

# E2C: A Visual Simulator for Heterogeneous Computing Systems

Ali Mokhtari, Mohsen Amini Salehi  
 High Performance Cloud Computing (HPCC) Laboratory,  
 School of Computing and Informatics,  
 University of Louisiana at Lafayette, LA 70503, USA  
 E-mail: {ali.mokhtari1, amini}@louisiana.edu



**Abstract**—Heterogeneity has been an indispensable aspect of distributed computing throughout the history of these systems. In particular, with the increasing prevalence of accelerator technologies (e.g., GPUs and TPUs) and the emergence of domain-specific computing via ASICs and FPGA, the matter of heterogeneity and harnessing it has become a more critical challenge than ever before. Harnessing system heterogeneity has been a longstanding challenge in distributed systems and has been investigated extensively in the past. Making use of real infrastructure (such as those offered by the public cloud providers) for benchmarking the performance of heterogeneous machines, for different applications, with respect to different objectives, and under various workload intensities is cost- and time-prohibitive. To mitigate this burden, we develop an open-source simulation tool, called E2C, that can help researchers and practitioners study any type of heterogeneous (or homogeneous) computing system and measure its performance under various system configurations. E2C has an intuitive graphical user interface (GUI) that enables its users to easily examine system-level solutions (scheduling, load balancing, scalability, etc.) in a controlled environment within a short time and at no cost. In particular, E2C offers the following features: (i) simulating a heterogeneous computing system; (ii) implementing a newly developed scheduling method and plugging it into the system, (iii) measuring energy consumption and other output-related metrics; and (iv) powerful visual aspects to ease the learning curve for students. Potential users of E2C can be undergraduate and graduate students in computer science/engineering, researchers, and practitioners.

## 1 INTRODUCTION

Heterogeneity has been an indispensable aspect of distributed computing throughout the history of these systems. In the modern era, as Moore’s law is losing momentum due to the power density and heat dissipation limitations [7], [14], heterogeneous computing systems have attracted even more attention to overcome the slowdown in Moore’s law and fulfilling the desire for higher performance in various types of distributed systems. In particular, with the increasing prevalence of accelerator technologies (e.g., GPUs and TPUs) and the emergence of domain-specific computing via ASICs [15] and FPGA [4], the matter of heterogeneity and harnessing it has become a more critical challenge than ever before to deal with.

Examples of heterogeneity can be found in any type of distributed system. Public cloud providers offer and operate

based on a wide variety of machine types. Hyperscalers such as AWS and Microsoft Azure provide computing services ranging from general-purpose X86-based and ARM-based machines to FPGAs and accelerators [1]. In the context of Edge computing, domain-specific accelerators (ASICs and FPGA) and general-purpose processors are commonly used together to perform near-data processing, thereby unlocking various real-time use cases (e.g., edge AI and AR/VR applications) [2], [3]. In the HPC context, deploying various machine types with different architectures on HPC boards to fulfill the power and performance requirements is becoming a trend [5].

Heterogeneity plays a key role in improving various performance objectives of distributed systems, such as cost, energy consumption, and QoS. That is why harnessing system heterogeneity has been a longstanding challenge in distributed systems and has been investigated extensively in the past (e.g., [6], [8]–[11], [13]). Making use of real infrastructure (such as those offered by the public cloud providers) for benchmarking the performance of heterogeneous machines, for different applications, with respect to different objectives, and under various workload intensities is cost- and time-prohibitive. As an example, consider an IoT-based system that offers multiple smart applications to its users (e.g., object detection, face recognition, speech recognition, etc.); there exists a wide range of machine types with different architectures (such as x-86 or ARM-based multi-core CPUs, different types of GPUs, FPGAs, and ASICs) that can process these services. To find an optimal configuration, ideally, all permutations of these configurations must be examined. Moreover, there can be multiple workload intensities and scheduling policies that can affect the performance of the system. In addition, the energy consumption of the system might be of interest, which in turn, it adds another dimension to the evaluation tests.

To mitigate the burden of examining all cases, we need simulation tools that can help us study the performance of various system configurations. To that end, in this tutorial, we introduce E2C that is a framework to simulate any type of heterogeneous computing system. By using E2C, the researchers can easily examine their system-level solutions (scheduling, load balancing, scalability, etc.) in a controlled

environment within a short time and at no cost. In particular, E2C offers the following features: (i) simulating a heterogeneous computing system; (ii) implementing a newly developed scheduling method and plugging it into the system, (iii) measuring power and other output-related things, and (iv) visual aspects to ease the learning curve for students. These features help researchers who study resource allocation solutions in distributed systems to test and evaluate their solutions easier and faster. Moreover, the graphical user interface would help students to gain a deeper knowledge of resource allocation procedures in distributed computing systems. In the following sections, we elaborate on these features in more detail and discuss why UCC is a good fit for this work.

## 2 SIMULATING A HETEROGENEOUS COMPUTING SYSTEM VIA E2C

Figure 1 shows the overview of the E2C. The major components in E2C are: (i) workload, (ii) batch queue (iii) scheduler, (iv) machine queue, and (v) machines. In addition, there are two more components that contain canceled and dropped tasks.

Note that the heterogeneity of the system is modeled by the Expected Execution Time (EET) matrix. As shown in Figure 2, the user has access to EET matrix by selecting the workload component. Users can either modify the EET matrix manually or load the desired one as a CSV file.

The arrival of a task is also generated in the workload component. As shown in Figure 2, the user can load the desired workload trace in this section.

Upon the arrival of a task, the simulator transfers the task to the batch queue. Next, based on the selected scheduling method, the scheduler will select a task from the batch queue. Figure 3 shows the scheduler options.

The user can modify the scheduling method by selecting the scheduler component. There exist two options for the selected tasks: (i) it might be canceled or (ii) it might be mapped to one of the available machines. The status of a canceled task is set to “canceled” and no more process is needed. The canceled tasks component shows the tasks have been canceled so far. In the case of mapping decisions, the task is appended to the local queue of the assigned machine. Tasks are executed on the assigned machine in a sequential manner by default. If a task missed its deadline while executing on the machine, it is dropped from the machine. As shown in Figure 4, the Missed Tasks component shows the tasks that missed their deadline.

The simulator can be used for instructional and research purposes. So far, we have used the E2C simulator for undergraduate and graduate Distributed Computing courses to examine the impact of different scheduling policies on homogeneous and heterogeneous systems with various workload intensities. Similarly, the simulator can be used for the operating system course at the undergraduate and graduate levels to teach students about scheduling methods. The simulator has also been used to implement two research works. In [12], we have used E2C to examine energy efficiency and fairness of scheduling methods on a heterogeneous edge. In [16], we extended E2C to simulate the memory allocation policies of multi-tenant applications on a homogeneous edge.

## 3 POTENTIAL USERS OF E2C SIMULATOR

Students learning distributed systems can benefit from E2C simulator. They can deeply investigate scheduling methods and gain insights into their advantages and disadvantages. They can also learn how heterogeneity can improve the performance of the system. Moreover, they can study the energy consumption of the system once a certain scheduling method is applied. Researchers in resource allocation area and cloud solution architects can employ E2C simulator to test their solution prior to implementation.

## 4 BIOGRAPHY OF THE INSTRUCTOR(S)

Dr. Mohsen Amini Salehi is an Associate Professor and Francis Patrick Clark/BORSF Endowed Professorship holder at the School of Computing and Informatics (CMIX), University of Louisiana Lafayette, USA. He is the director of High Performance and Cloud Computing (HPCC) Laboratory where several graduate and undergraduate students research on various aspects of Distributed and Cloud computing. Dr. Amini is an NSF CAREER Awardee and, so far, he has had 11 research projects funded by National Science Foundation (NSF) and Board of Regents of Louisiana. He has also received five awards and certificates from University of Louisiana at Lafayette in recognition of his innovative research. Dr. Amini has been an active researcher in Distributed and Cloud computing research areas since 2004. He has received several awards in recognition of his research, including the “Best Intern Award” from Infosys Ltd., in 2012, and the IEEE award for “Applied Research”, in 2009. His research interests are in building smart systems across edge-to-cloud continuum, virtualization, resource allocation, energy-efficiency, heterogeneity, and trustworthy in Distributed, Edge, and Cloud computing systems.

Ali Mokhtari is currently a Ph.D. candidate in computer science at the University of Louisiana at Lafayette. Ali works as a research assistant at HighPerformance Cloud Computing (HPCC) Lab in the computer science department. His research interest includes using reinforcement learning for efficient resource allocation in heterogeneous edge computing systems. He earned an M.Sc. in aerospace engineering from the Sharif University of Technology in Iran and a B.Sc. in mechanical engineering from Shiraz University in Iran.

## 5 E2C CODE AND RESOURCE AVAILABILITY

E2C core is available for download at the following address:

<https://github.com/hpcclab/E2C-Sim>

The manual document on how to run E2C and its options and full documentations are available here:

<https://hpcclab.github.io/E2C-Sim-docs/>

The video resources for E2C are in this [YouTube page](#).

## 6 ACKNOWLEDGEMENT

Development of E2C was made possible by the funding support provided by National Science Foundation (NSF) under awards# CNS-2007209 and CNS-2117785 (NSF CAREER Award).

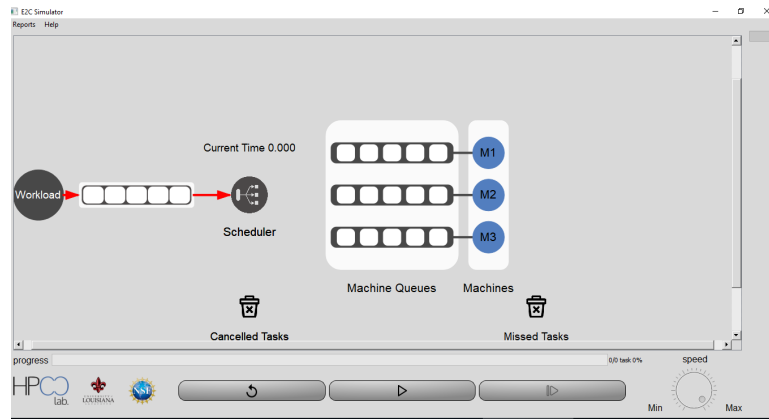


Fig. 1. Overview of the E2C Simulator

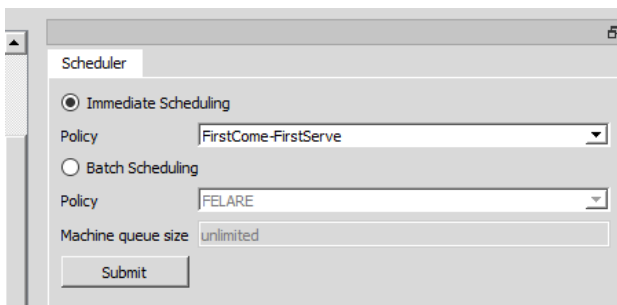


Fig. 2. Workload component. The user can modify the EET matrix and arrival times of the task via the workload component.




Fig. 3. The scheduler options. The user can change the scheduling method here. The user should also set the machine queue size for batch scheduling methods.

	Task ID	Type	Assigned Machine	Arrival Time	Start Time	Missed Time
1	11	T3	m3	12.90	42.21	57.90
2	14	T4	m3	15.96	57.90	60.96
3	20	T4	m3	23.18	65.92	68.18
4	25	T3	m2	31.01	70.54	76.01
5	27	T3	m1	31.03	74.07	76.03

Fig. 4. Missed Tasks component shows tasks that missed their deadline.

## REFERENCES

- [1] Amazon sagemaker. <https://aws.amazon.com/sagemaker/>.
- [2] Glass enterprise edition 2. <https://www.google.com/glass/tech-specs/>, note = Accessed: September 2023 .
- [3] Qualcomm reveals the world's first dedicated xr platform. <https://www.qualcomm.com/news/releases/2018/05/qualcomm-reveals-worlds-first-dedicated-xr-platform>, note = Posted on: May 28, 2018.
- [4] Christophe Bobda, Joel Mandebi Mbongue, Paul Chow, Mohammad Ewais, Naif Tarafdar, Juan Camilo Vega, Ken Eguro, Dirk Koch, Suranga Handagala, Miriam Leeser, et al. The future of fpga acceleration in datacenters and the cloud. *ACM Transactions on Reconfigurable Technology and Systems (TRETS)*, 15(3):1–42, 2022.
- [5] Suma George Cardwell, Craig Vineyard, Willam Severa, Frances S Chance, Frederick Rothganger, Felix Wang, Srideep Musuvathy, Corinne Teeter, and James B Aimone. Truly heterogeneous hpc: Co-design to achieve what science needs from hpc. In *Smoky Mountains Computational Sciences and Engineering Conference*, pages 349–365. Springer, 2020.
- [6] Chavit Denninnart, James Gentry, Ali Mokhtari, and Mohsen Amini Salehi. Efficient task pruning mechanism to improve robustness of heterogeneous computing systems. *Journal of Parallel and Distributed Computing (JPDC)*, 142:46–61, 2020.
- [7] Hadi Esmailzadeh, Emily Blem, Renee St. Amant, Karthikeyan Sankaralingam, and Doug Burger. Dark silicon and the end of multicore scaling. In *Proceedings of the 38th annual international symposium on Computer architecture*, pages 365–376, 2011.
- [8] James Gentry, Chavit Denninnart, and Mohsen Amini Salehi. Robust dynamic resource allocation via probabilistic task pruning in heterogeneous computing systems. In *Proceedings of the 33rd IEEE International Parallel & Distributed Processing Symposium, IPDPS '19*, May 2019.
- [9] Xiangbo Li, Mohsen Amini Salehi, Yamini Joshi, Mahmoud Darwich, Landreneau Brad, and Magdi Bayoumi. Performance analysis and modelling of video stream transcoding using heterogeneous cloud services. *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, 30(4):910–922, Sep. 2018.
- [10] Xiangbo Li, Mohsen Amini Salehi, Magdy Bayoumi, Nian-Feng Tzeng, and Rajkumar Buyya. Cost-efficient and robust on-demand video stream transcoding using heterogeneous cloud services. *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, 29(3):556–571, Mar. 2018.
- [11] Ali Mokhtari, Chavit Denninnart, and Mohsen Amini Salehi. Autonomous task dropping mechanism to achieve robustness in heterogeneous computing systems. In *Proceedings of 34th IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, pages 17–26, 2020.
- [12] Ali Mokhtari, MD Abir Hossen, Pooyan Jamshidi, and Mohsen Amini Salehi. FELARE: Fair Scheduling of Machine Learning Applications on Heterogeneous Edge Systems. In *Proceedings of the 15th IEEE International Conference on Cloud Computing, IEEE Cloud '22*, 2022.
- [13] Mohsen Amini Salehi, Jay Smith, Anthony A Maciejewski, Howard Jay Siegel, Edwin KP Chong, Jonathan Apodaca, Luis D Briceño, Timothy Renner, Vladimir Shestak, Joshua Ladd, et al. Stochastic-based robust dynamic resource allocation for independent tasks in a heterogeneous

- computing system. *Journal of Parallel and Distributed Computing (JPDC)*, 97:96–111, Nov. 2016.
- [14] Michael B Taylor. Is dark silicon useful? harnessing the four horsemen of the coming dark silicon apocalypse. In *DAC Design Automation Conference 2012*, pages 1131–1136. IEEE, 2012.
  - [15] Michael Bedford Taylor, Luis Vega, Moein Khazraee, Ikuo Magaki, Scott Davidson, and Dustin Richmond. Asic clouds: Specializing the datacenter for planet-scale applications. *Communications of the ACM*, 63(7):103–109, 2020.
  - [16] SM Zobaed, Ali Mokhtari, Jaya Prakash Champati, Mathieu Kourouma, and Mohsen Amini Salehi. Edge-MultiAI: Multi-Tenancy of Latency-Sensitive Deep Learning Applications on Edge. In *Proceedings of 15th IEEE/ACM International Conference on Utility and Cloud Computing, UCC '22*, Dec. 2022.