SPIRT: A Fault-Tolerant and Reliable Peer-to-Peer Serverless ML Training Architecture

Amine BARRAK*, Mayssa JAZIRI*, Ranim TRABELSI*, Fehmi JAAFAR*, Fabio PETRILLO†

*Department of Computer Science and Mathematics, University of Quebec at Chicoutimi, UQAC, Saguenay, Canada
Email: {mabarrak, mjaziri, ranim.trabelsi1, fehmi.jaafar}@uqac.ca
†Département de génie logiciel, École de Technologie Supérieure, ÉTS, Montreal, QC
Email: fabio.petrillo@etsmtl.ca

Abstract—The advent of serverless computing has ushered in notable advancements in distributed machine learning, particularly within parameter server-based architectures. Yet, the integration of serverless features within peer-to-peer (P2P) distributed networks remains largely uncharted. In this paper, we introduce SPIRT, a fault-tolerant, reliable, and secure serverless P2P ML training architecture. designed to bridge this existing gap.

Capitalizing on the inherent robustness and reliability innate to P2P systems, SPIRT employs RedisAI for in-database operations, leading to an 82% reduction in the time required for model updates and gradient averaging across a variety of models and batch sizes. This architecture showcases resilience against peer failures and adeptly manages the integration of new peers, thereby highlighting its fault-tolerant characteristics and scalability. Furthermore, SPIRT ensures secure communication between peers, enhancing the reliability of distributed machine learning tasks. Even in the face of Byzantine attacks, the system's robust aggregation algorithms maintain high levels of accuracy. These findings illuminate the promising potential of serverless architectures in P2P distributed machine learning, offering a significant stride towards the development of more efficient, scalable, and resilient applications.

Keywords–Distributed Machine Learning, Peer-to-Peer (P2P), Serverless Computing, Fault Tolerance, Robust Aggregation.

I. INTRODUCTION

Distributed machine learning (ML) has become increasingly important in tackling the rising complexity and large amounts of data associated with ML models, which traditional single-machine learning frameworks struggle to manage. This approach harnesses multiple computational nodes for simultaneous processing, greatly reducing the time necessary to train larger, more intricate models[1].

Multiple distributed ML architectures have been introduced over the years, many of which are fundamentally based on the structures of the Parameter Server (PS) and Peer-to-Peer (P2P) architectures [2]. Each of these represents a unique methodology for orchestrating the management and distribution of tasks and data across nodes in a distributed system, with their own set of benefits and challenges [3], [4].

In the parameter server architecture, for instance, the worker nodes perform computations on their respective data partitions and communicate with the parameter server (PS) to update the global model [5]. In contrast, peer-to-peer (P2P) architectures distribute the model parameters and computation across all nodes in the network, eliminating the need for a central coordinator [6].

In the pursuit of optimizing distributed ML architectures, serverless computing has emerged as a revolutionary development. Today, the growing trend of serverless computing, including platforms such as Amazon Lambda[7], Google Cloud Functions[8], and Azure Functions[9], offers flexibility and cost-effective solutions for distributed ML systems[10], [11], [12]. A common trait among existing serverless architectures—whether based on Parameter Server or P2P setups—is a heavy reliance on databases as a communication channel. This necessity arises from the stateless nature of serverless architectures and limitations in directly transmitting large data sizes[13], [14]. As a result, communication latency can arise, especially during the iterative process of model update and gradient aggregation[14], [15]. In fact, in these architectures, the model parameters and computed gradients must be fetched from the database, processed, and re-uploaded in each iteration. This recurring cycle imposes a significant burden, leading to additional computational and communication costs that can significantly impact overall efficiency.

Notably, while the usage of serverless computing within parameter server architecture has been extensively studied and proven to yield benefits like cost reduction [16], scalability [17], [18], and performance efficiency [19]. Surprisingly, however, the potential of serverless computing within peer-to-peer (P2P) architectures remains less explored, despite the immense advantages P2P offers in terms of data privacy, load balancing, and fault tolerance, as a single node's failure does not disrupt the system [20].

In relation to this, our preceding work [21] explored the integration of serverless computing within VM-based P2P architecture, aiming to offload excessive gradient computation when faced with resource constraints. Despite the benefits of this strategy, it represents a piecemeal approach that does not fully harness the potential of a completely serverless system. Adding to these challenges, security is a significant concern in distributed Machine Learning (ML) environments, particularly in Peer-to-Peer (P2P) based architectures [22]. One issue is the absence of stringent mechanisms to authenticate new peers

joining the network [23], [24]. Unauthenticated or malicious peers could join the network, potentially causing significant disruption to the training process.

Beyond the threat of unauthenticated new peers, even existing, trusted peers can become compromised, introducing and contributing malicious gradients intended to sway the model training in harmful directions, severely compromising the model's integrity and performance [25], [23], [26]. This risk emphasizes the importance of robust aggregation methods [27], [28], [29], which can help mitigate the effect of such outliers on the overall model's performance.

Addressing the outlined challenges, we introduce the Serverless Peer Integrated for Robust Training (SPIRT) Architecture. In this setup, each peer is identified by its associated database, e.g., Redis, and executes shard-based gradient computation, aggregation, and model convergence checks. With orchestrated workflow coordination, e.g., AWS Step Functions, SPIRT effectively overcomes the stateless limitations inherent in serverless computing. Consequently, it takes full advantage of the serverless architecture within a peer-to-peer setup for machine learning training.

In this endeavor, we propose a reliable peer to peer, serverless ML training architecture with the following contributions:

1) Automated ML training Workflow: We propose an automated ML training serverless workflow within a P2P architecture, maximizing efficiency and scalability while reducing operational complexities.
2) Serverless-adapted Redis Enhancement: We modify RedisAI for serverless ML training systems that rely on databases, introducing in-database model updates to eliminate the traditional fetch-process-reupload cycle. This enhancement streamlines communication and reduces overhead, facilitating more efficient serverless ML operations.
3) Enhanced Fault Tolerance: We integrate secure, scalable peer-to-peer communication, new participant integration, and robust aggregation mechanisms to establish a fault-tolerant environment. This approach ensures reliable distributed ML training, even in the presence of potential disruptions or adversarial actors.

Once the paper is accepted, a replication package will be made available for further research.

II. Related work

This section reviews literature on distributed machine learning and serverless computing, focusing on decentralized ML on P2P paradigm and fault tolerance mechanisms.

1. Distributed Machine Learning Architectures

Distributed machine learning has gained prominence due to the intense computational requirements and extensive data involved in training sophisticated ML models [30], [2], [31]. The choice of topology, either Centralized, Tree-like or P2P, is a critical factor influencing performance, scalability, reliability, and security [32], [2], [33], [34]. P2P architectures, for instance, offer scalability and reliability by eliminating the "Single Point of Failure". Recent studies have begun to explore serverless computing for machine learning, a promising yet nascent field [10], [12], [35], [17], [13]. While many works [36], [13], [37], [14], [18], [38] are based on the PS paradigm, others have started to present decentralized learning solutions using serverless computing [21].

Unlike previous studies, we propose a fully automated ML training serverless workflow within a P2P architecture.

2. Communication and Model Updating in Serverless based Distributed ML

Serverless computing coupled with database systems has seen increased application in the field of distributed learning. Works such as LambdaML [14], SMLT [17], and MLLess [16] utilize a database as a communication channel, collecting intermediate local statistics from all workers through a designated communication pattern. Other works propose in-database distributed machine learning solutions to reduce communication overhead [39]. Additionally, libraries such as RedisAI [40] enable seamless inference on trained models within Redis, and other solutions like OML4R [41] and MADlib [42] offer SQL-based functions for prediction and inference within database environments.

Differentiating from these works, our architecture utilizes Redis, a high-performance in-memory data structure store, serving as both a communication channel and an automatic model update tool through RedisAI.

3. Security and Robustness in Distributed ML

Fault tolerance in distributed ML is facilitated by techniques like Availability Zones, retries, architectural decisions, and checkpointing mechanisms [43], [44], [45], [13], [14]. The heartbeat technique is also used for failure detection [46]. Several studies also propose securing communication within distributed systems to prevent data leakage [18], [47], [23], [48]. In addition, robust aggregation techniques such as KRUM [27], MULTI-KRUM [27], GeoMed [28], MarMed [28], and ZENO [29] are used to address Byzantine faults [49]. These robust aggregations rules have been adapted to P2P architectures in works like the BRIDGE framework [26] and a blockchain-based solution by Xu *et al.* [23].

We expand on this body of work by implementing a secure, scalable peer-to-peer communication, new participant integration into the ML training network, and incorporating robust aggregation mechanisms for reliable distributed ML.

III. Design Architecture of Logical Peer to Peer training ML

In this section, we delve into the design architecture of a serverless, peer-to-peer machine learning training system. The discussion encompasses a broad overview of the proposed architecture, an in-depth examination of the core components, and an exploration of the operational dynamics driving the system's overall functionality and performance.
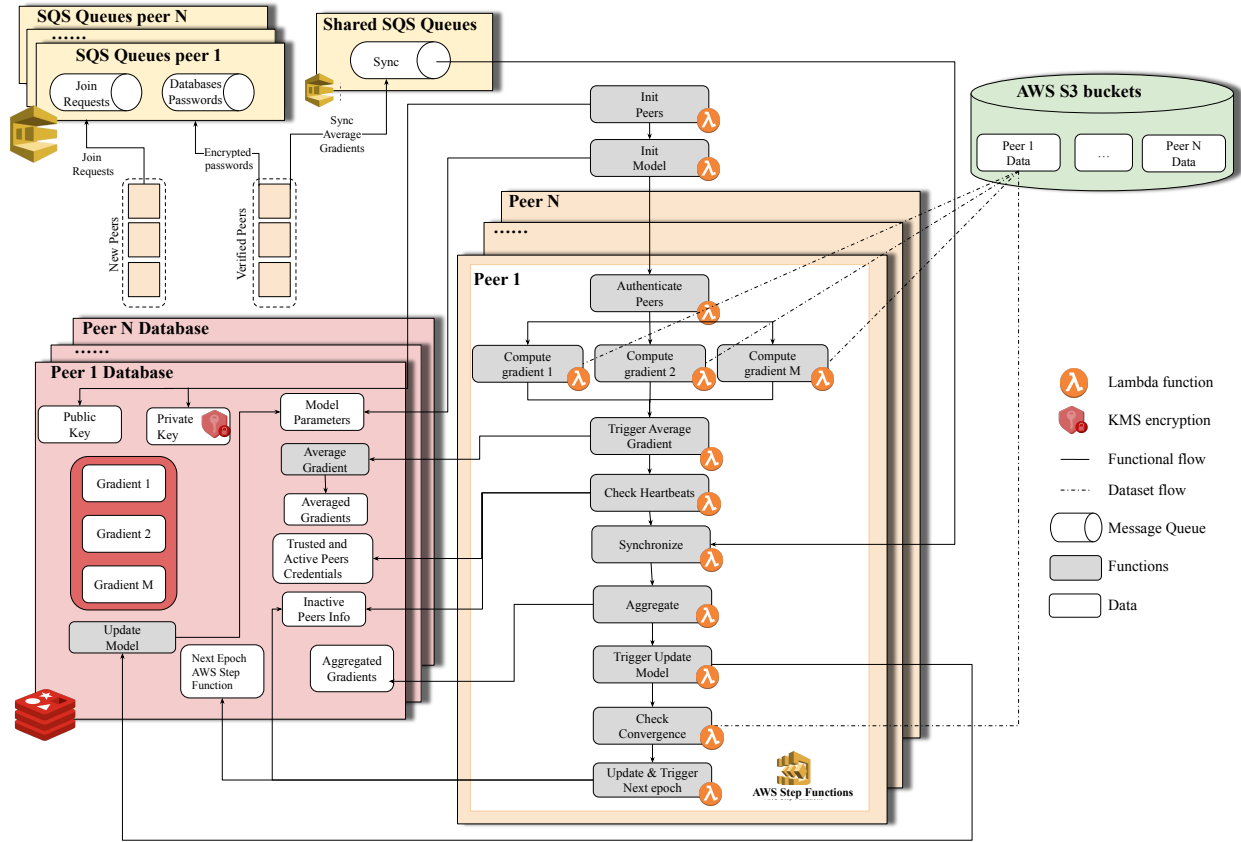
Figure 1. Overview of the proposed Peer To Peer training based on Serverless computing

1. Comprehensive Overview of the Proposed Architecture

Our proposed serverless P2P distributed training architecture, as illustrated in Figure 1, kicks off with system initialization, followed by the authentication of new peers, if any. In this configuration, every peer in the network is uniquely identified by the IP address and port tied to its corresponding stateful component - a dedicated Redis database. A heartbeat monitoring system ensures the constant availability of all peers.

Within the system, the assigned dataset of each peer is divided into smaller shards (batches). The peer computes gradient for each shard, averages them, and stores the result in its Redis database. These averaged gradients are then collectively aggregated among all peers, filtering out any outliers in the process by applying robust aggregation. The resultant set of trusted, aggregated gradients is used to update the model parameters. The system performs periodic checks for model convergence *i.e.,*every ten epochs, assuring optimal progress during the learning phase. Data integrity and confidentiality are ensured through secure peer communication.

The entire process is coordinated using AWS Step Functions, which seamlessly orchestrates the flow of each epoch within the training process of each peer.

This architecture is deployed on Amazon AWS for its unique benefits like the 15-minute timeout and 10GB RAM from AWS Lambda [50]. Notably, comparable services on platforms like Google Cloud, Azure, and IBM Cloud enable possible

architecture replication.

2. Deep Dive into Core Architectural Components

Following the initial overview, we will now go over each key facet of our proposed Peer-to-Peer (P2P) serverless architecture.

1) Training Dataset Management and Partitioning

Individual peers have the ability to pull data from multiple distributed storage systems, including Amazon S3. The specific data each peer is responsible for is determined by its unique rank. This data is then divided into smaller units, or shards, to enable batch processing.

2) Leveraging Serverless Computing Across Peer Training Tasks

The cornerstone of our architecture is serverless computing, embodied by Amazon Lambda functions. Incorporated throughout the peer training workflow—from peer authentication to model updates—it offers benefits such as isolation for uninterrupted operations, capacity for managing compute-intensive tasks like gradient computations, and scalability for workload fluctuations.

3) Training Workflow Orchestration

Due to the serverless structure of our system, where functions operate independently, their orchestration is crucial.

To manage this, we use AWS Step Functions, a powerful serverless workflow service. It coordinates the entire machine learning training process within each peer during each epoch, including tasks like peers authentication, gradient computation and averaging, model updates, and convergence assessment. Notably, our architecture integrate continuous invocation of step functions for each epoch, which effectively mitigates AWS Lambda's cold start delays, thereby enhancing overall performance and minimizing latency during ML training.

4) State Management and Processing in Database

In our architecture, we use Redis, an open-source in-memory data store, for quick access to machine learning artifacts such as model parameters and gradients - a key requirement for any stateless distributed ML system.

Beyond a simple key-value store, we utilize the RedisAI module which supports various deep learning backends, enabling in-database ML operations, and minimizing data transfer latency. RedisAI is especially efficient at serving models at scale and in real-time. Unique to our architecture is the extension of RedisAI's capabilities to directly modify model parameters within the database, eliminating the traditional process of external processing, our routine performs these operations inside the database itself.

5) Synchronization between Peers

Within our proposed architecture, achieving synchronization amongst peers is paramount for ensuring the correctness of the distributed training process. To manage this aspect of distributed computation, we employ the AWS Simple Queue Service (SQS).

Once a peer completes gradient computation for its data shards and averages local gradients, it sends a notification message to a designated synchronization queue, the "Sync Queue", signifying the task completion. If a peer doesn't respond or acknowledge within a designated timeout period, others proceed without waiting indefinitely. The unresponsive peer is identified as a failed node in the next epoch by our heartbeat monitoring system.

We note that the messages inside the "sync queue" will be deleted by any peer in initialisation phase.

6) Secure Communication: Safeguarding Data Integrity and Confidentiality

Our architecture employs stringent secure communication protocols to ensure data integrity and confidentiality during inter-peer interactions, accomplished through the RSA algorithm for asymmetric encryption, with unique public and private keys for each peer. Beyond encryption, unique digital signatures derived from private keys authenticate sender identity and verify data integrity.

Each peer's private key is safeguarded by encryption using a unique key from AWS Key Management Service (KMS), with access strictly limited to few authorized services (Lambda functions), enhancing security against unauthorized access.

3. Operational Dynamics of the Proposed Architecture

Within this subsection, we take a closer look on the operational dynamics that power our proposed architecture. Moving beyond the standalone examination of the key components, to their interactions, collaborative processes, and the mechanisms that drive the overall system performance and functionality.

1) Peers Initialization and Authentication

Our architecture relies on Amazon SQS for two main operations: *peers' initialization* and *new peers' integration*. Each peer has two distinct SQS queues - the first for join requests and the second for receiving encrypted passwords of other peers' databases. Every peer also maintains an AWS Key Management Service (KMS) encryption key, securing its private key within its database, and ensuring exclusive key access.
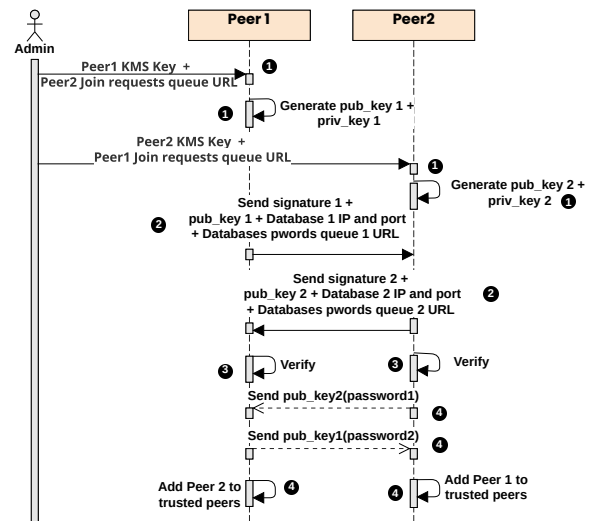
Peers Initialisation

Figure 2. Sequence Diagram Illustrating the Initialization Process of Peers

This phase involves the setup of individual peers. The process below outlines an example of initializing a two-peer system, and it's also illustrated in Figure 2:

1) At the onset, the admin initiates the peers, by providing each their own KMS encryption key, URLs of their neighboring peers' "join requests" SQS queues, and a unique rank. Each peer then generates private and public keys. The public key is stored in its plain form, while the private key is encrypted using the KMS encryption key before storage in the respective database.

2) Each peer then generates a digital signature and broadcasts this, along with its public key, database IP address and port, and the URL of its "databases passwords" queue inside other peer's "join requests" queue.

3) Verification ensues as each peer validates the other's signature

4) Upon successful verification, peers mutually exchange their encrypted database passwords. Furthermore, they save each other details, including the rank, into their respective databases.

Novel peer integration In this phase, we illustrate the process of integrating a new peer into an existing training network of two peers, as shown in Figure 3. The steps are as follows:

1) The process is initiated when the admin provisions the new peer (Peer 3) with the URLs of the "join requests" SQS queues of existing peers and their corresponding public keys, and their ranks. Peer 3, in response, generates a pair of public and private keys. These keys are stored in its database, with the private key encrypted using the designated KMS encryption key.
2) The new peer (Peer 3) generates a digital signature, broadcasting it and its public key, database port and IP address, URL of its "databases passwords" queue, and rank. It sends its password, encrypted with the recipient's public key, via their "join request" queues.
3) Once the broadcasting phase is completed, Peer 3 waits for validation from the existing peers. These peers undertake the task of validating Peer 3's authenticity. They accomplish this by rigorously comparing the signature provided by Peer 3 against the information contained in its public key, ensuring a match.
4) After Peer 3's successful validation, Peers 1 and 2 send back their individual signatures, and their databases passwords encrypted with Peer 3's public key, into Peer 3's "databases passwords" queue. Furthermore, they incorporate Peer 3's details into their respective databases.
5) Finally, Peer 3 validates the sender peers based on their signatures and public keys. Upon successful validation, Peer 3 records the details of Peers 1 and 2 into its own database.
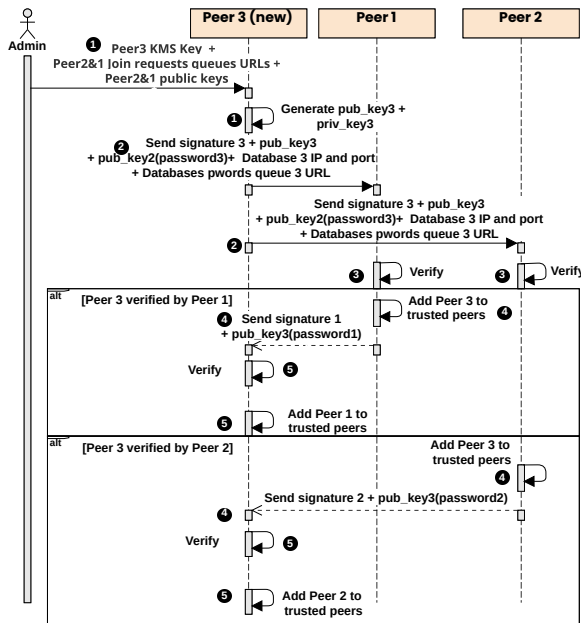

Figure 3. Sequence Diagram of New Peer Integration into Training Network

2) Model Initialisation

The model initialization phase involves the establishment of a unified model to serve all initialized peers. This model can be initialized with random parameters or pre-trained ones, depending on the requirements. The chosen model, which could be a specific ML or deep learning model, is then stored in each peer's Redis database using RedisAI. This ensures a consistent starting point for distributed learning.

3) Distributed Gradient Computation

Leveraging the serverless concurrent abilities of AWS Lambda, we implement a parallel gradient computation within our architecture. Each peer partakes in this distributed process by calculating gradients on assigned data batches and storing these computed gradients in its local Redis database. To expedite the gradient computation, data, segmented into smaller batches, is fetched from S3 storage, and model parameters are retrieved from Redis.

4) Averaged Gradient Computation

Once the gradients have been computed, an embedded Lua script calculates the gradients' average within the Redis database environment, capitalizing on in-database programming. This approach eliminates costly external data transfers. Once local averaging is complete, peers send a completion message to the "sync queue", notifying others of the task's conclusion.

5) Heartbeat Monitoring

Our system incorporates a 'heartbeat' mechanism designed to be triggered every epoch, where each peer checks the operational status of other peers databases.
Peers send a signal, waiting for responses to confirm others' activeness. Failure to respond within a set timeframe and a number of trials denotes a peer as "inactive", subsequently removed from the trusted peers list and added to the "inactive peers info" list. This continuous health check is vital for network integrity and uninterrupted peer-to-peer communication.

6) Peers Synchronization

To collate the individually computed average gradients from each participating peer and derive the aggregated gradient, a specially designated Lambda function "synchronize" is instantiated. This function serves as a synchronization barrier that waits until the count of messages in the queue equals the current number of active peers, those determined by the preceding heartbeat check.

7) Gradient Aggregation

Once all peers are synchronized, the gradient aggregation phase begins. Each peer fetches the average gradients from the databases of all active peers in the network. An aggregation function then amasses these gradients, utilizing robust algorithms to discard outlier gradients. The final aggregated gradient is then stored in each peer's Redis database. This approach ensures the integrity and accuracy of our gradient aggregation process.

8) Model Update

Utilizing the in-database programming feature of RedisAI, we directly update each peer's model parameters stored in the Redis database using the aggregate gradient, thereby bypassing the conventional read-process-write cycle.

9) Convergence Checking

Upon the completion of model parameter updates, a Lambda function is intermittently called (e.g., after every tenth iteration) to check model convergence. This approach reduces unnecessary function calls, as significant model changes aren't anticipated after each iteration.

10) Update and Trigger new epoch

In the context of our AWS Step Function based workflow, each step function orchestrates the process for a single epoch. The stateful nature of this service necessitates the instantiation of a new Step Function at the end of each epoch to maintain the continuity of training inside each peer. To facilitate this, we employ a dedicated Lambda function that is responsible for spawning the new Step Function. This Lambda function initializes the new Step Function with the correct inputs, including the subsequent epoch number, the degree of parallelism for gradient computation tasks, and a flag indicating whether a convergence check is required (based on the current epoch number). This process ensures an updated step function based on current needs. The ARN (Amazon Resource Name) of the new function is stored in the Redis database, which is reused in subsequent epochs if no new requirements arise.

The same Lambda function also undertakes the responsibility of updating the list of inactive peers through a consensus-driven approach. To ensure the integrity of this process, the list of inactive peers for each node is cross-validated against corresponding lists from all other nodes within the network. An inclusive agreement is applied, meaning a peer is only marked as inactive if it is listed as such in every peer's record.

11) Fault Tolerance and Peer Data Redistribution

A core strength of our architecture lies in its fault-tolerance mechanisms. This resilience is primarily manifested through our approach to handle instances of peer inactivity or failure. When one or many peers become inactive, our dedicated Lambda function first identifies these instances using the consensus-based approach described previously. Once inactive peers are identified, the data originally assigned to the downed peers is segmented and distributed among the active peers based on a predefined ranking system. Here, each peer, according to their rank, inherits a corresponding portion of the data from the inactive peer.

Following this data reassignment, the "Update and Trigger new epoch" Lambda function adjusts the configuration of the new Step Function to account for the change in data distribution and the increased workload for each active peer. This adjustment may include tuning the degree of parallelism for gradient computation tasks to accommodate the additional data batches.

IV. EXPERIMENTAL SETUP

In order to evaluate the efficacy of the Serverless Peer Integrated for Robust Training (SPIRT) architecture, we design a series of experiments to evaluate the performance of various CNN models across different datasets on the proposed architectures.

1. Datasets

We utilized two public datasets for our experiments:
MNIST: The MNIST Handwritten Digit Collection [51] consists of 60,000 handwritten digit samples, each belonging to one of ten classes.

2. Model Architectures and Hyperparameters

The experiments involve three different CNN models:
MobileNet V3 Small: A lightweight CNN developed for mobile and edge devices, it features inverted residual blocks, linear bottlenecks, and squeeze-and-excitation modules, with roughly 2.5 million trainable parameters [52].
ResNet-18: A deep learning CNN model with approximately 11.7 million parameters, featuring 18 layers and using "skip connections" to aid training of deeper networks [53].
DenseNet-121: A uniquely structured CNN with about 8 million parameters and 121 layers. It leverages dense connections, where each layer is connected to all other layers, enhancing learning efficiency [54].

3. Redis and RedisAI Configuration

We utilized Redis and RedisAI for in-database model updates. Each peer's Redis was deployed within an Amazon EC2 R6a.large instance. To access the database hosted on EC2, we utilized SSH tunnels to forward the port and enable access from public services such as AWS Lambda.

4. AWS Lambda Configuration

Several AWS Lambda functions, discussed in Section III, were set up for the training procedure. Dependencies such as PyTorch, NumPy, Redis, RedisAI, and sshtunnel were necessary. Managing these dependencies posed a challenge due to AWS Lambda's deployment package size limit. The zipped package must be less than 50MB, and the unzipped files cannot exceed 250MB.

We utilized a Virtual Private Cloud (VPC) to minimize the need for multiple SSH connections between compute gradients Lambda functions and the database EC2 instance.

V. EFFICIENCY AND PERFORMANCE IN SERVERLESS P2P ARCHITECTURE

1. Motivation:

Our innovative serverless peer-to-peer architecture marks the advent of a decentralized era in model training. Peers are enabled to concurrently compute gradients in parallel based on their assigned datasets, exchange these gradients for robust aggregation, and update their model parameters. Yet, a theoretically sound architecture is merely the initial phase; the critical next step is to assess its performance under various
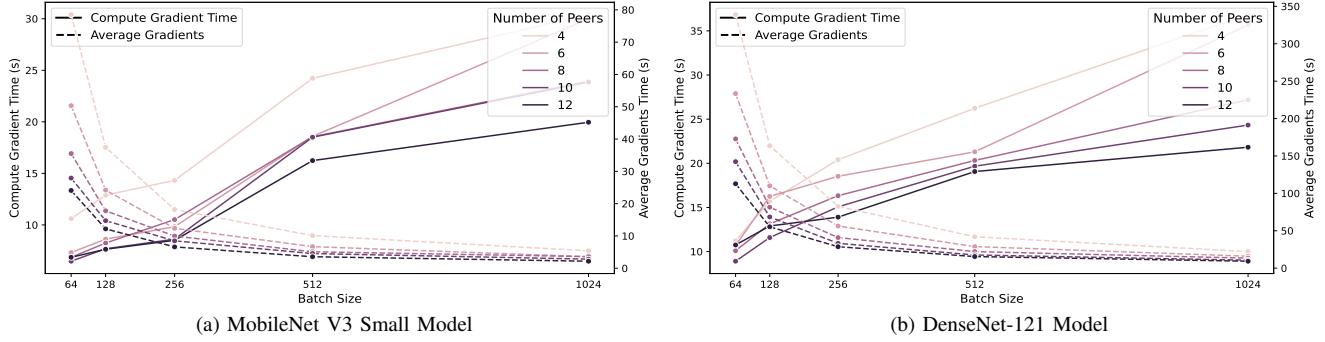
Figure 4. Comparison of Compute Gradient Time and Average Gradients Across Varying Batch Sizes for Different Number of Peers

operational conditions using multiple deep learning models of different complexity and computational demands.

Moreover, it's crucial to explore scalability in two significant aspects: *(1) Intra-peer scalability:* The ability of our architecture to manage parallel gradient computations within a single peer, and *(2) Inter-peer scalability:* The adaptability of our architecture to changing numbers of peers in the system.

Our motivation lies in thoroughly understanding our architecture's scalability under these two critical dimensions. By altering the number of peers and harnessing each peer's parallelization capabilities, we can observe their impacts on performance metrics like gradient computation time, aggregation time, and overall training time per epoch. This approach gives us a holistic view of our architecture's performance under varied operational circumstances.
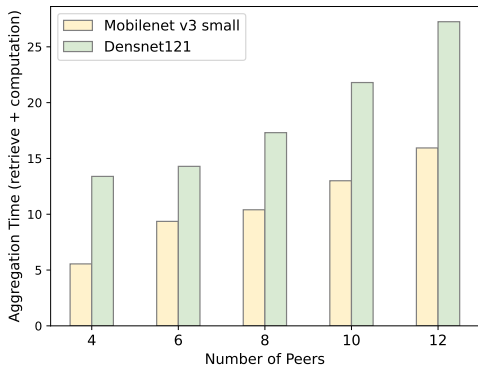

Figure 5. Comparison of Average Aggregation Times for Mobilenet v3 small and Densnet121 Models Against # of Peer

2. Approach:

For our experiments, we utilized two distinct deep learning models, Densenet121 and MobileNet V3 Small, to represent a range of model complexities, using the standard MNIST dataset for training.

Intra-peer scalability was explored by adjusting the batch size, thereby determining the number of concurrent gradient computations within a peer. With batch sizes of 64, 128, 256, 512, and 1024, we traced the impact on gradient computation and averaging times.

Inter-peer scalability was assessed by varying the number of peers, starting from 4 and incrementing by 2, up to 12. This

helped us gauge the system's adaptability and its effect on parameters like aggregation time and training time per epoch.

3. results:

As we observed in Figure 4, an increase in batch size yields an increased time to compute gradients. This is expected, as a larger batch size implies more data to process at once per peer, which intuitively would increase the time needed for computation. We observed this increase in computation time consistently across both Densenet121 and MobileNet V3 Small models, signifying that the effect is model-agnostic. Additionally, this trend persists across different numbers of peers. This indicates that the increase in computational burden due to a larger batch size cannot be offset simply by adding more peers, as the computation time per gradient computation still increases.

Table I
TRAINING TIME PER EPOCH ACROSS DIFFERENT BATCH SIZES AND PEER COUNTS FOR MOBILENET V3 SMALL AND DENSENET121 MODELS

# Peers	Mobilenet v3 small (Batch Size)					Densenet121 (Batch Size)				
	64	128	256	512	1024	64	128	256	512	1024
4	96.27	57.46	39.61	41.41	42.85	375.82	202.62	125.7	90.67	81.17
6	68.55	43.69	33.05	36.11	44.39	268.13	150.16	98.41	73.39	73.77
8	54.42	38.05	32.5	35.64	39.48	209.79	121.41	83.65	67.44	65.25
10	48.97	36.97	31.65	37.56	41.17	182.32	110.95	79	67.02	63.89
12	48.43	37.3	32.57	37.21	39.53	169.12	104.39	74.41	65.19	62.03

On the other hand, as we decrease the batch size, we enable more parallel gradient computations due to the availability of more smaller-sized batches. Although this parallelization could potentially lead to faster computation times, we interestingly notice an increase in the average computation time within each peer's database due to the overhead associated with averaging multiple gradients.

Further dissection of the number of peers' impact on performance, independent of batch size, reveals an interesting trend. As illustrated in Figure 5, the aggregation time, which includes both the computation of the aggregated gradient and the time to retrieve locally averaged gradients from other peers, decreases when the number of peers is reduced (as shown in the bar plot). This suggests an inherent computational overhead linked with gradient aggregation over larger networks of peers,
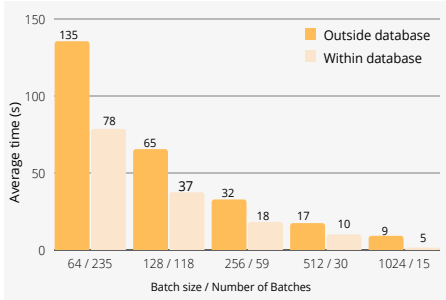
(a) MobileNet V3 Small Model
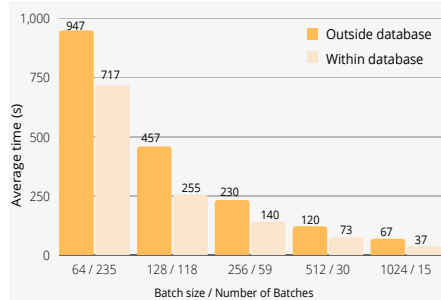
(b) Resnet-18 Model

Figure 6. Time taken for calculating gradient averages within and outside the database
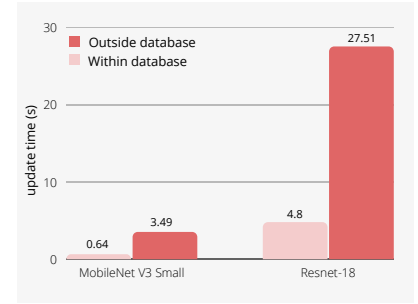

Figure 7. Time taken for Model Update within and outside the database

a factor that requires careful consideration when designing and scaling such decentralized learning systems.

This network size impact is further reflected in Table I, which present the total training time per epoch for both Densenet121 and MobileNet V3 Small models, respectively. We observe an overall decrease in training time per epoch with a larger batch size and a higher number of peers. Nonetheless, these tables also highlight the nonlinear relationship and the trade-offs between batch size and peers network scale.

VI. OPTIMIZING COMMUNICATION OVERHEAD

1. Motivation:

As we delve deeper into the intricate world of distributed serverless environments, characterized by numerous autonomous and stateless services, the challenges associated with frequent database retrieval loads during training phases emerge. While works such as previous work [21], LambdaML [14], SMLT [17], and MLLess [16] utilize a database as a communication channel, storing and retrieving model parameters as necessary, the exploration of the communication overhead these operations create remains largely unexplored. Such conditions often precipitate a significant increase in communication overhead, consequently undermining the overall training performance. However, within our unique architectural framework, we incorporate a customized Redis. This component allows us to perform average and update operations directly within Redis. Through this investigation, we aim to shed light on how our approach can reduce communication overhead and subsequently enhance training performance.

2. Approach:

In our experiment, we first explore the efficiency gains that can be realized by conducting model updates directly within RedisAI. Following this, we delve into the benefits of calculating gradient averages within Redis. For a comprehensive evaluation that considers the impact of model size on the overhead, we will use two distinct models - MobileNetV3 Small and ResNet18 - and run these models on the MNIST dataset. The insights from these tests will then be compared with the traditional method of iterative fetch-update-store operations, typically used with standard Redis, as outlined in previous studies.

3. Results:

An in-depth analysis of the experimental results provides persuasive evidence of the significant efficiency improvements facilitated by in-database operations, as vividly depicted in Figures 6 and 7.

Figure 6 concisely illustrates the substantial reduction in time for gradient averaging calculations within the database, as opposed to outside, for both the MobileNetV3 Small and the ResNet-18 models. As the plot underscores, the in-database approach yields consistently lower computation times across the spectrum of batch sizes. The MobileNetV3 Small model, for example, experiences an impressive decrease in computation time from 135.29 seconds outside the database to 78.52 seconds within, for a batch size of 64. The time efficiency continues to escalate as the batch size increases, with the largest batch size of 1024 resulting in a striking 82% improvement, requiring only 5.4 seconds for in-database computations.

A similar efficiency advantage is observed with the larger ResNet-18 model, where gradient averaging computations conducted within the database halved the processing time to 37.41 seconds, a stark reduction from the 67.32 seconds required when computed outside the database, for the batch size of 1024. This trend equates to a remarkable 69% improvement in processing efficiency, proving that our approach's benefits are applicable even for larger, more demanding models.

Figure 7 delves into the time taken for model updates within and outside the database. It echoes the efficiency improvements demonstrated by the previous figure, particularly the sharp decrease in model update times when processed within the database. MobileNetV3 Small model updates saw an 82% reduction in time when executed within the database, shrinking from 3.49 seconds outside the database to a mere 0.64 seconds within. Likewise, ResNet-18 updates experienced an approximately 83% improvement, with update times reducing dramatically from 27.5 seconds outside the database to 4.8 seconds within.

VII. SCALABILITY AND FAULT TOLERANCE EVALUATION

1. Motivation:

In the dynamic realm of distributed machine learning (ML) peer-to-peer training, disruptions, such as peer failures, new

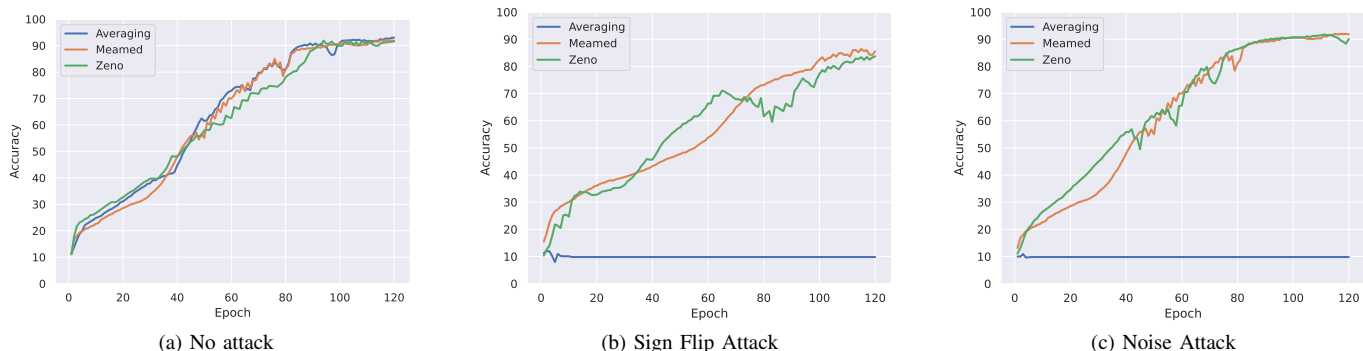(a) No attack   (b) Sign Flip Attack   (c) Noise Attack

Figure 8. Evaluation of network accuracy using Averaging, Zeno, and Meamed aggregation methods under various conditions: (a) Normal training with no attack, (b) Training under sign flip attack, where gradients are manipulated, and (c) Training amidst a noise attack, where random Gaussian noise is added to local updates.

peers joining, and potential Byzantine behaviors, are inevitable, and if not effectively managed, can impact the reliability and performance of the ML training. This provides the impetus for our experiment, which is meticulously designed to rigorously evaluate the robustness, scalability, and fault tolerance of a serverless, peer-to-peer ML training architecture. In this architecture, peers are carefully identified to form a secure training network, and its capacity to handle peer failures, seamlessly integrate new peers, and resist Byzantine attacks will be critically evaluated to ensure long-running distributed ML training sessions can continue to deliver accurate results amidst these disruptions.

2. Approach:

In our experimental approach, we center on the evaluation of three crucial aspects using the Mobilenet V3 Small model trained on the MNIST dataset: *peer failures*, the *addition of new peers*, and *Byzantine attacks*.

We initiate with the simulation of ***peer failures*** where we begin training with four peers, each processing 15 compute gradients. A peer failure is then artificially induced, leading us to measure the time taken by the remaining peers to identify the failure. In response to this failure, these remaining peers incorporate the data of the failed peer into their following epochs, augmenting the compute gradients to 20. This simulation aids in testing the system's resilience and ability to adapt in response to computational loss during peer failures.

To assess the system's scalability, we execute the ***addition of new peers*** scenario where a new peer is introduced into the network. We proceed to calculate the time taken by the existing peers to recognize and integrate this newcomer into the ongoing training process.

astly, our approach includes countering potential ***Byzantine attacks*** using robust aggregation algorithms, Zeno and Meamed. Zeno [29] ensures resilience by using a validation set to score and exclude any potentially adversarial local updates. The medians-based approach, or Meamed [28], combats Byzantine attacks by creating a vector that minimizes the overall distance to all local updates, thereby mitigating the influence of adversarial elements. By simulating adversarial scenarios of

one malicious peer, such as a sign flipping attack [55] where the malicious peer inverts and amplifies its local gradient, and a noise attack [56] where the malicious peer introduces Gaussian noise to its local updates, we track the training progression to convergence. This comprehensive approach enables us to gauge the architecture's resilience and robustness against adversarial behavior.

3. results:

1) *Peer Failure*

The flow of the experiment are depicted in Figure 9 for a more visual understanding. Initially, we began with four peers, each designed to process 15 batches per epoch. The first epoch proceeded unimpeded, with the total training time recorded at 52.6 seconds.
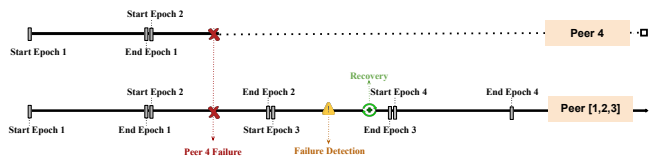

Figure 9. Recovery Time when a Peer Fails

A simulated peer failure has been introduced at the beginning of the second epoch, immediately after a health check had validated the 'failed' peer as operational. This timing represented a worst-case scenario, as it extended the detection period. Nonetheless, the remaining peers were in recognizing the failed peer within the ongoing epoch, taking a total 50.9 seconds to align with the next epoch's heartbeat step.

Post this single-peer-level detection, the remaining peers reached a consensus on the failed peer within 9.66 seconds. This aggregated the total detection time to 61.56s seconds of the failed peer. we note that, despite the deferred detection of the simulated failure in the third epoch, the training process was not interrupted, thus maintaining the progression towards model convergence. Upon consensus on the peer failure, the recovery process was triggered. This involved creating a new AWS Step Function to redistribute the failed peer's

9

computational workload, which increased the remaining peers' load from 15 to 20 batches per epoch. The complete recovery process, encapsulating the creation, deployment, and database entry of the new AWS Step Function, required an extra 1.94 seconds. After recovery, the remaining peers continued training in the 4th epoch with an adjusted workload to account for the lost peer. By the end of the fourth epoch, the total training time post-recovery was registered as 55,06 seconds. Our analysis shows a slight increase in total training time from 52.6 to 55.06 seconds, as the system transitions from four to three peers. This can be attributed to the increased number of gradients being averaged due to the redistribution of the failed peer's workload.

2) *Adding a new peer*

We found that in the scenario of *adding a new peer*, it took approximately 7.2 seconds for the existing three peers to recognize and integrate the new peer into their list of trusted peers.

3) *Tolerating Byzantine attacks*

The adoption of Meamed and Zeno aggregation algorithms introduced a computational overhead, leading to an approximate increase of 8.2 and 5.9 times in computational time, respectively, in contrast to the average aggregation method.

In a scenario without adversarial attacks, all three aggregation methods – Averaging, Zeno, and Meamed – achieved an accuracy above 90% within approximately 100 epochs. During a sign flip attack, the robust aggregations, Zeno and Meamed, managed to converge to almost 85%, while the normal averaging method did not. During a noise attack scenario, both Zeno and Meamed reached convergence above 90% after nearly 90 epochs, but the normal averaging method remained divergent.

VIII. Discussions

This section critically examines the results of our study and the performance of SPIRT, identifying its implications for peer-to-peer serverless ML training and potential directions for future research.

1. Serverless P2P for ML training

In our research, we created a distributed training workflow utilizing serverless computing, customized to inherit characteristics from a peer-to-peer (P2P) system, such as robustness, fault tolerance, and scalability. This design led to the creation of logical peers, each focusing on training a distinct part of the dataset. The statelessness of serverless computing necessitates the use of redis databases into our architecture for maintaining state information and enabling inter-component communication. To reduce communication overhead, we incorporated a modified version of RedisAI into our architecture, enabling immediate model updates within the database. Our proposed architecture paves the way for further exploration into the enhancement of in-database computing in distributed serverless ML systems.

2. Security, reliability, and Fault Tolerance in Machine Learning Architecture

Exploring scalability, a cornerstone of distributed systems, we delved into balancing inter and intra-peer scalability. While handling multiple data batches within a single peer increases efficiency, it also risks a single point of failure. Conversely, extending processing power across multiple peers introduces challenges related to gradient averaging and multiple connections. By testing our system with varied model and batch sizes, and different numbers of peers, while further exploration is required to fully comprehend these trade-offs, our efforts lay a significant foundation towards improving scalability in distributed ML systems.

In terms of fault tolerance, our design incorporates a robust heartbeat system to monitor peer health. Although this method is generally effective, it can introduce minor delays in detecting failures due to its periodic checking nature. To enhance detection speed, more frequent heartbeats could be implemented. As for fault recovery, our architecture ensures swift recovery time either by intra-peer scalability (adding a new peer) or by redistributing the load among existing peers (inter-peer scalability).

Aiming to augment the security of our architecture, we've implemented sophisticated cryptographic mechanisms which provide robust safeguards for data integrity, authenticity, and confidentiality. These measures ensure that data shared between two peers can only be understood by them. In addition, we've adopted robust aggregation that securely consolidates gradients from diverse peers. While the robust aggregation procedure may extend the aggregation time - for instance, take longer than the average aggregation - its capacity to safeguard against Byzantine attacks and guarantee model convergence underscores its importance.

IX. Conclusion

In this study, we present SPIRT, a serverless machine learning (ML) architecture that streamlines training in distributed settings. By minimizing communication overhead, demonstrating resilience to adversarial attacks, and seamlessly integrating new peers.

We found that SPIRT's use of in-database operations in RedisAI resulted in remarkable computational and communication efficiency, improving processing times by up to 82%. Moreover, SPIRT exhibited a high degree of resilience amidst disruptions, rapidly adapting to peer failures with minimal impact on training times and integrating new peers swiftly into the training network. When subjected to adversarial scenarios, the robustness of the architecture became more apparent. Through the use of Byzantine-resistant aggregation methods, Zeno and Meamed, high accuracy levels above 85% and 90% were maintained during sign flip and noise attacks, respectively.

These findings provide a robust foundation for the evolution of distributed ML training frameworks, paving the way for advancements in efficiency, resilience, and scalability in serverless environments.

REFERENCES

[1] B. Yuan, C. R. Wolfe, C. Dun, Y. Tang, A. Kyrillidis, and C. Jermaine, "Distributed learning of fully connected neural networks using independent subnet training," *Proc. VLDB Endow.*, vol. 15, no. 8, p. 1581–1590, apr 2022. [Online]. Available: https://doi.org/10.14778/3529337.3529343

[2] J. Verbraeken, M. Wolting, J. Katzy, J. Kloppenburg, T. Verbelen, and J. S. Rellermeyer, "A survey on distributed machine learning," *Acm computing surveys (csur)*, vol. 53, no. 2, pp. 1–33, 2020.

[3] S. Alqahtani and M. Demirbas, "Performance analysis and comparison of distributed machine learning systems," *arXiv preprint arXiv:1909.02061*, 2019.

[4] T. Sun, D. Li, and B. Wang, "Decentralized federated averaging," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.

[5] M. Li, L. Zhou, Z. Yang, A. Li, F. Xia, D. G. Andersen, and A. Smola, "Parameter server for distributed machine learning," in *Big learning NIPS workshop*, vol. 6, no. 2, 2013.

[6] R. Šajina, N. Tanković, and I. Ipšić, "Peer-to-peer deep learning with non-iid data," *Expert Systems with Applications*, vol. 214, p. 119159, 2023.

[7] "Serverless computing - aws lambda - amazon web services," https://aws.amazon.com/lambda/, (Accessed on 04/20/2023).

[8] "Cloud functions | google cloud," https://cloud.google.com/functions, (Accessed on 01/26/2022).

[9] "Cloud computing services | microsoft azure," https://azure.microsoft.com/en-us/, (Accessed on 01/26/2022).

[10] A. Barrak, F. Petrillo, and F. Jaafar, "Serverless on machine learning: A systematic mapping study," *IEEE Access*, vol. 10, pp. 99 337–99 352, 2022.

[11] Y. Yang, L. Zhao, Y. Li, H. Zhang, J. Li, M. Zhao, X. Chen, and K. Li, "Infless: a native serverless system for low-latency, high-throughput inference," in *Proceedings of the 27th ACM International Conference on Architectural Support for Programming Languages and Operating Systems*, 2022, pp. 768–781.

[12] A. Bhattacharjee, Y. Barve, S. Khare, S. Bao, A. Gokhale, and T. Damiano, "Stratum: A serverless framework for the lifecycle management of machine learning-based data analytics tasks," in *2019 USENIX Conference on Operational Machine Learning (OpML 19)*. Santa Clara, CA: USENIX Association, May 2019, pp. 59–61. [Online]. Available: https://www.usenix.org/conference/opml19/presentation/bhattacharjee

[13] P. G. Sarroca and M. Sánchez-Artigas, "Mlless: Achieving cost efficiency in serverless machine learning training," *arXiv preprint arXiv:2206.05786*, 2022.

[14] J. Jiang, S. Gan, Y. Liu, F. Wang, G. Alonso, A. Klimovic, A. Singla, W. Wu, and C. Zhang, "Towards demystifying serverless machine learning training," in *Proceedings of the 2021 International Conference on Management of Data*, 2021, pp. 857–871.

[15] D. Chahal, M. Mishra, S. C. Palepu, R. K. Singh, and R. Singhal, "Pay-as-you-train: Efficient ways of serverless training," in *2022 IEEE International Conference on Cloud Engineering (IC2E)*. IEEE, 2022, pp. 116–125.

[16] P. G. Sarroca and M. Sánchez-Artigas, "Mlless: Achieving cost efficiency in serverless machine learning training," *arXiv preprint arXiv:2206.05786*, 2022.

[17] A. Ali, S. Zawad, P. Aditya, I. E. Akkus, R. Chen, and F. Yan, "Smlt: A serverless framework for scalable and adaptive machine learning design and training," *arXiv preprint arXiv:2205.01853*, 2022.

[18] A. Grafberger, M. Chadha, A. Jindal, J. Gu, and M. Gerndt, "Fedless: Secure and scalable federated learning using serverless computing," *arXiv preprint arXiv:2111.03396*, 2021.

[19] J. Sampé, G. Vernik, M. Sánchez-Artigas, and P. García-López, "Serverless data analytics in the ibm cloud," in *Proceedings of the 19th International Middleware Conference Industry*, 2018, pp. 1–8.

[20] R. Anthony, *Systems programming: designing and developing distributed applications*. Morgan Kaufmann, 2015.

[21] A. Barrak, R. Trabelssi, F. Jaafar, and F. Petrillo, "Exploring the impact of serverless computing on peer to peer training machine learning," *International Conference on Cloud Engineering*, 2023.

[22] H. Wang, L. Muñoz-González, M. Z. Hameed, D. Eklund, and S. Raza, "Sparsfa: Towards robust and communication-efficient peer-to-peer federated learning," *Computers & Security*, vol. 129, p. 103182, 2023.

[23] M. Xu, Z. Zou, Y. Cheng, Q. Hu, D. Yu, and X. Cheng, "Spdl: A blockchain-enabled secure and privacy-preserving decentralized learning system," *IEEE Transactions on Computers*, 2022.

[24] M. Shayan, C. Fung, C. J. Yoon, and I. Beschastnikh, "Biscotti: A ledger for private and secure peer-to-peer machine learning," *arXiv preprint arXiv:1811.09904*, 2018.

[25] R. Guerraoui, A. Guirguis, J. Plassmann, A. Ragot, and S. Rouault, "Garfield: System support for byzantine machine learning (regular paper)," in *2021 51st Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. IEEE, 2021, pp. 39–51.

[26] C. Fang, Z. Yang, and W. U. Bajwa, "Bridge: Byzantine-resilient decentralized gradient descent," *arXiv preprint arXiv:1908.08098*, 2019.

[27] P. Blanchard, E. M. El Mhamdi, R. Guerraoui, and J. Stainer, "Machine learning with adversaries: Byzantine tolerant gradient descent," *Advances in neural information processing systems*, vol. 30, 2017.

[28] C. Xie, O. Koyejo, and I. Gupta, "Generalized byzantine-tolerant sgd," *arXiv preprint arXiv:1802.10116*, 2018.

[29] C. Xie, S. Koyejo, and I. Gupta, "Zeno: Distributed stochastic gradient descent with suspicion-based fault-

tolerance," in *International Conference on Machine Learning*. PMLR, 2019, pp. 6893–6901.

[30] E. Haußmann, "Accelerating i/o bound deep learning on shared storage," 2018.

[31] M. Zaharia, M. Chowdhury, M. J. Franklin, S. Shenker, I. Stoica *et al.*, "Spark: Cluster computing with working sets." *HotCloud*, vol. 10, no. 10-10, p. 95, 2010.

[32] P. Baran, "On distributed communications networks," *IEEE transactions on Communications Systems*, vol. 12, no. 1, pp. 1–9, 1964.

[33] A. Agarwal, O. Chapelle, M. Dudík, and J. Langford, "A reliable effective terascale linear learning system," *The Journal of Machine Learning Research*, vol. 15, no. 1, pp. 1111–1133, 2014.

[34] J. Wei, W. Dai, A. Qiao, Q. Ho, H. Cui, G. R. Ganger, P. B. Gibbons, G. A. Gibson, and E. P. Xing, "Managed communication and consistency for fast data-parallel iterative analytics," in *Proceedings of the Sixth ACM Symposium on Cloud Computing*, 2015, pp. 381–394.

[35] N. Shahidi, J. R. Gunasekaran, and M. T. Kandemir, "Cross-platform performance evaluation of stateful serverless workflows," in *2021 IEEE International Symposium on Workload Characterization (IISWC)*, 2021, pp. 63–73.

[36] D. Barcelona-Pons, P. Sutra, M. Sánchez-Artigas, G. París, and P. García-López, "Stateful serverless computing with crucial," *ACM Transactions on Software Engineering and Methodology (TOSEM)*, vol. 31, no. 3, pp. 1–38, 2022.

[37] F. Xu, Y. Qin, L. Chen, Z. Zhou, and F. Liu, "λ dnn : Achieving predictable distributed dnn training with serverless architectures," *IEEE Transactions on Computers*, 2021.

[38] M. Sánchez-Artigas and P. G. Sarroca, "Experience paper: Towards enhancing cost efficiency in serverless machine learning training," in *Proceedings of the 22nd International Middleware Conference*, 2021, pp. 210–222.

[39] S. S. Sandha, W. Cabrera, M. Al-Kateb, S. Nair, and M. Srivastava, "In-database distributed machine learning: Demonstration using teradata sql engine," *Proc. VLDB Endow.*, vol. 12, no. 12, p. 1854–1857, aug 2019. [Online]. Available: https://doi.org/10.14778/3352063.3352083

[40] "Redisai - a server for machine and deep learning models," https://oss.redis.com/redisai/#quick-links, (Accessed on 06/20/2023).

[41] "Oracle report scaling r to the enterprise," https://www.oracle.com/a/otn/docs/bringing-r-to-the-enterprise.pdf, July 2019, (Accessed on 06/20/2023).

[42] J. Hellerstein, C. Ré, F. Schoppmann, D. Z. Wang, E. Fratkin, A. Gorajek, K. S. Ng, C. Welton, X. Feng, K. Li *et al.*, "The madlib analytics library or mad skills, the sql," *arXiv preprint arXiv:1208.4165*, 2012.

[43] "Cross-region replication in azure | microsoft learn," https://learn.microsoft.com/en-us/azure/reliability/cross-region-replication-azure, (Accessed on 05/29/2023).

[44] "Configuring a lambda function to access resources in a vpc - aws lambda," https://docs.aws.amazon.com/lambda/latest/dg/configuration-vpc.html, (Accessed on 05/29/2023).

[45] A. Qiao, B. Aragam, B. Zhang, and E. Xing, "Fault tolerance in iterative-convergent machine learning," in *International Conference on Machine Learning*. PMLR, 2019, pp. 5220–5230.

[46] Y. Bouizem, N. Parlavantzas, D. Dib, and C. Morin, "Active-standby for high-availability in faas," in *Proceedings of the 2020 Sixth International Workshop on Serverless Computing*, 2020, pp. 31–36.

[47] T. Wink and Z. Nochta, "An approach for peer-to-peer federated learning," in *2021 51st Annual IEEE/IFIP International Conference on Dependable Systems and Networks Workshops (DSN-W)*. IEEE, 2021, pp. 150–157.

[48] H. Bourreau, E. Guichet, A. Barrak, B. Simon, and F. Jaafar, "On securing the communication in iot infrastructure using elliptic curve cryptography," in *2022 IEEE 22nd International Conference on Software Quality, Reliability, and Security Companion (QRS-C)*. IEEE, 2022, pp. 758–759.

[49] L. Lamport, R. Shostak, and M. Pease, "The byzantine generals problem," in *Concurrency: the works of leslie lamport*, 2019, pp. 203–226.

[50] "Lambda quotas - aws lambda," https://docs.aws.amazon.com/lambda/latest/dg/gettingstarted-limits.html, (Accessed on 06/24/2023).

[51] L. Deng, "The mnist database of handwritten digit images for machine learning research [best of the web]," *IEEE signal processing magazine*, vol. 29, no. 6, pp. 141–142, 2012.

[52] B. Koonce and B. Koonce, "Mobilenetv3," *Convolutional Neural Networks with Swift for Tensorflow: Image Recognition and Dataset Categorization*, pp. 125–144, 2021.

[53] 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*, 2016, pp. 770–778.

[54] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 4700–4708.

[55] L. Li, W. Xu, T. Chen, G. B. Giannakis, and Q. Ling, "Rsa: Byzantine-robust stochastic aggregation methods for distributed learning from heterogeneous datasets," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 33, no. 01, 2019, pp. 1544–1551.

[56] S. Li, Y. Cheng, W. Wang, Y. Liu, and T. Chen, "Learning to detect malicious clients for robust federated learning," *arXiv preprint arXiv:2002.00211*, 2020.