

Walid A. Hanafy

PhD Candidate

Manning College of Information and Computer Sciences
University of Massachusetts Amherst
☎ (+1) 413 345-9453
✉ whanafy@cs.umass.edu
🏠 WebPage 📺 washraf in whanafy

Summary

I am a 5th-year PhD student at UMass Amherst with extensive experience developing large-scale and sustainable distributed systems. My recent research has focused on reducing cloud computing systems' carbon footprint and designing resource-efficient machine learning systems for the edge.

My work has been published in top-tier conferences, e.g., *ASPLOS*, *INFOCOM*, *SIGMETRICS*, *E-Energy*, *IoTDI*, *SEC*, and *IWQOS*.

Education

2019–Now MS'23/PhD, Computer Science, University of Massachusetts Amherst, MA, USA.

Cloud and Edge Computing, Green Computing, Energy Informatics, and Machine Learning Systems.

Supervisor Prof. Prashant Shenoy, LASS research lab.

2015–2018 Masters of Computer Engineering, Faculty of Engineering, Helwan University, Egypt.

Cloud Computing, Elastic Infrastructures, Containers

Supervisors Prof. Sameh Salem, Prof. Amr Mohamed.

Thesis Title Container Type Virtualization Management in Cloud Computing.

2012 Diploma in Information Technology, Major in System Development, Minor in Software Architecture, Information Technology Institute (ITI), Egypt.

2006–2011 BS of Computer Engineering (Class Valedictorian), Helwan University, Egypt.

Thesis Assistive Drive System: Driver Behavior Monitoring using computer vision and ultrasonic sensors.

Research Experience

University of Massachusetts Amherst, Research Assistant

2021–Now Carbon-aware/Sustainable Computing.

Papers: ASPLOS'23, HotCarbon'23, SIGMETRICS/Performance 2024, ASPLOS'24, E-Energy'24

Designed and Implemented Carbon and Energy aware system monitoring and management solutions that help reduce the carbon footprint of computing.

2020–Now Model-Serving on The Edge, Papers: WEEE'21, TAAS'23, IoTDI'23, SEC'23, MilCom'23.

Designed and Implemented Model-Serving platforms that adapt to workload and environment dynamics while providing fairness, resiliency, and latency guarantees.

Nokia Bell Labs, Network Systems Research Department, Summer Intern

2022 Low-Latency Model Serving, Papers: IWQOS'23.

Designed a low latency model serving framework using RDMA and GPUDirect RDMA.

2021 Telemetry system, Patent: US20230246931A1, Papers: INFOCOM'24.

Designed an End-to-End telemetry system for RDMA network traffic

Public University of Navarra, Spain, Research Assistant (Funded by ERASMUS+)

2018 Solar Panels anomaly detection, Papers: ICENCO'19.

Cleanliness detection of Solar Panels using computer vision and machine learning.

Helwan University, Research Assistant

2015-2019 *Containers Clusters Management, Papers: ICENCO'17, ICCES'17, IEEE Access'19.*
Developed elasticity control/load balancing algorithms for containers as a service cloud.

Publications and Patents

- 2024 **Walid A. Hanafy**, Qianlin Liang, Noman Bashir, Abel Souza, David Irwin, and Prashant Shenoy. Going Green for Less Green: Optimizing the Cost of Reducing Cloud Carbon Emissions. In *29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 3 (ASPLOS '24)*, April 27-May 1, 2024, La Jolla, CA, USA, 2024.
- 2024 Hyunseok Chang, **Walid A. Hanafy**, Sarit Mukherjee, and Limin Wang. INSERT: In-Network Stateful End-to-End RDMA Telemetry. In *IEEE Conference on Computer Communications, INFOCOM'24*, 2024.
- 2024 Roozbeh Bostandoost, Adam Lechowicz, **Walid A. Hanafy**, Noman Bashir, Prashant Shenoy, and Mohammad Hajiesmaili. LACS: Learning-Augmented Algorithms for Carbon-Aware Resource Scaling with Uncertain Demand. In *The 15th ACM International Conference on Future and Sustainable Energy Systems (E-Energy '24)*, June 4–7, 2024, Singapore, Singapore, 2024.
- 2023 **Walid A. Hanafy**, Li Wu, Tarek Abdelzaher, Suhas Diggavi, and Prashant Shenoy. Failure-Resilient ML Inference at the Edge through Graceful Service Degradation. In *Proceedings of the 41st IEEE Military Communications Conference (MILCOM) workshop on Internet of Things for Adversarial Environments*, 10 2023.
- 2023 **Walid A. Hanafy**, Limin Wang, Hyunseok Chang, Sarit Mukherjee, T. V. Lakshman, and Prashant Shenoy. Understanding the Benefits of Hardware-Accelerated Communication in Model-Serving Applications. In *2023 IEEE/ACM 31st International Symposium on Quality of Service (IWQoS)*, IWQoS'23, pages 1–10, 2023.
- 2023 **Walid A. Hanafy**, Qianlin Liang, Noman Bashir, David Irwin, and Prashant Shenoy. CarbonScaler: Leveraging Cloud Workload Elasticity for Optimizing Carbon-Efficiency. *Proc. ACM Meas. Anal. Comput. Syst.*, volume 7. Association for Computing Machinery, dec 2023.
- 2023 **Walid A. Hanafy**, Roozbeh Bostandoost, Noman Bashir, David Irwin, Mohammad Hajiesmaili, and Prashant Shenoy. The War of the Efficiencies: Understanding the Tension between Carbon and Energy Optimization. In *Proc. 2nd ACM Workshop on Hot Topics in Sustainable Computing Systems (HotCarbon'23)*, July 2023.
- 2023 Abel Souza, Noman Bashir, Jorge Murillo, **Walid A. Hanafy**, Qianlin Liang, David Irwin, and Prashant Shenoy. Ecovisor: A Virtual Energy System for Carbon-Efficient Applications. In *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems*, volume 2 of *ASPLOS 2023*, page 252–265, 2023.
- 2023 Qianlin Liang, **Walid A. Hanafy**, Noman Bashir, David Irwin, and Prashant Shenoy. Energy Time Fairness: Balancing Fair Allocation of Energy and Time for GPU Workloads. In *Proceedings of the 8th ACM/IEEE Symposium on Edge Computing (SEC)*, 12 2023.
- 2023 Qianlin Liang, **Walid A. Hanafy**, Noman Bashir, Ahmed Ali-Eldin, David Irwin, and Prashant Shenoy. Dēlen: Enabling Flexible and Adaptive Model-serving for Multi-tenant Edge AI. In *Proceedings of IEEE/ACM Eighth International Conference on Internet-of-Things Design and Implementation (IoTDI)*, San Antonio, 5 2023.
- 2023 Qianlin Liang, **Walid A. Hanafy**, Ahmed Ali-Eldin, and Prashant Shenoy. Model-driven Cluster Resource Management for AI Workloads in Edge Clouds. *ACM Transactions on Autonomous and Adaptive Systems*, volume 18, pages 1–26, 2023.
- 2023 Hyunseok Chang, Limin Wang, Sarit Mukherjee, and **Walid Abdelrahman**. End-to-End RDMA Telemetry System, 2023. Patent US20230246931A1.

- 2021 **Walid A. Hanafy**, Tergel Molom-Ochir, and Rohan Shenoy. Design Considerations for Energy-Efficient Inference on Edge Devices. In *Proceedings of International Workshop on Energy-Efficient Learning at the Edge*, WEEE'21, page 302–308, 2021.
- 2019 **Walid A. Hanafy**, Alfredo Pina, and Sameh A. Salem. Machine Learning Approach for Photovoltaic Panels Cleanliness Detection. In *2019 15th International Computer Engineering Conference (ICENCO)*, pages 72–77, 2019.
- 2019 **Walid A. Hanafy**, Amr E. Mohamed, and Sameh A. Salem. A New Infrastructure Elasticity Control Algorithm for Containerized Cloud. *IEEE Access*, volume 7, pages 39731–39741, 2019.
- 2017 **Walid A. Hanafy**, Amr E. Mohamed, and Sameh A. Salem. Novel Selection Policies for Container-based Cloud Deployment Models. In *2017 13th International Computer Engineering Conference (ICENCO)*, pages 237–242, 2017.
- 2017 **Walid A. Hanafy**, Amr E. Mohamed, and Sameh A. Salem. A Load Balancing with Power Optimization Algorithm for Container-based Infrastructure Management. In *2017 12th International Conference on Computer Engineering and Systems (ICCES)*, pages 161–166, 2017.

Development Experience

2017-2019 Testicide, Egypt, Research Engineer

Designed a microservices system that implements the startup novel “codeless” GUI testing framework.

Tools: Azure Service Fabric, CosmosDB, AngularJS

2016-2017 IT-Bits, UAE, Software Engineer

Led the successful development of software solutions for multiple pharmaceutical companies.

Tools: Asp.net Core, MSSQL server, Angular

2016-2017 Consukorra-PES, Egypt, Software Engineer

Design and Implement a GIS-based data collection platform, leading to a 200% speedup in the collection process

Tools: QGIS, PostgreSQL, .net framework

Teaching Experience

University of Massachusetts Amherst, Teaching Assistant

Spring 2023 COMPSCI 677: Distributed and Operating Systems

Spring 2020 COMPSCI 577: OS Design and Implementation

Information Technology Institute (ITI), Visiting Lecturer

2015-2019 Design and teach courses in Software Design, Cloud Computing, and Software Architecture

Technologies: Microsoft Stack, Microsoft Azure, and Docker multiple recurrent courses

Helwan University, Teaching Assistant

2015-2019 Software Engineering, Database Systems, Operating Systems, and Distributed Systems

Achievements & Recognitions

2021 **Outstanding Innovation Award** for my internship from *Nokia Bell Labs*, USA.

2012 Winner of **ICT4Change** for my ITI graduation project from *ICT Ministry*, Egypt.

2011 **Young Innovators Award** for my BS graduation project from *Nahdet El Mahrosa*, Egypt.

Technical skills

Programming Python, C/C++, C#, Java

Web HTML 5, .net Stack, AngularJS

Database MSSQL, PostgreSQL, MySQL, COSMOSDB
Cloud VMware, Docker, and K8s
Testing XUnit, Specflow, and Selenium
ML Python Stack and Pytorch

Services

2020-2024 **UMass Muslim Student Association**, *Grduates Branch*, Treasurer.

2021 **Artifcat Evaluation Committe Memeber**, *Eurosys 2021*.

2017 **Docker Egypt Meetup**, *Tutorials on containers*, Funded by Microsoft and Docker, Egypt.