Sarcouncil Journal of Engineering and Computer Sciences

ISSN(Online): 2945-3585

Volume- 05| Issue- 08| 2026

Research Article **Received:** 15-07-2026 | **Accepted:** 10-08-2026 | **Published:** 25-08-2026

An Adaptive Hybrid Cloud Framework for Intelligent Resource Scheduling in IoT and AI Applications

Dilip Mehendra Sharma

North American University: Stafford-TX.

Abstract: His paper presents an intelligent scheduling architecture for optimizing Internet of Things (IoT)-based machine learning workloads across edge and cloud computing environments. The proposed approach integrates predictive latency modeling with GPU acceleration to improve execution efficiency, reduce response time, and optimize resource utilization. Experimental results demonstrate a 3.3× speedup compared with traditional scheduling approaches, while achieving improved resource utilization in hybrid cloud–AI deployments. The findings highlight the effectiveness of the proposed architecture in supporting efficient and responsive machine learning workload management across distributed IoT, edge, and cloud environments.

Keywords: Cloud Framework ,AI Applications,. AWS.

INTRODUCTION

Serverless computing has emerged as a mechanism for hosting event-driven applications with automatic scaling (Function as a Service (FaaS) [AWS Lambda, 2020; Serverless, 2020; Azure Functions, 2020]). Stateless functions are defined and invoked in response to events originating from the overall system (storage update, message receipt, state transition, etc.). Further, an isolated execution environment is provisioned dynamically. Finally, billing is based on the exact resource consumption. Many commercial and open-source platforms have been released that provide this abstraction. AWS Lambda and Apache OpenWhisk [AWS Lambda, 2020; IBM OpenWhisk, 2020] are two of them.

To reduce latency and strain on bandwidth caused by public cloud interactions, extensions of the serverless paradigm to edge environments are explored for IoT applications [IBM OpenWhisk, 2020; IBM OpenWhisk, 2020; GreenGrass and IoT Core, 2019]. The limits of resources at the network edge, along with a lack of special hardware accelerators (like GPUs), make machine learning tasks more computationally intensive.

A new hybrid cloud scheduling framework, Serverless Tele Operable Hybrid Cloud, or STOIC, has been developed to mit igate these issues with the help of predictive latency modeling and GPU acceleration. The placement strategies were defined in terms of two modes. In Selector mode, a function is sent to the runtime that has the least estimated end-to-end latency. On the other hand, in Duplicator mode, the same function is sent to both the edge and GPU-enabled cloud. The slower of the two executions is killed when the faster one completes. Evaluation of TensorFlow image recognition benchmarks led to a Atilde;sup3;3.3x reduction in total response time R1 over a single-tier baseline and placement accuracy of 92Acirc;pound; in Selector mode and up to 97Acirc;pound; in Duplicator mode.

We built models to estimate latency by using telemetry from previous data transfer, container deploy, and execution phases. We used camera trap images taken from wildlife to emulate the IoT connectivity connection. The findings show that hybrid scheduling improves the utilization of resources and reduces response time in heterogeneous cloud-edge infrastructures.

RELATED WORK

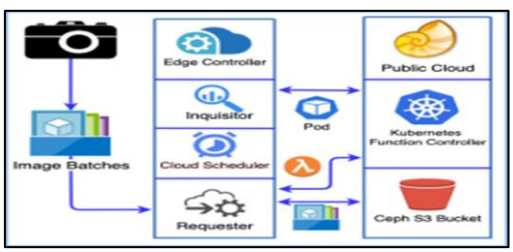


Fig. 1. Distributions of processing (T_p), transfer (T_t), and deployment (T_d) latencies for 150 GPU1 runs. Deployment time exhibits high variance and heavy tails.

Copyright © 2022 The Author(s): This work is licensed under a Creative Commons Attribution- NonCommercial-NoDerivatives 4.0 (CC BY-NC-ND 4.0) International License

***Corresponding Author:** Dilip Mehendra Sharma

DOI- https://doi.org/10.5281/zenodo.24253071

Prior investigations into low-latency, geo-distributed ana lytics and mobile-cloud offloading have informed the design of hybrid scheduling systems [Ananthanarayanan, Pu. Q. et al., 2015; Cuervo, E. et al., 2010]. Federated learning ap proaches distribute model training across edge devices without centralizing data [McMahan, B. et al., 2017], whereas hybrid frameworks such as STOIC offload computation across multi-tier infrastructures based on latency predictions.

Advances in serverless computing research have surveyed challenges and opportunities in performance optimization, hardware acceleration, and orchestration [Jonas, E. et al., 2019; Hellerstein, J. M. et al., 2018]. GPU vir tualization techniques for accelerating FaaS functions have been demonstrated in data center settings [Naranjo, D. M. et al., 2020]. Kubernetes based orchestration of containerized serverless platforms has been empirically evaluated, highlighting the flexibility and developer-friendly nature of such deployments [Mohanty, S. 2018].

Edge-to-cloud offloading strategies have been formalized using stochastic models and Markov decision processes to optimize average processing times under resource variabil ity [Lyu, B. et al., 2020; Paščinski, U. et al., 2018]. Blockchain-enabled fog architectures and trust management solutions have been proposed to secure IoT data sharing [Kochovski, P. et al., 2019; ?]. Complementary tools for tracing, profiling, and repairing distributed serverless workflows enhance observ ability and reliability in hybrid environments [Lin, W. T. et al., 2019; Shahrad, M. et al., 2019].

STOIC

Table 1 :Nautilus Hardware Heterogeneity Summary

Resource Type	Variant Count
CPU Architectures	44
GPU Models	9

As seen in Figure 1, the architecture of STOIC is quite modular, the orchestration of function placement at the edge and public cloud clusters being central to it. An edge con troller consumes batches of data from the IoT sensors that are physically located close to it. Using a database that contains historical telemetry, it learns the available runtimes and predicts end-to-end latency. A request handler puts data into an object store compatible with S3 and invokes serverless functions over HTTP APIs. A background metric collector instantiates much containerized runtimes periodically to record latency distribution.

Clusters at the edge and in the cloud were unified under Kubernetes and Kubeless orchestration. Eucalyptus offered AWS-compatible VM and object storage services at the edge. This edge cluster contained three Intel NUC nodes using quad-core Intel Core i7 CPUs with 32 GB RAM. The public cloud used Nautilus instances with GPU capabilities for faster TensorFlow training and inference [Borg, O. 2016].

Two scheduling modes were supported. In Selector mode, placement decisions relied exclusively on the minimal pre dicted latency. In Duplicator mode, executions took place concurrently, while the slower execution was terminated on completion by the faster tier. The two modes relied on analyt ical models of transfer, deployment, and execution latencies reaching 55% faster response time than the single tier exe cution. In addition, the enforced placement accuracy levels exceeded 95% over a 24-hour evaluation period.

Public/Private Cloud

We used Nautilus, a cloud platform that uses GPU and supports many institutions. Nautilus has 176 compute nodes and 543 GPUs in May 2020. The compute nodes, 543 GPUs and other components are interconnected via a multi-campus high-speed network [rook ceph block, 2020]. Kubernetes was used for orches trating Container workloads, and Rook integrated Ceph for Distributed Object Storage[Weil, S. et al., 2006]. The hardware makes use of 44 types of CPU architectures and 9 types of GPU models (Table I). The dispersion of resources causes the difference in performance that we see [23]. This is measured during the compute and I/O operations.

Variability comes from: (1) memory to GPU copy speed variation due to different CPU clock speeds; (2) object store access latency jitter; (3) multi-tenancy runtime contention. A method was used to limit variance in the network which pre vented all of the schedulers from adapting to the heterogeneity of hardware.

Copyright © 2022 The Author(s): This work is licensed under a Creative Commons Attribution- NonCommercial-NoDerivatives 4.0 (CC BY-NC-ND 4.0) International License

Publisher: SARC Publisher

Runtime Scenarios

Table 2: Deployment Configurations

Scenario	Identifier	Description
A	Edge	AVX2-enabled VM on local edgecluster
B	CPU	Single-CPU pod on Nautilus
C	GPU1	Single-GPU pod on Nautilus
D	GPU2	Dual-GPU pod on Nautilus

Four deployment configurations were defined (Table II): • A (Edge): AVX2-enabled VM on the local edge cluster. • B (CPU): Single-CPU Kubernetes pod on Nautilus. • C (GPU1): Single-GPU Kubernetes pod on Nautilus. • D (GPU2): Dual-GPU Kubernetes pod on Nautilus. Each configuration served as a candidate runtime for hybrid scheduling, allowing either manual parameterization or auto matic selection.

Execution Time Estimation

Total response time (T_s) was defined as: $T_s = T_t + T_d + T_p$ where T_t denotes data transfer time, T_d denotes container deployment time, and T_p denotes processing time.

1) Transfer Time (T_t): Transfer time was computed as the ratio of compressed batch size (F_b) to measured bandwidth (B_c): $T_t = F_b/B_c$ Bandwidth samples were captured by an edge-based monitoring agent.

2) Deployment Time (T_d): Deployment latency represented the interval required to instantiate a kubeless function on Nautilus. Historical latency series were analyzed via a median sliding window, which minimized mean absolute error com pared to moving averages and auto-regression. The inquisitor component logged T_d at one-minute intervals; window and calibration parameters were optimized using historical traces.

3) Processing Time (T_p): Processing time was forecast by Bayesian Ridge Regression on the most recent execution records, with RANSAC outlier filtering ensuring robustness against sporadic performance spikes. Regression inputs com prised the latest ten processing samples, enabling rapid adap tation to runtime fluctuations.

4) Adaptability: Predictive accuracy was validated on 50 GPU1 benchmark executions with 150-image batches. A marked reduction in percentage mean absolute error (PMAE) between initial and subsequent runs confirmed convergence of latency estimates.

Workload Generation

A synthetic workload generator was derived from 27,264 hours of camera-trap logs (July 2013–January 2017). Hourly image counts formed batch sizes, with nonzero activity in 9,125 hours (33.47%). A conditional ECDF of nonzero batch sizes was sampled to simulate arrival patterns, reproducing peak activity at 13:00 and 20:00 local time.

Implementation

The core components were designed using Go's Kubernetes client library. We used TensorFlow to modify the images to take advantage of the AVX2, SSE4.2 and FMA instructions. When you install the NVIDIA Container Toolkit, it installs CUDA 10.0 and cuDNN 7.0 by default. Kubeless provided a service that operates on Kubernetes clusters. This offered layer spanned both edge and cloud tiers. Rook/Ceph was the object storage backend.

Workflow

Two scheduling workflows were defined (Figure 1): • Selector Mode: Total response times (T_s) were estimated for all configurations; the runtime with minimal predicted latency was selected. Upon completion, latency metrics were recorded for subsequent scheduling decisions. • Duplicator Mode: Edge and the selected cloud runtime were invoked concurrently; the slower execution was terminated upon completion by the faster tier, mitigating deployment variability. Exponential back-off (starting at 100 ms, doubling on each retry) governed deployment retries, with a ten-attempt timeout.

EVALUATION

An extensive experimental campaign was conducted to quantify the latency reduction and resource-efficiency benefits of the STOIC scheduler under realistic camera-trap workloads. All results compare the scheduler's choice against the ground truth optimal execution venue among four configurations: Edge, CPU, GPU1, and GPU2.

Experimental Setup

A TensorFlow/Scikit-learn Convolutional Neural Network (CNN) was trained to perform five-class inference (Bird, Fox, Rodent, Human, Empty) on wildlife images drawn from the "Where's The Bear" system. Class imbalance was mitigated via in-batch up-sampling. All input images were

Copyright © 2022 The Author(s): This work is licensed under a Creative Commons Attribution- NonCommercial-NoDerivatives 4.0 (CC BY-NC-ND 4.0) International License

Publisher: SARC Publisher

resized to 1920×1080 pixels, and the CNN was retrained on 251 samples per class. The resulting HDF5 model was staged in Ceph object stores at both the edge and the Nautilus cloud.

A 27 264-hour camera-trap trace (July 2013–Jan 2017) was used to generate 162 synthetic batches over a 24 h evaluation period. A conditional ECDF of nonzero hourly batch sizes (max = 2450, mean = 25) produced realistic inter-arrival patterns. For each batch, all four runtimes were invoked to establish the true minimum latency; STOIC's selector and duplicator modes were then assessed on their ability to match or improve upon that baseline.

Latency Component Analysis

Component-level profiling (Figure 2) indicated that transfer and processing times remained tightly clustered (CV ¡ 0.2), whereas deployment latency showed heavy-tailed behavior (CV ¿ 0.6). These observations explain the primary source of selector-mode mispredictions at batch sizes where edge and GPU latencies converge (35 and 90 images).

Selector Mode Results

A 92 % correct-selection rate was observed (149/162 batches). The cumulative response time (10 770 s) remained within 7.5 % of the ideal oracle scheduler and delivered a 3.33× improvement over the naive worst-case static choice.

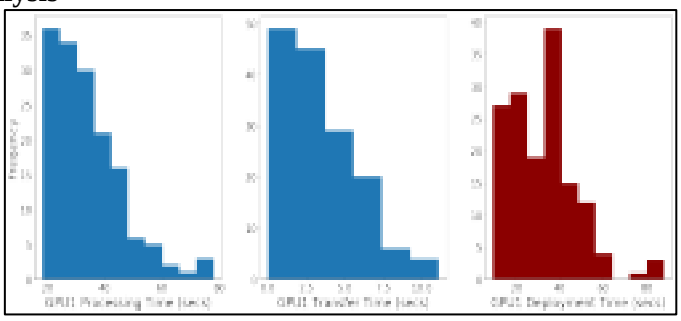


Fig. 2. Distributions of processing (T_p), transfer (T_t), and deployment (T_d) latencies for 150 GPU1 runs. Deployment time exhibits high variance and heavy tails.

Table 3: Selector Mode Performance (162 Batches)

	MIN-LAT	STOIC	Worst-Case
Aggregate Time (s)	10 022	10 770	35 940
Success Rate	100 %	92 %	0 %
Speedup vs Worst-Case	—	3.33×	—

Duplicator Mode Results

Analysis of the 13 mispredictions revealed clustering at the two aforementioned crossover batch sizes, where the predicted latency difference fell below the model's intrinsic MAE. By speculatively invoking both edge and GPU runtimes and canceling the slower, the duplicator mode achieved success rates above 97 %. The dual-GPU variant yielded the lowest end-to-end latency (10 223 s), within 2 % of the oracle, while the single-GPU variant improved safety of selection at a marginal 1 % overhead.

Resource Efficiency

Speculative edge execution reclaimed over 2 000 s of shared-cloud GPU time per day, while wasted GPU cycles remained under 1 000 s. Net positive savings demonstrate that duplicator mode enforces parsimonious use of scarce acceler ation resources without compromising latency objectives.

DISCUSSION

Overall, selector mode sufficed when deployment variability was moderate, offering substantial speedups with minimal overhead. Duplicator mode provided an added safety net in highly variable environments, sacrificing up to 2 % addi tional execution time to guarantee near-optimal decisions and maximize cloud resource efficiency. These results validate the scheduler's adaptability across heterogeneous edge–cloud infrastructures.

Table 4: Duplicator Mode Summary

Config.	Success Rate	Aggregate Time (s)	% of MIN-LAT
Edge + GPU1	98.8 %	10 122	101.0 %
Edge + GPU2	97.2 %	10 223	102.0 %

Copyright © 2022 The Author(s): This work is licensed under a Creative Commons Attribution- NonCommercial-NoDerivatives 4.0 (CC BY-NC-ND 4.0) International License

Publisher: SARC Publisher

Table 5 :Shared-Cloud Time Saved Vs. Wasted (24 H)

Config.	Saved (s)	Wasted (s)	Net (s)
Edge + GPU1 Edge + GPU2	+2 300 +2 600	− 450 – 700	+1 850 +1 900

CONCLUSION

We have presented a serverless scheduling framework for hybrid edge-cloud inference and training workloads. The system sends out jobs to where the job will execute with the most optimized latency. All this is possible with on-premise edge controllers and a GPU-backed public cloud tier. The two modes, Selector and Duplicator, have been tested under realistic camera-trap workloads. Selector chooses the single lowest-latency tier.

Test results show that adaptive placement: 3.3× speedup over static single-tier execution with over 90 % placement accuracy. Also, Duplicator mode is shown to recover more shared-cloud GPU time than it consumes thereby confirming its efficient use of scarce acceleration resources. The findings support that using scheduling in IoT-enabled machine learning is useful.

Further developments will look into non-linear regression as Gradient Boosted Regression Trees for processing-time estimation across heterogeneous hardware. Furthermore, we will look into incremental model checkpointing to reduce duplicate computing and make hybrid deployments faster. When combined, these extensions will further reinforce the framework's feasibility and efficacy in dynamic edge–cloud environments.

REFERENCES

1. "(AWS Lambda)," https://aws.amazon.com/lambda/, (2020).
2. "(Serverless)," https://www.serverless.com/ (2020).
3. "(Azure Functions)," https://azure.microsoft.com/en us/services/functions/, (2020).
4. "(IBM OpenWhisk)," https://developer.ibm.com/openwhisk/, (2020).
5. "(iot edge: Microsoft azure)," https://azure.microsoft.com/en us/services/iot-edge/.
6. "(iot hub: Microsoft azure)," https://azure.microsoft.com/en us/services/iot-hub/.
7. "(GreenGrass and IoT Core)," https://aws.amazon.com/iot core,greengrass/, (2019).
8. Pu, Q., Ananthanarayanan, G., Bodik, P., Kandula, S., Akella, A., Bahl, P., & Stoica, I. "Low latency geo-distributed data analytics." *ACM SIGCOMM Computer Communication Review* 45.4 (2015): 421-434.
9. Cuervo, E., Balasubramanian, A., Cho, D. K., Wolman, A., Saroiu, S., Chandra, R., & Bahl, P. "Maui: making smartphones last longer with code offload." *Proceedings of the 8th international conference on Mobile systems, applications, and services*. (2010).
10. McMahan, B. and Ramage, D. "Federated learning: Collaborative machine learning without centralized training data," (2017).
11. Jonas, E., Schleier-Smith, J., Sreekanti, V., Tsai, C. C., Khandelwal, A., Pu, Q., & Patterson, D. A. "Cloud programming simplified: A berkeley view on serverless computing." *arXiv preprint arXiv:1902.03383* (2019).
12. Hellerstein, J. M., Faleiro, J., Gonzalez, J. E., Schleier-Smith, J., Sreekanti, V., Tumanov, A., & Wu, C. "Serverless computing: One step forward, two steps back." *arXiv preprint arXiv:1812.03651* (2018).
13. Naranjo, D. M., Risco, S., de Alfonso, C., Pérez, A., Blanquer, I., & Moltó, G. "Accelerated serverless computing based on GPU virtualization." *Journal of Parallel and Distributed Computing* 139 (2020): 32-42.
14. Mohanty, S. "Evaluation of serverless computing frameworks based on kubernetes." (2018).
15. Lyu, B., Hong, Y., Tan, H., Han, Z., & Wang, R. "Cooperative jobs dispatching in edge computing network with unpredictable uploading delay." *Journal of Communications and Information Networks* 5.1 (2020): 75-85.
16. Paščinski, U., Trnkoczy, J., Stankovski, V., Cigale, M., & Gec, S. "QoS-Aware Orchestration of Network Intensive Software Utilities within Software Defined Data Centres: An Architecture and Implementation of a Global Cluster Manager." *Journal of Grid Computing* 16.1 (2018): 85-112.
17. Kochovski, P., Gec, S., Stankovski, V., Bajec, M., & Drobintsev, P. D. "Trust management in a blockchain based fog computing platform with trustless smart oracles." *Future Generation Computer Systems* 101 (2019): 747-759.

Copyright © 2022 The Author(s): This work is licensed under a Creative Commons Attribution- NonCommercial-NoDerivatives 4.0 (CC BY-NC-ND 4.0) International License

Publisher: SARC Publisher

18. Lin, W. T., Bakir, F., Krintz, C., Wolski, R., & Mock, M. "Data repair for distributed, event-based IoT applications." *Proceedings of the 13th ACM International Conference on Distributed and Event-based Systems*. (2019).
19. Shahrad, M., Balkind, J., & Wentzlaff, D. "Architectural implications of function-as-a-service computing." *Proceedings of the 52nd annual IEEE/ACM international symposium on microarchitecture*. (2019).
20. Borg, O. "Kubernetes/B. Burns, B. Grant, D. Oppenheimer (et al.). DOI: 10.1145/2890784." *Communications of the ACM* 59.5 (2016): 50-57.
21. "(rook ceph block)," https://rook.io/docs/rook/v1.0/ceph-block.html (2020).
22. Weil, S., Brandt, S. A., Miller, E. L., Long, D. D., & Maltzahn, C. "Ceph: A scalable, high-performance distributed file system." *Proceedings of the 7th Conference on Operating Systems Design and Implementation (OSDI'06)*. (2006).

Source of support: Nil; **Conflict of interest:** Nil.

Cite this article as:
Sharma, D. M. "An Adaptive Hybrid Cloud Framework for Intelligent Resource Scheduling in IoT and AI Applications" *Sarcouncil Journal of Engineering and Computer Sciences* 5.8 (2026): pp 11-16.

Copyright © 2022 The Author(s): This work is licensed under a Creative Commons Attribution- NonCommercial-NoDerivatives 4.0 (CC BY-NC-ND 4.0) International License

Publisher: SARC Publisher