# Laércio LIMA PILLA

Laboratoire Bordelais de Recherche en Informatique (LaBRI)
33405 TALENCE, France
☎ +33 5 24 57 41 55 • ✉ laercio.lima-pilla@labri.fr
www.labri.fr/perso/llimapilla/

*Researcher in High-Performance Computing*

Curriculum Vitae: November, 2024

## Education

**Univ. Grenoble Alpes & Federal University of Rio Grande do Sul** — **FR & BR**
*PhD in Computer Science* — *2010–2014*
Thesis Advisers: Prof. Jean-François Méhaut & Prof. Philippe O. A. Navaux

**Federal University of Rio Grande do Sul** — **BR**
*Bachelor in Computer Science* — *2005–2009*

## Work Experience

**CNRS** — **Talence, FR**
*Research Scientist,UMR5800 - LaBRI, Topal team (prev. Storm team)* — *2021–*

**Orsay, FR**
*Research Scientist,UMR8623 - Laboratoire de Recherche en Informatique, ParSys team* — *2018–2020*

**Inria Grenoble - Rhône-Alpes** — **Grenoble, FR**
*Postdoctoral Researcher* — *2018–2018*

**Federal University of Santa Catarina** — **Florianópolis, BR**
*Associate Professor* — *2014–2017*

**LIG & Parallel and Distributed Processing Group** — **Grenoble, FR & Porto Alegre, BR**
*Assistant Researcher* — *2010–2014*

**Parallel and Distributed Processing Group** — **Porto Alegre, BR**
*Junior Researcher, Microsoft Interop-Lab UFRGS* — *2008–2009*

**Parallel and Distributed Processing Group** — **Porto Alegre, BR**
*Junior Researcher* — *2007–2008*

## Honors & Awards

**2019**: Best Paper Award for "*On server-side file access pattern matching*", International Conference on High Performance Computing & Simulation (HPCS)

**2017**: 3rd Place Award on Multi-Deck Stand. Cell Legalization, ICCAD CAD Contest

**2015**: Best Paper Award for "*An Efficient Algorithm for Communication-Based Task Mapping*", 23rd Euromicro International Conference on Parallel, Distributed, and Network-Based Processing (PDP)

**2014**: Best PhD Thesis in the Computer Architecture and High Performance Computing Contest (WSCAD-CTD)

**2010**: 2nd place, 5th International Marathon of Parallel Programming at SBAC-PAD

**2009**: 1st place, 4th International Marathon of Parallel Programming at SBAC-PAD

## Publications and Scientific Content

Technical Reports & Pre-prints

[T3] O. Beaumont, R. Bouzel, L. Eyraud-Dubois, E. Korkmaz, L. L. Pilla, A. van Kempen, "*Approximation Algorithms for Scheduling with/without Deadline Constraints where Rejection Costs are Proportional to Processing Times*, October 2024. https://hal.science/hal-04745701/

[T2] D. Orhan, L. L. Pilla, D. Barthou, A. Cassagne, O. Aumage, R. Tajan, C. Jégo, C. Leroux, "*OTAC: Optimal Scheduling for Pipelined and Replicated Task Chains for Software-Defined Radio*", October 2023. https://hal.science/hal-04228117v1

[T1] L. L. Pilla, "*Optimal Workload Scheduling Algorithm for Data-Parallel Applications on Heterogeneous Platforms based on Dynamic Programming*", October 2022. https://hal.inria.fr/hal-03776372

Journals

[J23] L. Scravaglieri, L. L. Pilla, A. Guermouche, O. Aumage, E. Saillard, M. Popov, "*Optimizing Performance and Energy Across Problem Sizes Through a Search Space Exploration and Machine Learning*", Journal of Parallel and Distributed Computing, vol. 180, October 2023. https://doi.org/10.1016/j.jpdc.2023.104720

[J22] L. L. Pilla, "*Scheduling Algorithms for Federated Learning with Minimal Energy Consumption*", in IEEE Transactions on Parallel and Distributed Systems, vol. 34, no. 4, April 2023. https://doi.org/10.1109/TPDS.2023.3240833

[J21] M. Regragui, B. Coye, L. L. Pilla, R. Namyst, D. Barthou, "*SimSGamE : Scheduling simulator for modern game engines*", in Journal of Open Source Software, vol. 7, no. 76, pp. 4592, August 2022. https://doi.org/10.21105/joss.04592

[J20] R. Netto, S. Fabre, T. A. Fontana, V. Livramento, L. L. Pilla, L. Behjat, J. L. Guntzel, "*Algorithm Selection Framework for Legalization Using Deep Convolutional Neural Networks and Transfer Learning*", in IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems, vol. 41, no. 5, pp. 1481-1494, May 2022. https://doi.org/10.1109/TCAD.2021.3079126

[J19] A. Santana, V. Freitas, M. Castro, L. L. Pilla and J.-F. Méhaut, "*ARTful: A specification for user-defined schedulers targeting multiple HPC runtime systems*", in Software: Practice and Experience, vol. 51, pp. 1622–1638. July 2021. https://dx.doi.org/10.1002/spe.2977

[J18] V. Freitas, L. L. Pilla, A. Santana, M. Castro and J. Cohen, "*PackStealLB: A Scalable Distributed Load Balancer based on Work Stealing and Workload Discretization*", in Journal of Parallel and Distributed Computing, vol. 150, pp. 35–45, April 2021. https://doi.org/10.1016/j.jpdc.2020.12.005

[J17] E. H. M. Cruz, M. Diener, L. L. Pilla, and P. O. A. Navaux, "*Online Thread and Data Mapping Using a Sharing-Aware Memory Management Unit*", in ACM Transactions on Modeling and Performance Evaluation of Computing Systems, vol. 5, no. 4, article 16, 28 pages, January 2021. https://doi.org/10.1145/3433687

[J16] E. H. M. Cruz, M. Diener, L. L. Pilla, and P. O. A. Navaux, "*EagerMap: A Task Mapping Algorithm to Improve Communication and Load Balancing in Clusters of Multicore Systems*", in ACM Transactions on Parallel Computing, vol. 5, no. 4, article 17, 24 pages, March 2019. https://doi.org/10.1145/3309711

[J15] M. Hinz, G. P. Koslovski, C. C. Miers, L. L. Pilla and M. A. Pillon, "*A Cost Model for IAAS Clouds based on Virtual Machine Energy Consumption*", in Journal of Grid Computing, vol. 16, no. 3, pp. 493–512, September 2018. https://doi.org/10.1007/s10723-018-9440-8

[J14] A. S. Furtado, L. L. Pilla and V. Bogorny, "*A Branch and Bound Strategy for Fast Trajectory Similarity Measuring*", in Data and Knowledge Engineering, vol. 115, pp. 16–31, May 2018. https://doi.org/10.1016/j.datak.2018.01.003

[J13] R. da R. Righi, R. de Q. Gomes, V. F. Rodrigues, C. A. da Costa, A. M. Alberti, L. L. Pilla and P. O. A. Navaux, "*MigPF: Towards on self-organizing process rescheduling of Bulk-Synchronous Parallel applications*", in Future Generation Computer Systems, vol. 78, part 1, pp. 272–286, January 2018. https://doi.org/10.1016/j.future.2016.05.004

[J12] E. H. M. Cruz, M. Diener, L. L. Pilla and P. O. A. Navaux, "*Hardware-Assisted Thread and Data Mapping in Hierarchical Multicore Architectures*", in ACM Transactions on Architecture and Code Optimization, vol. 13, no. 3, article 28, 28 pages, September 2016. https://doi.org/10.1145/2975587

[J11] E. H. M. Cruz, M. Diener, M. A. Z. Alves, L. L. Pilla and P. O. A. Navaux, "*LAPT: A locality-aware page table for thread and data mapping*", in Parallel Computing, vol. 54, pp. 59–71, May 2016. https://doi.org/10.1016/j.parco.2015.12.001

[J10] D. A. G. de Oliveira, L. L. Pilla, T. Santini and P. Rech, "*Evaluation and Mitigation of Radiation-Induced Soft Errors in Graphics Processing Units*", in IEEE Transactions on Computers, vol. 65, no. 3, pp. 791–804, March 2016. https://doi.org/10.1109/TC.2015.2444855

[J9] L. L. Pilla, D. A. G. Oliveira, C. Lunardi, P. O. A. Navaux, L. Carro and P. Rech, "*Memory Access Time and Input Size Effects on Parallel Processors Reliability*", in IEEE Transactions on Nuclear Science, vol. 62, no. 6, pp. 2627–2634, December 2015. https://doi.org/10.1109/TNS.2015.2496381

[J8] L. L. Pilla, T. C. Bozzetti, M. Castro, P. O. A. Navaux and J.-F. Méhaut, "*ComprehensiveBench: a Benchmark for the Extensive Evaluation of Global Scheduling Algorithms*", in Journal of Physics: Conference Series, vol. 649, conf. 1, 12 pages, October 2015. https://doi.org/10.1088/1742-6596/649/1/012007

[J7] M. Diener, E. H. M. Cruz, L. L. Pilla, F. Dupros and P. O. A. Navaux, "*Characterizing communication and page usage of parallel applications for thread and data mapping*", in Performance Evaluation, vol. 88-89, pp. 18–36, June 2015. https://doi.org/10.1016/j.peva.2015.03.001

[J6] E. L. Padoin, L. L. Pilla, M. Castro, F. Z. Boito, P. O. A. Navaux and J.-F. Méhaut, "*Performance/energy trade-off in scientific computing: the case of ARM big.LITTLE and Intel Sandy Bridge*", in IET Computers & Digital Techniques, vol. 9, no. 1, pp. 27–35, January 2015. https://doi.org/10.1049/iet-cdt.2014.0074

[J5] L. L. Pilla, P. Rech, F. Silvestri, C. Frost, P. O. A. Navaux, M. Sonza Reorda and L. Carro, "*Software-Based Hardening Strategies for Neutron Sensitive FFT Algorithms on GPUs*", in IEEE Transactions on Nuclear Science, vol. 61, no. 4, pp. 1874–1880, August 2014. https://doi.org/10.1109/TNS.2014.2301768

[J4] L. L. Pilla, C. P. Ribeiro, P. Coucheney, F. Broquedis, B. Gaujal, P. O. A. Navaux and J.-F. Méhaut, "*A topology-aware load balancing algorithm for clustered hierarchical multi-core machines*", in Future Generation Computer Systems, vol. 30, pp. 191–201, January 2014. urlhttps://doi.org/10.1016/j.future.2013.06.023

[J3] E. L. Padoin, L. L. Pilla, F. Z. Boito, R. V. Kassick, P. Velho and P. O. A. Navaux, "*Evaluating application performance and energy consumption on hybrid CPU+GPU architecture*", in Cluster Computing, vol. 16, no. 3, pp 511–525, September 2013. https://doi.org/10.1007/s10586-012-0219-6

[J2] C. Osthoff, F. Z. Boito, R. V. Kassick, L. L. Pilla, P. O. A. Navaux, C. Schepke, J. Panetta, P. J. Grunmann, N. Maillard, P. L. S. Dias and R. L. Walko, "*Atmospheric models hybrid OpenMP/MPI*

*implementation multicore cluster evaluation*", in International Journal of Information Technology, Communications and Convergence, vol. 2, no. 3, pp. 212–233, November 2012. http://doi.org/10.1504/ijitcc.2012.050411

[J1] R. da R. Righi, L. L. Pilla, N. Maillard, A. Carissimi and P. O. A. Navaux, "*Observing the impact of multiple metrics and runtime adaptations on BSP process rescheduling*", in Parallel Processing Letters, vol. 20, no. 2, pp. 123–144, June 2010. https://doi.org/10.1142/S0129626410000107

Conferences

[C38] A. L. Nunes, C. Boeres, L. M. A. Drummond, L. L. Pilla, "*Optimal Time and Energy-Aware Client Selection Algorithms for Federated Learning on Heterogeneous Resources*", to be published in Proceedings of the 36th International Symposium on Computer Architecture and High Performance Computing (SBAC-PAD), November 2024. https://hal.science/hal-04690494/

[C37] O. Beaumont, R. Bouzel, L. Eyraud-Dubois, E. Korkmaz, L. L. Pilla, A. van Kempen, "*A $1.25(1+\epsilon)$-Approximation Algorithm for Scheduling with Rejection Costs Proportional to Processing Times*", in European Conference on Parallel Processing (Euro-Par), pp. 225–238, August 2024. https://doi.org/10.1007/978-3-031-69577-3_16

[C36] M. Regragui, B. Coye, L. L. Pilla, R. Namyst, D. Barthou, "*Exploring scheduling algorithms for parallel task graphs: a modern game engine case study*", in European Conference on Parallel Processing (Euro-Par), pp. 103–118, August 2022. https://doi.org/10.1007/978-3-031-12597-3_7

[C35] L. L. Pilla, "*Optimal Task Assignment for Heterogeneous Federated Learning Devices*", in 2021 IEEE International Parallel and Distributed Processing Symposium (IPDPS), pp. 661–670, May 2021. https://doi.org/10.1109/IPDPS49936.2021.00074

[C34] J. H. M. Korndörfer, M. Bielert, L. L. Pilla and F. M. Ciorba, "*Mapping Matters: Application Process Mapping on 3-D Processor Topologies*", to be published in the 2020 International Conference on High Performance Computing & Simulation (HPCS) 2021. https://hal.archives-ouvertes.fr/hal-02617127v1

[C33] A. V. C. R. Oikawa, V. Freitas, M. Castro and L. L. Pilla, "*Adaptive Load Balancing based on Machine Learning for Iterative Parallel Applications*", in 28th Euromicro International Conference on Parallel, Distributed, and Network-Based Processing (PDP), pp. 94–101, March 2020. https://doi.org/10.1109/PDP50117.2020.00021

[C32] V. Freitas, A. Santana, M. Castro and L. L. Pilla, "*Distributed Memory Graph Representation for Load Balancing Data: Accelerating Data Structure Generation for Decentralized Scheduling*", in International Conference on High Performance Computing & Simulation (HPCS), pp. 787–794, July 2019. https://doi.org/10.1109/HPCS48598.2019.9188134

[C31] F. Z. Boito, R. Nou, L. L. Pilla, J. L. Bez, J.-F. Méhaut, T. Cortes and P. O. A. Navaux, "*On server-side file access pattern matching*", in International Conference on High Performance Computing & Simulation (HPCS), pp. 217–224, July 2019. https://doi.org/10.1109/HPCS48598.2019.9188092

[C30] R. Netto, S. Fabre, T. Fontana, V. Livramento, L. L. Pilla and J. L. Güntzel, "*How Deep Learning Can Drive Physical Synthesis Towards More Predictable Legalization*", in Proceedings of the 2019 International Symposium on Physical Design (ISPD), pp 3–10, April 2019. https://doi.org/10.1145/3299902.3309754

[C29] V. Freitas, A. L. Santana, M. Castro and L. L. Pilla, "*A Batch Task Migration Approach for Decentralized Global Rescheduling*", in Proceedings of the 30th International Symposium on Computer Architecture

and High Performance Computing (SBAC-PAD), pp. 49–56, September 2018. https://doi.org/10.1109/CAHPC.2018.8645953

[C28] A. Santana, V. Freitas, L. L. Pilla, M. Castro and J.-F. Méhaut, "*Reducing Global Schedulers' Complexity Through Runtime System Decoupling*", in Proceedings of the 2018 Symposium on High Performance Computing Systems (WSCAD), pp.38–44, September 2018. https://doi.org/10.1109/WSCAD.2018.00016

[C27] S. Fabre, J. L. Güntzel, L. L. Pilla, R. Netto, T. Fontana and V. Livramento, "*Enhancing Multi-Threaded Legalization Through k-d Tree Circuit Partitioning*", in Proceedings of the 31st Symposium on Integrated Circuits and Systems Design (SBCCI), pp. 1–6, August 2018. https://doi.org/10.1109/SBCCI.2018.8533264

[C26] E. Cruz, M. Diener, M. Serpa, P. Navaux, L. L. Pilla and I. Koren, "*Improving Communication and Load Balancing with Thread Mapping in Manycore Systems*", in 26th Euromicro International Conference on Parallel, Distributed, and Network-Based Processing (PDP), pp. 93–100, March 2018. https://doi.org/10.1109/PDP2018.2018.00021

[C25] D. Oliveira, L. Pilla, N. DeBardeleben, S. Blanchard, H. Quinn, I. Koren, P. Navaux and Paolo Rech, "*Experimental and analytical study of Xeon Phi reliability*", in Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC), article 28, 12 pages, November 2017. https://doi.org/10.1145/3126908.3126960

[C24] T. A. Fontana, S. Almeida, R. Netto, V. Livramento, C. Guth, L. Pilla and J. L. Güntzel, "*Exploiting cache locality to speedup register clustering*", in Proceedings of the 30th Symposium on Integrated Circuits and Systems Design: Chip on the Sands (SBCCI), pp. 191–197, September 2017. https://doi.org/10.1145/3109984.3110005

[C23] T. Fontana, R. Netto, V. Livramento, C. Guth, S. Almeida, L. Pilla and J. L. Güntzel, "*How Game Engines Can Inspire EDA Tools Development: A use case for an open-source physical design library*", in Proceedings of the 2017 ACM on International Symposium on Physical Design (ISPD), pp. 25–31, March 2017. https://doi.org/10.1145/3036669.3038248

[C22] D. A. G. de Oliveira, L. L. Pilla, M. Hanzich, V. Fratin, F. Fernandes, C. Lunardi, J. M. Cela, P. O. A. Navaux, L. Carro and P. Rech, "*Radiation-Induced Error Criticality in Modern HPC Parallel Accelerators*", in IEEE International Symposium on High Performance Computer Architecture (HPCA), pp. 577–588, February 2017. https://doi.org/10.1109/HPCA.2017.41

[C21] R. Netto, C. Guth, V. Livramento, M. Castro, L. L. Pilla and J. L. Güntzel, "*Exploiting parallelism to speed up circuit legalization*", in IEEE International Conference on Electronics, Circuits and Systems (ICECS), pp. 624–627, December 2016. https://doi.org/10.1109/ICECS.2016.7841279

[C20] R. C. De Moura, G. O. Torres, M. L. Pilla, L. L. Pilla, A. T. Da Costa and F. M. G. França, "*Value Reuse Potential in ARM Architectures*", in 28th International Symposium on Computer Architecture and High Performance Computing (SBAC-PAD), pp. 174–181, October 2016. https://doi.org/10.1109/SBAC-PAD.2016.30

[C19] E. L. Padoin, L. L. Pilla, M. Castro, P. O. A. Navaux and J.-F. Méhaut, "*Exploration of Load Balancing Thresholds to Save Energy on Iterative Applications*", in Third Latin American Conference in High Performance Computing (CARLA), pp. 76–88, September 2016. https://doi.org/10.1007/978-3-319-57972-6_6

[C18] E. H. M. Cruz, M. Diener, L. L. Pilla and P. O. A. Navaux, "*A Sharing-Aware Memory Management Unit for Online Mapping in Multi-core Architectures*", in 22nd International Conference on Parallel

and Distributed Computing (Euro-Par), pp. 490–501, August 2016. https://doi.org/10.1007/978-3-319-43659-3_36

[C17] D. A. G. de Oliveira, L. Pilla, C. Lunardi, L. Carro, P. O. A. Navaux and Paolo Rech, "*The Path to Exascale: Code Optimizations and Hardening Solutions Reliability*", in Proceedings of the 5th Workshop on Fault Tolerance for HPC at eXtreme Scale (FTXS), pp. 55–62, June 2015. https://doi.org/10.1145/2751504.2751510

[C16] E. H. M. Cruz, M. Diener, L. L. Pilla and P. O. A. Navaux, "*An Efficient Algorithm for Communication-Based Task Mapping*", in 23rd Euromicro International Conference on Parallel, Distributed, and Network-Based Processing (PDP), pp. 207–214, March 2015. https://doi.org/10.1109/PDP.2015.25

[C15] E. L. Padoin, M. Castro, L. L. Pilla, P. O. A. Navaux and J. F. Méhaut, "*Saving energy by exploiting residual imbalances on iterative applications*", in 21st International Conference on High Performance Computing (HiPC), 10 pages, December 2014. https://doi.org/10.1109/HiPC.2014.7116895

[C14] D. A. G. Oliveira, P. Rech, L. L. Pilla, P. O. A. Navaux and L. Carro, "*GPGPUs ECC efficiency and efficacy*", in IEEE International Symposium on Defect and Fault Tolerance in VLSI and Nanotechnology Systems (DFT), pp. 209–215, October 2014. https://doi.org/10.1109/DFT.2014.6962085

[C13] E. H. M. Cruz, M. Diener, M. A. Z. Alves, L. L. Pilla and P. O. A. Navaux, "*Optimizing Memory Locality Using a Locality-Aware Page Table*", in IEEE 26th International Symposium on Computer Architecture and High Performance Computing (SBAC-PAD), pp. 198–205, October 2014. https://doi.org/10.1109/SBAC-PAD.2014.22

[C12] D. A. G. Oliveira, C. B. Lunardi, L. L. Pilla, P. Rech, P. O. A. Navaux and L. Carro, "*Radiation Sensitivity of High Performance Computing Applications on Kepler-Based GPGPUs*", in 44th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN), pp. 732–737, June 2014. https://doi.org/10.1109/DSN.2014.75

[C11] P. Rech, L. L. Pilla, P. O. A. Navaux and L. Carro, "*Impact of GPUs Parallelism Management on Safety-Critical and HPC Applications Reliability*", in 44th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN), pp. 455–466, June 2014. https://doi.org/10.1109/DSN.2014.49

[C10] R. K. Tesser, L. L. Pilla, F. Dupros, P. O. A. Navaux, J. F. Méhaut and C. Mendes, "*Improving the Performance of Seismic Wave Simulations with Dynamic Load Balancing*", 22nd Euromicro International Conference on Parallel, Distributed, and Network-Based Processing (PDP), pp. 196–203, February 2014. https://doi.org/10.1109/PDP.2014.37

[C9] L. L. Pilla, P. O. A. Navaux, C. P. Ribeiro, P. Coucheney, F. Broquedis, B. Gaujal and J.-F. Méhaut, "*Asymptotically Optimal Load Balancing for Hierarchical Multi-Core Systems*", in IEEE 18th International Conference on Parallel and Distributed Systems (ICPADS), pp. 236–243, December 2012. https://doi.org/10.1109/ICPADS.2012.41

[C8] L. L. Pilla, C. P. Ribeiro, D. Cordeiro, C. Mei, A. Bhatele, P. O. A. Navaux, F. Broquedis, J.-F. Méhaut, L. V. Kale "*A Hierarchical Approach for Load Balancing on Parallel Multi-core Systems*", in 41st International Conference on Parallel Processing (ICPP), pp. 118–127, September 2012. https://doi.org/10.1109/ICPP.2012.9

[C7] R. da R. Righi, L. Graebin, R. B. Avila, P. O. A. Navaux and L. L. Pilla, "*Combining Multiple Metrics to Control BSP Process Rescheduling in Response to Resource and Application Dynamics*", in IEEE 17th International Conference on Parallel and Distributed Systems (ICPADS), pp. 72–79, December 2011. https://doi.org/10.1109/ICPADS.2011.44

[C6] C. Osthoff, P. Grunmann, F. Boito, R. Kassick, L. Pilla, P. Navaux, C. Schepke, J. Panetta, N. Maillard, P. L. S. Dias and R. Walko, "*Improving Performance on Atmospheric Models through a Hybrid OpenMP/MPI Implementation*", in IEEE Ninth International Symposium on Parallel and Distributed Processing with Applications (ISPA), pp. 69–74, May 2011. https://doi.org/10.1109/ISPA.2011.56

[C5] R. da R. Righi, L. L. Pilla, A. Carissimi, P. O. A. Navaux, H.-U. Heiss, "*Applying Process Migration on a BSP-Based LU Decomposition Application*", in 9th International Conference on High Performance Computing for Computational Science (VECPAR), pp. 314–326, June 2010. https://doi.org/10.1007/978-3-642-19328-6_30

[C4] R. da R. Righi, L. L. Pilla, A. Carissimi, P. Navaux and H. U. Heiss, "*MigBSP: A Novel Migration Model for Bulk-Synchronous Parallel Processes Rescheduling*", in 11th IEEE International Conference on High Performance Computing and Communications (HPCC), pp. 585–590, June 2009. https://doi.org/10.1109/HPCC.2009.54

[C3] R. da R. Righi, L. L. Pilla, A. S. Carissimi, P. O. A. Navaux and H.-U. Heiss, "*Applying Processes Rescheduling over Irregular BSP Application*", in 9th International Conference on Computational Science (ICCS), pp. 213–223, May 2009. https://doi.org/10.1007/978-3-642-01970-8_22

[C2] R. da R. Righi, L. L. Pilla, A. Carissimi and P. O. A. Navaux, "*Controlling Processes Reassignment in BSP Applications*", in 20th International Symposium on Computer Architecture and High Performance Computing (SBAC-PAD), pp. 37–44, November 2008. https://doi.org/10.1109/SBAC-PAD.2008.16

[C1] R. Righi, L. Pilla, A. Carissimi, N. Maillard and P. Navaux, "*ICE: Managing Multiple Clusters Using Web Services*", 11th IEEE International Conference on Computational Science and Engineering (CSE), pp. 64–71, July 2008. https://doi.org/10.1109/CSE.2008.61

Book Chapters

[B5] A. V. Schmidt, F. Z. Boito and L. L. Pilla, "*Fundamentos de Estatística para Análise de Desempenho*", in Anais da XVII Escola Regional de Alto Desempenho (ERAD/RS), editors E. L. Padoin, M. C. Cera and A. S. Charao, vol. 17, chapter 1, pp. 3-24, Sociedade Brasileira de Computação, 2017.

[B4] L. L. Pilla and E. Meneses, "*Programação Paralela em Charm++*", in Anais da XV Escola Regional de Alto Desempenho (ERAD/RS), editors C. de Rose, M. Pasin and L. M. Schnorr, vol. 15, chapter 2, pp. 25-44, Sociedade Brasileira de Computação, 2015.

[B3] C. Osthoff, R. P. Souto, F. Vilasbôas, P. Grunmann, P. L. S. Dias, F. Z. Boito, R. V. Kassick, L. L. Pilla, P. O. A. Navaux, C. Schepke, N. Maillard, J. Panetta, P. P. Lopes, R. Walko, "*Improving Atmospheric Model Performance on a Multi-Core Cluster System*", in Atmospheric Model Applications, editor I. Yucel, vol. 1, chapter 1, pp. 1-24, InTech Europe, 2012.

[B2] P. O. A. Navaux, C. de Rose and L. L. Pilla, "*Fundamentos das Arquiteturas para Processamento Paralelo e Distribuído*", in Anais da XI Escola Regional de Alto Desempenho (ERAD/RS), editors N. Maillard, C. de Rose, A. M. S. Adario, B. G. Apel and S. D. K Mor, vol. 11, chapter 2, pp. 22-59, Sociedade Brasileira de Computação, 2011.

[B1] R. da R. Righi, L. L. Pilla, N. Maillard, A. Carrisimi and P. O. A. Navaux, "*Process Rescheduling: Enabling Performance by Applying Multiple Metrics and Efficient Adaptations*", in Future Manufacturing Systems, editor Tauseef Aized, vol. 1, chapter 3, pp. 39-64, Sycio, 2010.

Datasets

[D4] O. Beaumont, L. Eyraud-Dubois, E. Korkmaz, L. L. Pilla, "*Experimental codes and results for the paper "A 5/4(1+eps)-Approximation Algorithm for Scheduling with Rejection Costs Proportional to Processing Times"*", March 2024. `https://inria.hal.science/hal-04517532/`

[D3] L. L. Pilla, "*llpilla/energy-optimal-federated-learning: Full set of algorithms: (MC)$^2$MKP, MarIn, MarCo, MarDec, and MarDecUn*", November 2022. `https://doi.org/10.5281/zenodo.7377419`

[D2] M. Regragui, B. Coye, L. L. Pilla, R. Namyst, D. Barthou, "*Performance results of different scheduling algorithms used in the simulation of a modern game engine*", May 2022. `https://doi.org/10.5281/zenodo.6532251`

[D1] L. L. Pilla and V. Freitas, "*Performance results of LeanMD on the Joliot-Curie supercomputer using different load balancing algorithms*", June 2020. `https://doi.org/10.5281/zenodo.3878953`

## Research Projects

**2024–2026**: Data movement, Energy COnsumption and performance in High-Performance Computing (DECoHPC), *Inria + LNCC, UFF, UFRGS, and CEFET-RJ joint team*

**2022–2026**: PUshing Low-carbon Services towards the Edge (PULSE), *Défi Commun Inria – Qarnot computing with the support of Ademe*

**2022–2024**: Maelstrom, *Inria + Simula joint team*

**2021–2022**: 3BEARS: Broad Bundle of BEnchmARks for Scheduling in HPC, Big Data, and ML, *PHC Germaine de Staël, (Project number 2021-13)*

**2019–2020**: Exploring the Trade-offs of Scheduling in Two Contexts, *LRI Internal Funding*

**2019–2019**: Mapping for Emerging Memory Models in HPC Systems, *CNRS INS2I PEPS*

**2017–2018**: Adaptive Global Scheduling for Scientific Applications, *CNPq Universal (grant 401266/2016-8)*

**2016–2017**: Energy-aware Scheduling and Fault Tolerance Techniques for the Exascale Era, *STIC-AmSud (grant 99999.007556/2015-02)*

## Professional Service

**PC Member**: CARLA (2015–2019), COMPAS (2021–2023), DRMEC (2019), ICCS (2019–2022), HCW (2020–2023), HPCEuropeLatAm (2022), ISC PhD Forum (2019, 2021), PAW-ATM (2022), SBAC-PAD (2014, 2015, 2019, 2023), SBESC (2017), SC (2020), SC Tutorials (2020), WCC (2023).

**Journal Reviewer**: 4OR, ACM Surveys, ACM TACO, ACM TCPS, CCPE, FGCS, IEEE TC, IEEE TNS, IEEE TPDS Reproducibility, IEEE TVLSI, JPDC, JSS, PARCO, SPE.

## Teaching: Long Courses

**Federal University of Santa Catarina, BR**: Advanced Topics in Computer Architecture, Compilers Construction, Computer Architecture and Organization, Computer Organization Laboratory, Introduction to Informatics, Parallel Computing (graduate course).

**INP Bordeaux, FR**: Algorithms for High-Performance Computing Platforms (*currently teaching*).

**Polytech Paris-Saclay, FR**: GPU Programming.

**Université Paris-Saclay, FR**: Scheduling and Runtime Systems.

## Teaching: Short Courses

**2017**: Adaptive MPI, Massively Parallel Programming Week, National Laboratory for Scientific Computing (LNCC), Brazil

**2016**: Adaptive MPI, Massively Parallel Programming Week, National Laboratory for Scientific Computing (LNCC), Brazil

**2015**: Programming with Parallel Objects: from MPI to Charm++, 27th International Symposium on Computer Architecture and High Performance Computing (SBAC-PAD)

**2015**: Parallel Programming with Charm++, 16th High Performance Computing Regional School

## Software

Self-assessment (*software* and *contribution*) based on Inria's criteria.

**2020–2022**: Unified Scheduling Simulator. *Software:* A-1-up, SO-3, SM-2-up, EM-2, SDL-4. *Contribution:* DA-4, CD-4, MS-4, TPM-4. Available at `https://github.com/llpilla/unified-scheduling-simulato`

**2019–2020**: Communication Statistics. *Software:* A-2, SO-2, SM-3, EM-2, SDL-4. *Contribution:* DA-4, CD-4, MS-4, TPM-4. Available at `https://github.com/llpilla/communication-statistics`

**2017–2020**: MOGSLib. *Software:* A-2-up, SO-3, SM-3-up, EM-3, SDL-4. *Contribution:* DA-2, CD-2, MS-2, TPM-4. Available at `https://github.com/ECLScheduling/MOGSLib`

**2016–2019**: Ophidian. *Software:* A-3, SO-4, SM-3-up, EM-3, SDL-4. *Contribution:* DA-1, CD-1, MS-1, TPM-2. Available at `https://github.com/eclufsc/ophidian`

**2011–2015**: HieSchella. *Software:* A-2, SO-4, SM-2, EM-1, SDL-4. *Contribution:* DA-4, CD-4, MS-4, TPM-4. Available at `https://github.com/llpilla/HieSchella`

**2008–2010**: hpcgpu. *Software:* A-3, SO-2, SM-2, EM-1, SDL-4. *Contribution:* DA-4, CD-4, MS-4, TPM-4. Archived at `https://archive.codeplex.com/?p=hpcgpu`

## Keywords

**Scientific**: Scheduling, Load Balancing, Computer Architectures, Fault Tolerance, Parallel Computing

**Technical**: Linux, C/C++, Bash, Python, R, MPI, OpenMP, Charm++, CUDA, Git, LaTeX

**Leisure**: Computer & Board Games, Traveling, Role Playing Games, Science Fiction Media