

Loïc Pottier

Research Scientist in High-Performance Computing

4676 Admiralty Way, Suite 1001
Marina Del Rey, CA 90292

✉ lpottier@isi.edu

🌐 pottier.github.io

French Nationality

Employments

- 2020 – Now **Computer Scientist**, *University of Southern California*, Los Angeles, USA. I am part of the Science Automation Technologies group at the Information Sciences Institute (ISI). I am working on scientific workflows management (scheduling and data management) for large-scale infrastructures, with a strong focus on HPC systems. I have also recently started exploring I/O optimizations for machine learning workloads at scale.
- 2019 – 2020 **Postdoctoral scholar – Research Associate**, *University of Southern California*, Los Angeles, USA. I have been working under the supervision of Dr Ewa Deelman in the Science Automation Technologies research group. I have been working on scientific workflows management (scheduling and data management) for large-scale infrastructures (in particular HPC system).
- 2015 – 2018 **Teaching Assistant in Computer Science**, *École Normale Supérieure de Lyon and University of Lyon*. I have been a teaching assistant for 3 years during my PhD. See below for the list of my classes.
- 2015 – 2018 **PhD in Computer Science**, *LIP laboratory – École Normale Supérieure de Lyon*. “Co-scheduling for large-scale applications: memory and resilience”, under the supervision of Anne Benoit and Yves Robert in the ROMA team, defended on September 18, 2018
- 2016 (3 months) **Research intern**, *Argonne National Laboratory*, Chicago, USA. I have been working with Swann Perarnau on scheduling and data management problems for the new many-core architectures that exhibit new memory hierarchies, such as Xeon Phi Knights Landing.

Education

- 2015 – 2018 **PhD in Computer Science**, *LIP laboratory – École Normale Supérieure de Lyon*. Advisors: Anne Benoit and Yves Robert, defended on September 18, 2018
- 2013 – 2015 **Master of Science**, *University of Versailles, with High Honors («Bien»)*. Major in Computer Science, specialized in High Performance Computing
- 2010 – 2013 **Bachelor of Science**, *University of Lower Normandy, with Honors («Assez Bien»)*. Major in Computer Science
- 2010 **High school diploma**, Caen, Lower Normandy, *with Honors («Assez Bien »)*. Major in Science, minor in Mathematics

Funding

- 2021 – 2023 **Scalability of Deep-Learning Methods on HPC Systems: An I/O-centric Approach**, *Principal Investigator*, \$175,000, NSF CISE Research Initiation Initiative (CRII). #2105044
- 2021 – 2024 **Simulation-driven Evaluation of Cyberinfrastructure Systems**, *Principal Investigator*, \$600,000, NSF Software Institutes: Cyberinfrastructure for Sustained Scientific Innovation (CSSI) Elements: Collaborative Research. #2103508, #2103489
- 2021 – 2024 **Simulation-driven runtime resource management for distributed workflow applications**, *Principal Investigator*, \$499,988, NSF CISE Core Program: Collaborative Research. #2106147, #2106059
- 2019 – 2022 **Integrating core CI literacy and skills into university curricula via simulation-driven activities**, *Principal Investigator*, \$500,000, NSF CyberTraining - Training-based: Collaborative Research: CyberTraining: Implementation: Small. #1923539, #1923621

Publications — For publications prefixed by * authors are listed in alphabetical order

Thesis

- [T1] L. Pottier. “Co-scheduling for large-scale applications : memory and resilience”. PhD thesis. Université de Lyon, Sept. 2018.

Book Chapters

- [*B1] G. Aupy, A. Benoit, L. Pottier, P. Raghavan, Y. Robert, and M. Shantharam. “Co-scheduling high-performance computing applications”. In: *Big Data Management and Processing*. Ed. by K.-C. Li, H. Jiang, and A. Zomaya. Chapman and Hall/CRC Press, 2017. Chap. 5. ISBN: 9781351650045.

Articles in International Refereed Journals

- [J1] T. Coleman, H. Casanova, L. Pottier, M. Kaushik, E. Deelman, and R. Ferreira da Silva. “WfCommons: A framework for enabling scientific workflow research and development”. In: *Future Generation Computer Systems* 128 (2022). Funding Acknowledgments: NSF 1923539, pp. 16–27. ISSN: 0167-739X. DOI: 10.1016/j.future.2021.09.043.
- [J2] T. M. A. Do, L. Pottier, S. Caíno-Lores, R. Ferreira da Silva, M. A. Cuendet, H. Weinstein, T. Estrada, M. Taufer, and E. Deelman. “A Lightweight Method for Evaluating In Situ Workflow Efficiency”. In: *Journal of Computational Science* 48 (2021). Funding Acknowledgments: NSF 1741040, DOE DE-SC0012636, p. 101259. DOI: 10.1016/j.jocs.2020.101259.
- [*J3] G. Aupy, A. Benoit, B. Goglin, L. Pottier, and Y. Robert. “Co-scheduling HPC workloads on cache-partitioned CMP platforms”. In: *International Journal of High Performance Computing Applications* (Apr. 2019). DOI: 10.1177/1094342019846956.
- [*J4] G. Aupy, A. Benoit, S. Dai, L. Pottier, P. Raghavan, Y. Robert, and M. Shantharam. “Co-scheduling Amdahl applications on cache-partitioned systems”. In: *International Journal of High Performance Computing and Applications* (2017). DOI: 10.1177/1094342017710806.
- [*J5] A. Benoit, L. Pottier, and Y. Robert. “Resilient co-scheduling of malleable applications”. In: *International Journal of High Performance Computing and Applications* (2017). DOI: 10.1177/1094342017704979.

Articles in International Refereed Conferences

- [C1] T. M. A. **Do**, L. **Pottier**, O. Yildiz, K. Vahi, P. Krawczuk, T. Peterka, and E. Deelman. “Accelerating Scientific Workflows on HPC Platforms with In Situ Processing”. In: *2022 IEEE/ACM 22nd International Symposium on Cluster, Cloud and Internet Computing (CCGrid)*. Funding Acknowledgments: NSF 1664162, DOE DE-AC02-06CH11357, DE-AC02-05CH11231, DE-SC0012636 and DE-SC0022328. IEEE. 2022, To appear. *The highlighted authors are joint first authors with equal contributions.*
- [C2] T. M. A. Do, L. Pottier, S. Thomas, R. Ferreira da Silva, M. A. Cuendet, H. Weinstein, T. Estrada, M. Taufer, and E. Deelman. “A Novel Metric to Evaluate In Situ Workflows”. In: *International Conference on Computational Science (ICCS)*. Funding Acknowledgments: NSF 1741040. 2020, pp. 538–553. DOI: 10.1007/978-3-030-50371-0_40.
- [C3] L. Pottier, R. Ferreira da Silva, H. Casanova, and E. Deelman. “Modeling the Performance of Scientific Workflow Executions on HPC Platforms with Burst Buffers”. In: *2020 IEEE International Conference on Cluster Computing (CLUSTER)*. Funding Acknowledgments: DOE DE-SC0012636, NSF 1664162, NSF 1741040, NSF 1923539, NSF 1923621. 2020, pp. 92–103. DOI: 10.1109/CLUSTER49012.2020.00019.
- [C4] E. Deelman et al. “Cyberinfrastructure Center of Excellence Pilot: Connecting Large Facilities Cyberinfrastructure”. In: *15th International Conference on eScience (eScience)*. Funding Acknowledgments: NSF 1842042. San Diego, CA, USA, 2019.
- [C5] S. Thomas, M. Wyatt, T. M. A. Do, L. Pottier, R. Ferreira da Silva, H. Weinstein, M. A. Cuendet, T. Estrada, E. Deelman, and M. Taufer. “Characterization of In Situ and In Transit Analytics of Molecular Dynamics Simulations for Next-generation Supercomputers”. In: *15th International Conference on eScience (eScience)*. Funding Acknowledgments: NSF 1741040. 2019, pp. 188–198. DOI: 10.1109/eScience.2019.00027.
- [*C6] G. Aupy, A. Benoit, B. Goglin, L. Pottier, and Y. Robert. “Co-scheduling HPC workloads on cache-partitioned CMP platforms”. In: *IEEE International Conference on Cluster Computing, CLUSTER 2018, Belfast, UK, September 10-13*. IEEE. Sept. 2018. DOI: 10.1109/CLUSTER.2018.00052.
- [*C7] A. Benoit, S. Perarnau, L. Pottier, and Y. Robert. “A performance model to execute workflows on high-bandwidth-memory architectures”. In: *47th International Conference on Parallel Processing, ICPP 2018, Eugene, USA, August 13-16*. Aug. 2018. DOI: 10.1145/3225058.3225110.
- [*C8] A. Benoit, L. Pottier, and Y. Robert. “Resilient application co-scheduling with processor redistribution”. In: *45th International Conference on Parallel Processing, ICPP 2016, Philadelphia, USA, August 16-19*. Aug. 2016. DOI: 10.1109/ICPP.2016.21.

Articles in International Refereed Workshops

- [W1] H. Casanova, Y. Ching Wong, L. Pottier, and R. Ferreira da Silva. “On the Feasibility of Simulation-driven Portfolio Scheduling for Cyberinfrastructure Runtime Systems”. In: *2022 Job Scheduling Strategies for Parallel Processing (JSSPP)*. Funding Acknowledgments: NSF 2106059 and 2106147, DOE DE-AC05-00OR22725. Springer Nature, 2022, To appear.

- [W2] T. M. A. Do, L. Pottier, R. Ferreira da Silva, S. Cano-Lores, M. Taufer, and E. Deelman. “Assessing Resource Provisioning and Allocation of Ensembles of In Situ Workflows”. In: *50th International Conference on Parallel Processing Workshop*. ICPP Workshops '21. Funding Acknowledgments: NSF 1741040, DOE SC0012636. Lemont, IL, USA: Association for Computing Machinery, 2021. ISBN: 9781450384414. DOI: 10.1145/3458744.3474051.
- [W3] R. Ferreira da Silva et al. “A Community Roadmap for Scientific Workflows Research and Development”. In: *2021 IEEE Workshop on Workflows in Support of Large-Scale Science (WORKS)*. 2021, pp. 81–90. DOI: 10.1109/WORKS54523.2021.00016.
- [W4] P. Krawczuk et al. “A Performance Characterization of Scientific Machine Learning Workflows”. In: *2021 IEEE/ACM Workflows in Support of Large-Scale Science (WORKS)*. Funding Acknowledgments: DOE DE-SC0012636, NSF 1664162. 2021, pp. 58–65. DOI: 10.1109/WORKS54523.2021.00013.
- [W5] R. Ferreira da Silva, L. Pottier, T. Coleman, E. Deelman, and H. Casanova. “WorkflowHub: Community Framework for Enabling Scientific Workflow Research and Development”. In: *2020 IEEE/ACM Workflows in Support of Large-Scale Science (WORKS)*. Funding Acknowledgments: NSF 2016619, DOE DE-SC0012636, NSF 1664162, NSF 1923539. 2020, pp. 49–56. DOI: 10.1109/WORKS51914.2020.00012.
- [W6] R. Mitchell, L. Pottier, S. Jacobs, R. Ferreira da Silva, M. Rynge, K. Vahi, and E. Deelman. “Exploration of Workflow Management Systems Emerging Features from Users Perspectives”. In: *First International Workshop on Big Data Tools, Methods, and Use Cases for Innovative Scientific Discovery (BTSD)*. Funding Acknowledgments: NSF 1842042. 2019.
- [*W7] G. Aupy, A. Benoit, L. Pottier, P. Raghavan, Y. Robert, and M. Shantharam. “Co-scheduling algorithms for cache-partitioned systems”. In: *19th Workshop on Advances in Parallel and Distributed Computational Models APDCM 2017*. IEEE Computer Society Press, 2017. DOI: 10.1109/IPDPSW.2017.60.

White Papers

- [R1] R. Ferreira da Silva et al. *Workflows Community Summit: Advancing the State-of-the-art of Scientific Workflows Management Systems Research and Development*. Tech. rep. May 2021. DOI: 10.5281/zenodo.4915801.
- [R2] R. Ferreira da Silva et al. *Workflows Community Summit: Bringing the Scientific Workflows Community Together*. Tech. rep. Mar. 2021. DOI: 10.5281/zenodo.4606958.

Teaching

- 2017 – 2018 Master – Parallel algorithms (22h, University of Lyon)
 Master – Distributed algorithms (10h, University of Lyon)
 Bachelor – Programming 1 (32h, ENS de Lyon)
- 2016 – 2017 Bachelor – ASR2: Advanced Computer Architecture and Network (32h, ENS de Lyon)
 Bachelor – Programming 1 (32h, ENS de Lyon)
- 2015 – 2016 Master – Image Processing and Computational Geometry (20h, ENS de Lyon)
 Bachelor – ASR1 : Computer Architecture and Network (6h, ENS de Lyon)

Bachelor – ALGO2 : Advanced Algorithms (32h, ENS de Lyon)

Languages skills

French **Native.**

English **Fluent.**

Computer Science skills

Programming C, C++, Python, R, \LaTeX .

Parallelism OpenMP, MPI, Parallel architectures.

Operating Systems Unix, System Programming.

Theory Scheduling, Performance Models.

Collectives responsibilities

Administrative

2017-2018 Elected representative for non-tenured members at the LIP (ENS Lyon computer science laboratory) council, co-organized a two-days seminar for PhD students.

Program Committee

2022 SuperComputing Technical Program and Undergraduate Posters, PPAM

2021 SuperComputing Technical Program, EuroPar, eScience

2020 SuperComputing Workshops, EuroPar

2019 PPAM, ICPP, eScience, EuroPar

Paper Refereeing

2021 SuperComputing, PPAM, ICPP, IJHPCA, EuroPar, FGCS, JPDC, TCC, TSC

2020 CCGrid, SuperComputing, EuroPar, CCPE, IJHPCA, JPDC, TPDS

2019 PPAM, ICPP, IJHPCA, Computing Journal, eScience, EuroPar

Panels Participation

National Science Foundation

US Department of Energy

References

Anne Benoit

Laboratoire d'Informatique du Parallélisme

ENS Lyon

46 allée d'Italie

69364 Lyon Cedex 07, France

✉ anne.benoit@ens-lyon.fr

Yves Robert

Laboratoire d'Informatique du Parallélisme

ENS Lyon

46 allée d'Italie

69364 Lyon Cedex 07, France

✉ yves.robert@inria.fr

Rafael Ferreira da Silva
Oak Ridge Leadership Computing
Facility
Oak Ridge National Laboratory
P.O. Box 2008
Oak Ridge, TN 37831
✉ silvarf@ornl.gov

Swann Perarnau
Mathematics and Computer Science
Argonne National Laboratory
9700 S. Cass Avenue
Argonne, IL 60439
✉ swann@anl.gov