

David Daharewa Gureya

SYSTEMS · RESEARCHER

INESC-ID, Universidade de Lisboa & KTH Royal Institute of Technology

☎ (+46) 72-3641828 | ✉ gureya.daharewa@gmail.com | 🏠 gureya.github.io | 📷 gureya | 🌐 gureya

Summary

Currently Systems Researcher, interested in building system software for multi-socket systems and resource management in the cloud. 10+ years of experience working in diverse environments with a breadth of programs and technologies. Experienced in object-oriented (C++ & Java) and low-level C programming. Quickly learn and master new technologies if the need arises.

Professional Experience

INESC-ID, Instituto Superior Técnico, Portugal & KTH Royal Institute of Technology, Sweden

Stockholm, Sweden

RESEARCH ASSISTANT

October 2016 - present

- Dissertation: Resource Allocation for Data-Intensive Services in the Cloud.
- OnlineElastMan: an "out-of-the-box" elasticity controller, which can be deployed and adopted by different storage systems without complicated tailoring/configuring efforts. The controller was implemented in Java & Matlab and was evaluated on a private cloud running OpenStack software stack and is openly available.
- BWAP: a novel bandwidth-aware page placement in NUMA systems. Implemented in C/C++ and open-sourced it.
- BALM: a QoS-aware memory bandwidth allocation technique for multi-socket architectures. Implementation done in C/C++ and will be readily available.

LOG Associates

Nairobi, Kenya

CONSULTANT

March 2016 - September 2016

- Lead consultant in the development of a Monitoring and Evaluation (M&E) System for the Kenya Cereal Enhancement Programme (KCEP).
- Designed and implemented a Management Information System (MIS) for storing/analysing data and information relating to the programme, using Java Spring Framework & MySQL.

University of Nairobi

Nairobi, Kenya

TUTORIAL FELLOW

August 2013 - September 2016

- Participation in the fields of teaching, research and development.
- Lecturing and project supervision for undergraduate and graduate students. Taught programming methodology and operating systems courses.

Swedish ICT

Stockholm, Sweden

RESEARCH ENGINEER

February 2015 - September 2015

- Conducted research on elasticity management for cloud storage services. Main research topics included: Cloud computing, Online machine learning, and Time series analysis.
- Gain skills in Java, Linux, Matlab, OpenStack, and Cassandra.

C4DLab, University of Nairobi

Nairobi, Kenya

SOFTWARE ENGINEER

April 2013 - September 2013

- Research, capacity building and training, systems design and development.
- Organized and participated in technological boot-camps, workshops, meetups, and seminars within Nairobi.
- Team Lead in developing system prototypes, notably among them were Games for Learning (G4L).

Cellulant Corporation

Nairobi, Kenya

SOFTWARE ENGINEER

August 2012 - March 2013

- Requirements elicitation and planning, Solution analysis and design, Products development and Maintenance.
- Developed web applications using Yii PHP Framework, MySQL (LAMP), HTML/CSS/JavaScript.
- Worked on USSD applications and developed Java daemons for a payments ecosystem.

Education

Instituto Superior Técnico, Portugal & KTH Royal Institute of Technology, Sweden

Portugal/Sweden

DOCTOR OF SCIENCE (PHD) IN COMPUTER SCIENCE

October 2016 - present

- Advisor(s): Prof. João Barreto & Prof. Vladimir Vlassov.
- Thesis Title: Resource Allocation for Data-Intensive Services in the Cloud

Instituto Superior Técnico, Portugal & KTH Royal Institute of Technology, Sweden

ERASMUS MUNDUS MASTER (MSc) IN DISTRIBUTED COMPUTING

- Advisor: Prof. Vladimir Vlassov.
- Grade: 17/20.
- Thesis Title: Self-trained Proactive Elasticity Manager for Cloud-based Storage Services [pdf]

Portugal/Sweden

September 2013 - July 2015

University of Nairobi

BACHELOR OF SCIENCE (BSc) IN COMPUTER SCIENCE

- Grade: First Class Honours
- Thesis Title: Automatic language identification using N-Gram-Based Text Categorization.

Kenya

October 2008 - August 2012

Skills

System research	Extensive experience in performing research on system development.
Back-end	Possess expertise in object-oriented (C++ & Java) and low-level C programming.
Distributed systems	A firm foundation of distributed computing concepts, including cloud computing.
Full-stack	Well-versed with full-stack knowledge including HTML/CSS/JavaScript, PHP, MySQL/MariaDB, NoSQL(Cassandra, Memcached); familiarity with scripting languages (Python, Shell).

Honors & Awards

INTERNATIONAL

2013	Recipient , Erasmus Mundus scholarship (academic-merit based, awarded each semester of MSc)	Stockholm, Sweden
2016	Recipient , Erasmus Mundus fellowship for Joint Doctorate in Distributed Computing (EMJD-DC)	Brussels, Belgium
2016	Finalist , Best Poster Award at 2016 IEEE International Conference on Cloud and Autonomic Computing (ICCAC)	Augsburg, Germany

Peer Reviewed Publications

2021	Generalizing QoS-Aware Memory Bandwidth Allocation to Multi-Socket Cloud Servers , David Gureya , João Barreto and Vladimir Vlassov	IEEE Cloud'21
2021	BALM: QoS-Aware Memory Bandwidth Partitioning for Multi-Socket Cloud Nodes , David Gureya , Vladimir Vlassov and João Barreto	SPAA'21
2020	BWAP: Bandwidth-Aware Page Placement in NUMA , David Gureya , Joao Neto, Reza Karimi, Rodrigo Rodrigues, Paolo Romano, Pramod Bhatotia, Vivien Quéma, João Barreto and Vladimir Vlassov	IPDPS'20
2018	Asymmetry-aware page placement for contemporary NUMA architectures , David Gureya , Joao Neto, Reza Karimi, Rodrigo Rodrigues, Paolo Romano, Pramod Bhatotia, Vivien Quéma, and João Barreto	SFMA'18
2018	Tackling memory asymmetry in NUMA architectures , David Gureya , Joao Neto, Reza Karimi, Rodrigo Rodrigues, Paolo Romano, Pramod Bhatotia, Vivien Quéma, and João Barreto	INForum'18
2017	OnlineElastMan: Self-Trained Proactive Elasticity Manager for Cloud-Based Storage Services , Ying Liu, Daharewa Gureya , Ahmad Al-Shishtawy and Vladimir Vlassov	Journal of Cluster Computing
2017	Profiling for Asymmetric NUMA Systems , David Gureya and João Barreto	EuroDW'17
2016	OnlineElastMan: Self-Trained Proactive Elasticity Manager for Cloud-Based Storage Services , Ying Liu, Daharewa Gureya , Ahmad Al-Shishtawy and Vladimir Vlassov	ICCAC'16

Projects

2020	BWAP , A library for dynamic page placement in NUMA nodes.	github.com/epeec/BWAP
2017	MemoryTracer , An ISA-portable PIN tool for tracing memory accesses.	MemoryTracer
2017	OnlineElastMan , A Self-Trained Proactive Elasticity Manager for Cloud-Based Storage Services.	OnlineElastMan
2016	Forecasting , A set of Prediction algorithms for forecasting time series data.	forecasting