Dept. Informática, Universidade do Minho
Rua da Universidade, Gualtar (4710-057)
Braga, Portugal
☎ (+351) 253 604 477
✉ ricardo.g.macedo@inesctec.pt

Ricardo Macedo

Personal Information

Name Ricardo Gonçalves Macedo
Inst. Page https://www.inesctec.pt/en/people/ricardo-goncalves-macedo
Pers. Page https://rgmacedo.github.io/

ORCID 0000-0003-4036-0126
Google Scholar 3QPFbOcAAAAJ
GitHub rgmacedo

Education

2017 – 2023 **MAP-i Doctoral Program in Computer Science.**
Universities of Minho, Aveiro, and Porto, Portugal.
Advisors: João Tiago Paulo and José Orlando Pereira.
Thesis: *User-level Software-Defined Storage Data Planes*.
Classification: Very good (UMinho's highest grade).

2011 – 2017 **Integrated Master's in Informatics Engineering.**
University of Minho, Portugal.
Advisors: João Tiago Paulo and Rui Carlos Oliveira.
Thesis: *Secure Computation in NoSQL Databases*.
Classification (thesis): 20 out of 20.

Experience

2023 – 2024 **Researcher (contract)**, *HASLab INESC TEC*.
Research in storage systems for modern I/O infrastructures, resource disaggregation, and energy-aware computing.

2017 – 2023 **Researcher (PhD student)**, *HASLab INESC TEC*.
Research in Software-Defined Storage, local and distributed storage, key-value storage, kernel-bypass storage.

2020 (summer) **Research Intern**, *AIST*.
Tsukuba, Japan (Remote)
Advisors: Jason Haga, Yusuke Tanimura
Investigating storage data planes for ensuring bandwidth guarantees in the ABCI supercomputer.

2016 – 2017 **Researcher (MSc student)**, *HASLab INESC TEC*.
Research in secure computation over NoSQL and SQL databases.

Publications

Conferences

CCGrid 2023 **Ricardo Macedo**, Mariana Miranda, Yusuke Tanimura, Jason Haga, Amit Ruhela, Stephen Lien Harrell, Richard Todd Evans, José Pereira, João Paulo. Taming Metadata-intensive HPC Jobs Through Dynamic, Application-agnostic QoS Control. In *23rd IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing*. 2023.

CCGrid 2022 Marco Dantas, Diogo Leitão, Peter Cui, **Ricardo Macedo**, Xinlian Liu, Weijia Xu, João Paulo. Accelerating Deep Learning Training Through Transparent Storage Tiering. In *22nd IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing*. 2022.

FAST 2022 **Ricardo Macedo**, Yusuke Tanimura, Jason Haga, Vijay Chidambaram, José Pereira, João Paulo. PAIO: General, Portable I/O Optimizations With Minor Application Modifications. In *20th USENIX Conference on File and Storage Technologies*. 2022.

SYSTOR 2021 Alberto Faria, **Ricardo Macedo**, José Pereira, João Paulo. BDUS: Implementing Block Device Drivers in User Space. In *14th ACM International System and Storage Conference*, 2021. ***Best paper runner-up.***

SRDS 2017 **Ricardo Macedo**, João Paulo, Rogério Pontes, Bernardo Portela, Tiago Oliveira, Miguel Matos, Rui Oliveira. A Practical Framework for Privacy-Preserving NoSQL Databases. In *36th International Symposium on Reliable Distributed Systems*, 2017.

Journals

CSUR 2020 **Ricardo Macedo**, João Paulo, José Pereira, Alysson Bessani. A Survey and Classification of Software-Defined Storage Systems. *ACM Computing Surveys 53, 3 (48)*, 2020.

Workshops

DCDS 2023 Tânia Esteves, **Ricardo Macedo**, Rui Oliveira, João Paulo. Diagnosing applications' I/O behavior through system call observability. In *5th Workshop on Data-Centric Dependability and Security*, co-located with IEEE/IFIP International Conference on Dependable Systems and Networks. 2023.

REX-IO 2022 **Ricardo Macedo**, Mariana Miranda, Yusuke Tanimura, Jason Haga, Amit Ruhela, Stephen Lien Harrell, Richard Todd Evans, João Paulo. Protecting Metadata Servers From Harm Through Application-level I/O Control. In *2nd Workshop on Re-envisioning Extreme-Scale I/O for Emerging Hybrid HPC Workloads*, co-located with IEEE International Conference in Cluster Computing. 2022.

WoC 2021 Alberto Faria, **Ricardo Macedo**, João Paulo. Pods-as-Volumes: Effortlessly Integrating Storage Systems and Middleware into Kubernetes. In *7th International Workshop on Container Technologies and Container Clouds*,co-located with ACM/IFIP International Middleware Conference. 2021.

REX-IO 2021 **Ricardo Macedo**, Cláudia Correia, Marco Dantas, Cláudia Brito, Weijia Xu, Yusuke Tanimura, Jason Haga, João Paulo. The Case for Storage Optimization Decoupling in Deep Learning Frameworks. In *1st Workshop on Re-envisioning Extreme-Scale I/O for Emerging Hybrid HPC Workloads*, co-located with IEEE International Conference in Cluster Computing. 2021.

REX-IO 2021 Marco Dantas, Diogo Leitão, Cláudia Correia, **Ricardo Macedo**, Weijia Xu, João Paulo. Monarch: Hierarchical Storage Management for Deep Learning Frameworks. In *1st Workshop on Re-envisioning Extreme-Scale I/O for Emerging Hybrid HPC Workloads*, co-located with IEEE International Conference in Cluster Computing. 2021.

SRDSW 2019 **Ricardo Macedo**, Alberto Faria, João Paulo, José Pereira. A Case for Dynamically Programmable Storage Background Tasks. In *38th International Symposium on Reliable Distributed Systems Workshops*, 2019.

SRDSW 2019 Tânia Esteves, **Ricardo Macedo**, Alberto Faria, Bernardo Portela, João Paulo, José Pereira, Danny Harnik. TrustFS: An SGX-enabled Stackable File System Framework. In *38th International Symposium on Reliable Distributed Systems Workshops*, 2019.

Preprint

2023 Tânia Esteves, **Ricardo Macedo**, Rui Oliveira, João Paulo. Diagnosing applications' I/O behavior through system call observability. *arXiv preprint arXiv:2304.08569*, 2023.

Ricardo Macedo, Mariana Miranda, Yusuke Tanimura, Jason Haga, Amit Ruhela, Stephen Lien Harrell, Richard Todd Evans, José Pereira, João Paulo. PADLL: Taming Metadata-intensive HPC Jobs Through Dynamic, Application-agnostic QoS Control. *arXiv preprint arXiv:2302.06418*, 2023.

2021 **Ricardo Macedo**, Yusuke Tanimura, Jason Haga, Vijay Chidambaram, José Pereira, João Paulo. PAIO: A Software-Defined Storage Data Plane Framework. *arXiv preprint arXiv:2106.03617*, 2021.

Theses

PhD Thesis **Ricardo Macedo**. User-level Software-Defined Storage Data Planes. *PhD Thesis, Universidades do Minho, Aveiro, e Porto*. Supervised by João Paulo and José Pereira. 2023.

MSc Thesis **Ricardo Macedo**. Computação Segura em Bases de Dados NoSQL. *Master's Thesis, Universidade do Minho*. Supervised by João Paulo and Rui Oliveira. 2017.

Projects

2020 – cur. **BigHPC: A Management Framework for Consolidated Big Data and HPC**.
POCI-01-0247-FEDER-045924

Research and development of Software-Defined Storage systems to improve performance isolation, QoS, and fairness in converged HPC and Big Data infrastructures.

2022 **LazyFS: Lazy File System Project**.
PP2022-0019

Consulting services with Jepsen LLC. Research and development of fault-injection file system.

2020 – 2021 **PAStor: Programmable and Adaptable Storage for AI-oriented HPC Ecosystems**.
UTA-EXPL/CA/0075/2019

Research and development of a Software-Defined Storage data plane for improving I/O performance of AI-powered applications in HPC infrastructures.

2018 – cur. **CENTRA: Efficient and Secure Data Management for HPC and Cloud Computing**.

Research and development of Software-Defined Storage systems to improve the programmability and configurability of HPC and Cloud Computing infrastructures.

2018 – 2019 **IBM Research Haifa: Joint Study Agreement**.

Research and development of a programmable and modular stackable file system framework for implementing secure content-aware storage functionalities over hardware-assisted trusted execution environments.

2016 – 2018 **SafeCloud**.
H2020-DS-2014-1/653884

Design, development, and evaluation of a privacy-preserving framework for NoSQL databases. Integration with a SQL query engine in order to provide secure computation over real-time SQL-based applications.

Advising

MSc Students

2022 – cur. **Maria Beatriz Moreira.** *I/O Optimizations for Distributed Deep Learning Training*. Co-advised with João Paulo and Cláudia Brito. University of Minho.

Maria Ramos. *Realistic Fault Assessment in Distributed Storage Systems*. Co-advised with João Paulo and Tânia Esteves. University of Minho.

Pedro Peixoto. *Programmable Caches with System-wide Visibility*. Co-advised with João Paulo. University of Minho.

Rúben Adão. *Co-designing LSM-based Key-Value Stores with Hybrid Storage for Optimal Cost-Performance Ratio*. Co-advised with João Paulo. University of Minho.

2021 – 2023 **Alexandre Miranda.** *Realistic Assessment of Failures in the SPDK Platform*. Co-advised with João Paulo. 18 out of 20. University of Minho.

2021 – 2022 **João Azevedo.** *LazyFS: A file system for assessing applications data durability*. Co-advised with João Paulo. 20 out of 20. University of Minho.

2020 – 2022 **Marco Dantas.** *Accelerating Deep Learning Training on High-Performance Computing with Storage Tiering*. Co-advised with João Paulo and Rui Oliveira. 20 out of 20. University of Minho.

2019 – 2021 **Cláudia Mendonça.** *PRISMA: A Prefetching Storage Middleware for Accelerating Deep Learning Frameworks*. Co-advised with João Paulo and António Sousa. 18 out of 20. University of Minho.

Diogo Leitão. *RSafeFS: A Modular File System for Remote Storage.* Co-advised with João Paulo and José Orlando Pereira. 17 out of 20. University of Minho.

Undergraduate Research Projects

2023 **Maria Beatriz Moreira**, **Pedro Peixoto**, and **Rúben Adão**. *Programmable and Adaptable Data Transformations.* Co-advised with João Paulo. University of Minho.

2022 **Susana Marques** and **Tomás Costa.** *A Key-Value Storage Benchmarking Tool.* Co-advised with João Paulo. University of Minho.

2021 **Alberto Faria.** *An Evaluation of Linux Interfaces for Storage I/O*. Co-advised with João Paulo. University of Minho.

Research Mentor

2022 – cur. **Guilherme Fernandes**. *Metadata trace replayer for HPC parallel file systems.*

Diana Rodrigues, **João Fernandes**, **Mariana Amorim** and **Sara Pereira**. *Exploring energy-aware policies for modern I/O infrastructures.*

2021 – 2022 **Maria Beatriz Moreira** and **Pedro Peixoto.** *Building transparent data transformations mechanisms (compression, encryption).*

Rúben Adão. *Storage tiering strategies for LSM-based key-value stores.*

2021 **Diogo Ribeiro.** *Storage tiering strategies in FUSE-based file systems.*

Teaching

2020 – 2023 **University of Minho**, *Distributed Systems.*
Invited Assistant Professor (25%).
MIEI course (3rd year), University of Minho.

2019 – 2023 **University of Minho**, *Operating Systems.*
Invited Assistant Professor (25%).
MIEI course (2nd year), University of Minho.

Service

Program Committee Member

2024 ASPLOS

2023 EuroSys ShadowPC, ESSA Workshop (IPDPS), Rex-IO Workshop (CLUSTER)

2022 Rex-IO Workshop (CLUSTER)

External Reviewer

2022 DSN

2021 DSN, ICDCS

2019 OPODIS, SRDS, DSN

2018 Middleware

Other Service

2018 **Volunteer.**
ACM European Conference on Computer Systems *(Porto, Portugal).*

Awards

2022 **"The Incredibles" Award.** *INESC TEC, Portugal.* Award given by INESC TEC and HASLab coordinators to João Paulo and Ricardo Macedo, due to their contribution to the field of storage, together with their respective teams of MSc students.

2021 **ATC'21 Student Grant.** *2021 USENIX Annual Technical Conference*, USENIX Association.

Best paper runner-up. *"BDUS: Implementing Block Device Drivers in User Space"*, *14th International System and Storage Conference*, Association for Computing Machinery.

2020 **OSDI'20 Student Grant.** *14th USENIX Symposium on Operating Systems Design and Implementation*, USENIX Association.

Merit Grant. *Direção-Geral do Ensino Superior (DGES), Portugal.*

2019 **PhD Grant.** *SFRH/BD/146059/2019, Fundação para a Ciência e a Tecnologia, Portugal.*

2018 **Excellence Scholarship.** *Universidade do Minho, Portugal.*

Talks

May 2023 *PADLL: Taming Metadata Burstiness of HPC Jobs Through Application-level QoS Control (Poster)* at *ACM EuroSys* (Rome, Italy)

May 2023 *Taming Metadata-intensive HPC Jobs Through Dynamic, Application-agnostic QoS Control* at *ACM/IEEE CCGrid* (Bangalore, India)

Jun 2022 *User-level Software-Defined Storage Data Planes* at *ENSD* (Évora, Portugal)

May 2022 *Building User-level Data Planes with PAIO* at *ADAC Seminar* (Virtual)

May 2022 *User-level Software-Defined Storage Data Planes* at *On the Road to HPC: Major Challenges and New Opportunities* (Webinar, Virtual)

Feb 2022 *PAIO: General, Portable I/O Optimizations With Minor Application Modifications* at *USENIX FAST* (Santa Clara, CA, USA)

Sep 2021 *The Case for Storage Optimization Decoupling in Deep Learning Frameworks* at *REX-IO* (Virtual)

Oct 2019 *A Case for Dynamically Programmable Storage Background Tasks* at *SRDSW* (Lyon, France)

Sep 2017 *A Practical Framework for Privacy-Preserving NoSQL Databases* at *SRDS* (Hong Kong)

References

Dr. João Tiago Paulo
INESC TEC & University of Minho
joao.t.paulo@inesctec.pt

Prof. José Orlando Pereira
University of Minho
jop@di.uminho.pt

Dr. Jason Haga
National Institute of Advanced Industrial Science and Technology (AIST)
jh.haga@aist.go.jp

Prof. Alysson Bessani
LASIGE & Faculdade de Ciências da Universidade de Lisboa
anbessani@fc.ul.pt

Prof. Vijay Chidambaram
University of Texas at Austin
vijayc@utexas.edu

Prof. Rui Oliveira
University of Minho
rco@di.uminho.pt