

Benchmarking Different Application Types across Heterogeneous Cloud Compute Services

Nivedhitha Duggi, Masoud Rafiei, Mohsen Amini Salehi

High Performance Cloud Computing (HPCC) Lab,
University of North Texas, Denton TX, USA
nivedhithaduggi@my.unt.edu and mohsen.aminisalehi@unt.edu

Abstract. Infrastructure as a Service (IaaS) clouds have become the predominant underlying infrastructure for the operation of modern and smart technology. IaaS clouds have proven to be useful for multiple reasons such as reduced costs, increased speed and efficiency, and better reliability and scalability. Compute services offered by such clouds are heterogeneous—they offer a set of architecturally diverse machines that fit efficiently executing different workloads. However, there has been little study to shed light on the performance of popular application types on these heterogeneous compute servers across different clouds. Such a study can help organizations to optimally (in terms of cost, latency, throughput, consumed energy, carbon footprint, etc.) employ cloud compute services. At HPCC lab, we have focused on such benchmarks in different research projects and, in this report, we curate those benchmarks in a single document to help other researchers in the community using them. Specifically, we introduce our benchmarks datasets for three application types in three different domains, namely: Deep Neural Networks (DNN) Inference for industrial applications, Machine Learning (ML) Inference for assistive technology applications, and video transcoding for multimedia use cases.

1 Overview

Heterogeneous computing systems (HCS) have become essential in overcoming computational limitations as Moore’s Law reaches its twilight, especially with the rising demand for scalable and efficient computational solutions for smart applications. These systems leverage architecturally diverse machines, enabling them to handle tasks with varying computational requirements. Heterogeneity is increasingly prevalent across cloud systems, including public cloud platforms like Amazon Web Services (AWS) and Microsoft Azure, which provide a mix of x86-based CPUs, ARM processors, GPUs, and domain-specific accelerators like FPGAs.

On the application side, cloud computing has become a cornerstone for deploying a wide range of applications, including Deep Neural Network (DNN) inference and training tasks, and more conventional workloads, such as video

processing, across various industries. These systems are highly diverse (aka application level heterogeneity) and have affinity with heterogeneous compute servers offered by cloud platforms.

This technical report delves into the benchmarking of different application types across different cloud providers with heterogeneity compute offerings. This benchmarking provides a base to evaluate and compare the performance of diverse workloads. In this paper, we focus on benchmarking applications for inference tasks across heterogeneous servers. Specifically, we analyze:

- Deep Neural Network (DNN) inference for time-critical applications in the Oil and Gas industry [1,2].
- Machine Learning inference tasks, including image classification, NLP, and speech recognition for assistive technology, particularly, for blind and visually impaired users [3].
- Video transcoding tasks involving resolution, frame rate, bit rate, and codec conversions [4,5].

This technical report introduces the methodologies, datasets, and cloud platforms used for benchmarking and provides insights into accessing these benchmarks across heterogeneous systems. The datasets of this report are all available on the High-Performance Cloud Computing Lab’s GitHub repositories.

1.1 Benchmarking Structure

Figure 1 provides an overview to the benchmarking we introduce in this technical report. According to the figure, our benchmarking encompasses two Cloud providers, namely AWS and Chameleon Cloud; each one with heterogeneous processors that are listed in the overview taxonomy. On the application side we consider three application domains, namely “Video Transcoding”, “Machine Learning” inference for assistive technology, and “Machine Learning” inference for industrial Use Cases in the context of Oil and Gas. Each application is benchmarked on each processor type in different clouds.

1.2 Application Types

The industrial use cases include applications like fire detection, oil spill detection, and acoustic impedance estimation in the Oil and Gas sector, which require real-time and highly accurate predictions to prevent catastrophic failures [?], [?].

Assistive Technology applications are in the context of supporting blind and visually impaired users, focusing on speech recognition and object (obstacle) detection to enhance accessibility and independence.

Video Processing application are for video transcoding and compression in multimedia streaming. The applications include changing resolution, bit rate, frame rate, and compression format.

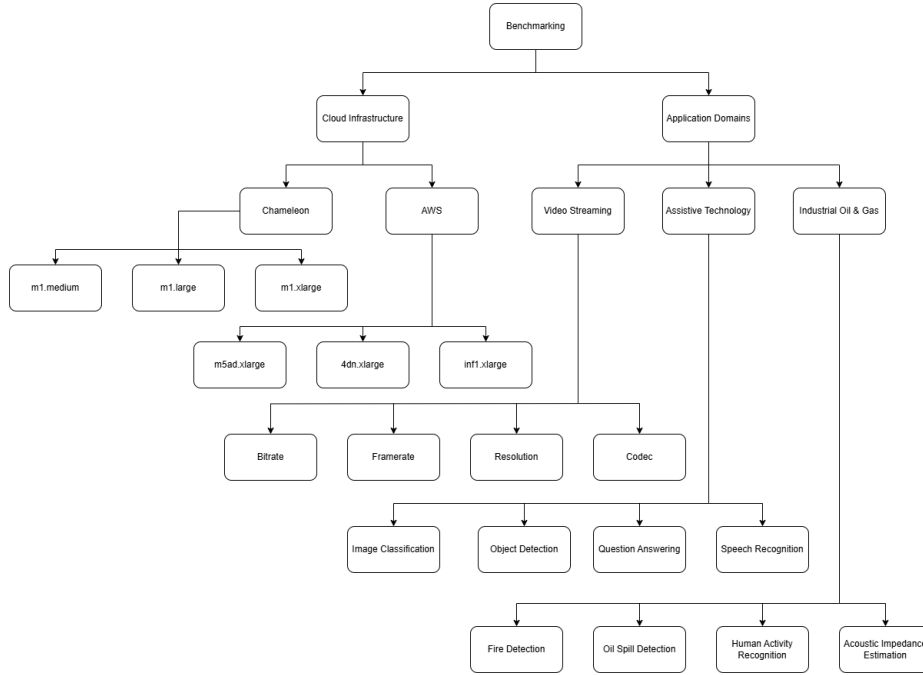


Fig. 1: Benchmarking structure

1.3 Heterogeneous Resources in Cloud Platforms

Leading providers like Amazon Web Services (AWS) and research-oriented platforms such as Chameleon Cloud, offering a range of virtual machine (VM) configurations to cater to different application needs with different prices.

In this benchmarking we use VMs configured with different processor types offered by these cloud providers.

Performance variability in heterogeneous systems poses significant challenges, especially for inference tasks where latency and throughput are critical. By analyzing execution times across varied machine types, we can figure out: (A) the randomness in application performance due to heterogeneity. (B) Comprehensive benchmarking results, including resource utilization and execution times. These results can be used by researchers to explore strategies to optimize resource allocation for cost-efficiency and scalability. (C) This benchmarking provides access to essential resources through curated links to GitHub repositories and datasets, enabling reproducibility and further research

In the rest of this technical report, we will explain details of benchmarking for industrial use case (Section 2), assistive technology for blind and visually-impaired users (Section 3), and video processing (Section 4).

2 Benchmark I: DNN Applications in Industrial (Oil & Gas) Use Case

2.1 Overview

In the context of O&G operations, several activities, such as fire detection, toxic gas monitoring, and spill detection, require accurate and timely processing. A failure to process these operations correctly and quickly can lead to disastrous consequences, including oil spills, explosions, and loss of life. Understanding the uncertainties in the execution times of different applications and properly modeling them is crucial for improving the safety, reliability, and efficiency of O&G operations in Industry 4.0.

We benchmark the inference time of four Deep Neural Networks (DNN) applications in the context of O&G industry on both AWS and Chameleon cloud, as explained in the next subsections. We note that more extensive results about this benchmark is available in [1]. This benchmark was also used in conducting the research and evaluation of [2].

2.2 DNN Inference Time for O&G Applications

1. **Fire Detection:** Utilizing the FireNet DNN model(which itself is based on Alexnet [6]), this application detects fire in real-time from video frames. FireNet integrates convolutional layers with max-pooling and normalization to optimize processing time and accuracy. The dataset consists of 240 videos processed on heterogeneous cloud systems, including AWS and Chameleon Cloud.
All the benchmarking results and scripts for the Fire Detection Application are available in [GitHub Repository for Fire Detection](#).
2. **Oil Spill Detection:** We utilize a detection system that operates based on the FCN-8 model [7]. The model contains five Fully Convolutional Network (FCN) blocks and two up-sampling blocks that collectively perform semantic segmentation (i.e., classifying every pixel) of an input image and output a labeled image. The FCN-8 model functions based on the satellite (a.k.a. SAR) [8] images. We configure the analysis to obtain the inference time of 110 SAR images collected by MKLab [9].
All the benchmarking results and scripts for the Fire Detection Application are available in [GitHub link for Oil Spill Detection](#).
3. **Human Activity Recognition (HAR):** Human Activity Recognition (HAR) systems are widely used in Industry 4.0 to ensure workers safety in hazardous zones. For this purpose, motion sensors, such as accelerometer and gyroscope, that are widely available on handheld PDA devices are utilized. The HAR system we use operates based on the sequential neural network model with four layers to identify the worker’s activities (namely, walking, walking upstairs, walking downstairs, sitting). For analysis, we use a dataset of UCI machine learning repository, known as Human Activity Recognition Using Smartphones [10].

All the benchmarking results and scripts for this Application are available in [GitHub Repository for Human Activity Recognition](#).

4. **Acoustic Impedance Estimation:** Estimating acoustic impedance (AI) from seismic data is an important step in O&G exploration. To estimate AI from seismic data, we utilize a solution functions based on the temporal convolutional network [11]. Marmousi 2 dataset [12] is used to estimate AI. All the benchmarking results and scripts for this Application are available in [GitHub Repository for Acoustic Impedance Estimation](#).

2.3 Cloud Platforms Used for Benchmarking the Industrial Applications

From **Amazon Web Services (AWS)**, we employed the following instance type virtual machines (VMs):

1. **General Purpose Instances (e.g., m5ad.xlarge):** Balances compute, memory, and networking for diverse workloads.
2. **GPU Instances (e.g., g4dn.xlarge):** Optimized for high-performance tasks like deep learning and video processing.
3. **Inference-Optimized Instances (e.g., inf1.xlarge):** Designed for machine learning inference at scale.
For configurations and raw data of the benchmarking under these AWS instances, please refer to [AWS results for industrial use cases](#).

Chameleon Cloud is a configurable experimental platform that enables extensive benchmarking research. It offers VM flavors such as `m1.medium`, `m1.large`, and `m1.xlarge`, catering to various resource needs [13].

For configurations and raw data of the benchmarking under these Chameleon Cloud instances, please refer to [Chameleon results for industrial use cases](#).

2.4 Benchmarking Methodology for Industrial Applications on Heterogeneous Clouds

To benchmark the DNN applications, the inference time is measured in AWS and Chameleon cloud providers with respect to different machine types they offer. That is, the aforementioned applications are deployed across different cloud instances mentioned earlier. The random nature of the inference time is captured by running each application multiple times under varying cloud conditions to obtain a comprehensive understanding of how the system’s heterogeneity influences the performance.

Detailed benchmarking data are available in: [Benchmarking Repository](#).

2.5 Analysis of Inference Time of the Applications

The inference time of each application will be analyzed using both AWS and Chameleon Cloud, and the times across the different machines will be compared. First, we will test for the normality of the datasets using the Shapiro-Wilk test, then we will utilize other statistical inference methods such as mean and standard deviations to compare the execution times on the different cloud services.

Shapiro-Wilk Test for Normality of the Data The Shapiro-Wilk test is a test that evaluates whether the data set is normally distributed [?]. A large p-value indicates that the dataset is normally distributed, while a small p-value indicates that the data is not normally distributed. The normality tests are used to understand the distribution of the collected data; therefore, this test is used.

Results of the AWS Cloud Shapiro-Wilk Test; Raw data can be found at: [Results of the AWS Cloud Shapiro-Wilk Test](#)

Results of the Chameleon Cloud Shapiro-Wilk Test; Raw data can be found at: [Results of the Chameleon Cloud Shapiro-Wilk Test](#)

Kolmogorov-Smirnoff Goodness of Fit Test Due to the lack of normality in several cases after using the Shapiro-Wilk test, we utilized the Kolmogorov-Smirnoff test to determine the best fitting distribution of the inference times. The Kolmogorov-Smirnoff Goodness of Fit test identifies whether a set of samples derived from a population fits to a specific distribution.

Results of the Chameleon Cloud Shapiro-Wilk Test; Raw data can be found at: [Results of the Chameleon Cloud Shapiro-Wilk Test](#)

Mean and Standard Deviation of Inference Executive Times The mean and standard deviation of the inference times summarizes the behavior of the observations in a single value.

Results of the AWS Mean and Standard Deviation; Raw data can be found at: [Results of the AWS Mean and Standard Deviation](#)

Results of the Chameleon Mean and Standard Deviation. Raw data can be found at: [Results of the Chameleon Mean and Standard Deviation](#)

For further analysis of these results, please check our paper dedicated to this benchmarking [1].

3 Benchmark II: Machine Learning Inference Benchmarking on Heterogeneous Cloud Resources for Assistive Technology

3.1 Overview

In the context of assistive technology, several activities, such as identifying surrounding object, obstacle detection, and interaction with the system (in form of speech recognition and question answering) require accurate and real-time processing. This benchmarking aims at Understanding the inference time of different applications for assistive technology of blind and visually impaired individuals.

We benchmark the inference time of four ML applications in the context of assistive technology on different VM types of AWS cloud, as explained in the next subsections. We note that this benchmark has been used and explained in our prior publication [3].

3.2 Task Types

In this section, we benchmark the execution times of four ML tasks used for assisting blind and visually impaired individuals, namely Image Classification, Object Detection, Question Answering, and Speech Recognition. To ensure consistency across the task types, they will all be converted to the ONNX (Open Neural Network Exchange) format using the Python ONNX converter. The workload of the tasks was also designed to be available for execution from the beginning.

The summary of the benchmarking of these applications on different AWS VM types are shown in Table 1.

Task Type	t2.large	c5.2xlarge	g4dn.xlarge
Image Classification [14]	74	27	12
Object Detection [15]	36	21	12
Question Answering [16]	36	16	4
Speech Recognition [17]	621	237	20

Table 1: Table showing the expected execution times for the different tasks across the different machines. The table is also available in the [ML Github repository](#).

Image Classification Image classification [14] is a supervised learning problem where a model is trained to identify specific target classes using labeled example images. The aim is to classify unseen images accurately by learning patterns from the training data. In this dataset, image classification is implemented using the ResNet50 model, which is a robust deep neural network for image recognition tasks. To measure performance, 1,000 sample images are processed on each machine (VM) type to evaluate execution times for the classification task.

The average execution (inference) time of this model on different AWS VM types are shown in Table 2. Raw data of this benchmark can be found at [Image Classification Github page](#), and figure 2 graphically shows these different inference times.

Machine Type	Average Inference Time
c5.2xlarge	24.929
g4dn.xlarge	8.663
t2.large	74.096

Table 2: Average inference times of the three machine types when performing image classification tasks.




Fig. 2: Chart comparing the Image Classification Inference times across the Machines

Object Detection Object detection [15] is to locate and identify instances of objects within images or videos. In this dataset, YOLOv5 model is utilized to perform object detection, a task run on 1,000 sample images ten times to measure the expected execution time for object recognition. The set of images used for this purpose are available [in this link](#).

The average execution (inference) time of this model on different AWS VM types are shown in Table 3. Raw data of this benchmark can be found at [Object Detection Github page](#), and figure 3 graphically shows these different inference times.

Machine	Average Inference Time
c2.2xlarge	21.483
g4dn.xlarge	14.149
t2.large	38.824

Table 3: The Average inference times of the three machines when performing the Object Detection.

Question Answering Question-Answering (QA) models are advanced machine and deep learning models designed to answer questions based on provided context, or in some cases, without context, as seen in open-domain QA tasks.

In this dataset, the QA task is implemented using the DistilBERT model [16], a lightweight variant of BERT. A sample context and question are provided as input, and the inference task is executed 1,000 times.

The average execution (inference) time of this model on different AWS VM types are shown in Table 4. Raw data of this benchmark can be found at [Question Answering Github page](#), and figure 4 graphically shows these different inference times.




Fig. 3: The average inference times for Object Detection across the different AWS VM types.

Machine Type	Average Inference Time
c2.2xlarge	16.995
g4dn.xlarge	4.035
t2.large	36.116

Table 4: Average inference time of performing the question-answering inference across heterogeneous AWS VM types.




Fig. 4: Average inference times of different machines performing questions answering tasks.

Speech Recognition Automatic speech recognition (ASR) is the technology that enables programs to convert spoken language into written text. It differs from voice recognition, which focuses solely on identifying individual speakers.

In this dataset, speech recognition tasks are performed using the Wav2vec2 model [17]. Four-second audio samples are executed across various machine types, and the average inference time is calculated to evaluate the system’s performance.

The average execution (inference) time of this model on different AWS VM types are shown in Table 5. Raw data of this benchmark can be found at [Speech Recognition Github page](#), and figure 5 graphically shows these different inference times.

Machine Types	Average Inference Times
c2.2xlarge	230.32
g4dn.xlarge	21.579
t2.large	621.614

Table 5: Average inference times of heterogeneous AWS VM types performing the speech recognition task.




Fig. 5: Average inference times of heterogeneous AWS machines performing speech recognition.

4 Benchmark III: Video Transcoding Benchmark on Heterogeneous Cloud Resources

4.1 Overview

In this section, we explain the benchmark we produced for video processing. The benchmarking focuses on the video transcoding operation, which CPU-intensive,

using FFMpeg, a popular tool for video processing. We note that FFMpeg is a powerful tool with the ability to concurrently process multiple operations. Therefore, the benchmarking includes multi-parameter processing information on the cloud, in addition to single parameter transcoding.

This benchmark has been used in [5] and [4] publications. Interested researchers can refer to these papers for more details on the specifications of the benchmark and how to use it. Here is the [Github repository](#) where all the benchmarking results and scripts are made available. This Github repository includes a [folder](#) where all the scripts for splitting and transcoding benchmark videos are included.

4.2 Video Transcoding Tasks

This repository includes 100 open-source videos that are gathered in a variety of content forms (i.e. fast/slow pace, action, scenery, animation, etc.). All 100 videos are available in the following formats

1. [In their original form](#). Also, in this [link](#) the meta-data (content type, codec, resolution, of each video is described in the CSV format.
2. [In form of 2-second split video segments](#).
3. The 2-second split video segments in [the uniform \(standard\) compression format](#). We define standard as to have a 720p resolution, 30 fps framerate, H.264 codec, and 4500 kb/s bitrate
4. the 2-second split video segments available with the [standard and HLS-compatible format](#).

The task types we considered for this benchmarking are the following three transcoding operations:

Bitrate Bitrate describes the rate at which bits are transferred from one location to another along a digital network. The average execution time is collected for five parameters: 384K, 512K, 768K, 1024K, and 1536K.

Frame rate Frame rate is the measurement of how quickly a number of frames appear in a second. The average execution time is collected for the five parameters: 10fps, 15fps, 20fps, 30fps, and 40fps.

Resolution Resolution is the total number of pixels in a video frame rate. The average execution time for the resolution tasks will be collected for the five parameters: 352×288, 680×320, 720×480, 1280×800, and 1920×1080.

4.3 Single Parameter Video Transcoding Benchmarking

To limit the degree of freedom in execution time, in this part, each task is configured to change only one specification of the videos in the benchmark dataset.

Single Parameter Transcoding folder include transcoding which changes one parameter of the standardized video segment at a time; the folder also includes the required scripts to perform such benchmarking. All the results of this benchmark are from processing the video dataset only on **Chameleon Cloud Small VM** type. Each experiment is run 30-times. Both individual transcoding time and the mean and standard deviation are reported in the benchmark.

4.4 Multi-Parameter Video Transcoding Benchmarking

Multiple Parameter Transcoding folder include transcoding which using FFMpeg, we change three parameters, namely Framerate, Resolution, Codec (e.g., 30, 240p, HEVC) of the standardized video segment in one transcoding operation (i.e., output one video with three parameters changed). This benchmarking was conducted both on the following VM types of Chameleon Cloud: **Medium, Large, XLarge, XXXLarge**; For AWS cloud we used **c5.2xlarge, g4dn.xlarge, m5.2xlarge, m5a.2xlarge, r5.2xlarge** VM types. More information about these resources are shown in the tables of Figure 6 These results are available in separate folders.

Chameleon Cloud				
Type	Abbr.	CPUs	RAM (GB)	Disk (GB)
medium	CC-Medium	2	4	40
large	CC-Large	4	8	40
xlarge	CC-XLarge	8	16	40
xxlarge	CC-2xLarge	16	32	40

Amazon Web Services					
Type	Family	Abbr.	vCPUs	Memory (GiB)	Instance Storage
c5.2xlarge	Compute optimized	AWS-CO	8	16	EBS only
g4dn.xlarge	GPU instances	AWS-GPU	4	16	1 x 125 (SSD)
m5.2xlarge	General Purpose	AWS-GPI	8 (Intel)	32	EBS only
m5a.2xlarge	General Purpose	AWS-GPA	8 (AMD)	32	EBS only
r5.2xlarge	Memory Optimized	AWS-MO	8	64	EBS only

Fig. 6: Different machine types used for benchmarking multi-parameter video transcoding. Image taken for the main report file available in [here](#).

The summary of analysis of multi-parameter transcoding for different codec types for AWS cloud is show in the table of Figure 7. Each entry of the table shows the minimum and maximum transcoding time of the resolution, frame rate, and/or bit rate change for different compression formats (codec).

Execution Time Range of Every Resolution and Framerate (Min-Max)				
	HEVC	VP9	MPEG4	H.264
CO	0.660 - 4.729	0.107 - 0.188	0.272 - 0.902	0.028 - 0.040
GPA	1.017 - 6.901	0.112 - 0.165	0.317 - 1.13	0.032 - 0.036
GPI	0.825 - 5.131	0.115 - 0.197	0.292 - 0.977	0.031 - 0.332
MO	0.741 - 5.076	0.114 - 0.198	0.299 - 0.983	0.031 - 0.036

Fig. 7: Min and Max of transcoding time for different machine types ins AWS. CO is AWS c5.2xlarge, GPA is AWS m5a.2xlarge with AMD processor, GPI is AWS m5.2xlarge with Intel processor, and MO is AWS r5.2xlarge . Image taken for the main report file available in [Page 45](#).

4.5 Merging Multiple Transcoding Benchmarking

FFMpeg has the ability to produce multiple outputs in one command (a.k.a. transcoding merging). Note that this is different with multi-parameter transcoding, because in merging, there are more than one output files. For instance, if we transcode frame rate, bit rate, and resolution of `video1`, and generate one output video file, this is a multi-parameter transcoding, whereas, changing the resolution of `video1` to 720p and 1280p and generate two separate video files is called merging (i.e., merging to different FFMpeg commands in one).

Early evaluation of the collected execution-time revealed a remarkable variation in the execution-time of some task types. Specifically, we noticed that codec execution-time is far beyond the other three task types. Accordingly, we categorize the tasks types into two groups: *First* group is called Video Information Conversion (*VIC*) that includes changing bit-rate, frame-rate, or resolution task types. Tasks of this group have a low variation in their execution-times, when processing different video segments on the same machine type. *Second* group is Video Compression Conversion that only includes the codec task type (hence, we call it the Codec group). In contrast to the first group, the codec execution-time (and subsequently its merge-saving) for different video segments varies remarkably even on the same machine.

This is the [link](#) to the merging benchmark and the paper elaborates on it is available here [5]. The above Github page contains three sub folder, namely scripts, Codec class, and (a.k.a. VIC) class. This dataset is prepared to show merge saving percentage (i.e., relative percentage between merge and non-merge counterpart).

The **Codec class** contains 3 subfolders which each of them have two Microsoft Excel files (xls format) that include the execution time of merged and non merged transcoding. Task merging in this case are performed with the same operation only. There is also “merge across operations.xls” that includes the data from where tasks from multiple operations within VIC class are merged together. The **VIC class** contains 3 subfolders where each one of them represents transcoding time of each of the codec (HEVC, mpeg, VP9). Both as a single task and a merged task.

The data collected in this benchmark are performed on our own local cloud which is a DELL PowerEdge R830 with 4×Intel Xeon E5-4628Lv4 processors with 112 homogeneous cores, 384 GB memory (24×16 GB DRAM), and RAID1 (2×900 GB HDD) storage.

5 Summary

This technical report serves as a handbook for different benchmarking we have done on various application types on different clouds with heterogeneous machines. We presented three main benchmarked applications in different domains:

1. Four DNN applications in industrial use case (Oil and Gas context). This benchmarking was conducted on Chameleon and AWS clouds, each one with heterogeneous VMs.
2. Four ML applications in assistive technology (blind and visually impaired) use case. This benchmarking was exclusively done on heterogeneous AWS VMs.
3. Four video transcoding operations using FFMpeg. This benchmarking was performed in three phases. First, single parameter transcoding on Chameleon Cloud. Second, Multi-parameter transcoding on both Chameleon and AWS clouds. Third, merging multiple transcoding operations in one FFMpeg command that was done on our local cluster.

All these benchmarks are publicly available through our Github and are reflected in our publications. We hope that other researchers in the community can make use of this benchmarks as well.

References

1. R. Hussain, M. Amini, and A. Pakravan, “Analyzing the performance of smart industry 4.0 applications on cloud computing systems,” in *Proceedings of IEEE 22nd International Conference on High Performance Computing and Communications (HPCC)*, 2020.
2. R. F. Hussain and M. A. Salehi, “Resource allocation of industry 4.0 micro-service applications across serverless fog federation,” *Future Generation Computer Systems*, vol. 154, pp. 479–490, 2024.
3. A. Mokhtari, S. Ghafouri, P. Jamshidi, and M. Amini Salehi, “Heet: A heterogeneity measure to quantify the difference across distributed computing systems,” in *Proceedings of the 17th IEEE/ACM International Conference on Utility and Cloud Computing, UCC '24*, 2024.
4. C. Denninnart and M. A. Salehi, “Harnessing the potential of function-reuse in multimedia cloud systems,” *IEEE Transactions on Parallel and Distributed Systems*, vol. 33, no. 3, pp. 617–629, 2022.
5. S. Wu, C. Denninnart, X. Li, Y. Wang, and M. A. Salehi, “Descriptive and predictive analysis of aggregating functions in serverless clouds: the case of video streaming,” in *Proceedings of the 22nd International Conference on High Performance Computing and Communications (HPCC '20)*, pp. 19–26, 2020.

6. M. Z. Alom, T. M. Taha, C. Yakopcic, S. Westberg, P. Sidike, M. S. Nasrin, B. C. Van Esesn, A. A. S. Awwal, and V. K. Asari, "The history began from alexnet: A comprehensive survey on deep learning approaches," *arXiv preprint arXiv:1803.01164*, 2018.
7. J. Long, E. Shelhamer, and T. Darrell, "Fully convolutional networks for semantic segmentation," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3431–3440, 2015.
8. Z. Huang, C. O. Dumitru, Z. Pan, B. Lei, and M. Datcu, "Classification of large-scale high-resolution sar images with deep transfer learning," *Journal of Geoscience and Remote Sensing Letters*, 2020.
9. M. Krestenitis, G. Orfanidis, K. Ioannidis, K. Avgerinakis, S. Vrochidis, and I. Kompatsiaris, "Oil spill identification from satellite images using deep neural networks," *Journal of Remote Sensing*, vol. 11, no. 15, p. 1762, 2019.
10. D. Anguita, A. Ghio, L. Oneto, X. Parra, and J. L. Reyes-Ortiz, "A public domain dataset for human activity recognition using smartphones.," in *Esann*, 2013.
11. A. Mustafa, M. Alfarraj, and G. AlRegib, "Estimation of acoustic impedance from seismic data using temporal convolutional network," *arXiv preprint arXiv:1906.02684*, 2019.
12. R. Versteeg, "The Marmousi experience; velocity model determination on a synthetic complex data set," *Journal of the Leading Edge*, vol. 13, pp. 927–936, 09 1994.
13. K. Keahey, P. Riteau, D. Stanzione, T. Cockerill, J. Mambretti, P. Rad, and P. Ruth, "Chameleon: a scalable production testbed for computer science research," in *Contemporary High Performance Computing: From Petascale toward Exascale*, vol. 3, ch. 5, pp. 123–148, May 2019.
14. K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
15. G. Jocher, "YOLOv5 by Ultralytics." <https://github.com/ultralytics/yolov5>, May 2020. (Accessed: Aug. 2, 2023).
16. V. Sanh, L. Debut, J. Chaumond, and T. Wolf, "Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter," *arXiv preprint arXiv:1910.01108*, 2019.
17. A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," *Advances in neural information processing systems*, vol. 33, pp. 12449–12460, 2020.