Giorgis Georgakoudis

Curriculum Vitae – March 2026

Pleasanton, CA, USA
+1-925-495-9544 • georgakoudis1@llnl.gov • ggeorgakoudis
ggeorgakoudis

Research Profile

Research scientist and technical lead working at the intersection of high-performance computing, compilers, runtime systems, and AI-assisted software optimization. My recent work focuses on OpenMP, GPU performance portability, autotuning, approximate computing, and LLM-driven tooling for scientific software development. I lead and co-lead multi-institution efforts spanning programming models, compiler/runtime co-design, and performance engineering for production HPC applications.

Appointments

Principal Computer Scientist (05/2025 – present)
Computer Scientist (02/2020 – 05/2025)
Postdoctoral Researcher (11/2018 – 02/2020)
Lawrence Livermore National Laboratory (USA) *11/2018 – present*

- Principal Investigator of LDRD 25-ERD-019, a $2.7M project on just-in-time optimization for HPC applications
- Technical Lead for SciDAC efforts in boundary plasma dynamics and OpenMP in Python, coordinating compiler, runtime, and application-facing research
- Co-Investigator on Ellora, an ASCR project applying LLMs to software development workflows for scientific computing
- Research spans OpenMP, GPU/offload performance, autotuning, approximate computing, and compiler-guided optimization

Research Fellow
Queen's University Belfast (UK) *09/2016 – 11/2018*
Worked on throughput optimizations in OpenMP runtimes, resilience-aware adaptation, and transprecision computing through the EPSRC SERT project and the EU H2020 OPRECOMP project.

Research Intern
Lawrence Livermore National Laboratory (USA) *11/2016 – 03/2017*
Developed compiler-based fault injection techniques in LLVM for the REFINE project.
Supervisors: Ignacio Laguna, Martin Schulz

Research Assistant
Queen's University Belfast (UK) *01/2013 – 08/2016*
Developed runtime and systems software for streaming analytics, energy-aware execution, and heterogeneous architectures through the NanoStreams and GEMSCLAIM projects.

HPC Engineer
Simtec Software and Services *05/2011 – 01/2013*
Procured and managed a $300K Linux cluster for CFD simulations

Research Assistant
Center for Research & Technology Hellas (CERTH) (GR) *11/2007 – 04/2011*
Developed middleware, simulation infrastructure, and embedded systems software for wireless sensor networks and wearable sensing platforms.

Education

University of Thessaly (GR) **GPA 10/10 (Excellent)**
PhD in Comp. Eng. *09/2010 – 03/2017*
Dissertation: *"Scheduling and Performance Characterization on Heterogeneous Computing Systems"*.
Supervisors: Prof. D. S. Nikolopoulos, Prof. S. Lalis, Prof. C. Antonopoulos

University of Thessaly (GR) **GPA 8.75/10 (Excellent)**
Master's Degree in Comp. Eng. *02/2008 – 06/2009*

University of Thessaly (GR) **GPA 7.26/10 (Very Good)**
Diploma in Comp. Eng. *02/2002 – 09/2007*

Funding

2024: LDRD Exploratory Research *Maximize Performance of HPC Applications Through Just-In-Time Optimization (PI)*, $2.7M/3Y

2020: LDRD Exploratory Research *Optimizing Parallelism Compilation*, $1.6M/3Y

2019: LDRD Feasibility Study *Optimizing Performance Through Parallelism-Aware Compilation*, $150K/1Y

Honors and Awards

2025: **Best paper** award, IWOMP 2025

2024: **LLNL Director's Excellence in Publication** award

2024: **LLNL Achievement Award** for substantial advancements and impressive efficiency in achieving technical goals.

2023: **Best paper finalist**, SC'23

2021: **Best reproducibility** award, SC'21

2021: **Best paper** award, IWOMP 2021

2020: **LLNL Achievement Award** for exemplary leadership and teamwork on the Apollo project

2020: **Best paper** award, IWOMP 2020

2020: **Best paper** award in track Test and Dependability, DATE 2020

2017: **Discretionary Award for Exceptional Performance**, Queen's University Belfast (UK)

2013: **Best paper** award, COSMIC'13 (CGO)

Invited Talks

HiPEAC'18: REFINE: A Compiler Based Tool to Simulate Faults at Scale, Post-Exascale Workshop, Manchester, UK

HiPEAC'14: Informed Dynamic Scheduling On Shared-ISA Heterogeneous MPSoCs, Computing Systems Week, Athens, Greece

Professional Service

Professional Societies: **IEEE** (Senior Member), **ACM** (Member),
Technical Chamber of Greece (Member)

Standards

OpenMP Python Language Subcommittee (Chair)

Committees and Boards

Journals: **IEEE IT Professional** (Associate Editor) (2023–2026); **IEEE Transactions on Parallel and Distributed Systems (TPDS)** (Review Board) (2023–present); **Frontiers in High Performance Computing** (Review Board) (2023–present)

Conferences: **SC24** (Program Committee) (2024), **CCGRID** (Program Committee) (2024), **HPC Asia** (Program Committee) (2022), **CF** (Program Committee) (2020), **ARCS** (Program Committee) (2020), **Euro-Par** (Program Committee) (2022, 2018), **ICS** (Program Committee) (2023), **IISWC** (Program Committee) (2024, 2022), **eScience** (Program Committee) (2024)

Workshops: **LLPP (ICPP)** (Co-chair) (2023, 2022, 2021), **HIPS (IPDPS)** (Program Committee) (2024, 2023), **AI4DEV (SC23)** (Co-chair) (2023)

Grants: **NSF Panel Reviewer** (Computer and Information Science and Engineering (CISE), March 2023)

Peer Reviewer

Journals: **IEEE TCC** (2023), **IEEE TPDS** (2024, 2022, 2020, 2014), **ACM TACO** (2024), **JPDC** (2015)

Conferences: **IPDPS** (2025, 2020, 2017, 2014, 2012), **ICPP** (2016), **ICS** (2017), **EuroPar** (2017 – 2014), **CF** (2017, 2016), **CLUSTER** (2017, 2015), **CCGRID** (2024, 2019, 2017), **ICPADS** (2016),

Workshops: **COSMIC (CGO)** (2015), **AsHES (IPDPS)** (2014),

Other

Google Summer of Code (GSoC) (Mentor) (2020), **CLUSTER** (Web Chair) (2022, 2018)

Selected Publications

Google Scholar: 807 citations, h-index: 17, i10-index: 28

Conferences

ICS'25 J. H. Davis, P. Sivaraman, J. Kitson, K. Parasyris, H. Menon, I. Minn, G. **Georgakoudis**, and A. Bhatele, "Taking GPU Programming Models to Task for Performance Portability," in *Proceedings of the 39th ACM International Conference on Supercomputing*, 2025, pp. 776–791. DOI: 10.1145/3721145.3730423

SSDBM'25 A. Lisan, T. Patki, S. Brink, K. Parasyris, B. Gunnarson, G. **Georgakoudis**, and H. Childs, "Enabling Lightweight Performance Analysis of Complex Scientific Workflows with PerfFlowAspect," in *Proceedings of the 37th International Conference on Scientific and Statistical Database Management*, 2025, 5:1–5:12. DOI: 10.1145/3733723.3733734

CGO'25 G. **Georgakoudis**, K. Parasyris, and D. Beckingsale, "Proteus: Portable Runtime Optimization of GPU Kernel Execution with Just-In-Time Compilation," in *Proceedings of the 23rd ACM/IEEE International Symposium on Code Generation and Optimization*, 2025, pp. 507–522. DOI: 10.1145/3696443.3708939

ISC'25 A. Tran, I. Laguna, K. Parasyris, G. **Georgakoudis**, and G. Gopalakrishnan, "GORC: A Graph Neural Network Based Static Data Race Checker for OpenMP," in *ISC High Performance 2025*, 2025, pp. 1–15.

SC'24 Z. Fink, K. Parasyris, P. Rathi, **Georgakoudis, Giorgis**, H. Menon, and P.-T. Bremer, "HPAC-ML: A Programming Model for Embedding ML Surrogates in Scientific Applications," in *SC24: International Conference for High Performance Computing, Networking, Storage and Analysis*, acceptance rate 21% (99/470), 2024, pp. 1–16. DOI: 10.1109/SC41406.2024.00078

CGO'24 A. Murtovi, G. **Georgakoudis**, K. Parasyris, C. Liao, I. Laguna, and B. Steffen, "Enhancing Performance Through Control-Flow Unmerging and Loop Unrolling on GPUs," in *2024 IEEE/ACM International Symposium on Code Generation and Optimization (CGO)*, acceptance rate 32% (37/114), 2024, pp. 106–118. DOI: 10.1109/CGO57630.2024.10444819

SC23 Z. Fink, K. Parasyris, G. **Georgakoudis**, and H. Menon, "HPAC-Offload: Accelerating HPC Applications with Portable Approximate Computing on the GPU," in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, ser. SC '23, acceptance rate 24% (90/376), Denver, CO, USA: Association for Computing Machinery, 2023. DOI: 10.1145/3581784.3607095

SC23 K. Parasyris, G. **Georgakoudis**, E. Rangel, I. Laguna, and J. Doerfert, "Scalable Tuning of (OpenMP) GPU Applications via Kernel Record and Replay," in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, ser. SC '23, **Best paper finalist**, acceptance rate 24% (90/376), Denver, CO, USA: Association for Computing Machinery, Nov. 2023. DOI: 10.1145/3581784.3607098

PACT'22 J. Doerfert, M. Jasper, J. Huber, K. Abdelaal, G. **Georgakoudis**, T. Scogland, and K. Parasyris, "Breaking the Vendor Lock: Performance Portable Programming through OpenMP as Target Independent Runtime Layer," in *Proceedings of the International Conference on Parallel Architectures and Compilation Techniques*, ser. PACT '22, (acceptance 34% 40/118), Chicago, Illinois: Association for Computing Machinery, Mar. 2023, 494–504. DOI: 10.1145/3559009.3569687

IPDPS'22 J. Doerfert, A. Patel, J. Huber, S. Tian, J. M. M. Diaz, B. Chapman, and G. **Georgakoudis**, "Co-Designing an OpenMP GPU Runtime and Optimizations for Near-Zero Overhead Execution," in *2022 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, (acceptance 26% 123/474), 2022, pp. 504–514. DOI: 10.1109/IPDPS53621.2022.00055

CGO'22 J. Huber, M. Cornelius, **Georgakoudis, Giorgis**, S. Tian, J. M. M. Diaz, K. Dinel, B. Chapman, and J. Doerfert, "Efficient Execution of OpenMP on GPUs," in *2022 IEEE/ACM International Symposium on Code Generation and Optimization (CGO)*, (acceptance 27% 27/99), 2022, pp. 41–52. DOI: 10.1109/CGO53902.2022.9741290

CCGrid'21 K. Parasyris, G. **Georgakoudis**, L. Bautista-Gomez, and I. Laguna, "Co-Designing Multi-Level Checkpoint Restart for MPI Applications," in *2021 IEEE/ACM 21st International Symposium on Cluster, Cloud and Internet Computing (CCGrid)*, (acceptance 26% 60/230), 2021, pp. 103–112. DOI: 10.1109/CCGrid51090.2021.00020

SC21 K. Parasyris, G. **Georgakoudis**, H. Menon, J. Diffenderfer, I. Laguna, D. Osei-Kuffuor, and M. Schordan, "HPAC: Evaluating Approximate Computing Techniques on HPC OpenMP Applications," in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, ser. SC '21, **Best reproducibility award** (acceptance 25.8% 98/379), St. Louis, Missouri: Association for Computing Machinery, 2021. DOI: 10.1145/3458817.3476216

DSN'21 A. Taherin, T. Patel, G. **Georgakoudis**, I. Laguna, and D. Tiwari, "Examining Failures and Repairs on Supercomputers with Multi-GPU Compute Nodes," in *2021 51st Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, (acceptance 16.3% 48/295), 2021, pp. 305–313. DOI: 10.1109/DSN48987.2021.00043

ISC'21 C. Wood, G. **Georgakoudis**, D. Beckingsale, D. Poliakoff, A. Gimenez, K. Huck, A. Malony, and T. Gamblin, "Artemis: Automatic Runtime Tuning Of Parallel Execution Parameters Using Machine Learning," in *High Performance Computing: 36th International Conference, ISC High Performance 2021, Virtual Event, June 24 – July 2, 2021, Proceedings*, (acceptance 32.4% 24/74), Berlin, Heidelberg: Springer-Verlag, 2021, 453–472. DOI: 10.1007/978-3-030-78713-4_24

IISWC'20 L. Guo, G. **Georgakoudis**, K. Parasyris, I. Laguna, and D. Li, "MATCH: An MPI Fault Tolerance Benchmark Suite," in *2020 IEEE International Symposium on Workload Characterization (IISWC)*, (acceptance 37% 26/70), 2020, pp. 60–71. DOI: 10.1109/IISWC50251.2020.00015

IISWC'20 K. Parasyris, I. Laguna, H. Menon, M. Schordan, D. Osei-Kuffuor, G. **Georgakoudis**, M. O. Lam, and T. Vanderbruggen, "HPC-MixPBench: An HPC Benchmark Suite for Mixed-Precision Analysis," in *2020 IEEE International Symposium on Workload Characterization (IISWC)*, (acceptance 37% 26/70), 2020, pp. 25–36. DOI: 10.1109/IISWC50251.2020.00012

SC20 B. Swain, Y. Li, P. Liu, I. Laguna, G. **Georgakoudis**, and J. Huang, "OMPRacer: A Scalable and Precise Static Race Detector for OpenMP Programs," in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, ser. SC '20, (acceptance 25% 95/378), Atlanta, Georgia: IEEE Press, 2020.

ISC'20 G. **Georgakoudis**, L. Guo, and I. Laguna, "Reinit^{++}: Evaluating the Performance of Global-Restart Recovery Methods for MPI Fault Tolerance," in *High Performance Computing*, P. Sadayappan, B. L. Chamberlain, G. Juckeland, and H. Ltaief, Eds., (acceptance 31% 27/87), Cham: Springer International Publishing, 2020, pp. 536–554.

DATE'20 I. Tsiokanos, L. Mukhanov, G. **Georgakoudis**, D. S. Nikolopoulos, and G. Karakonstantis, "DEFCON: Generating and Detecting Failure-prone Instruction Sequences via Stochastic Search," in *2020 Design, Automation Test in Europe Conference Exhibition (DATE)*, **Best paper award**, (acceptance 26% 194/748), 2020, pp. 1121–1126.

HiPC'19 G. **Georgakoudis**, N. Jain, T. Ono, K. Inoue, S. Miwa, and A. Bhatele, "Evaluating the Impact of Energy Efficient Networks on HPC Workloads," in *2019 IEEE 26th International Conference on High Performance Computing, Data, and Analytics (HiPC)*, (acceptance 23% 39/171), 2019, pp. 301–310.

IPDPS'19 G. **Georgakoudis**, I. Laguna, H. Vandierendonck, D. S. Nikolopoulos, and M. Schulz, "SAFIRE: Scalable and Accurate Fault Injection for Parallel Multithreaded Applications," in *2019 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, (acceptance 28% 103/372), May 2019, pp. 890–899. DOI: 10.1109/IPDPS.2019.00097

SC17 G. **Georgakoudis**, I. Laguna, D. S. Nikolopoulos, and M. Schulz, "REFINE: Realistic Fault Injection via Compiler-based Instrumentation for Accuracy, Portability and Speed," in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, ser. SC17, (acceptance 19% 61/327), Denver, Colorado: ACM, 2017, 29:1–29:14. DOI: 10.1145/3126908.3126972

ARCS'16 M. Marcu, O. Boncalo, M. Ghenea, A. Amaricai, J. Weinstock, R. Leupers, Z. Wang, G. **Georgakoudis**, D. S. Nikolopoulos, C. Cernazanu-Glavan, L. Bara, and M. Ionascu, "Low-Cost Hardware Infrastructure for Runtime Thread Level Energy Accounting," in *Architecture of Computing Systems – ARCS 2016: 29th International Conference, Nuremberg, Germany, April 4–7, 2016, Proceedings*, (acceptance 33% 29/87), Cham: Springer International Publishing, 2016, pp. 277–289. DOI: 10.1007/978-3-319-30695-7_21

SAMOS'16 G. **Georgakoudis**, C. Gillan, A. Hassan, U. I. Minhas, I. Spence, G. Tzenakis, H. Vandierendonck, R. Woods, D. S. Nikolopoulos, M. Shyamsundar, P. Barber, M. Russell, A. Bilas, S. Kaloutsakis, H. Giefers, P. Staar, C. Bekas, N. Horlock, R. Faloon, and C. Pattison, "NanoStreams: Codesigned microservers for edge analytics in real time," in *2016 International Conference on Embedded Computer Systems: Architectures, Modeling and Simulation (SAMOS)*, (acceptance 48% 25/52), Jul. 2016, pp. 180–187. DOI: 10.1109/SAMOS.2016.7818346

SAMOS'14 G. **Georgakoudis**, D. S. Nikolopoulos, H. Vandierendonck, and S. Lalis, "Fast Dynamic Binary Rewriting for flexible thread migration on shared-ISA heterogeneous MPSoCs," in *Embedded Computer Systems: Architectures, Modeling, and Simulation (SAMOS XIV), 2014 International Conference on*, (acceptance: 40% 37/92), Jul. 2014, pp. 156–163. DOI: 10.1109/SAMOS.2014.6893207

Journals

IJHPCA'24 D. E. Bernholdt, G. Bosilca, A. Bouteiller, R. Brightwell, J. Ciesko, M. G. Dosanjh, G. **Georgakoudis**, I. Laguna, S. Levy, T. Naughton, S. L. Olivier, H. P. Pritchard, W. Schonbein, J. Schuchart, and A. Shehata, "Taking the MPI standard and the open MPI library to exascale," *The International Journal of High Performance Computing Applications*, vol. 0, no. 0, p. 10 943 420 241 265 936, 2024, (JCR IF 3.5). DOI: 10.1177/10943420241265936

PARCO'24 D. Miao, I. Laguna, G. **Georgakoudis**, K. Parasyris, and C. Rubio-González, "An automated OpenMP mutation testing framework for performance optimization," *Parallel Computing*, vol. 121, p. 103 097, 2024. DOI: https://doi.org/10.1016/j.parco.2024.103097

IEEE ITPro'23 H. Menon, J. Diffenderfer, G. **Georgakoudis**, I. Laguna, M. O. Lam, D. Osei-Kuffuor, K. Parasyris, and J. Vanover, "Approximate High-Performance Computing: A Fast and Energy-Efficient Computing Paradigm in the Post-Moore Era," *IT Professional*, vol. 25, no. 2, pp. 7–15, 2023. DOI: 10.1109/MITP.2023.3254642

IEEE TC'22 I. Tsiokanos, S. Tompazi, G. **Georgakoudis**, L. Mukhanov, and G. Karakonstantis, "ARETE: Accurate Error Assessment via Machine Learning-Guided Dynamic-Timing Analysis," *IEEE Transactions on Computers*, pp. 1–14, 2022. DOI: 10.1109/TC.2022.3191966

CSE'21 T. G. Mattson, T. A. Anderson, and G. **Georgakoudis**, "PyOMP: Multithreaded Parallel Programming in Python," *Computing in Science Engineering*, vol. 23, no. 6, pp. 77–80, 2021, (JCR IF 2.080). DOI: 10.1109/MCSE.2021.3128806

IJHPCA'18 C. Chalios, G. **Georgakoudis**, K. Tovletoglou, G. Karakonstantis, H. Vandierendonck, and D. S. Nikolopoulos, "DARE: Data-Access Aware Refresh via spatial-temporal application resilience on commodity servers," *The International Journal of High Performance Computing Applications*, vol. 32, no. 1, pp. 74–88, 2018, (JCR IF 1.956). DOI: 10.1177/1094342017718612

IEEE TMCS'17 U. I. Minhas, M. Russell, S. Kaloutsakis, P. Barber, R. Woods, G. **Georgakoudis**, C. Gillan, D. S. Nikolopoulos, and A. Bilas, "NanoStreams: A Microserver Architecture for Real-time Analytics on Fast Data Streams," *IEEE Transactions on Multi-Scale Computing Systems*, vol. PP, no. 99, pp. 1–1, 2017. DOI: 10.1109/TMSCS.2017.2764087

TACO'17 G. **Georgakoudis**, H. Vandierendonck, P. Thoman, B. R. D. Supinski, T. Fahringer, and D. S. Nikolopoulos, "SCALO: Scalability-Aware Parallelism Orchestration for Multi-Threaded Workloads," *ACM Trans. Archit. Code Optim.*, vol. 14, no. 4, 54:1–54:25, Dec. 2017, (JCR IF 1.166). DOI: 10.1145/3158643

CCPE'16 G. **Georgakoudis**, C. J. Gillan, A. Sayed, I. Spence, R. Faloon, and D. S. Nikolopoulos, "Methods and metrics for fair server assessment under real-time financial workloads," *Concurrency and Computation: Practice and Experience*, vol. 28, no. 3, pp. 916–928, 2016, (JCR IF 1.167). DOI: 10.1002/cpe.3704

PPL'15 G. **Georgakoudis**, C. Gillan, A. Sayed, I. Spence, R. Faloon, and D. S. Nikolopoulos, "Iso-Quality of Service: Fairly Ranking Servers for Real-Time Data Analytics," *Parallel Processing Letters*, vol. 25, no. 03, p. 1 541 004, 2015. DOI: 10.1142/S0129626415410042

Workshops and Posters

ISCW'25 K. Teranishi, H. Menon, W. F. Godoy, P. Balaprakash, D. Bau, T. Ben-Nun, A. Bhatele, F. Franchetti, M. E. Franusich, T. Gamblin, G. **Georgakoudis**, T. Goldstein, A. Guha, S. E. Hahn, C. Iancu, Z. Jin, T. R. Jones, T. M. Low, H. Mankad, N. R. Miniskar, M. A. H. Monil, D. Nichols, K. Parasyris, S. Pophale, P. Valero-Lara, J. S. Vetter, S. Williams, and A. R. Young, "Leveraging AI for Productive and Trustworthy HPC Software: Challenges and Research Directions," in *High Performance Computing - ISC High Performance 2025 International Workshops, Revised Selected Papers*, ser. Lecture Notes in Computer Science, vol. 16091, Springer Science and Business Media Deutschland GmbH, 2026, pp. 615–625. DOI: 10.1007/978-3-032-07612-0_47

IWOMP'25 G. **Georgakoudis**, T. A. Anderson, S. Archibald, B. R. de Supinski, and T. G. Mattson, "Programming GPUs with OpenMP and Python," in *OpenMP: Balancing Productivity and Performance Portability*, ser. Lecture Notes in Computer Science, **Best paper award**, vol. 16123, Springer, Cham, 2026, pp. 212–226. DOI: 10.1007/978-3-032-06343-4_14

AI4Sys'25 G. Bolet, G. **Georgakoudis**, H. Menon, K. Parasyris, N. Hasabnis, H. Estes, K. W. Cameron, and G. Oren, "Can Large Language Models Predict Parallel Code Performance?" In *Workshop on AI For Systems*, 2025. DOI: 10.1145/3731545.3743645

SCW'25 J. Bowen, K. Parasyris, D. Beckingsale, T. Ben-Nun, T. Stitt, and G. **Georgakoudis**, "Extending RAJA Parallel Programming Abstractions with Just-In-Time Optimization," in *SC25-W: Workshops of the International Conference for High Performance Computing, Networking, Storage and Analysis*, 2025, pp. 1159–1171.

IPDPSW'25 J. Wu, J. Ren, S. Yang, K. Parasyris, G. **Georgakoudis**, I. Laguna, and D. Li, "LM-Offload: Performance Model-Guided Generative Inference of Large Language Models with Parallelism Control," in *2025 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, 2025, pp. 840–849. DOI: 10.1109/IPDPSW66978.2025.00134

IPDPSW'24 G. Bolet, G. **Georgakoudis**, K. Parasyris, K. W. Cameron, D. Beckingsale, and T. Gamblin, "An Exploration of Global Optimization Strategies for Autotuning OpenMP-based Codes," in *2024 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, 2024, pp. 741–750. DOI: 10.1109/IPDPSW63119.2024.00138

PMAM'24 D. Miao, I. Laguna, G. **Georgakoudis**, K. Parasyris, and C. Rubio-González, "MUPPET: Optimizing Performance in OpenMP via Mutation Testing," in *Proceedings of the 15th International Workshop on Programming Models and Applications for Multicores and Manycores*, ser. PMAM '24, Edinburgh, United Kingdom: Association for Computing Machinery, 2024, 22–31. DOI: 10.1145/3649169.3649246

WACCPD'21 C. Liao, A. Wang, G. **Georgakoudis**, B. R. de Supinski, Y. Yan, D. Beckingsale, and T. Gamblin, "Extending OpenMP for Machine Learning-Driven Adaptation," in *Accelerator Programming Using Directives*, S. Bhalachandra, C. Daley, and V. Melesse Vergara, Eds., Cham: Springer International Publishing, 2022, pp. 49–69.

CANOPIE-HPC'22 D. J. Milroy, C. Misale, G. **Georgakoudis**, T. Elengikal, A. Sarkar, M. Drocco, T. Patki, J.-S. Yeom, C. E. A. Gutierrez, D. H. Ahn, and Y. Park, "One Step Closer to Converged Computing: Achieving Scalability with Cloud-Native HPC," in *2022 IEEE/ACM 4th International Workshop on Containers and New Orchestration Paradigms for Isolated Environments in HPC (CANOPIE-HPC)*, 2022, pp. 57–70. DOI: 10.1109/CANOPIE-HPC56864.2022.00011

P3HPC'22 K. Parasyris, G. **Georgakoudis**, J. Doerfert, I. Laguna, and T. R. Scogland, "Piper: Pipelining OpenMP Offloading Execution Through Compiler Optimization For Performance," in *2022 IEEE/ACM International Workshop on Performance, Portability and Productivity in HPC (P3HPC)*, 2022, pp. 100–110. DOI: 10.1109/P3HPC56579.2022.00015

IWOMP'22 G. **Georgakoudis**, T. R. W. Scogland, C. Liao, and B. R. de Supinski, "Extending OpenMP to Support Automated Function Specialization Across Translation Units," in *OpenMP in a Modern World: From Multi-device Support to Meta Programming*, M. Klemm, B. R. de Supinski, J. Klinkenberg, and B. Neth, Eds., Cham: Springer International Publishing, 2022, pp. 159–173.

IWOMP'21 J. Huber, W. Wei, G. **Georgakoudis**, J. Doerfert, and O. Hernandez, "A Case Study of LLVM-Based Analysis for Optimizing SIMD Code Generation," in *OpenMP: Enabling Massive Node-Level Parallelism*, S. McIntosh-Smith, B. R. de Supinski, and J. Klinkenberg, Eds., **Best paper award**, Cham: Springer International Publishing, 2021, pp. 142–155.

ICPPW'21 T. Jayatilaka, H. Ueno, G. **Georgakoudis**, E. Park, and J. Doerfert, "Towards Compile-Time-Reducing Compiler Optimization Selection via Machine Learning," in *50th International Conference on Parallel Processing Workshop*. New York, NY, USA: Association for Computing Machinery, 2021.

IWOMP 2020 G. **Georgakoudis**, J. Doerfert, I. Laguna, and T. R. W. Scogland, "FAROS: A Framework to Analyze OpenMP Compilation Through Benchmarking and Compiler Optimization Analysis," in *OpenMP: Portable Multi-Level Parallelism on Modern Systems*, K. Milfeld, B. R. de Supinski, L. Koesterke, and J. Klinkenberg, Eds., **Best paper award**, Cham: Springer International Publishing, 2020, pp. 3–17.

IWOMP 2015 F. Alessi, P. Thoman, G. **Georgakoudis**, T. Fahringer, and D. S. Nikolopoulos, "Application-Level Energy Awareness for OpenMP," in *OpenMP: Heterogenous Execution and Data Movements: 11th International Workshop on OpenMP, IWOMP 2015, Aachen, Germany, October 1-2, 2015, Proceedings*, Cham: Springer International Publishing, 2015, pp. 219–232. DOI: 10.1007/978-3-319-24595-9_16

WHPCF'14 (SC14) C. J. Gillan, D. S. Nikolopoulos, G. **Georgakoudis**, R. Faloon, G. Tzenakis, and I. Spence, "On the Viability of Microservers for Financial Analytics," in *High Performance Computational Finance (WHPCF), 2014 Seventh Workshop on*, Nov. 2014, pp. 29–36. DOI: 10.1109/WHPCF.2014.11

COSMIC'13 G. **Georgakoudis**, D. S. Nikolopoulos, and S. Lalis, "Fast Dynamic Binary Rewriting to Support Thread Migration in shared-ISA Asymmetric Multicores," in *Proceedings of the First International Workshop on Code OptimiSation for MultI and Many Cores (CGO)*, ser. COSMIC'13, **Best paper award**, Shenzhen, China: ACM, 2013, pp. 1–10. DOI: 10.1145/2446920.2446924

HPDC'12 (Poster) G. **Georgakoudis**, S. Lalis, and D. S. Nikolopoulos, "Summary: Dynamic Binary Rewriting and Migration for shared-ISA Asymmetric, Multicore Processors," in *Proceedings of the 21st International Symposium on High-Performance Parallel and Distributed Computing*, ser. HPDC'12, Delft, The Netherlands: ACM, 2012, pp. 127–128. DOI: 10.1145/2287076.2287096

Preprints

CoRR'25 G. Bolet, G. **Georgakoudis**, K. Parasyris, H. Menon, N. Hasabnis, K. W. Cameron, and G. Oren, "Counting Without Running: Evaluating LLMs' Reasoning About Code Complexity," *CoRR*, vol. abs/2512.04355, 2025. DOI: 10.48550/arXiv.2512.04355

CoRR'25 D. Nichols, K. Parasyris, H. Menon, B. R. Bartoldson, G. **Georgakoudis**, T. Ben-Nun, and A. Bhatele, "Modeling Code: Is Text All You Need?" *CoRR*, vol. abs/2507.11467, 2025. DOI: 10.48550/arXiv.2507.11467