the Scala Virtualized Compiler[3].

Saiph applications are compiled at the front end with the LMS and the Saiph implementation together using the Scala Virtualized Compiler. At the middle end, the output of the first phase is compiled using our embedded compiler; the domain specific optimizations implementations are applied at this point and the LMS generates the corresponding IR nodes. Finally, the output of the embedded compiler (C++ file) is compiled and linked with our low-level C++ library that handles the numerical and parallelization issues, producing a binary ready to be executed in parallel. Figure 1 shows the compilation process and the internal design structured by layers.

Saiph has two main layers: the Scala compiler and the C++ library. This separation eases the DSL development, as in each of them the efforts are devoted to the developments naturally belonging to the layer. In that way, we can take advantage of each tool being used at its natural layer. High-level domain-oriented syntax and domain specific optimizations are thus implemented at the Scala compiler, while MPI and OpenMP libraries and auxiliary, mainly numerical methods, are integrated into the C++ library.

B. Saiph as a language

Saiph offers a high level syntax to unambiguously define a complete system of PDE's used for the characterization of a physical phenomena. The main components are presented in this section.

1) *Units*: The basic component of Saiph, they represent a physical magnitude through a value and its dimensionality information. *Units* can be combined and compared, and are internally used to check the consistency of the equations.

```
def L = 1 * Meters
```

2) *Mesh*: Saiph works with Cartesian meshes that can be defined in 1D, 2D or 3D. The sizes are specified by the users using physical continuous space.

```
val mesh = CartesianMesh(L, L)
```

3) *Terms and Consterms*: Saiph offers these components to represent dimensional variables and constants for the problem.

```
val T1 = Term(Temperature)("Temp1", mesh, 300 * Kelvins)
```

4) *Operators*: There is a complete set of numerical operators available to combine the different terms in order to build the system of discretized equations.

5) *Equations*: Defining equations in Saiph involves declaring terms and combine them through operators. An equation is formed by the left-hand side and the right-hand side expressions. Consequently, the units of both sides must match, otherwise, Saiph emits and error.

```
val eq1 = Equation(lhs_expr, rhs_expr)
```

C. Internal features

Internal features are completely transparent to the user.

1) *Numerical methods*: Numerical evaluations require the discretization of continuous functions, models, and equations which are time-space dependent. The spatial discretization follows explicit high-order schemes based on the finite difference method for which different operators are available, while the time marching is based on high-order Runge Kutta methods. Saiph uses non-uniform structured meshes.

2) *Exploiting parallelism*: Inter and intra-node parallelism are harmoniously combined. For the inter-node parallelism, the mesh is partitioned by the last dimension and a similar workload is distributed across the available MPI processes. Each process solve its part of the mesh for the whole simulation. Computations at each time-step are parallel and dependence's free but, after each of them, each MPI process has to exchange its boundaries with its neighbors in order to correctly update all the values to be used for the spatial derivatives of future computational steps. Regarding intra-node parallelism, each equation can be integrated in parallel at any time-step. The iteration space of the loop traversing the mesh is distributed across available OpenMP threads and executed in parallel. Each equation is solved in parallel for all the points of the mesh, one equation after the other.

3) *Domain specific optimizations*: Saiph has features and components which does not change the user external interface while changing the Saiph internal behaviour. Those optimizations are specific for the resolution of PDEs systems and even more specific for the resolution of CFD problems. As an example of those optimizations, we considered the *advection term* of a convection-diffusion-reaction (CDR) equation. When identified, this term can be treated with different operators depending on the type of problem. For instance, for convection-dominated flows, a low-dissipation *central scheme* can be used, although *upwind differentiation* can be activated when there is no sufficient resolution to capture the sharp of the gradients.




Fig. 2. Saiph - Two-dimensional advection-diffusion simulation

III. APPLICATION EXAMPLE

We present an advection-diffusion problem, represented by the following partial differential equation.

$$\frac{\partial T}{\partial t} = -\mathbf{u} \cdot \nabla T + \nabla \cdot (k \nabla T) \quad (1)$$

A small hot 2D cube is being transported within a periodic domain. The following snippet code corresponds to the transcription of this equation into Saiph code.

```
val eq = Equation(dt(T), -u*grad(T) + div(k*grad(T)))
```

Figure 2 shows the results from this advection-diffusion Saiph application.

IV. CONCLUSION

Saiph appears to be a powerful tool that can be advantageously used by scientists without knowledge on numerical methods and High-Performance Computing, while internally providing the advantages of such expertise. Several domain specific optimizations can be implemented as well as new suitable numerical methods and parallelization strategies boosting the efficiency, flexibility and genericity while maintaining its usability.

REFERENCES

- [1] M. Odersky and al., "An Overview of the Scala Programming Language," EPFL, Lausanne, Switzerland, Tech. Rep. IC/2004/64, 2004.
- [2] T. Rompf and M. Odersky, "Lightweight modular staging: a pragmatic approach to runtime code generation and compiled dsls," ser. GPCE '10. New York, NY, USA: ACM, 2010, pp. 127–136. [Online]. Available: <http://doi.acm.org/10.1145/1868294.1868314>
- [3] A. Moors, T. Rompf, P. Haller, and M. Odersky, "Scala-virtualized," ser. PEPM '12. New York, NY, USA: ACM, 2012, pp. 117–120. [Online]. Available: <http://doi.acm.org/10.1145/2103746.2103769>



Sandra Macià studied physics at the UB. After obtaining her degree she enrolled the Master in Innovations and Research in Informatics, on the High Performance Computing specialization, MSc at UPC-FIB, where she obtained the *Severo-Ochoa MSc scholarship*. For her master thesis she joined the *Barcelona Supercomputing Center*, Computer Science Dpt., where she started to work on Domain Specific Languages for the resolution of systems of Partial Differential Equations. Currently she is developing her PhD studies as a Severo Ochoa PhD student, targeting the same subject and with focus on Computational Fluid Dynamics applications.

Excited state gradients within a polarizable QM/MM formulation

Maximilian F.S.J. Menger*, Stefano Caprasecca, and Benedetta Mennucci

Dipartimento di Chimica e Chimica Industriale, University of Pisa*m.menger@studenti.unipi.it

Abstract—Multiscale approaches that partition the system into an active site (where the electronic process under study occurs) and a remaining region, the environment, have proven to be good strategies for the computation of electronic excitations in complex systems. In this work the implementation of a polarizable QM/MM scheme for the computation of excited state gradients is presented and are applied to a test case.

1 INTRODUCTION

In standard QM/MM models, the environment is described in terms of fixed charges (or multipoles), that remain unchanged during the calculation. In our polarizable QM/MM approach (QM/MMPol), on the other hand, polarizabilities on the MM sites are included additionally (see Fig. 1). Therefore the environment is allowed to respond to the electron density of the QM subsystem or its changes upon excitation by polarizing. The polarization of the environment is represented in terms of classical dipoles, induced by the QM electric field F_p . The dipoles are obtained by

$$\mathbf{F}_p = (\mathbf{D}\boldsymbol{\mu})_p$$

where \mathbf{D} is the MMPol matrix, that only depends on the geometrical parameters of the MM sites and the atomic polarizabilities.

The effective interaction energy between QM and MM region of the system is given by:

$$\begin{aligned} E^{es} &= U_q^{es} + U_n^{es} + U_e^{es} \\ E^{pol} &= \frac{1}{2} [U_{en}^{pol} + U_{ne}^{pol} + U_{ee}^{pol} + U_{nn}^{pol}] \end{aligned}$$

Hereby is U_i^{es} the electrostatic interaction energy between MM charges and the potential by the other MM charges ($i = q$), the nuclei ($i = n$) and the electron density ($i = e$). E^{pol} denotes the interaction energy between the induced dipoles and the electric fields generated by charges, nuclei, electrons and the other dipoles.




Fig. 1. Schematic representation of the polarizable QM/MM approach.

2 POLARIZABLE TDDFT/MM GRADIENTS

The energy derivative of a system represented by its wavefunction Ψ and its corresponding Hamiltonian \mathbf{H} with respect to a generic parameter λ can in general be written as:

$$E^\lambda = \left\langle \Psi \left| \frac{\partial \mathbf{H}}{\partial \lambda} \right| \Psi \right\rangle + 2 \left\langle \frac{\partial \Psi}{\partial \lambda} \left| \mathbf{H} \right| \Psi \right\rangle \quad (1)$$

The first term is the so-called Hellmann Feynman force, while the second term is the Pulay force. In the case of a purely variational energy functional the Pulay force vanishes. This is also desirable in the case for non variational energy functionals, as the computation of the molecular orbital response with respect to the perturbation is very costly. A better ansatz is to introduce a so-called relaxed density to account for the wavefunction response. This avoids the computation of the perturbed basis.

In our approach for the analytic gradients of the polarizable TDDFT/MM approach, we use an ansatz that was already applied by Furche and Ahlrichs [1] to pure TDDFT and later expanded by Scalmani et al. [2] on TDDFT gradients in combination with the polarizable continuum model [3] (PCM). The starting point for this approach is a fully variational Lagrangian \mathbf{L} , defined as following:

$$\begin{aligned} \mathbf{L}[\mathbf{X}, \mathbf{Y}, \Omega, \mathbf{C}, \mathbf{Z}, \mathbf{W}] &= G[\mathbf{X}, \mathbf{Y}, \Omega] + \sum_{ia} \mathbf{Z}_{ia} \mathbf{F}_{ia} \\ &\quad - \sum_{pq, p \leq q} \mathbf{W}_{pq} (\mathbf{S}_{pq} - \delta_{pq}) \end{aligned}$$

The derivation of analytic gradients from \mathbf{L} is straightforward and the resulting Z-Vector method is computational advantageous with respect to others.

The equation for the analytic derivatives can be written in a compact form, by expressing it in the AO basis:

$$\begin{aligned} \Omega^\xi &= \sum_{\mu\nu} h_{\mu\nu}^\xi P_{\mu\nu}^\Delta + \sum_{\mu\nu} S_{\mu\nu}^\xi W_{\mu\nu} + \sum_{\mu\nu\kappa\lambda} \langle \mu\nu | \kappa\lambda \rangle^\xi \Gamma_{\mu\nu\kappa\lambda} \\ &\quad + \omega^{xc, \xi} + \omega^{MMPol, \xi} \end{aligned}$$

The exponent ξ denotes derivatives with respect to a general variable. P^Δ and Γ are the one- and two-particle density matrices. $\omega^{xc, \xi}$ is the derivative of the DFT exchange-

correlation energy and $\omega^{MMPol,\xi}$ of the QM-MMPol interaction energy. The exact expressions and a detailed discussion can be found in the articles of Furche [1] and Scalmani [2] for the case of pure DFT and DFT/PCM.

3 IMPLEMENTATION AND APPLICATION

3.1 Implementation

The analytic gradients are implemented in a local modified version of Gaussian 09 [4]. Analytical Forces for QM and MM sites in excited state calculations at (TDA)-TDDFT/TDHF level of theory for a polarizable embedding are available within the ONIOM scheme. The implementation is an extension of the work of Caprasecca et al. [5], who implemented the corresponding ground state QM/MMPol gradients.

3.2 Application

We applied the model to study the intercalation of the doubly protonated fluorescent stain 4',6-diamidino-2-phenylindole (DAPI, shown in Fig. 2) within two pairs of CG DNA bases. DAPI is a commonly used fluorescent marker that can interact with DNA both by minor groove binding and intercalation. In this test study we will only focus on the latter. The initial structure of the system




Fig. 2. Structure of the doubly protonated 4',6-diamidino-2-phenylindole (DAPI).

was taken from the article of Biancardi et al. [6] and the conformer 1 of the intercalated structures was used. Both ground state (GS) and excited state (ES) geometry optimizations were performed and the QM method of choice was DFT/TDA-TDDFT. To assess the effect of polarization on the optimized structures, three different models were introduced: (i) a full-QM model, where both DAPI and DNA are treated at QM level (f-QM); (ii) a non polarizable QM/MM model, where the DAPI is treated at QM level and the DNA is treated at MM level; (iii) a polarizable QM/MM model (QM/MMPol), where the DNA is now represented as a polarizable environment. Computations of the isolated DAPI molecule were performed (g-QM), additionally. In all calculations the DFT functional m06-2X and the 6-31G(d) basis set was used for the DAPI molecule, while 6-31G was used on the DNA in the f-QM calculations. The excited state of interest is the lowest state localized on the DAPI molecule, characterized by a HOMO-LUMO transition in the g-QM model. This state corresponds to the first excited state in the QM/MM and QM/MMPol approaches and the third for f-QM, as two charge transfer state appear at lower energy.

To assess the structural changes induced by the environment on ground and excited state geometries, as well as the structural rearrangement upon excitation, eleven internal

bonds (marked in red in Fig. 2) were chosen. From these it is possible to define the bond length alternation (BLA), which is often used as a measure of conjugation. The bond difference between GS and ES calculations are shown in Fig. 3. It is evident from the alternating pattern of the first 8 bond length that, upon excitation, the single (double) bonds tend to become shorter (longer), with an increase in conjugation as consequence. The comparison between different models for the inclusion of the environment do not show significant differences. In several cases, however, the optimization carried out with the newly implemented QM/MMPol model is closer to the more expensive f-QM optimizations than the other models. This seems to point out that the inclusion of the polarisation is particularly useful in describing accurately the environment response to electron density redistribution.




Fig. 3. Difference in bond length of eleven internal bonds and the corresponding BLA in the DAPI molecule (marked in Fig. 2) between ES and GS optimized structures at different levels of theory.

ACKNOWLEDGMENTS

Financial support of the ITN-EJD-TCCM PhD Fellowship as part of the Marie Skłodowska-Curie Actions is greatly acknowledged.

REFERENCES

- [1] F. Furche, R. Ahlrichs: "Adiabatic time-dependent density functional methods for excited state properties", *J. Chem. Phys.* **2002**, 117, 7433; Erratum *J. Chem. Phys.* **2004**, 121, 12772.
- [2] G. Scalmani, M. J. Frisch, B. Mennucci, J. Tomasi, R. Cammi, V. Barone: "Excited States in the Gas Phase and in Solution. Theory and Application of a Time-Dependent DFT Polarizable Continuum Model", *J. Chem. Phys.* **2006**, 124, 094107.
- [3] J. Tomasi, B. Mennucci and R. Cammi: "Quantum Mechanical Continuum Solvation Models", *Chem. Rev.* **2005**, 105, 2999.
- [4] M. J. Frisch et al.: *Gaussian 09*, Gaussian Inc. Wallingford CT **2009**.
- [5] S. Caprasecca, S. Jurinovich, L. Viani, C. Curutchet, B. Mennucci: "Geometry Optimization in Polarizable QM/MM Models: The Induced Dipole Formulation", *J. Chem. Theory Comput.* **2014**, 10, 1588.
- [6] A. Biancardi, T. Biver, F. Secco and B. Mennucci: "An investigation of the photophysical properties of minor groove bound and intercalated DAPI through quantum-mechanical and spectroscopy tools", *Phys. Chem. Chem. Phys.* **2013**, 15, 4596.
- [7] J. Wang, P. Cieplak, J. Li, T. Hou, R. Luo, Y. Duan: "Development of Polarizable Models for Molecular Mechanical Calculations I: Parameterization of Atomic Polarizability" *J. Phys. Chem. B* **2011**, 115, 3091.
- [8] M. Caricato, B. Mennucci, J. Tomasi, F. Ingrosso, R. Cammi, S. Corni, G. Scalmani: "Formation and relaxation of excited states in solution: A new time dependent polarizable continuum model based on time dependent density functional theory" *J. Chem. Phys.* **2006**, 124, 124520.

Maximilian F.S.J. Menger Bachelor (2013) and Master degree (2015) in Chemistry from the Ruprecht-Karls University Heidelberg with a dissertation on ultrafast electron dynamics in phenylalanine, under the supervision of PD Alexander Kuleff. Currently, he works as a PhD student within the EJD-ITN-TCCM Marie Skłodowska-Curie actions in the groups of Prof. Benedetta Mennucci (University of Pisa) and Prof. Leticia González (University of Vienna), with a project on The modelling of light-harvesting processes in multichromophoric systems. His research interest include QM/classical models as well as diabatic and non adiabatic molecular dynamics.

Impact of Aerosol Microphysical Properties on Mass Scattering Cross Sections

Vincenzo Obiso*, Oriol Jorba*, Marco Pandolfi†, Marina Ealo†

*Barcelona Supercomputing Center (BSC-ES), c/ Jordi Girona 29, 08034, Barcelona, Spain.

Email: vincenzo.obiso@bsc.es (V. Obiso); oriol.jorba@bsc.es (O. Jorba).

†Institute of Environmental Assessment and Water Research (IDAEA-CSIC), c/ Jordi Girona 18-26, 08034, Barcelona, Spain.

Email: marco.pandolfi@idaea.csic.es (M. Pandolfi).

Abstract—We assessed the sensitivity of simulated mass scattering cross sections (α_{λ}^{sca} [m^2/g]) of three aerosol species to perturbed particle microphysical properties and derived constraints on these microphysical properties, suitable for the north-western Mediterranean basin, from a comparison between code calculations and observations. In detail, we calculated α_{λ}^{sca} of mineral dust, organic carbon and sulfate at three wavelengths in the visible range with a T-matrix optical code, considering $\pm 20\%$ perturbations on size distribution, refractive index and mass density, and spheroids with two different axial ratios as shape perturbations. Then, we compared the simulation results with a set of observed α_{λ}^{sca} of mineral dust, aged organics and ammonium sulfate sources provided by the Institute of Environmental Assessment and Water Research (IDAEA-CSIC) and representative of the north-western Mediterranean basin.

I. INTRODUCTION

Atmospheric aerosols can scatter and absorb electromagnetic radiation, causing a redistribution of the radiative energy in the atmosphere [1]. The Aerosol-Radiation Interaction (ARI) radiative forcing still contributes to dominate the uncertainty associated with the anthropogenic contribution to the climate change [2]. Also the role of the natural aerosols in affecting the Earth's radiative balance through ARI is poorly constrained [3]. The ARI parameterization mainly consists in the characterization of the aerosol optical properties, which in turn depend on the microphysical properties of the particles [1], [4]. Hence, errors in the particle microphysical assumptions can affect the calculation of the optical properties and so the assessments of the ARI radiative effects [5], [6]. For this reason, we consider a study on the relationship between microphysical and optical properties a recommendable first step to better parameterize the ARI in the atmospheric models.

II. T-MATRIX CODE

The T-matrix method is a powerful exact technique for computing light scattering by non-spherical particles [7]. For this work, the T-matrix code by [7] has been used. It allows calculating integrated optical properties of poly-disperse, randomly oriented, rotationally symmetric, homogeneous particles.

III. OBSERVATIONS

We used a set of observed aerosol α_{λ}^{sca} provided by the IDAEA-CSIC. [8] collected measurements of dry ($RH < 40\%$) aerosol PM_{10} mass concentrations (gravimetric masses from 24h filters) and optical properties (scattering coefficients with Nephelometer AURORA 3000, Ecotech), in the period 2010-2014, at the Montseny regional background station (middle altitude emplacement within the Montseny Natural Park, Spain: $41^{\circ}46'45, 63^{\circ}N - 02^{\circ}21'28, 92^{\circ}E$;

$720m$ a.s.l.). From these data, [9] derived α_{λ}^{sca} at 3 visible wavelengths ($\lambda_1 = 0,450\mu m$; $\lambda_2 = 0,525\mu m$; $\lambda_3 = 0,635\mu m$) through Multilinear Regression (MLR) analysis applied to aerosol sources [10] detected in the PM_{10} mass chemical speciated data [11] through the application of the Positive Matrix Factorization (PMF) model [12]. In our work we used three of the seven sources detected at the Montseny station: mineral dust (typical crustal elements), aged organics (mainly organic carbon from biogenic sources but with a significant contribution from biomass burning) and ammonium sulfate (mainly SO_4^{2-} and NH_4^+) [10]. These data can be considered representative of the regional background of the north-western Mediterranean basin.

IV. EXPERIMENT SETUP

We defined the reference microphysical properties following the aerosol parameterization used in the NMMB-MONARCH model (formerly known as NMMB/BSC-CTM) [13], [14], [15], which in turn is based on the OPAC database [16], but with some deviations. We performed calculations of integrated α_{λ}^{sca} of mineral dust, organic carbon and sulfate. The calculations have been performed at the same three visible wavelengths of the observations ($\lambda_1 = 0,450\mu m$; $\lambda_2 = 0,525\mu m$; $\lambda_3 = 0,635\mu m$). Then, in order to generate variability in the microphysical assumptions for the particles, we gave as inputs to the T-Matrix code the reference microphysical properties independently perturbed. In particular, for each aerosol species, we considered perturbations of $\pm 20\%$ on size distribution (separately for geometric radius r_g and standard deviation σ_g), refractive index (separately for real n_R and imaginary n_I parts) and mass density ρ ; as perturbations for the spherical shape we considered two types of spheroid: moderate and extreme, averaging for each axial ratio χ the oblate and the prolate options.

V. DATA ANALYSIS

We compared simulated and observed α_{λ}^{sca} associating mineral dust with the mineral dust source, organic carbon with the aged organics source and sulfate with the ammonium sulfate source. At first, we evaluated the sensitivity of α_{λ}^{sca} values and spectral dependence to the different perturbed microphysical properties. Regarding the α_{λ}^{sca} values, for each species and microphysical property we estimated the extent of the variability range (generated by the perturbations) of the simulated value at $\lambda_2 = 0,525\mu m$ respect to the uncertainty affecting the observed value at the same wavelength. Regarding the α_{λ}^{sca} spectral dependence, instead, we estimated the extent of the variability range of the log-linear spectral dependence slope of the simulated values respect to the uncertainty affecting the slope of the observed values. The spectral dependence slopes have been derived by performing linear fits of the logarithmic spectral dependence of both simulated and observed α_{λ}^{sca} , taking into account

Table I

PRESCRIPTIONS DERIVED FROM OUR ANALYSIS FOR SIZE DISTRIBUTION (SD), REFRACTIVE INDEX (RI), MASS DENSITY (DN) AND SHAPE (SH) OF MINERAL DUST (DU), ORGANIC CARBON (OC) AND SULFATE (SU). THE VALUES REPORTED FOR THE MICROPHYSICAL PARAMETERS MARKED BY * ARE REFERENCE VALUES, FOUND TO BE SUITABLE FOR THE CORRESPONDENT SPECIES IN THE GEOGRAPHICAL AREA UNDER STUDY.

	DU	OC	SU
SD: $r_g(\mu m)$	$3,583 \cdot 10^{-1}$	$2,544 \cdot 10^{-2}$	$8,340 \cdot 10^{-2}$
σ_g	1,600	1,760	1,624
$r_{eff}(\mu m)$	$6,221 \cdot 10^{-1}$	$5,656 \cdot 10^{-2}$	$1,501 \cdot 10^{-1}$
RI: n_R	(1,530; 1,530; 1,530)(*)	(1,576; 1,576; 1,576)	(1,547; 1,545; 1,543)
n_I^*	$(8,500; 6,650; 4,500) \cdot 10^{-3}$	$(1,730; 1,250; 0,696) \cdot 10^{-2}$	$(1,000; 1,000; 1,610) \cdot 10^{-8}$
DN: $\rho^*(kg/m^3)$	$2,506 \cdot 10^3$	$1,800 \cdot 10^3$	$1,700 \cdot 10^3$
SH: χ^*	1,000	1,000	= 1,000

the measurement errors as weights for the observation fits. In order to complete the comparison analysis and to constrain the perturbed microphysical properties, then, we performed a compatibility test on the best fit parameters. The test has been performed, for each aerosol species, on the best fit parameters of all the perturbed simulations respect to those of the correspondent observations.

VI. RESULTS

We found that the mineral dust α_{λ}^{sca} values are only affected by the size distribution and, with a lower impact, by the mass density perturbations. On the other hand, no microphysical properties seem to have any impact on the α_{λ}^{sca} spectral dependence. So, it appears that, due mainly to the size of the particles bigger than the visible wavelengths, the dust α_{λ}^{sca} are quite stable respect to the microphysical perturbations (in the spectral range of the experiment). The organic carbon α_{λ}^{sca} values are affected mainly by refractive index (real part), size distribution and, with a lower impact, by mass density and shape perturbations. The α_{λ}^{sca} spectral dependence, instead, is significantly affected only by the size distribution perturbations. So, the organic carbon α_{λ}^{sca} appear less stable than the mineral dust ones respect to the microphysical perturbations, due mainly to the size of the particles smaller than the visible wavelengths. The sulfate α_{λ}^{sca} values are mainly affected by the refractive index (real part) and, with a lower impact, by size distribution, mass density and shape perturbations. On the other hand, the α_{λ}^{sca} spectral dependence is only affected by size distribution and refractive index perturbations. So, it seems that the sulfate α_{λ}^{sca} are the most unstable respect to the microphysical perturbations, due to the particle size approximately comparable with the visible wavelengths. The microphysical prescriptions derived from our analysis for the three aerosol species, suitable for the north-western Mediterranean Basin, are reported in Table I.

ACKNOWLEDGEMENTS

This work has been presented in the paper "Impact of aerosol microphysical properties on mass scattering cross sections" (Obiso et al. [2017]), accepted for publication by *Journal of Aerosol Science*.

REFERENCES

- [1] O. Boucher et al., "Clouds and Aerosols," 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*, T. F. Stocker et al., Ed. Cambridge, United Kingdom and New York, NY, USA: Cambridge University Press, 2013.
- [2] G. Myhre et al., "Anthropogenic and Natural Radiative Forcing," 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*, T. F. Stocker et al., Ed. Cambridge, United Kingdom and New York, NY, USA: Cambridge University Press, 2013.
- [3] A. Rap et al., "Natural aerosol direct and indirect radiative effects," *Geophys. Res. Lett.*, vol. 40, pp. 3297–3301, 2013.

- [4] J. L. Hand and W. C. Malm, "Review of aerosol mass scattering efficiencies from ground-based measurements since 1990," *J. Geophys. Res.*, vol. 112, D16203, 2007.
- [5] H. Yu et al., "A review of measurement-based assessments of the aerosol direct radiative effect and forcing," *Atmos. Chem. Phys.*, vol. 6, pp. 613–666, 2006.
- [6] Z. Zhang et al., "Shortwave direct radiative effects of above-cloud aerosols over global oceans derived from 8 years of CALIOP and MODIS observations," *Atmos. Chem. Phys.*, vol. 16, pp. 2877–2900, 2016.
- [7] M. I. Mishchenko and L. D. Travis, "Capabilities and limitations of a current FORTRAN implementation of the T-matrix method for randomly oriented, rotationally symmetric scatterers," *J. Quant. Spectrosc. Radiat. Transfer*, vol. 60, pp. 309–324, 1998.
- [8] M. Pandolfi et al., "Variability of aerosol optical properties in the Western Mediterranean Basin," *Atmos. Chem. Phys.*, vol. 11, pp. 8189–8203, 2011.
- [9] M. Ealo et al., "From air quality to climate; Impact of aerosol sources on optical properties at urban, regional and continental levels in the north-western Mediterranean," *Atmos. Chem. Phys.*, vol. (in preparation), unpublished results.
- [10] M. Pandolfi et al., "Trends analysis of PM source contributions and chemical tracers in NE Spain during 2004–2014: A multi-exponential approach," *Atmos. Chem. Phys.*, vol. 16, pp. 11787–11805, 2016.
- [11] X. Querol et al., "Variability in regional background aerosols within the Mediterranean," *Atmos. Chem. Phys.*, vol. 9, pp. 4575–4591, 2009.
- [12] P. Paatero, "Least squares formulation of robust non-negative factor analysis," *Chemometrics and Intelligent Laboratory Systems*, vol. 37, pp. 23–35, 1997.
- [13] O. Jorba et al., "Potential significance of photoexcited NO₂ on global air quality with the NMMB/BSC chemical transport model," *J. Geophys. Res.*, vol. 117, D13301, 2012.
- [14] C. Pérez et al., "Atmospheric dust modeling from meso to global scales with the online NMMB/BSC-Dust model - Part 1: Model description, annual simulations and evaluation," *Atmos. Chem. Phys.*, vol. 11, pp. 13001–13027, 2011.
- [15] M. Spada et al., "Modeling and evaluation of the global sea-salt aerosol distribution: sensitivity to size-resolved and sea-surface temperature dependent emission schemes," *Atmos. Chem. Phys.*, vol. 13, pp. 11735–11755, 2013.
- [16] M. Hess et al., "Optical Properties of Aerosols and Clouds: The Software Package OPAC," *Bull. Am. Met. Soc.*, vol. 79, pp. 831–844, 1998.



Vincenzo Obiso Vincenzo Obiso studied Theoretical Physics at the University of Rome (Italy). He is a Ph.D. student at the Earth Sciences Department of the Barcelona Supercomputing Center (BSC-ES, Barcelona, Spain) since 2013. His research plan focuses on the implementation of the dynamic aerosol-radiation interaction within the NMMB-MONARCH model, an online integrated meteorology-chemistry model under development at the BSC-ES, and on the evaluation of its impact on the meteorological forecasts.

Stabilization of Microturbulence by Fast Ions in ASDEX Upgrade

F.N. de Oliveira^{*1}, H. Doerk², M. J. Mantsinen^{1,3}, C. Angioni², R. Bilato², D. Gallart¹, A. Gutiérrez-Milla¹, P. Mantica⁴, T. Odstrčil², X. Sáez¹, G. Tardini², and the ASDEX Upgrade and EUROfusion MST1 Teams^{*}

¹Barcelona Supercomputing Center, Barcelona, Spain

²Max-Planck-Institut für Plasmaphysik, Garching, Germany

³ICREA, Barcelona, Spain

⁴Instituto di Fisica del Plasma “P. Caldirola”, CNR, Milano, Italy

^{*}see H. Meyer et. Al. Nuclear Fusion FEC 2016 Special issue (2017)

*felipe.deoliveira@bsc.es

Keywords— Nuclear Fusion, ICRF, microturbulence.

EXTENDED ABSTRACT

Tokamaks are devices used to confine a fusion fuel mixture heated to high temperatures in the range of hundreds of millions of degrees, arriving in a state of matter called plasma. Several auxiliary systems are used for this purpose. They include Neutral Beam Injection (NBI) and electromagnetic waves. The largest tokamak up-to-date, aiming to achieve self-sustained nuclear fusion, is the International Thermonuclear Experimental Reactor (ITER) [1] shown in Fig. 1. It is presently under construction in Cadarache, France.




Fig. 1 Cross section of the ITER tokamak [1]. Its reactor core consists of a vessel used to contain the gas and several magnetic coils used to generate the magnetic field that confines the plasma.

A. Introduction

It has been experimentally found that transport of particles and heat in tokamaks is larger than expected when compared to neoclassical theoretical predictions, *i.e.* based on Coulomb collisions with the addition of toroidal geometry. It is thought that this is due to the presence of microinstabilities which are self-organized small-scale turbulent plasma fluctuations that can lead to transport of particles and heat. Given that a steady state is obtained when sources and transport are in balance, microturbulence is believed to be a critical player in the quality of the confinement. In order to improve the confinement of the plasma and avoid losses of particles and heat to the vessel, mechanisms of stabilization of microturbulence such as those studied in this work are of particular interest.

B. Ion Cyclotron Resonance Heating (ICRH)

One of the main auxiliary heating systems foreseen for ITER is the Ion Cyclotron Resonance Heating (ICRH). In this system, powerful electromagnetic wave antennas are used to transfer energy to the plasma.

It has been suggested [2,7-8] that the use of ICRH has helped to increase the ion-bulk temperature in the Joint

European Torus (JET) tokamak by stabilizing microturbulent effects. In that case, a relaxation of ion pressure profile stiffness is believed to be induced by microturbulence stabilization, arising in turn from fast ion suprathreshold pressure gradients.

The goal of the present work is to study the impact of fast ions, generated by ICRH, on microturbulence. Our methodology is similar to that in [2] and is based on extensive simulations of discharges carried out on the ASDEX Upgrade (AUG) tokamak, in Germany using state-of-art fusion simulations codes.

C. Method

The three main codes used in the work are PION [3], FIDO [4] and GENE [5]. The Finite Ion Drift Orbit (FIDO) code and the PION code are used to compute Ion Cyclotron Resonance Heating in tokamaks. They are used to study the waves, interaction with the plasma and the profile of ICRH power deposition in the plasma. The output of these codes is used as input in GENE, enabling us to quantify the impact of the presence of fast ions in the stabilization of microturbulences. The Gyrokinetic Electromagnetic Numerical Experiment (GENE) code is used to compute gyroradius-scale fluctuations, including microturbulence, and the resulting transport coefficients.

GENE solves for the distribution function f in (x,v) space and time. Integrating f over velocity space yields the equation

$$\frac{\partial n}{\partial t} + v_E \cdot \nabla n_0 + v_E \nabla \cdot n + n_0 \nabla \cdot v_E + \frac{2}{m\Omega B^2} B \times \nabla B \cdot \nabla p = 0 \quad (1)$$

for the time evolution of the density n . Here n_0 stands for the initial density, v_E for the ExB drift velocity, m for the species mass, Ω for the cyclotron frequency, B for the magnetic field and p for the normalized pressure.

In the present work, we solve the linearized version of Eq. (1), which is less computationally expensive. This study yields the growth rate and frequency of different unstable plasma waves for a broad parameter range consistent with the experimental conditions of the AUG discharges under consideration.

D. Results

The two AUG discharges we consider in this work are discharges 31563 and 31562 with combined NBI and ICRH heating. They were prepared in the same way except for the on-axis magnetic field to vary the ICRH resonance position and the concentration of He-3 ions resonant with ICRH waves.

Figure 2 shows the measured ion temperature T_i for these discharges together with T_i for a reference discharge 31555 without ICRH heating. The results show an increase from 3 keV to 5.5 keV in central T_i with ICRH as compared to 31555 without ICRH.




Fig. 2 Ion temperature T_i profiles for AUG discharges 31562, 31563 and 31555 [6]. T_i in discharges 31562 and 31563 is higher than what is expected from bulk ion heating alone. $1 \text{ keV} = 11604525 \text{ K}$.

Figure 3 shows the fast ion pressure profile as calculated with FIDO. Discharge 31563 has a steeper fast ion pressure profile although the heating power is the same for both cases, $P_{\text{ICRF}} = 3.5 \text{ MW}$.




Fig. 3 Pressure profile of ^3He ions for discharges 31562 and 31563 calculated by FIDO.

A linear analysis done with GENE showed that fast ions have a strong impact on the instability growth rate of Ion Temperature Gradient (ITG) instability, as it can be seen in figure 4. In particular, our results indicate a dual impact of fast ions in the instability growth rate, affecting both ITG and Kinetic Ballooning Mode (KBM). The presence of fast ions (red and black curves) not only diminishes the value of the instability growth rate at low betas, but it also delays the onset of the rapid increase observed at higher betas, associated with the KBM instability, which is related to fluid dynamics activity. Here, the beta is the plasma pressure normalized to the magnetic pressure. As expected from KBM theory, a large impact of the safety factor q is also found. The parameter q measures the ratio of the toroidal and poloidal magnetic fields used to confine the plasma. In order to make accurate predictions, this parameter needs to be accurately measured.

E. Summary

In this work, we have studied the effect of fast ions on microturbulence stabilization for two specific AUG discharges with combined NBI and ICRH heating. According to the GENE simulations in linear mode, the presence of fast ions has a stabilizing effect on microinstabilities and could be responsible for the ion stiffness reduction associated with suprathermal pressure gradients.

Further work involves performing nonlinear simulations and comparing the simulated values of particle and heat fluxes with the ones found in reference [6].




Fig. 4 Instability growth rate of discharge 31563 with and without fast ions as a function of beta. The KBM threshold is characterized by the sharp upturn in the growth rate as a function of beta. The vertical line represents the experimental thermal beta. Red and black curves represent simulations with fast ions for safety factor q equals to 1.6 and 2.0, blue and pink are simulations without fast ions for q equals to 1.6 and 2.0


F. ACKNOWLEDGEMENT

This work has been carried out within the framework of the EUROfusion Consortium and has received funding from the Euratom research and training programme 2014-2018 under the grant agreement No 633053. The views and opinion expressed herein do not necessarily reflect those of the European Commission. We would also like to thank Red Española de Supercomputación (RES) for the resources granted.

References

- [1] www.iter.org
- [2] J. Citrin et al., Nucl. Fusion 54 (2014) 023008.
- [3] L.-G. Eriksson and T. Hellsten, Phys. Scripta 55 (1995) 70.
- [4] J. Carlsson, T. Hellsten and L.-G. Eriksson, Proc. Joint Varenna Lausanne Workshop 'Theory of Fusion Plasmas' 1994, p. 351.
- [5] F. Jenko, W. Dorland, M. Kotschenreuther, and B. N. Rogers, Physics of Plasmas 7 (2000) 1904.
- [6] M.J. Mantinen, et al., RADIO FREQUENCY POWER IN PLASMAS: Proceedings of the 21st Topical Conference, AIP Conference Proceedings, vol. 1689. pp. 030005:1-8, 2015.
- [7] J. Citrin et al., Phys. Rev. Letters, 111 (2013) 155001.
- [8] J. Garcia et al., Nucl. Fusion 55 (2015) 053007.

Author biography



Felipe Nathan de Oliveira was born in Brasilia, Brasil, in 1993. Currently Intern at the Barcelona Supercomputing Center, Nuclear Science and Engineering master candidate at the Polytechnic University of Catalonia, Spain, and master in Theoretical Plasma Physics from the University of Brasilia, Brazil. Graduate Physicist from the University of Brasilia, has interests in Theoretical Plasma Physics, Plasma Astrophysics and High Performance Computing. Attended to a two-years Internship at the Brazilian Space Agency and a one-year internship at the University of Brasilia Plasma Physics Laboratory.

Vortex Induced Vibration (VIV) of circular cylinders at high Reynolds numbers

D. Pastrana*

daniel.pastrana@bsc.es,

J. C. Cajas*, O. Lehmkuhl*, I. Rodriguez** and G. Houzeaux*

Abstract—In this contribution, preliminary results on the study of the VIV at high Reynolds numbers are presented. First, a validation of the code and the turbulent LES WALE formulation in conjunction with a low-dissipative discretization is presented by means of the study of the smooth static cylinder at sub-critical and critical regimes: $Re = 1.1 \times 10^4$ and $Re = 3 \times 10^5$. Finally, some results of the forced damped oscillating cylinder for the $Re = 1.1 \times 10^4$ case are shown. One of the main objectives of this project is to analyse how the movement of the cylinder affects the sub-critical and critical regimes and the influence of this changes in the aerodynamics forces at high Reynolds numbers present in industrial applications.

Index Terms—Fluid-Structure Interaction, Turbulent Flow, Vortex Induced Vibrations.

1 INTRODUCTION

THE so-called Vortex Induced Vibration phenomenon (VIV) has been of great interest not only in flow control, for reducing oscillations structures that may cause fatigue and degradation on the components of engineering structures, but also for energy harvesting. The canonical problem system, a rigid circular cylinder both, elastically mounted and forced to oscillate at low Reynolds numbers, has been extensively studied and a good understanding of the physics that drives the phenomenon has been achieved. However, some important debates and unresolved problems are still open. The reader is referred to the work of Williamson and Govardhan [1] for a well-explained and comprehensive overview of the phenomenon. On the other hand, there are still few works neither experimental nor numerical simulations at high Reynolds numbers due to measurement difficulties in the near-wall region for the former and the massive computational resources and extra modeling efforts required for the later. Here, we succinctly describe the methodology implemented and some result obtained in this last direction.

2 METODOLOGY

A state-of-the-art turbulence methodology was recently developed and implemented in Alya (BSC simulation code for multi-physics problems). This methodology combines a low dissipation finite element (FE) scheme with the wall-adapting local eddy-viscosity (WALE) sub-grid scale (SGS) model. The low dissipative discretisation is based on the same principles proposed by Verstappen and Veldman [3], generalised for Jofre *et al.* [4] and Trias *et al.* [7] for unstructured finite volumes and extended to FE by Lehmkuhl *et al.* [5]. The set of equations is time integrated using an energy conserving Runge-Kutta explicit method lately proposed by Cappuano *et al.* [6] combine with an eigenvalue base

time step estimator [7]. Finally, the pressure stabilisation is achieved by means of a non-incremental fractional step.

The cylinder is treated as a rigid body allows to translate in two degrees of freedom, the in-line (x-axis) and the cross-flow (y-axis) directions govern by the forced damped oscillator equation in non-dimensional form:

$$\frac{\partial^2 \tilde{x}_i}{\partial \tilde{t}^2} + 4\pi \tilde{f}_n \zeta \frac{\partial \tilde{x}_i}{\partial \tilde{t}} + (2\pi \tilde{f}_n)^2 \tilde{x}_i = \frac{2C_i}{\pi \tilde{m}} \quad (1)$$

with $i = 1, 2$, \tilde{m} the mass ratio, \tilde{f}_n dimensionless natural frequency ζ the damping ratio and C_i the forces coefficients.

3 RESULTS

In table 1 and Fig. 1 common physical statistical quantities and the near wake structure are presented for the fixed cylinder at $Re = 1.1 \times 10^4$ and compared with data available in the literature [2]. A good agreement of the statistical quantities and an accurate reproduction of the vortex formation physics of the present LES simulations with reported data can be seen.

TABLE 1

Physical statistical quantities for the fixed cylinder flow at $Re = 1.1 \times 10^4$; drag coefficient C_D , base pressure coefficient $-C_{pb}$, Strouhal number St , and root mean square of the lift coefficient C_L .

	C_D	$-C_{pb}$	St	C_L
LES $Re = 1.1 \times 10^4$	1.228	1.318	0.215	0.581
Dong <i>et al.</i> 2006 $Re = 10^4$ [2]	1.208	1.201	0.200	0.547

In Fig. 2 (a) the instantaneous velocity field at two span-wise locations are shown. Clearly, an asymmetry is observed not only in 2D, with a subcritical like separation in the bottom and more supercritical in the top, but also with a 3D structure which is in good agreement with recent experimental visualisation. Figures 2 (b) and (c) show the average drag force coefficient and Strouhal number. General speaking fairly good agreement is observed for these values compared with experimental and numerical data available.

• * Barcelona Super Computing Center.
• ** Universitat Politècnica de Catalunya BarcelonaTech.




Fig. 1. Contours of mean stream-wise velocity and normalized mean span-wise vorticity for experimental data at $Re = 10^4$ (upper row images), and LES at $Re = 1.1 \times 10^4$ for the present study (bottom row images). Top images were taken from [2].

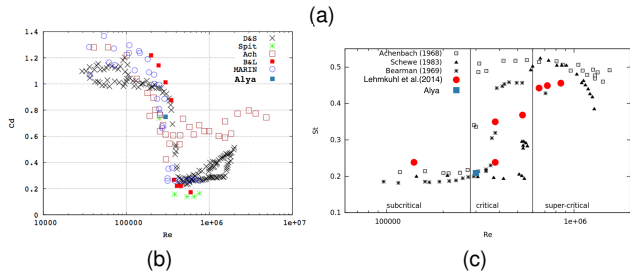



Fig. 2. (a) Instantaneous flow velocity field at different span-wise locations. (b) Drag coefficient (c) Strouhal number as a function the Reynolds numbers where the results for the present simulations at $Re = 3 \times 10^5$ have been included (blue rectangle).

Finally, a comparison between the damped and undamped oscillating cylinder is presented. The time history of force coefficients and cross-flow direction displacement, as well as the main components of their power spectra for the oscillating cylinder at $Re = 1.1 \times 10^4$ are presented. Besides of being clear the effect of the structural damping on the amplitude of the oscillations, it is also evident that the motion of the body with structural damping is also driven by a frequency close to the main frequency of the lift force, absent in the non-damped case.

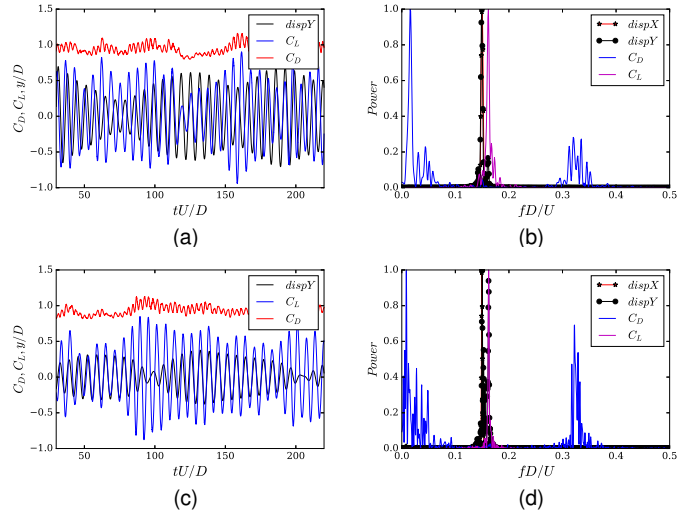


Fig. 3. Time history of lift and drag coefficients, and cross-flow direction displacement in flow past a circular cylinder at $Re = 1.1 \times 10^4$ and $f_N = 0.15$ (left). Power spectra of the lift and drag coefficients and body displacement in both directions (right). (Top) Non-damped. (bottom) Damping, structural damping $\xi = 0.5 \times 10^{-2}$.

4 CONCLUSION

A highly optimal and efficient turbulence methodology, recently developed and implemented in the Alya simulation code, has been tested and validated for the flow around a fixed circular cylinder at sub-critical and critical regimes. The methodology has been proved of being capable to accurately reproduce the physics present in both regimes and it has been applied to a more complex FSI scenario. From this experience first results for the flow around a rigid circular cylinder elastically mounted were presented and briefly discussed.

ACKNOWLEDGMENTS

The first author thanks to the Consejo Nacional de Ciencia y Tecnologia (CONACyT) for the SENER-CONACyT Ph.D. scholarship.

REFERENCES

- [1] C. H. K. Williamson and R. N. Govardhan, (2004) *Vortex-induced vibrations*, Annu. Rev. Fluid Mech. 36, 413–455.
- [2] S. Dong and G. E. Karniadakis, (2005) *DNS of flow past a stationary and oscillating cylinder at $Re = 10000$* , J Fluids Struct., 20, 519–531.
- [3] R. Verstappen and A. E. P. Veldman, (2003) *Symmetry-preserving discretization of turbulent flow*, Journal of Computational Physics, 187, 343–368.
- [4] L. Jofre, O. Lehmkühl, J. Ventosa, F. X. Trias and A. Oliva, (2014) *Conservation properties of unstructured finite-volume mesh schemes for the Navier-Stokes equations*, Numerical Heat Transfer, Part B: Fundamentals, 65, 1, 53–79.
- [5] O. Lehmkühl, G. Houzeaux, M. Avila, H. Owen, M. Vasquez and D. Mira, (2017) *A low dissipation finite element scheme for the large eddy simulation on complex geometries*, 19th International Conference on Finite Elements in Flow Problems - FEF 2017.
- [6] F. Cappuano, G. Coppola, L. Rndez and L. de Luca, (2017) *Explicit Runge-Kutta schemes for incompressible flow with improved energy-conservation properties*, Journal of Computational Physics, 328, 86–94.
- [7] F. X. Trias and O. Lehmkühl, (2011) *A self-adaptive strategy for the time integration of Navier-Stokes equations*, Numerical Heat Transfer, Part B: Fundamentals, 60, 2, 116–134.



Daniel Pastrana was born in Mexico City. He holds a bachelor's degree in physics (2012) and a master's degree in Ocean and Coastal Engineering (2015) from the Universidad Nacional Autónoma de México (UNAM). In early 2016, he was awarded by the Mexican government with SENER- CONACyT scholarship to study abroad. Nowadays, he is a PhD student in the Computational and Applied Physics program at the Universitat Politècnica de Catalunya - BarcelonaTech (UPC) and works at the Department of Com-

puter Applications in Science and Engineering of Barcelona Supercomputer Center. Daniel Pastrana has been involved in projects related to the study of heat transfer problems, fluid-structure interaction simulations and its applications to coastal engineering and the study of the energy and intensity decay of tropical cyclones when they make landfall. From this projects some important results have been published in international journals. Currently, his main research interests are High-Performance Computing, Fluid-Structure Interactions and Turbulence models and its industrial applications.

Application of Reduced order Modelling in Geophysics

Prattya Datta^{#1}, Josep De La Puente^{*2}, David Modesto^{#3}

[#]Universitat Politècnica de Catalunya - Campus Nord, Calle Jordi Girona, 1-3, 08034, Barcelona, España

¹prattya.datta@bsc.es, ³david.modesto@bsc.es

^{*}CASE, Barcelona Supercomputing Center-CNS, 08034, Barcelona, España

²josep.delapuente@bsc.es

Keywords— PGD, dimensionality, Geophysics

ABSTRACT

Simulations in science and engineering, and more specifically geophysics require huge resources, both in terms of processors and time. To tackle the problem of dealing with a large complex model reduced order methods have been devised to solve the model in less number of state variables. Proper Generalized Decomposition or PGD in short, is one such *a priori* technique to solve a high dimensional model in a reduced basis of low dimensions. PGD works by way of separated formulation of the original wave field. In this work, a detailed description of PGD is shown. PGD is applied to wave and heat models to successfully reduce higher dimensional problems and generate solution over the complete parameter range of interest in the offline phase. In the online phase accessing of solution inside the parameter range for any given value is efficient and fast.

A. Objectives

The main objective of this work is to effectively design an *a priori* Reduce Order Method solver to solve parametric high dimensional geophysical problems in a cost effective and fast manner.

In the field of geophysics, seismic inversion is an important process to identify geological parameters. Since, the geophysical data obtained (ground/sea) surveys does not produce the information of model of earth directly, inversion methods are applied, example being reverse time migration, Kirchhoff migration etc.[1, 3] The general technique for that is to find a representative smooth model, which in fact is a forward problem using the transient wave equation. Once, the forward problem is generated for all the parameters sufficiently the inversion process is performed by minimizing a suitable function characteristic of representative parameters [2]. In this case of geophysics with standard evaluation of forward models, this a cumbersome and time taking process considering the number of parameters involved (anisotropy, density etc.). To elaborate, each new parametric value requires complete calculation of overall 3D spatial model over the frequencies involved. The aim of this research is to minimize the time taken during such process effectively using Reduced Order Method (ROM) techniques especially *a priori* method and hence aid in inversion processes.

B. Methodology

ROM modeling or Model Order Reduction (MOR) works by way of looking for a generalized solution characteristic of the complete model [4]. It is classified into two types depending on whether it is an *a posteriori* method like, POD/SVD (Proper Orthogonal Decomposition/Singular Value Decomposition) or an *a priori* method like PGD (Proper Generalized Decomposition).

In the context of the problem statement being investigated, PGD is used.

C. Research Problem

Two problems are being investigated to find the effectiveness of PGD algorithm. The first is the heat equation to check the accuracy of the PGD and second is the simulation of a Helmholtz problem which requires more work because of its inherent difficulty in finding a good solution.

- The generic heat equation model is defined as follows;


$$\frac{\partial u}{\partial t} - k \cdot \Delta u = f. \text{ Where, } k = \text{conductivity.}$$

And, $u = u(x, k, t)$, $f = f(x, k, t)$, $x \in [0,1]$, $t \in [0,0.1]$ and $k \in [1, 5]$.


- The Helmholtz equation with 1D in space and 1D in frequency;

$$\nabla \cdot \nabla u + k^2 u = f. \text{ Where, } k = \text{wave number, } u = u(x, \omega) \text{ and, } f = f(x, \omega)$$

D. Results



The problem of heat equation is solved through finite difference scheme. Mode number represents the number of mode required to reach the final solution (here, final solution is the solution generated through direct solution of heat equation). It perfectly shows the error by the PGD formulation has reached until 10⁻⁷ after 10 nodes.



Similarly for the Helmholtz problem with a frequency range of 1-5Hz, the error corresponding to the direct finite difference formulation is shown. It has taken 50 modes to reach the desired level of error.

E. Conclusion and Future Enhancement

The problems solved within the purview of the PGD shows it is efficient and is suitable for solving further higher dimensional Helmholtz problem in geophysics. Currently,

PGD application in 2D spatial Helmholtz equation is being worked upon with few successes. Also, implementation of perfect absorbing layers, variable source location and different velocity regions is being processed.

F. ACKNOWLEDGEMENT

The author would like to show his gratitude towards BSC for providing the necessary resources and also a word of thanks to *La-Caixa* for helping me sustain myself while doing my research in Spain.

References

- [1] A. Tarantola. Inversion of seismic reflection data in the acoustic approximation. *Geophysics* 49, 8 (8 Sept. 1984), 1259–1266.
- [2] H. Talezer, D. Kosloff, and Z. Koren. An accurate scheme for seismic forward modelling. *Geophysical prospecting* 35, 5 (June 1987), 479–490.
- [3] J. Virieux and S. Operto. An overview of full waveform inversion in exploration geophysics. *Geophysics* 74, 6 (2009), WCC1–WCC26.
- [4] W. H. A. Schilders, H.A. van der Vorst, and J. Rommes. *Model Order Reduction: Theory, Research Aspects and Applications*, vol. 13. Springer, 2008.

Author biography



Prattya Datta was born in Kolkata, India, in 1989. He received the B.E. degree in Chemical engineering from the Jadavpur University, Kolkata, India, in 2011, and the M.Tech. degree in Petroleum engineering from the Indian Institute of Technology (IIT) Madras, Chennai, India, in 2014. He has also obtained degree of M.S. in Geophysics from Institut de Physics du Globe de Paris (IPGP) in 2013.

Since October 2015, he has been with the Department of CASE, BSC-CNS, where he is registered as Pre-Doctoral student. He is associated with La-caixa fellowship. Prior to joining he has also worked as business analyst in India for a period of 1 year. His current research interests include geophysics, reduced order modeling, data science and machine learning.

Simulations of Alfvénic Modes in TJ-II Stellarator

Allah Rakha^{1†} M. J. Mantsinen^{1,2} A. López-Fraguas³ F. Castejón³ J.L. de Pablos³
A.V. Melnikov⁴ S. E. Sharapov⁵ D. A. Spong⁶

¹Barcelona Supercomputing Center, Spain, ²ICREA, Barcelona, Spain, ³Fusion National Laboratory, CIEMAT, 28040, Madrid, Spain, ⁴Institute of Tokamak Physics, RRC ‘Kurchatov Institute’, 123182, Moscow, Russia, ⁵CCFE, Culham Science Centre, OX14 3DB, UK, ⁶Oak Ridge National Laboratory, TN, USA

†allah.rakha@bsc.es

Keywords— Alfvén Eigenmodes, MHD instabilities, Simulations

ABSTRACT

Alfvénic modes are one of the subclass of instabilities prevalent in burning plasmas due to interaction of energetic particles with background plasma. In this paper we investigate the properties of these modes with 3D simulations using modeling tools STELLGAP [1] and AE3D [2] of Neutral Beam Injection (NBI) heated H-plasmas in TJ-II low-magnetic-shear flexible heliac ($B_0 = 0.95$ T, $\langle R \rangle = 1.5$ m, $\langle a \rangle = 0.22$ m). These simulations focus on modelling the experimental observations [3] for prominent modes in TJ-II plasmas. Our simulations show consistency in frequency and radial location with the measured Alfvén Eigenmodes [3]. These simulations are performed for chirping and steady modes in TJ-II discharge # 29839 at $t = 1150$ and 1160 ms respectively.

INTRODUCTION

A. Importance of study

Alfvén Eigenmodes are of importance in fusion plasmas in view of next step burning plasma experiments. This is because they are likely to be excited by energetic particles and may ultimately cause losses of energetic fusion-born alpha particles in future fusion reactors. These instabilities can affect the efficiency of alpha-particle heating as well as cause damage to the first wall of the machine [4]. One of the challenges for a fusion reactor is how to contain the alpha particles in the vessel long enough for the particles to efficiently heat the plasma. Since these alpha particles can escape the fusion chamber prematurely by exciting high frequency Alfvén waves and scattering of these waves to the vessel walls. The study of the confinement properties of energetic particles, and, in particular their interaction with the background plasma, is therefore a key element in the preparation for ITER and DEMO operations [4, 5]. Fortunately, effective methods exist to produce significant fractions of energetic ions already in present experiments. Although these populations differ qualitatively in their energy and pitch-angle spectra from those of fusion-alfas, they still offer an excellent opportunity to test theoretical models and codes and to develop and benchmark diagnostic techniques [4]. The first observations of Alfvén Eigenmodes (AE) in TJ-II [6, 7] demand detailed investigations of these structures for future 3D fusion devices. In this work, the properties of energetic particle driven shear Alfvén Eigenmodes (AE) measured in [3] are studied for nominal plasma parameters of Flexible Helic TJ-II.

B. State of the art Modelling Tools

There are a variety of modelling tools available in the scientific community for different scales and devices. In this study, the following set of tools is used.

- I VMEC is a 3D numerical model used to calculate the equilibria in 3D fusion devices, using an energy minimizing principle.
- II STELLGAP [1] is a code for the calculation of shear Alfvén continua and gap formations in 3D toroidal / helical geometries.
- III AE3D [2] is a code for the evaluation of Alfvén spectral properties, and mode structure in toroidal devices.

C. Modelling of Alfvén Eigenmodes in 3D geometry

A reduced MHD shear Alfvén model [8] is used here for the modelling of Alfvénic instabilities in TJ-II plasmas. The continuum structures are obtained from the STELLGAP code [1] and Alfvén mode structures are calculated using the AE3D code [2]. The mode structure includes the effects of couplings from the 3D equilibrium for both the Alfvén continuum and Alfvén Eigenmode calculations. The main eigen-value equation used in these tools is;

$$\omega^2 \nabla \cdot \left(\frac{1}{V_A^2} \nabla \phi \right) + (\mathbf{B} \cdot \nabla) \left[\frac{1}{B} \nabla^2 \left(\frac{\mathbf{B}}{B} \cdot \nabla \phi \right) \right] + \nabla \zeta \times \nabla \left(\frac{\mathbf{B}}{B} \cdot \nabla \phi \right) \cdot \nabla \frac{J_{\parallel 0}}{B} = 0$$

Here \mathbf{B} is the equilibrium magnetic field; B is its magnitude, $J_{\parallel 0}$ is the equilibrium parallel current, Φ is the electrostatic potential, while ζ is the toroidal angle coordinate, $V_A^2 = B^2 / \mu_0 \rho_m$ is Alfvén speed and $\rho_m = n_{ion} m_{ion}$ is ion mass density. Setting coefficient of highest order radial (ρ) derivatives to zero leads to the continuum equation. These are the continua of shear Alfvén problem [1] and provide an initial guess for the location and frequency gaps where eigen functions of the Alfvénic instabilities exist.

D. Alfvén Eigenmodes in TJ-II Stellarator

TJ-II discharge # 29839 is interesting due to its distinct feature, specifically chirping and steady-frequency modes coexisting at the same time in this discharge. The experimental observations of Alfvén mode structures in this discharge are presented in [3]; in this work modelling of chirping and steady modes is presented. Experimental results confirm the signatures of these modes because of the stronger coherence of Heavy Ion Beam Probe (HIBP) and Mirnov Probe (MP) data as shown in Fig (1) left panel and the zoom of the prominent mode in right panel. The first results of 3D modelling of these modes along generation of gap structures are presented in Section E.




Fig. 1 Experimental observations of Alfvén Eigenmodes used in these simulations for TJ-II Stellarator Discharge # 29839 [3].

E. Results from STELLGAP simulations

Separate simulations of these experimental results for chirping and steady modes at $t = 1150$ and 1160 ms respectively have been performed by using the simulation tools discussed in Section B. Corresponding simulated Alfvén continuum gap structures for chirping and steady modes are plotted in Fig. 2 and 3 respectively.




Fig. 2 Alfvén continuum gap structures for chirping phase of mode at $t = 1150$ ms.




Fig. 3 Alfvén continuum gap structures for steady phase of mode at $t = 1160$ ms.

The prominent toroidal mode numbers are shown distinctly with color coding in the graphs.

F. Results from AE3D simulations

To investigate the mode structures and their profiles, AE3D simulations have been performed for the cases presented in section E. Corresponding results are presented in Figs 4 and 5.




Fig. 4 Alfvén Eigenmodes for chirping modes.

In both the cases of chirping and steady modes, there is single prominent mode. In the chirping phase, the modes ($m = 10$ and $n = -17$) have two different frequencies of 289 and 316

kHz, located at radial position $\rho = 0.7$. While in the steady phase, the same mode ($m = 10$ and $n = -17$) appears with a lower frequency of 254 kHz at smaller radial location $\rho = 0.55$. It can be seen that the mode ($m = 10$ and $n = -17$) is prominent in both cases and is comparable with experimental results, although there is a decrease in frequency from 289 to 254 kHz.




Fig. 5 Alfvén Eigenmode structure for the steady phase of the mode.

TABLE I
SUMMARY OF THESE MODES

Mode type	Prominent mode numbers (m, n)	Frequency (kHz)	Radial location (ρ)
Chirping	(10, -17)	289 316	0.70
Steady	(10, -17)	254	0.55

G. Summary and Future Enhancement

It can be inferred from these simulations that the modes move radially inward when there is a transition from chirping to steady frequencies. Moreover, the frequencies of modes are also decreasing significantly. Interestingly the mode ($m = 10$ and $n = -17$) remains persistent during this transition.

This work will be extended to investigate other properties and types of Alfvén eigenmodes in TJ-II plasmas. Active interaction of the energetic particles with these modes and resonance calculations for these interactions will be explored. Furthermore, the comparison of these calculations with the experimental data will be addressed.

References

- [1] D. A. Spong, R. Sanchez and A. Weller. Phys. Plasmas 10, 3217, (2003).
- [2] D. A. Spong, D’Azevedo and Y. Todo. Phys. Plasmas 17, 022106, (2010).
- [3] A.V. Melnikov et. al. Nucl. Fusion. 56, 076001 (2016).
- [4] N.N. Gorelenkov et al, Nucl. Fusion 54, 125001 (2014).
- [5] L. Chen and F. Zonca. Nucl. Fusion 47, S727-34 (2007)
- [6] C. Alejaldre et al. Nucl. Fusion 41, 1449 (2001).
- [7] R. Jiménez-Gómez et al. , Nucl. Fusion 51,033001 (2011).
- [8] S. E. Kruger, C. C. Hegna and J. D. Callen, Phys. Plasmas 5, 4169 (1998).

Author biography



Allah Rakha was born in Jhang, Pakistan, in 1985. He received the M.Phil. degree in Physics from the Pakistan Institute of Engineering & Applied Sciences (PIEAS), Islamabad in 2008, and the M.S. degree in Nuclear Fusion & Engineering Physics from the Ghent University, Belgium, in 2015.

Since October 2008, he has been working as Lecturer in Physics at Department of Physics & Applied Mathematics (DPAM), PIEAS Islamabad. He joined Barcelona Supercomputing Center (BSC) as PhD researcher in Fusion group on January 2016. His current research interests include plasma physics, MHD of fusion plasmas, and Alfvénic instabilities in Stellarator devices.

Leveraging FPGAs to Accelerate the Query Processing of SQL-Based DataBases

Behzad Salami

Computer Science Department

Barcelona Supercomputing Center (BSC) Barcelona Supercomputing Center (BSC) Barcelona Supercomputing Center (BSC)

Barcelona, Spain

Email: behzad.salami@bsc.es

Osman Unsal

Computer Science Department

Barcelona, Spain

Email: osman.unsal@bsc.es

Adrian Cristal Kestelman

Computer Science Department

Barcelona, Spain

Email: adrian.cristal@bsc.es

Abstract—With the rise of Big Data, providing high-performance query processing capabilities through the acceleration of database analytics has gained significant attention. Leveraging Field Programmable Gate Array (FPGA) technology, this approach can lead to clear benefits. In this work, we briefly introduce the design and implementation of an FPGA-based platform that enables fast query processing for database systems by melding novel database-specific accelerators with commercial-off-the-shelf (COTS) storage using modern interfaces, in a novel, unified, and a programmable environment. The proposed engine can perform a large subset of SQL queries through its set of instructions that can map compute-intensive database operations, such as filter, arithmetic, aggregate, group by, table join, or sort, on to the specialized high-throughput accelerators. To minimize the amount of SSD I/O operations required, it also supports hardware MinMax indexing for databases. We evaluated our query processing engine with five decision support queries from the TPC-H benchmark suite and achieved a speedup from 1.8X to 34.2X, in comparison to the state-of-the-art DBMS, i.e., PostgreSQL and MonetDB.

I. INTRODUCTION

FPGAs provide a unique opportunity to build an efficient query processing platform, by constructing a high-throughput execution engine with the additional aim of minimizing overheads of data movement [1]. It is mainly the consequence of; (i) the inherent characteristics of massively parallel and configurable architecture of FPGAs, suitable for data streaming in deep pipelined-style execution (ii) the rise of High-Level Synthesis (HLS) technology, which makes FPGA applications relatively easier to develop compared to low-level languages such as VHDL or Verilog, and (iii) the availability of soft cores that implement modern interfaces, such as PCIe 3.0 (Peripheral Component Interconnect Express) or SATA-3 (Serial AT Attachment).

For the query processing, FPGAs have been utilized in two distinct approaches: (i) traditional data offloading mechanisms, where data in the host-attached storage is offloaded towards external processing units or accelerators implemented in the FPGA [2], or (ii) placing the processing units directly in the data path between the host machine and the main storage units [3]. The first approach could incur overheads stemming from the additional data movement since data needs to be offloaded through the Operating System (OS) and device driver layers. In contrast, the second approach allows processing units to




Fig. 1. Overall architecture of the proposed engine with its major components. (1) software extensions for DBMS in the host, (2) Data storage units and device controllers, (3) Query accelerators that are organized in a clockwise unidirectional ring bus (RBAA), (4) Programmable Interconnection Unit (PIU) to manage the accesses to the off-chip data storage units, in a fully flexible fashion, (5) Data and Process Controller (DPC) to orchestrate the involved modules of the engine to process the SQL queries.

get direct access to the data blocks and facilitates low-latency data transmission. In addition, it can correspond to significant speedups in the query execution. In this work, we follow the second approach to minimize the data transmission overheads.

In a nutshell, the main objectives of this work as a novel query processing engine, are listed as below:

- to provide an infrastructure that sits between the host and the data storage in SSD, and utilizes PCIe-3 and SATA-3 interfaces to work directly with blocks of database columns
- to design a set of efficient query accelerators inside such an infrastructure that can facilitate the query processing in a fully pipelined fashion
- to allow the DBMS to utilize these accelerators by issuing query-specific instructions, exposing flexibility in the data movement and in enabling accelerators

II. THE OVERALL ARCHITECTURE

In the proposed query processing engine, the host communicates with the FPGA through an Application Program Interface (API), to transfer data and instructions using the PCIe-3 interface. When the host initiates the query execution, the query plan needs to be converted into our specific instructions. Inside the FPGA, these instructions are managed and executed by the Data and Process Controller (DPC), which orchestrates the movement of data blocks between the SSD, DDR-3, host, and

Accelerators. The query is effectively executed by streaming blocks of data, from the storage, through the accelerators, and back. Finally, the result of the query is returned to the software or stored back into the SSD.

The architecture of the proposed query processing engine is composed of five major components: **(1)** software extensions for DBMS in the host, including the Data Address Table (DAT) and the CStore Foreign Data Wrapper (FDW) extension of PostgreSQL, to manage the transfer of instructions and data, respectively, **(2)** data storage units, i.e., SSD, DDR-3, and host that are used as the primary or secondary database storage units and device controller cores, i.e., PCIe-3, DDR-3, SATA-3 to manage the data transfer to/from storage units, **(3)** a set of efficient database accelerators, i.e., filter, arithmetic, aggregation, group by, hash probe, hash build, and sort that are organized in a unidirectional ring bus, which is called RingBus of Accelerators (RBAA), **(4)** Programmable Interconnection Unit (PIU) to set up a path to transmit the data in a fully flexible fashion. It is composed of *i*) a 4-port bidirectional programmable data connection switch (PDCS) to exchange the data among SSD, DDR-3, host, and RBAA, *ii*) an arbiter to manage the DDR-3 concurrent requests, and *iii*) a set of synchronizing First In First Out (FIFO) modules for each individual port, separately for read and write directions, to cross the different clock domains, **(5)** Data and Process Controller (DPC) that is composed of an Instruction Cache (IC) to locate the instruction set and an execute Finite State Machine (FSM) *i*) to manage the accesses to the off-chip data sources and *ii*) to control the accelerators to execute the corresponding query, by issuing the appropriate control signals to the PIU and to the RBAA, respectively. These signals are generated by translating our query-specific instructions.

III. EXPERIMENTAL RESULTS

We developed our engine on a VC709 FPGA development board with an XC7VX690T FPGA and 4GB of DDR-3 RAM. It accesses a Crucial M4-256GB SSD through a customized version of an SATA-3 controller, based on Groundhog [XX]. We evaluated our engine against the query processing engines of several state-of-the-art software DBMS: *(i)* MonetDB 11.21 as a popular column-oriented database system, *(ii)* PostgreSQL 9.5 (PGSQL) as a popular object-relational row-oriented database system, and *(iii)* CStore as the PostgreSQL's column-oriented data store extension. We evaluated our engine with five decision-support TPC-H queries, under various conditions. The studied queries are Q01, Q03, Q04, Q06, and Q14, which heavily utilize and stress the various hardware accelerators.

- **process-intensive queries** (Q01), as it can be seen in Figure 2(a), the I/O time of SSD is negligible and the execution time is dominating. In process-intensive workloads, the performance gain of the proposed FPGA-based engine is mainly the consequence of exploiting highly efficient query accelerators, in a deeply pipelined fashion.
- **I/O-process-balanced queries** (Q03, Q04), the improvement of the proposed engine is the consequence of both

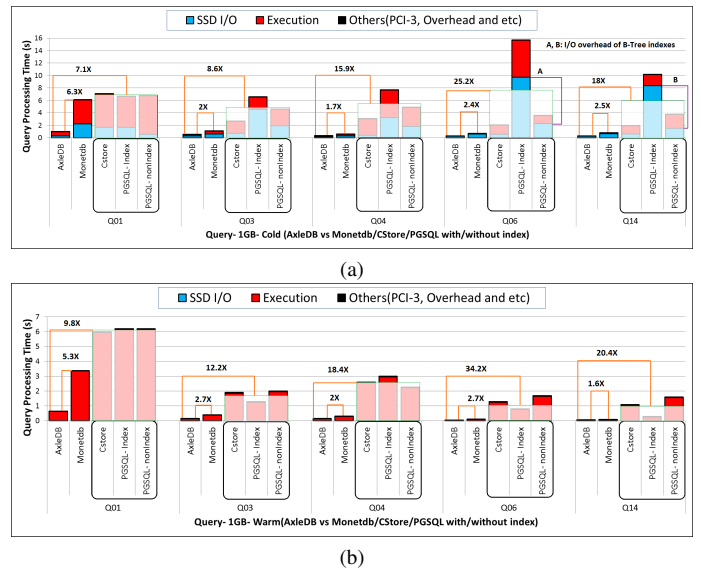


Fig. 2. Total query processing time of the studied benchmarks in cold (a) and warm (b) modes, comparing the proposed FPGA-based engine (AxleDB) vs. MonetDB, CStore, and PostgreSQL in 1GB scale. Lower is better.

I/O efficiency and faster execution. For instance, the FPGA-based engine reduces SSD I/O time by 17.9X and execution time by 31.5X for Q04, in comparison to the index-enabled PostgreSQL, which leads to a total of 23.4X speedup for this particular case.

- **I/O-intensive benchmarks** (Q06, Q14), SSD I/O time is dominating. Comparing PostgreSQL with index-enabled vs. non-index versions, we unexpectedly observed a significant overhead of B-Tree indexes, which causes substantial performance degradation. In contrast, our FPGA-Based engine, CStore, and MonetDB, thanks to their column-oriented data storage, significantly reduce SSD I/O transfers. The results demonstrate that the FPGA-Based engine can process these queries, on average 2.4X faster than MonetDB, and 21.6X faster than the average of different versions of PostgreSQL.

In summary, the significant speedup of the proposed FPGA-based engine against software-based comparison cases is the consequence of two optimization points: *i*) offloading the query processing onto the FPGA and following the streamline data flow execution model and *ii*) optimized accesses to the SSD (tightly coupled to the processing units -accelerators- in the FPGA). As mentioned earlier, these points have proportionally affected for each individual query,

REFERENCES

- [1] Jun, Sang-Woo and others. BlueDBM: an appliance for big data analytics. Proceedings of ISCA, pages=1–13, 2015, ACM.
- [2] Casper, Jared and Olukotun, Kunle. Hardware acceleration of database operations. Proceedings of FPGA, pages=151–160, 2014, ACM.
- [3] Woods, Louis and others. Ibox: an intelligent storage engine with support for advanced SQL offloading. Proceedings of the VLDB Endowment, volume=7, number=11, pages=963–974, 2014.

Time-Predictable Parallel Programming Models

Maria A. Serrano

Barcelona Supercomputing Center (BSC) and
 Technical University of Catalonia (UPC), Barcelona, Spain
 Email: maria.serranogracia@bsc.es

Eduardo Quinones

Barcelona Supercomputing Center (BSC), Barcelona, Spain
 Email: eduardo.quinones@bsc.es

Abstract—Embedded Computing (EC) systems are increasingly concerned with providing higher performance in real-time while HPC applications require huge amounts of information to be processed within a bounded amount of time. Addressing this convergence and mixed set of requirements needs suitable programming methodologies to exploit the massively parallel computation capabilities of the available platforms in a predictable way. OpenMP has evolved to deal with the programmability of heterogeneous many-cores, with mature support for fine-grained task parallelism. Unfortunately, while these features are very relevant for EC heterogeneous systems, often modeled as periodic task graphs, both the OpenMP programming interface and the execution model are completely agnostic to any timing requirement that the target applications may have. The goal of our work is to enable the use of the OpenMP parallel programming model in real-time embedded systems, such that many-cores architectures can be adopted in critical real-time embedded systems. To do so, it is required to guarantee the timing behavior of OpenMP applications.

I. INTRODUCTION

High performance computing (HPC) has been for a long time the realm of a specific community within academia and specialized industries. Similarly, embedded computing (EC) has also focused mainly on specific systems with specialized and fixed functionalities for which timing requirements were considered more important than performance requirements. However, with the ever-increasing availability of more powerful processing platforms, alongside affordable and scalable software solutions, both HPC and EC are extending to other sectors and application domains.

As a result, a new type of applications is crossing the boundaries between the HPC and the EC domains. For such applications, the correctness of the result is dependent on both performance and timing requirements, and the failure to meet either of them is critical to the functioning of the system. In this context, it is essential to guarantee the timing predictability of the performed computations.

The use of parallel programming models is fundamental to exploit the performance of current and future many-core architectures, while providing good programmability (and so productivity) of high performance systems. Among the different models, OpenMP [1] has become one of the most used parallel programming models due to its simplicity and scalability in shared memory and heterogeneous systems. The latest specifications of OpenMP incorporate a tasking model that enables very sophisticated types of fine-grained and irregular parallelism, in which the programmer may define explicit tasks and their related data dependencies, as well as an advanced

accelerator execution model to deal with data movement and efficient computation in heterogeneous architectures.

Unfortunately, OpenMP tasking and accelerator models were created for a very different purpose than describing real-time applications modeled as task graphs. However, its syntax and execution model retain certain similarities to that formalism that could make it a good candidate to fill the existing gap between: (i) a convenient programming model for heterogeneous many-cores and (ii) state-of-the-art techniques for scheduling with timing guarantees.

Our work focuses on adopting the current OpenMP v4.5 specification to provide timing guarantees so that it can be used in critical real-time embedded systems.

II. OPENMP TIMING CHARACTERIZATION

The first specifications of OpenMP (up to version 2.5) were focused on a thread-centric model to exploit massively data-parallel and loop-intensive types of applications. The latest specifications of OpenMP (versions 3.0, 4.0 and 4.5) have evolved to a task-centric model which enables very sophisticated types of fine-grained and irregular parallelism, including support for heterogeneous computing. This new model, known as *tasking model*, provides a very convenient abstraction of parallelism, being the run-time in charge of scheduling tasks to threads. Despite this model lacks any notion of real-time scheduling semantics, such as deadline, period or WCET, its structure and syntax have certain similarities with the Directed Acyclic Graph (DAG) real-time scheduling model [2] used to analyse the timing behavior of parallel execution in real-time. As an example, Figure 1a shows an OpenMP program composed of five tasks and Figure 1b shows the corresponding OpenMP-DAG, see more details in [3].

This tight correspondence between the structure and syntax of an OpenMP program and the DAG model makes OpenMP a firm candidate to be adopted in future real-time systems.

III. RESPONSE TIME ANALYSIS OF OPENMP DAGS


In real-time systems the correctness of each computation depends not only on its logical result but also on the time instant at which such result is produced. Failure to respond within a specified time interval is considered as a faulty response. Thus, a key property of a real-time system is time predictability. The sporadic DAG scheduling model is currently under investigation in the real-time community to address time predictability of parallel computation. In the sporadic DAG model, each real-time task (called DAG-task) is represented with a directed acyclic graph (DAG).

```

1 #pragma omp parallel {
2 #pragma omp single { //  $\tau_1$ 
3   p1,1
4 #pragma omp task { //  $\tau_2$ 
5   p2,1
6   #pragma omp task //  $\tau_3$ 
7   { p3,1 }
8   p2,2
9 #pragma omp taskwait
10  p2,3
11 } p1,2
12 #pragma omp task depend (out:a)
13 { p4,1 } //  $\tau_4$ 
14 p1,3
15 #pragma omp task depend (in:a)
16 { p5,1 } //  $\tau_5$ 
17 }

```

(a) OpenMP source code.



(b) OpenMP DAG.

Fig. 1: Example of an OpenMP program composed of tied tasks (a) and its corresponding OpenMP-DAG (b).

In the following subsections we briefly describe our research on computing the response time analysis of real-time OpenMP applications represented as DAGs. The response time analysis computes the worse-case response time a system may take, in this case, when being scheduled in a parallel processor with a given number of cores. Thus, when compared with the application deadline we can provide OpenMP applications with real-time guarantees.

A. Response Time Analysis of an OpenMP task

Despite being OpenMP a convenient candidate to be adopted in future real-time systems, it incorporates features that limit its practical usability in real-time systems. The most notable example is the distinction between tied and untied tasks. Tied tasks force all parts of a task to be executed on the same thread that started the execution, whereas a suspended untied task is allowed to resume execution on a different thread. The execution model of tied tasks has serious implications on the response time analysis of OpenMP applications, making difficult to adopt it in real-time environments.

In [6] we analyze, from a timing perspective, the two tasking existing models in OpenMP: tied and untied. The considerations drawn in this work suggest that using tied tasks inside time-critical applications is not recommendable because of the inherent pessimism that underlies the timing analysis of such tasks and the conceptual difficulties behind the construction of an accurate schedulability test. We also show that a simple schedulability analysis of OpenMP programs is possible whenever untied tasks are involved. This suggests that the use of untied tasks would be preferable for parallel applications in a real-time context, since it would permit to exploit a parallel execution model in a predictable way.

B. Response Time Analysis of a set of OpenMP tasks

When considering several OpenMP applications with different priorities, preemption is a key concept since it allows the operating system allocate the core to tasks requiring urgent service. The limited preemption (LP) approach has been proposed in the literature to reduce the run-time overhead

due to preemptions and still preserve the schedulability of the system. According to this approach, a task implicitly executes in non-preemptive mode and preemption is allowed only at predefined locations inside the code, called preemption points. In this way, a task is divided into a number of non-preemptive chunks (also called sub-tasks); if a higher priority task arrives between two preemption points of the running task, preemption is postponed until the next preemption point.

Interestingly, the LP approach with fixed preemption points resembles the OpenMP task execution model [3]. Therefore, in [4] and [5] we evaluate the LP strategy for DAG-based task-sets; we show the necessary conditions under which DAG tasks may experience lower and higher priority task interference for different LP sub-approaches.

IV. FUTURE WORK

Our future work focuses on the schedulability analysis of the OpenMP accelerator model for heterogeneous architectures. We will investigate novel time predictable scheduling solutions for the OpenMP accelerator model in current many-core heterogeneous architectures. In these architectures, the main problems lie in: (1) the proper scheduling of data transfers and computation on acceleration devices, and (2) the characterization of the multiple interferences and inter-dependencies that may arise among the simultaneous access of various accelerator devices in a many-core system.

V. ACKNOWLEDGMENT

This work was funded by the EU project P-SOCRATES (FP7-ICT-2013-10) and the Spanish Ministry of Science and Innovation under contract TIN2015-65316-P.

REFERENCES

- [1] OpenMP Application Program Interface, Version 4.5. November 2015.
- [2] S. Baruah, V. Bonifaci, A. Marchetti-Spaccamela, L. Stougie, and A. Wiese. A generalized parallel task model for recurrent real-time processes. In *RTSS*, 2012.
- [3] R. Vargas, et. al. OpenMP and Timing Predictability: A Possible Union? In *18th Design, Automation and Test in Europe Conference (DATE)*, 2015.
- [4] M. A. Serrano, A. Melani, M. Bertogna, and E. Quinones. Response-time analysis of DAG tasks under fixed priority scheduling with limited preemptions. In *DATE*, March 2016.
- [5] M. A. Serrano, A. Melani, S. Kehr, M. Bertogna, and E. Quinones. An analysis of lazy and eager limited preemption approaches under dag-based global fixed priority scheduling. In *ISORC*, May 2017 (to appear).
- [6] M. A. Serrano, A. Melani, R. Vargas, A. Marongiu, M. Bertogna, and E. Quinones. Timing characterization of OpenMP4 tasking model. In *CASES*, 2015.



Maria A. Serrano is a PhD. student in the Department of Computer Science at the Barcelona Supercomputing Center (BSC), Spain, since 2014. She studied a B.S / M.S in Computer Science Engineering (2006-2013) and a M.S. in Systems Engineering and Computer Science (20013-2014) at the University of Zaragoza, Spain. Her research interests focus on Real-Time systems, specifically on many-cores embedded processors and on providing parallel programming models with timing guarantees.

On the suitability of Time-Randomized Processors for Secure and Reliable High-Performance Computing

David Trilla^{†,‡}, Carles Hernandez[†], Jaume Abella[†], Francisco J. Cazorla^{†,*}

[†] Barcelona Supercomputing Center (BSC), Barcelona, Spain

[‡] Universitat Politècnica de Catalunya (UPC), Barcelona, Spain

* Spanish National Research Council (IIIA-CSIC), Barcelona, Spain.

Abstract—Time-randomized processor (TRP) architectures have been shown as one of the most promising approaches to deal with the overwhelming complexity of the timing analysis of high complex processor architectures for safety-related real-time systems. With TRPs the timing analysis step mainly relies on collecting measurements of the task under analysis rather than on complex timing models of the processor. Additionally, randomization techniques applied in TRPs provide increased reliability and security features. In this thesis, we elaborate on the reliability and security properties of TRPs and the suitability of extending this processor architecture design paradigm to the high-performance computing domain.

I. INTRODUCTION

Several years ago probabilistic timing analysis (PTA) [1] arised as a new timing analysis paradigm with the aim to facilitate the timing-verification of complex processors running safety-critical applications. Amongst other approaches, PTA proposes the utilization of time-randomized processors (TRPs) as a way to enable the derivation of timing bounds using probabilistic methods. Typically, complex commercial off-the-shelf (COTS) processor architectures are not amenable for timing-critical applications since they include hardware features such as caches, branch predictors, and multicore technology that severely complicates the derivation of execution time-bounds for the programs running on top of these processors. PTA applied on top of TRPs allows to reduce the complexity of the timing analysis process. To do so, TRPs employ hardware randomization techniques to make the latency of jittery resources to exhibit a probabilistic nature and thus, to allow the application of extreme value theory (EVT) [13] to bound program's execution time. Basically TRPs simplify the timing verification process.

Typically processors in each domain (i.e. personal devices, supercomputers, etc...) have been subject to different requirements in terms of area, power consumption, temperature, performance, reliability and security. For instance, in the high-performance domain reliability has been a second order concern for many years in comparison with performance first and power/temperature later. Conversely, in other domains like safety-critical systems, reliability and security have been primary concerns due to their potential threads affecting human lives.

However, the need to push technology scalability limits further every generation in the high-performance domain has

made faults to be more frequent and thus, reliability has become also a primary concern in this domain. Additionally, high-performance processors used in data-centers and servers have to meet strong security requirements to avoid malicious attacks.

The properties of TRPs to enable the utilization of high-performance processors in time-critical applications at a reasonable cost have been deeply analyzed [12]. Our focus is to highlight the reliability and security properties TRPs can offer to both high-performance and safety-critical domains.

II. ACHIEVING PREDICTABILITY WITH NON-REPRODUCIBLE TIMING BEHAVIOR

The main difference between TRPs and conventional (time-deterministic) processor designs resides on the way hardware resources exhibiting jitter, i.e. what factors trigger different latencies for those resources that do not exhibit a constant latency.

TRPs provide hardware support so that execution time measurements collected during analysis match or upperbound tightly those during operation. To upper-bound the jitter probabilistically, randomization (RND) techniques are required to make the jitter have a probabilistic nature during both analysis and operation phases.

RND techniques have been applied satisfactorily to caches and shared resource arbitration. Regarding caches, random replacement and random placement policies have been proposed to match TRP requirements [8]. Figure 1 shows and example of random modulo architecture. For the arbitration of shared resources the lottery bus [14] has been shown suitable for TRPs, while other new approaches like random permutations [10] have been shown also suitable and offer improved performance.

When all jittery resources are properly handled, EVT can be applied to the measurements collected for the programs executed in TRPs. The timing analysis methodology for which TRPs were proposed is called Measurement-Based Probabilistic Timing Analysis (MBPTA) [3], [2]. Time-measurements collected on top TRPs follow a probabilistic nature, are independent and identically distributed, and represent an upper-bound of the worst events due to the interactions of the different jittery resources that can occur in the processor. Once appropriate tests are passed, those measurements are used




Fig. 1. Random modulo architecture.

as input for EVT, which is a powerful statistical method to approximate the tail of a distribution that represents the high execution times. Then, the probabilistic WCET (pWCET) is the execution time value of the obtained distribution whose risk of being exceeded is upper-bounded by an arbitrarily low probability to be neglected (i.e. residual risk [9]). Figure 2 shows a pWCET curve with a cutoff probability of 10^{-14} .




Fig. 2. Example pWCET obtained with a TRP.

III. RELIABILITY PROPERTIES OF TIME-RANDOMIZED PROCESSORS

The majority of current processors, regardless of the application domain, are provided with some form of fault tolerance. We analyze how TRPs help maximizing the effectiveness of existing protection mechanisms and also provide enhanced robustness capabilities.

A. Improving graceful performance degradation

Fault-tolerant mechanisms are able to keep system functioning despite the presence of faults but sometimes this comes at the expense of a reduction in performance. For example, when one or several cache lines are permanently damaged, protection mechanisms disabling faulty lines allow the processor to operate correctly enlarging its lifetime. However, depending on the particular location of those faulty lines in cache, the performance provided by the processor can vary significantly [4]. Random cache policies included in TRPs make this degradation to occur more gracefully [16]. Figure 3 shows how the execution time for two benchmarks behaves as we decrease the cache space.




Fig. 3. Execution time (normalized) as cache size decreases for COTS and TRP designs in DL1 caches.

B. Reducing the bias in resource utilization

With conventional placement algorithms, cache sets access distribution is completely program dependent. On the contrary, with random placement algorithms [8], [11] accesses to the cache sets are randomly distributed since random placement algorithms employ a combination of address tags with random bits from a random seed to generate the cache index so, for every run, a different set is accessed for any given address. Having a highly biased cache set utilization is expected to lead to higher degradation since the most used sets are more exposed to hot carrier injection (HCI) [6] among other sources of transistor degradation. In this context, having set access distributions as uniform as possible is very convenient to mitigate those aging effects [7], [17]. Figure 4 shows that almost perfect set access distributions can be achieved with TRPs.

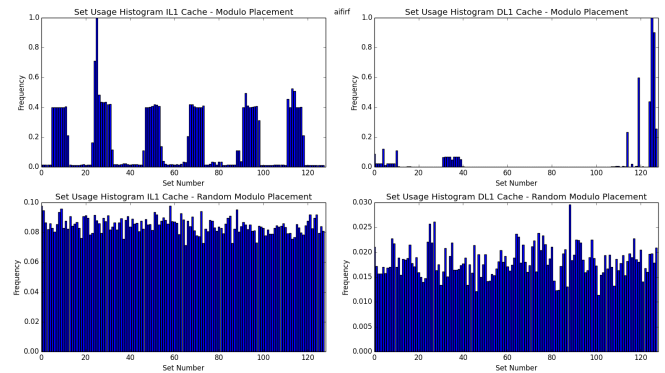


Fig. 4. Cache utilization across sets with and without random placement.

C. Power stability and voltage noise resilience

Randomization removes systematic pathological scenarios that can lead to corner situations with significantly bad performance. High demanding power events that occur at a given frequency and align with the resonance frequency of the power network distribution will significantly amplify voltage noise fluctuations [5]. Pathological cases occur, for example, due to systematic cache conflicts.

TRPs break systematic alignments of the events. Thus, TRPs can diminish the impact of voltage noise effects. Figure 5 shows the fast fourier transform of the power consumption

resulting from the execution of a LRU pathological case in a time-deterministic processor (a) and in a TRP (b).




Fig. 5. Fast fourier transform of the power consumption of pathological patterns in a TRP (a) and in an time-deterministic one (b).

IV. SECURITY PROPERTIES OF TIME-RANDOMIZED PROCESSORS

Side-channel attacks extract secret key data by exploiting the information leakage resulting from the physical implementation of the system. TRPs can mitigate these effects.

A. Cache side-channel attacks

Cache side-channel attacks are characterized by the abuse of its controllable and known behavior. Cache misses and hits are reflected on the execution time, thus they can be measured. This hits and misses are dependent on the placement of the data used, a pattern of execution times may indicate that certain data is being used and through this knowledge, cryptographic keys can be inferred. Layout randomization has been shown to be an effective mechanism to protect against contention-based attacks [18]. Since TRPs are based on the utilization of random cache designs [8], these processor designs are inherently protected.




Fig. 6. Power dissipation variability in a TRP.

B. Power analysis side channels

Power dissipation can also leak cryptographic information. When instructions are executed with fixed-time repetitive executions, they provide similar power profiles. Therefore, since cryptographic algorithms use multiple iterations for a given secret key, attackers can match the similar power profiles obtained to infer the cryptographic data. Randomizing the

execution time delay to achieve protection against power analysis attacks has been proved useful. TRPs provide time randomization by default by randomizing the existing processor jitter and thus, avoid detection of the execution of cryptographic algorithms. Figure 6 shows the power variability resulting from running 1000 times an encryption algorithm in a TRP.

V. TOWARDS SECURE AND RELIABLE HIGH PERFORMANCE TIME RANDOMISED PROCESSORS: CONCLUSIONS AND FUTURE WORK

A. Further Hardware Randomization and other EVT applications

In order for TRPs to be fully embraced some research is needed on unexplored features of EVT and hardware randomization. Some are listed below.

Hardware Randomization. Two important improvements to processors that due to being too complex to time analyze are not suitable for critical real-time systems are Data prefetchers and Out-of-Order execution. Both are able to partially hide latency of the program execution. At the same time they require randomization so systematic behaviors are avoided and probabilities attached to events. Furthermore randomizing prefetching policies can also enhance TRPs protection against side-channel attacks [15].

Extended EVT applications. EVT is an important tool that can be applied in other important aspects different from execution time bounding. Voltage Noise and Maximum Power Density are important factors when designing systems since reaching unsuspected limits on those domains can severely affect the behavior of the processor. Applying EVT on top of TRPs might provide safe bounds to power dissipation and can help in mitigating overdimensioning the safety margins.

B. Conclusions

As shown, TRPs provide unique properties that make them an ideal baseline to combat some of the reliability and security challenges that high-performance processors have to face these days. However, future work is needed to enhance these features. Despite this TRPs seem promisingly suitable to be used in markets other than safety-critical systems.

REFERENCES

- [1] F. Cazorla et al. PROARTIS: Probabilistically analysable real-time systems. *ACM TECS*, 2012.
- [2] F. Cazorla et al. Upper-bounding program execution time with extreme value theory. In *WCET Workshop*, 2013.
- [3] L. Cucu-Grosjean et al. Measurement-based probabilistic timing analysis for multi-path programs. In *ECRTS*, 2012.
- [4] J. A. et. al. Low vccmin fault-tolerant cache with highly predictable performance. *MICRO 42*, pages 111–121, New York, NY, USA, 2009. ACM.
- [5] R. B. et. al. Voltage noise in multi-core processors: Empirical characterization and optimization opportunities. In *MICRO '14*, pages 368–380, Cambridge, UK, December 2014. IEEE Computer Society.
- [6] V. H. et al. Managing sram reliability from bitcell to library level. In *Reliability Physics Symposium (IRPS), 2010 IEEE International*, pages 655–664, May 2010.

- [7] E. Gunadi, A. A. Sinkar, N. S. Kim, and M. H. Lipasti. Combating aging with the colt duty cycle equalizer. In *Microarchitecture (MICRO), 2010 43rd Annual IEEE/ACM International Symposium on*, pages 103–114, Dec 2010.
- [8] C. Hernandez et al. Random modulo: a new processor cache design for real-time critical systems. In *DAC*, 2016.
- [9] International Organization for Standardization. *ISO/DIS 26262. Road Vehicles – Functional Safety*, 2009.
- [10] J. Jalle, L. Kosmidis, J. Abella, E. Quinones, and F. Cazorla. Bus designs for time-probabilistic multicore processors. In *DATE*, 2014.
- [11] L. Kosmidis et al. A cache design for probabilistically analysable real-time systems. In *DATE*, 2013.
- [12] L. Kosmidis et al. Probabilistic timing analysis and its impact on processor architecture. In *DSD*, 2014.
- [13] S. Kotz and S. Nadarajah. *Extreme value distributions: theory and applications*. World Scientific, 2000.
- [14] K. Lahiri, A. Raghunathan, and G. Lakshminarayana. LOTTERYBUS: a new high-performance communication architecture for system-on-chip designs. *DAC '01*, pages 15–20, 2001.
- [15] F. Liu and R. B. Lee. Random fill cache architecture. In *Proceedings of the 47th Annual IEEE/ACM International Symposium on Microarchitecture, MICRO-47*, pages 203–215, Washington, DC, USA, 2014. IEEE Computer Society.
- [16] M. Slijepcevic et al. Timing verification of fault-tolerant chips for safety-critical applications in harsh environments. *IEEE Micro - Special Series on Harsh Chips*, 34(6), 2014.
- [17] D. Trilla, C. Hernandez, J. Abella, and F. Cazorla. Resilient random modulo cache memories for probabilistically-analyzable real-time systems. In *IOLTS*, 2016.
- [18] Z. Wang and R. B. Lee. A novel cache architecture with enhanced performance and security. In *2008 41st IEEE/ACM International Symposium on Microarchitecture*, pages 83–93, Nov 2008.



David Trilla is a PhD. Student for the CAOS group at BSC. He obtained his M.S. degree in 2016 and graduated in Informatics Engineering in 2014, both titles obtained from the Universitat Politècnica de Catalunya. He enrolled BSC in 2014 and has participated in the European project ESA-HAIR. He has been working on timing prediction models for real-time software for multicore during early design stages and his current research focuses on the effects on energy consumption behavior and security on randomized architectures.

Main Memory in HPC: Do We Need More or Could We Live with Less?

Darko Zivanovic^{*‡}, Petar Radojković^{*}, Eduard Ayguadé^{*‡}

^{*}Barcelona Supercomputing Center (BSC), [‡]Universitat Politècnica de Catalunya

darko.zivanovic@bsc.es, petar.radojkovic@bsc.es, eduard.ayguade@bsc.es

Keywords— High-performance computing, Memory capacity requirements, Production HPC applications, HPL, HPCG, Large-memory nodes, Energy-efficiency, Scaling-in.

EXTENDED ABSTRACT

An important aspect of High-performance Computing (HPC) system design is the choice of main memory capacity. This choice becomes increasingly important now that 3D-stacked memories are entering the market. Compared with conventional DIMMs, 3D memory chiplets provide better performance and energy efficiency but lower memory capacities. Hybrid memory systems, that combine 3D-stacked DRAM with standard DIMMs, should bring the best of two worlds — the bandwidth, latency and energy-efficiency of 3D-stacked DRAM together with the capacity of DIMMs. However, they are still difficult to manage, so 3D-memories will only be employed in HPC if enough applications have sufficiently small memory footprints to fit inside 3D memories exclusively.

This study analyzes the memory capacity requirements of important HPC benchmarks and applications. We find that the High Performance Conjugate Gradients benchmark could be an important success story for 3D-stacked memories in HPC, but High-performance Linpack is likely to be constrained by 3D memory capacity.

The study also emphasizes that the analysis of memory footprints of production HPC applications is complex and that it requires an understanding of application scalability and target category, i.e., whether the users target capability or capacity computing. In HPC, capability computing refers to using large-scale HPC installations to solve a single, highly complex problem in the shortest possible time, while capacity computing refers to optimizing system efficiency to solve as many mid-size or smaller problems as possible at the same time at the lowest possible cost.

The results show that most of the HPC applications under study have per-core memory footprints in the range of hundreds of megabytes, and these applications represent use cases in HPC that require memory capacities that could be provided solely by 3D memories, which is a first step toward their adoption in HPC.

We also detect applications and use cases in capacity computing that still require gigabytes of memory per core, and for these use cases we propose scaling-in, i.e. reducing the number of nodes for the execution. We show that scaling-in leads to significant energy savings and we propose upgrading the memory capacity which enables greater degree of scaling-in. We show that additional energy savings, of up to 52%, mean that in many cases the investment in upgrading the memory system would be recovered in a typical system lifetime of less than five years.

A. ACKNOWLEDGEMENT

This work is published on the International Symposium in Memory Systems (MEMSYS) [1] and in ACM Transactions on Architecture and Code Optimization (TACO) [2].

References

- [1] D. Zivanovic, M. Radulovic, G. Llort, D. Zaragoza, J. Strassburg, P. M. Carpenter, P. Radojković, E. Ayguade, “Large-Memory Nodes for Energy Efficient High-Performance Computing,” in Proceedings of the International Symposium on Memory Systems (MEMSYS), Pages 3-9, October 2016.
- [2] D. Zivanovic, M. Pavlovic, M. Radulovic, H. Shin, J. Son, S. A. McKee, P. M. Carpenter, P. Radojković, E. Ayguade, “Main Memory in HPC: Do We Need More or Could We Live with Less?,” ACM Transactions on Architecture and Code Optimization (TACO), Volume 14 Issue 1, Article No. 3, March 2017.

Author biography



Darko Zivanovic is a PhD candidate at Barcelona Supercomputing Center and Polytechnic University of Catalonia. His research is focused on Memory Systems for High-performance Computing. He obtained his B.Sc. and M.Sc. degrees from the School of Electrical Engineering at the University of Belgrade in 2008 and 2010, respectively. During his Master he joined the Institute Mihajlo Pupin in Belgrade, where he worked from 2009 to 2011 as Embedded System Developer. From 2011 to 2013 he was part of Architecture and Compilers (ARCO) research group in Barcelona, and in 2013 he joined Barcelona Supercomputing Center in a pursue for his PhD degree.



**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

/BSCCNS 

@BSC_CNS 

/BSCCNS 

bsc.es/linkedin 