# FEDDES: A DISCRETE-EVENT SIMULATOR FOR LARGE-SCALE FEDERATED LEARNING

Anonymous authors

Paper under double-blind review

#### Abstract

We introduce FedDES, a performance simulator for Federated Learning (FL) that leverages Discrete Event Simulation (DES) techniques to model key events-such as client updates, communication delays, and aggregation operations—as discrete occurrences in time. This approach accurately captures the runtime features of FL systems, providing a high-fidelity simulation environment that closely mirrors real-world deployments. FedDES incorporates all three known aggregation settings: Synchronous (e.g., FedAvg and FedProx), Asynchronous (e.g., FedAsync and FedFa), and Semi-Asynchronous (e.g., FedBuff and FedCompass). Designed to be framework-, dataset-, and model-agnostic, FedDES allows researchers and developers to explore various configurations without restrictions. Our evaluations involving over 1,000 clients with heterogeneous computation and communication characteristics demonstrate that FedDES accurately models event distribution and delivers performance estimates within 2% error of real-world measurements. While real-world workloads often take hours to evaluate, FedDES generates detailed, timestamped event logs in just few seconds. As a result, FedDES can significantly accelerate FL developing and debugging cycles, enabling developers to rapidly prototype and evaluate algorithms and system designs, bypassing the need for costly, time-consuming real-world deployments. It offers valuable performance insights—such as identifying bottlenecks, stragglers, fault-tolerance mechanisms, and edge-case scenarios—facilitating the optimization of FL systems for efficiency, scalability, and resilience.

031 032

033

004

010 011

012

013

014

015

016

017

018

019

021

023

024

025

026

027

028

029

#### 1 INTRODUCTION

Federated Learning (FL) has emerged as a key paradigm in machine learning, primarily driven by growing concerns over privacy and the dwindling availability of publicly accessible datasets McMahan et al. (2017); Kairouz et al. (2021). Unlike traditional centralized learning approaches, FL allows multiple clients to collaboratively train a shared model while keeping their local data decentralized, addressing critical issues related to privacy and data sovereignty Yang et al. (2019). This decentralized learning approach is gaining widespread adoption across various domains, including healthcare, finance, and mobile applications, where data privacy and security are of paramount importance Bonawitz et al. (2019).

042 The field of FL is evolving rapidly, particularly in aggregation strategies. Aggregation strategies 043 are being developed to tackle challenges such as data heterogeneity, fairness, privacy, and robust-044 ness, creating specialized aggregation techniques Li et al. (2020; 2021a). For instance, algorithms like FedAvg Karimireddy et al. (2020) and FedProx Li et al. (2020) have been introduced to address issues arising from non-IID data distributions and to improve robustness against stragglers and 046 malicious clients. At the same time, scheduling algorithms are evolving toward semi-synchronous 047 paradigms, which strike a balance between synchronous and asynchronous aggregations Li et al. 048 (2023b); Chen et al. (2020). These semi-synchronous approaches optimize training speed and model accuracy while maintaining system stability. However, integrating such sophisticated scheduling mechanisms can fundamentally alter an FL framework's workflow, making it challenging to assess 051 the effectiveness of new designs without comprehensive real-world evaluations. 052

In response to the rapidly changing landscape of FL, researchers are increasingly adopting fastprototyping methodologies to iterate and refine their algorithms Li et al. (2019; 2020; 2021b).

Therefore, there is a growing recognition of the need for efficient and extensive simulation tools 055 that can rigorously evaluate the robustness and efficacy of FL algorithms, particularly in areas such 056 as fault tolerance and scalability in large-scale deployments Zhao et al. (2018). To meet this need, 057 we propose FedDES: A Discrete-Event Simulator for Federated Learning, which applies Dis-058 crete Event Simulation (DES) principles to accurately model and analyze the intricate dynamics of FL systems Banks et al. (2010). FedDES captures the temporal evolution of FL processes by modeling client selection, local training, communication delays, and model aggregation as discrete events 060 within a simulated timeline. This event-driven approach allows for precise temporal coordination 061 and resource allocation, providing a detailed understanding of system behavior under various condi-062 tions. FedDES offers significant performance advantages inherent to DES, including high efficiency 063 and scalability, enabling rapid simulation of FL workflows with thousands of clients exhibiting het-064 erogeneous computational and communication capabilities. The simulator is framework-agnostic, 065 allowing modeling and simulating distinct aggregation strategies across all known settings. 066

Moreover, FedDES is designed to be independent of specific datasets and models, allowing re-067 searchers to integrate their own data distributions and model architectures seamlessly without mod-068 ifying the core simulation framework. We extensively evaluated FedDES on the NCSA Delta super-069 computer, scaling real-world experimental events to over 1,000 clients with heterogeneous computational and communication settings. By comparing these results with our simulations, we demon-071 strated that FedDES achieves highly accurate event distribution and performance estimates, with 072 an error margin of less than 2%. This underscores FedDES's ability to provide precise, large-scale 073 event logs, facilitating debugging, bottleneck analysis, and rapid prototyping of new FL algorithms 074 and system designs.

075 076

077

078 079

081

082

083

#### 2 RELATED WORK

#### 2.1 FL AGGREGATION STRATEGIES

Aggregation mechanisms critically influence the performance of FL systems. These mechanisms are typically classified into three categories:

Synchronous Aggregation: requires the server to wait for updates from all selected clients before aggregating and updating the global model. This ensures consistency but often leads to delays, especially when slow clients (stragglers) are involved. The most prominent synchronous algorithm, FedAvg McMahan et al. (2017); Karimireddy et al. (2020), aggregates client updates by computing a weighted average based on the number of samples each client holds. While effective for homogeneous training environments where each client's communication and communication resources are similar, FedAvg suffers from inefficiencies in environments with heterogeneous client resources, as the slowest client limits the overall training time.

Asynchronous Aggregation: update the global model without waiting for all clients, allowing 092 faster updates but risking using stale client models. FedAsync Xie et al. (2019) addresses this 093 by applying a staleness factor to penalize outdated client updates during aggregation. This method 094 reduces waiting time, improving resource utilization. However, staleness penalties can cause the global model to drift away from the local data of slower clients, particularly in non-IID settings, 096 affecting accuracy. FedFa Xu et al. (2024) takes asynchronous updates further by using immediate 097 global updates without waiting for any client group. It mitigates the effects of stale updates with a 098 sliding window aggregation technique, allowing the most recent updates to carry more weight. This approach boosts training speed significantly in heterogeneous environments but requires careful 099 handling of staleness to maintain accuracy in non-IID settings. 100

Semi-Asynchronous Aggregation: combine aspects of both synchronous and asynchronous methods by grouping clients based on their computational capacity and scheduling updates within these groups. FedBuff Nguyen et al. (2022) introduces a buffering mechanism to temporarily store client updates and aggregate them once conditions are met, balancing update frequency and reducing the staleness of updates. However, managing the buffer requires heuristic tuning of the buffer size. On the other hand, FedCompass Li et al. (2023b) uses a computing power-aware scheduler to dynamically assign local steps to clients so that updates are received near synchrony within each group. This approach mitigates model staleness while maintaining efficiency, allowing FedCompass to achieve faster convergence and higher accuracy in heterogeneous and non-IID settings than fully asynchronous methods.

110 111 112

2.2 FL SIMULATORS

113 FL simulators Beutel et al. (2020); Ryu et al. (2022); fls; flm; Li et al. (2023a); He et al. (2020); Sun et al. (2019); Ekaireb et al. (2022); Ro et al. (2021); Mugunthan et al. (2020) facilitate the 114 validation of theoretical FL research. For example, Flower Beutel et al. (2020) integrates Apache 115 Ray Moritz et al. (2018) to generate virtual clients, allowing for straightforward configuration of 116 client resources like GPU memory ratios and CPU core allocations. Researchers can deploy cus-117 tom aggregation strategies using selected models and datasets to evaluate statistical metrics such 118 as convergence and accuracy. Furthermore, ns3-fl Ekaireb et al. (2022) combines FLsim with ns3, 119 providing network configuration options that enable researchers to explore how network conditions 120 affect FL performance. Our work complements existing FL simulators; FedDES is particularly ef-121 fective for evaluating scheduling algorithms and delivering detailed performance metrics, including 122 latency and straggler effects in dynamic environments.

123 124

125

2.3 DISCRETE EVENT-DRIVEN SIMULATION (DES)

126 DES models a system as a series of discrete events, each triggering state changes at specific 127 times Banks (2005); Banks et al. (2010). DES has been widely applied in network simulations 128 and distributed systems, where the timing and order of events are critical. Tools like ROSS Pearce 129 (2002) and SimGrid Casanova (2001) demonstrate how DES can scale to millions of events and entities, providing insights into system performance in large-scale distributed environments. Our Fed-130 131 DES offers a significantly more scalable alternative by modeling FL workflows as discrete events. By simulating client-server activities and communication as state changes between events, FedDES 132 scales efficiently to tens of thousands of clients while capturing key performance metrics, making it 133 a robust tool for large-scale FL simulations. 134

- <sup>135</sup> 3 PROPOSED FEDDES
- 136 137

138

#### 3.1 COMMUNICATION-CENTRIC STATE MACHINE FOR FL EVENT MANAGEMENT

As illustrated in Figure II, FedDES models the FL system as a state machine, where states are de-139 fined by computational workloads (e.g., local training, aggregation) and communication workloads 140 (e.g., model parameter transmission). Client states track local models, computational capacity, and 141 training progress, while the server maintains the global model and monitors client updates. Network 142 states reflect link bandwidth, latency, and system topology, simulating real-world communication 143 delays. State transitions are primarily driven by communication events (i.e., red actions in Figure 1, 144 there are other transition events like the blue action, which depends on aggregation strategies), which 145 act as synchronization points between clients and the server. For instance, after receiving the global 146 model, a client transitions into the training state and, upon completion, triggers a communication event to send the updated model back to the server. Once sufficient updates are received upon 147 criteria of different aggregation paradigms, the server transitions into the aggregation state. This 148 communication-centric approach captures FL's asynchronous, distributed nature, where model ex-149 changes rather than client computations govern state changes. 150

In FedDES, clients and servers transition between states such as idle, training, communicating, and aggregating, with communication events marking key interactions. This communication-driven design is well-suited for distributed asynchronous environments like FL for several reasons. First, it optimizes computational efficiency by focusing on client-server interactions, bypassing idle times. Second, communication is a natural trigger for state transitions, eliminating the need for global synchronization. Finally, it enables accurate simulation of network effects, such as bandwidth constraints and delays, which are critical to FL performance.

Event execution in FedDES is managed using a priority queue, advancing the simulation clock based
on communication events. Key event types include *Client-Server Communication*, which simulates
model transmission and accounts for network delays; *Aggregation*, where the server updates the
global model; and *Client Computation*, which indirectly triggers communication events. By structuring the simulation around communication, FedDES efficiently captures the dynamics of large-



	Algorithm 1: FedDES-Avg (Exemplifies Synchronous FL)
ī	<b>Input:</b> Global model $W_i$ Number of clients $n$ Number of rounds $T$ Communication cost comm cost
	Training cost train_cost, Aggregation cost agg_cost
(	Output: Final global model after T rounds
]	Function Server ( $W_t$ , $n$ , $T$ , $dataloader\_cost$ , $agg\_cost$ , $comm\_cost$ ):
	for $t = 1$ to T do
;	for client set $c_i$ from 0 to $n-1$ do
L	$time + = \text{CommTime}(W_t, c_i, comm\_cost)  (Simulate sending global model)$
	<b>Event 1:</b> Distribute the global model to $c_i$
	$\begin{array}{c} \mathbf{end} \mathbf{10r} \\ \mathbf{for} \ a, \ from 0, \ to \ n, \ 1, \ \mathbf{do} \end{array}$
	$\begin{bmatrix} \text{time}_{i} & \text{from } 0 & $
	<b>Example</b> 1 Collect the local model from $c_1$
	end for
	$time + = \text{AggregateTime}(n. aag\_cost)$ (Simulate aggregation)
	end for
]	Function Client (ci, n, T, dataloader_cost, train_cost, comm_cost):
	for each epoch $t = 1$ to T do
	$time + = \text{CommTime}(W_t, \text{client}, comm\_cost)$ (Simulate receive global model)
	<b>Event 2:</b> $c_i$ receives the global model from the server.
	$time + = \text{ComputeTime}(W_t, train\_cost)$ (Simulate local training)
	$time + = \text{CommTime}(W_{t+1}^{(t)}, \text{server}, comm\_cost)$ (Simulate send local model)
	<b>Event 3:</b> $c_i$ sends the updated local model back to the server.
	end for
	Argonullin 2: reades-Async (exemplines Asynchronous PL)
]	<b>Input:</b> Global model $W_t$ , Number of clients n, Communication cost $comm\_cost$ , Training cost
]	<b>Input:</b> Global model $W_t$ , Number of clients <i>n</i> , Communication cost <i>comm_cost</i> , Training cost $train\_cost$ , Aggregation cost $agg\_cost$
]	<b>Input:</b> Global model $W_t$ , Number of clients <i>n</i> , Communication cost $comm\_cost$ , Training cost $train\_cost$ , Aggregation cost $agg\_cost$ <b>Output:</b> Final global model after all updates
	Argorithm 2: FedDES-Async (Exemplifies Asynchronous FL)         Input: Global model $W_t$ , Number of clients n, Communication cost $comm\_cost$ , Training cost $train\_cost$ , Aggregation cost $agg\_cost$ Output: Final global model after all updates         Server( $W_t$ , n, $agg\_cost$ , $comm\_cost$ )
] ( ( ( (	<b>Examplifies</b> Asynchronous FL) <b>Input:</b> Global model $W_t$ , Number of clients $n$ , Communication cost $comm\_cost$ , Training cost $train\_cost$ , Aggregation cost $agg\_cost$ <b>Output:</b> Final global model after all updates <b>Server</b> ( $W_t$ , $n$ , $agg\_cost$ , $comm\_cost$ ) <b>for</b> $c_i$ from 0 to $n - 1$ <b>do</b> $t_i$ time ( $W_i$ = comm $T$ ime)( $W_i$ = comm cost) (Sund clobal model)
1	Argorithm 2: FedDES-Async (Exemplifies Asynchronous FL)         Input: Global model $W_t$ , Number of clients $n$ , Communication cost $comm\_cost$ , Training cost $train\_cost$ , Aggregation cost $agg\_cost$ Output: Final global model after all updates         Server( $W_t$ , $n$ , $agg\_cost$ , $comm\_cost$ )         for $c_i$ from 0 to $n - 1$ do         time+ = CommTime( $W_t$ , $c_i$ , comm\_cost)         (Send global model)         Final 1: Sond the global model to client $a_i$
1	Argorithm 2: FedDES-Async (Exemplifies Asynchronous FL)         Input: Global model $W_t$ , Number of clients $n$ , Communication cost $comm\_cost$ , Training cost $train\_cost$ , Aggregation cost $agg\_cost$ Output: Final global model after all updates         Server( $W_t$ , $n$ , $agg\_cost$ , $comm\_cost$ )         for $c_i$ from 0 to $n - 1$ do $time+ = \text{CommTime}(W_t, c_i, comm\_cost)$ (Send global model)         Event 1: Send the global model to client $c_i$ .
1	Argorithm 2: FedDES-Async (Exemplifies Asynchronous FL)         Input: Global model $W_t$ , Number of clients $n$ , Communication cost $comm\_cost$ , Training cost $train\_cost$ , Aggregation cost $agg\_cost$ Output: Final global model after all updates         Server( $W_t$ , $n$ , $agg\_cost$ , $comm\_cost$ )         for $c_i$ from 0 to $n - 1$ do $time+ = \text{CommTime}(W_t, c_i, comm\_cost)$ (Send global model)         Event 1: Send the global model to client $c_i$ .         end for         for $t = 1$ to $n \times T$ do
	Argor tunn 2: FedDES-Async (Exemplifies Asynchronous FL)         Input: Global model $W_t$ , Number of clients $n$ , Communication cost $comm\_cost$ , Training cost $train\_cost$ , Aggregation cost $agg\_cost$ Output: Final global model after all updates         Server( $W_t$ , $n$ , $agg\_cost$ , $comm\_cost$ )         for $c_i$ from 0 to $n - 1$ do $time+ = \text{CommTime}(W_t, c_i, comm\_cost)$ (Send global model)         Event 1: Send the global model to client $c_i$ .         end for         for $t = 1$ to $n \times T$ do         if local model is received from any client $c_i$ then
1 1 1 1	Argorithm 2: FedDES-Async (Exemplifies Asynchronous FL)         Input: Global model $W_t$ , Number of clients $n$ , Communication cost $comm\_cost$ , Training cost $train\_cost$ , Aggregation cost $agg\_cost$ Output: Final global model after all updates         Server( $W_t$ , $n$ , $agg\_cost$ , $comm\_cost$ )         for $c_i$ from 0 to $n - 1$ do $time+ = \text{CommTime}(W_t, c_i, comm\_cost)$ (Send global model)         Event 1: Send the global model to client $c_i$ .         end for         for $t = 1$ to $n \times T$ do         if local model is received from any client $c_i$ then $time+ = \text{CommTime}(W_{t+1}^{(i)})$ , server, $comm\_cost$ ) (Receive local model)
	Argorithm 2: FedDES-Async (Exemplifies Asynchronous FL)         Input: Global model $W_t$ , Number of clients $n$ , Communication cost $comm\_cost$ , Training cost $train\_cost$ , Aggregation cost $agg\_cost$ Output: Final global model after all updates         Server( $W_t$ , $n$ , $agg\_cost$ , $comm\_cost$ )         for $c_i$ from 0 to $n - 1$ do $time+ = \text{CommTime}(W_t, c_i, comm\_cost)$ (Send global model)         Event 1: Send the global model to client $c_i$ .         end for         for $t = 1$ to $n \times T$ do         if local model is received from any client $c_i$ then $time+ = \text{CommTime}(W_{t+1}^{(i)}, \text{server}, comm\_cost)$ (Receive local model)         Event 4: Receive and aggregate the local model from client $c_i$ .
1 0 1 1 1	Argorithm 2: FedDES-Async (Exemplifies Asynchronous FL)         Input: Global model $W_t$ , Number of clients $n$ , Communication cost $comm\_cost$ , Training cost $train\_cost$ , Aggregation cost $agg\_cost$ Output: Final global model after all updates         Server( $W_t$ , $n$ , $agg\_cost$ , $comm\_cost$ )         for $c_i$ from 0 to $n - 1$ do $time+ = \text{CommTime}(W_t, c_i, comm\_cost)$ (Send global model)         Event 1: Send the global model to client $c_i$ .         end for         for $t = 1$ to $n \times T$ do         if local model is received from any client $c_i$ then $time+ = \text{CommTime}(W_{t+1}^{(i)}, \text{server}, comm\_cost)$ (Receive local model)         Event 4: Receive and aggregate the local model from client $c_i$ . $time+ = \text{AggregateTime}(agg\_cost)$ (Aggregation)
	Argorithm 2: FedDES-Async (Exemplifies Asynchronous FL)         Input: Global model $W_t$ , Number of clients $n$ , Communication cost $comm\_cost$ , Training cost $train\_cost$ , Aggregation cost $agg\_cost$ Output: Final global model after all updates         Server( $W_t$ , $n$ , $agg\_cost$ , $comm\_cost$ )         for $c_i$ from 0 to $n - 1$ do $time+ = \text{CommTime}(W_t, c_i, comm\_cost)$ (Send global model)         Event 1: Send the global model to client $c_i$ .         end for         for $t = 1$ to $n \times T$ do         if local model is received from any client $c_i$ then $time+ = \text{CommTime}(W_{t+1}^{(i)}, \text{server}, comm\_cost)$ (Receive local model)         Event 4: Receive and aggregate the local model from client $c_i$ . $time+ = \text{AggregateTime}(agg\_cost)$ (Aggregation)         Event 1: Send updated global model back to client $c_i$ .
	Argor tum 2: FedDES-Async (Exemplifies Asynchronous FL)         Input: Global model $W_t$ , Number of clients $n$ , Communication cost $comm\_cost$ , Training cost $train\_cost$ , Aggregation cost $agg\_cost$ Output: Final global model after all updates         Server( $W_t$ , $n$ , $agg\_cost$ , $comm\_cost$ )         for $c_i$ from 0 to $n - 1$ do $time+ = \text{CommTime}(W_t, c_i, comm\_cost)$ (Send global model)         Event 1: Send the global model to client $c_i$ .         end for         for $t = 1$ to $n \times T$ do         if local model is received from any client $c_i$ then $time+ = \text{CommTime}(W_{t+1}^{(i)}, \text{server}, comm\_cost)$ (Receive local model)         Event 4: Receive and aggregate the local model from client $c_i$ . $time+ = \text{AggregateTime}(agg\_cost)$ (Aggregation)         Event 1: Send updated global model back to client $c_i$ . $time+ = \text{CommTime}(W_{t+1}^{(i)}, c_i, comm\_cost)$ ) (Send updated model)
	Argor tum 2: FedDES-Async (Exemplifies Asynchronous FL)         Input: Global model $W_t$ , Number of clients $n$ , Communication cost $comm\_cost$ , Training cost $train\_cost$ , Aggregation cost $agg\_cost$ Output: Final global model after all updates         Server( $W_t$ , $n$ , $agg\_cost$ , $comm\_cost$ )         for $c_i$ from 0 to $n - 1$ do $time+ = \text{CommTime}(W_t, c_i, comm\_cost)$ (Send global model)         Event 1: Send the global model to client $c_i$ .         end for         for $t = 1$ to $n \times T$ do         if local model is received from any client $c_i$ then $time+ = \text{CommTime}(W_{t+1}^{(i)}, \text{server}, comm\_cost)$ (Receive local model)         Event 4: Receive and aggregate the local model from client $c_i$ . $time+ = \text{AggregateTime}(agg\_cost)$ (Aggregation)         Event 1: Send updated global model back to client $c_i$ . $time+ = \text{CommTime}(W_{t+1}^{(i)}, c_i, comm\_cost)$ (Send updated model)         event 1: Send updated global model back to client $c_i$ . $time+ = \text{CommTime}(W_{t+1}^{(i)}, c_i, comm\_cost)$ (Send updated model)         end if
	Argorithm 2: FedDES-Async (Exemplifies Asynchronous FL)         Input: Global model $W_t$ , Number of clients $n$ , Communication cost $comm\_cost$ , Training cost $train\_cost$ , Aggregation cost $agg\_cost$ Output: Final global model after all updates         Server( $W_t$ , $n$ , $agg\_cost$ , $comm\_cost$ )         for $c_i$ from 0 to $n - 1$ do $time+ = \text{CommTime}(W_t, c_i, comm\_cost)$ (Send global model)         Event 1: Send the global model to client $c_i$ .         end for         for $t = 1$ to $n \times T$ do         if local model is received from any client $c_i$ then $time+ = \text{CommTime}(W_{t+1}^{(i)}, \text{server}, comm\_cost)$ (Receive local model)         Event 4: Receive and aggregate the local model from client $c_i$ . $time+ = \text{AggregateTime}(agg\_cost)$ (Aggregation)         Event 1: Send updated global model back to client $c_i$ . $time+ = \text{CommTime}(W_{t+1}^{(i)}, c_i, comm\_cost)$ ) (Send updated model)         event 1: Send updated global model back to client $c_i$ . $time+ = \text{CommTime}(W_{t+1}^{(i)}, c_i, comm\_cost)$ ) (Send updated model)         end if
	Argor tunn 2: FedDES-Async (Exemplifies Asynchronous FL)         Input: Global model $W_t$ , Number of clients $n$ , Communication cost $comm\_cost$ , Training cost $train\_cost$ , Aggregation cost $agg\_cost$ Output: Final global model after all updates         Server( $W_t$ , $n$ , $agg\_cost$ , $comm\_cost$ )         for $c_i$ from 0 to $n - 1$ do $time+ = \text{CommTime}(W_t, c_i, comm\_cost)$ (Send global model)         Event 1: Send the global model to client $c_i$ .         end for         for $t = 1$ to $n \times T$ do         if local model is received from any client $c_i$ then $time+ = \text{CommTime}(W_{t+1}^{(i)}, \text{server}, comm\_cost)$ (Receive local model)         Event 4: Receive and aggregate the local model from client $c_i$ . $time+ = \text{AggregateTime}(agg\_cost)$ (Aggregation)         Event 1: Send updated global model back to client $c_i$ . $time+ = \text{CommTime}(W_{t+1}^{(i)}, c_i, comm\_cost)$ ) (Send updated model)         end for         for $c_i$ from remaining unterminated clients do
	Argorithm 2: FeadDES-Async (Exemplifies Asynchronous FL)         Input: Global model $W_t$ , Number of clients $n$ , Communication cost $comm\_cost$ , Training cost $train\_cost$ , Aggregation cost $agg\_cost$ Output: Final global model after all updates         Server( $W_t$ , $n$ , $agg\_cost$ , $comm\_cost$ )         for $c_i$ from 0 to $n - 1$ do $time+ = \text{CommTime}(W_t, c_i, comm\_cost)$ (Send global model)         Event 1: Send the global model to client $c_i$ .         end for         for $t = 1$ to $n \times T$ do         if local model is received from any client $c_i$ then $time+ = \text{CommTime}(W_{t+1}^{(i)}, \text{server}, comm\_cost)$ (Receive local model)         Event 4: Receive and aggregate the local model from client $c_i$ . $time+ = \text{AggregateTime}(agg\_cost)$ (Aggregation)         Event 1: Send updated global model back to client $c_i$ . $time+ = \text{CommTime}(W_{t+1}^{(i)}, c_i, comm\_cost))$ (Send updated model)         end for         for $c_i$ from remaining unterminated clients do         time+ = CommTime(Termination Signal, $c_i, comm\_cost$ ) $time+ = \text{CommTime}(\text{Termination Signal}, c_i, comm\_cost)$
	Augorithm 2: FedDES-Async (Exemplifies Asynchronous FL)Input: Global model $W_t$ , Number of clients $n$ , Communication cost $comm\_cost$ , Training cost $train\_cost$ , Aggregation cost $agg\_cost$ Output: Final global model after all updatesServer( $W_t$ , $n$ , $agg\_cost$ , $comm\_cost$ )for $c_i$ from 0 to $n - 1$ do $time+ = \text{CommTime}(W_t, c_i, comm\_cost)$ (Send global model)Event 1: Send the global model to client $c_i$ .end forfor $t = 1$ to $n \times T$ doif local model is received from any client $c_i$ then $time+ = \text{CommTime}(W_{t+1}^{(i)}, \text{server}, comm\_cost)$ (Receive local model)Event 4: Receive and aggregate the local model from client $c_i$ . $time+ = \text{AggregateTime}(agg\_cost)$ (Aggregation)Event 1: Send updated global model back to client $c_i$ . $time+ = \text{CommTime}(W_{t+1}^{(i)}, c_i, comm\_cost)$ (Send updated model)end forfor $c_i$ from remaining unterminated clients dotime+ = CommTime((Termination Signal, $c_i$ , comm\_cost)Event 5: Send termination signal to each client.met 4: Receive form mignal to each client.met 4: Receive form mignation signal to each client.met 4: Receive form remaining unterminated clients dotime+ = CommTime(Termination signal to each client.met 5: Send termination signal to each client.
	Augorithm 2: FedDES-Async (Exemptifies Asynchronous FL)         Input: Global model $W_t$ , Number of clients $n$ , Communication cost $comm\_cost$ , Training cost $train\_cost$ , Aggregation cost $agg\_cost$ Output: Final global model after all updates         Server( $W_t$ , $n$ , $agg\_cost$ , $comm\_cost$ )         for $c_i$ from 0 to $n - 1$ do $time+ = \text{CommTime}(W_t, c_i, comm\_cost)$ (Send global model)         Event 1: Send the global model to client $c_i$ .         end for         for $t = 1$ to $n \times T$ do         if local model is received from any client $c_i$ then $time+ = \text{CommTime}(W_{t+1}^{(i)}, \text{server}, comm\_cost)$ (Receive local model)         Event 4: Receive and aggregate the local model from client $c_i$ . $time+ = \text{AggregateTime}(agg\_cost)$ (Aggregation)         Event 1: Send updated global model back to client $c_i$ . $time+ = \text{CommTime}(W_{t+1}^{(i)}, c_i, comm\_cost)$ (Send updated model)         end for         for $c_i$ from remaining unterminated clients do         time+ = CommTime(Termination Signal, $c_i$ , comm\_cost)         Event 5: Send termination signal to each client.         end for         Collect (c: train cost comm cost)
	Augorithm 2: FEDES-Async (EXEMPTINES Asynchronous FL)         Input: Global model $W_t$ , Number of clients $n$ , Communication cost $comm\_cost$ , Training cost $train\_cost$ , Aggregation cost $agg\_cost$ Output: Final global model after all updates         Server( $W_t$ , $n$ , $agg\_cost$ , $comm\_cost$ )         for $c_i$ from 0 to $n - 1$ do $time+ = \text{CommTime}(W_t, c_i, comm\_cost)$ (Send global model)         Event 1: Send the global model to client $c_i$ .         end for         for $t = 1$ to $n \times T$ do         if local model is received from any client $c_i$ then $time+ = \text{CommTime}(W_{t+1}^{(i)}, \text{server}, comm\_cost)$ (Receive local model)         Event 4: Receive and aggregate the local model from client $c_i$ . $time+ = \text{AggregateTime}(agg\_cost)$ (Aggregation)         Event 1: Send updated global model back to client $c_i$ . $time+ = \text{CommTime}(W_{t+1}^{(i)}, c_i, comm\_cost)$ (Send updated model)         end for         for $c_i$ from remaining unterminated clients do         time+ = CommTime(Termination Signal, $c_i, comm\_cost$ )         end for         for $c_i$ from remaining unterminated clients do         time+ = Send termination signal to each client.         end for         Client( $c_i$ , train\_cost, comm\_cost)         while termination signal not received do
	Augorithm 2: FedDES-Async (Exemplifies Asynchronous FL)         Input: Global model $W_t$ , Number of clients $n$ , Communication cost $comm\_cost$ , Training cost $train\_cost$ , Aggregation cost $agg\_cost$ Output: Final global model after all updates         Server( $W_t$ , $n$ , $agg\_cost$ , $comm\_cost$ )         for $c_i$ , from 0 to $n - 1$ do $time+ = \text{CommTime}(W_t, c_i, comm\_cost)$ (Send global model)         Event 1: Send the global model to client $c_i$ .         end for         for $t = 1$ to $n \times T$ do         if local model is received from any client $c_i$ then $time+ = \text{CommTime}(W_{t+1}^{(i)}, \text{server}, comm\_cost)$ (Receive local model)         Event 4: Receive and aggregate the local model from client $c_i$ . $time+ = \text{AggregateTime}(agg\_cost)$ (Aggregation)         Event 1: Send updated global model back to client $c_i$ . $time+ = \text{CommTime}(W_{t+1}^{(i)}, c_i, comm\_cost)$ (Send updated model)         end for         for $c_i$ from remaining unterminated clients do $time+ = \text{CommTime}(\text{Termination Signal, c_i, comm\_cost)         Event 5: Send termination signal to each client.         end for         Client(c_i, train\_cost, comm\_cost)         while termination signal not received do         time+ = \text{CommTime}(W_{t+1}^{(i)}, c_i, comm\_cost) (Receive global model)   $
	Augorithm 2: FedDES-Async (Exemptines Asynchronous FL)         Input: Global model $W_t$ , Number of clients $n_i$ Communication cost $comm\_cost$ , Training cost $train\_cost$ , Aggregation cost $agg\_cost$ Output: Final global model after all updates         Server( $W_t$ , $n_i$ , $agg\_cost$ , $comm\_cost$ )         for $c_i$ from 0 to $n - 1$ do $time+ = \text{CommTime}(W_t, c_i, comm\_cost)$ (Send global model)         Event 1: Send the global model to client $c_i$ .         end for         for $t = 1$ to $n \times T$ do         if local model is received from any client $c_i$ then $time+ = \text{CommTime}(W_{t+1}^{(i)}, server, comm\_cost)$ (Receive local model)         Event 4: Receive and aggregate the local model from client $c_i$ . $time+ = \text{AggregateTime}(agg\_cost)$ (Aggregation)         Event 1: Send updated global model back to client $c_i$ . $time+ = \text{CommTime}(W_{t+1}^{(i)}, c_i, comm\_cost)$ (Send updated model)         end if         end for         for $c_i$ from remaining unterminated clients do $time+ = \text{CommTime}(Termination Signal, c_i, comm\_cost)         Event 5: Send termination signal to each client.         end for         Client(c_i, train\_cost, comm\_cost)         while termination signal not received do         time+ = CommTime(W_t, c_i, comm\_cost) (Receive global model)         Event 5: Send termination signal not$
	Argor truin 2: redDES-Asyne (Exemplifies Asynemotous PL)Input: Global model $W_t$ , Number of clients n, Communication cost $comm\_cost$ , Training cost $train\_cost$ , Aggregation cost $agg\_cost$ Output: Final global model after all updatesServer( $W_t$ , n, $agg\_cost$ , $comm\_cost$ )for $c_i$ from 0 to $n - 1$ dotime+ = CommTime( $W_t$ , $c_i$ , $comm\_cost$ ) (Send global model)Event 1: Send the global model to client $c_i$ .end forfor $t = 1$ to $n \times T$ doif local model is received from any client $c_i$ thentime+ = CommTime( $W_{t+1}^{(i)}$ , server, $comm\_cost$ ) (Receive local model)Event 4: Receive and aggregate the local model from client $c_i$ .time+ = AggregateTime( $agg\_cost$ ) (Aggregation)Event 1: Send updated global model back to client $c_i$ .time+ = CommTime( $W_{t+1}^{(i)}$ , $c_i$ , $comm\_cost$ ) (Send updated model)end forfor $c_i$ from remaining unterminated clients dotime+ = CommTime(Termination Signal, $c_i$ , $comm\_cost$ )Event 5: Send termination signal to each client.end forClient( $c_i$ , train\_cost, comm\_cost) (Receive global model)Event 5: Send termination signal to each client.end forClient( $c_i$ , train\_cost, comm\_cost) (Receive global model)Event 5: Send termination signal to each client.end forClient( $c_i$ , train\_cos
	Argon tume 2: real DES-Asyne (Exemplifies Asynemotous PL)Input: Global model $W_i$ , Number of clients $n$ , Communication cost $comm\_cost$ , Training cost $train\_cost$ , Aggregation cost $agg\_cost$ Output: Final global model after all updatesServer( $W_i$ , $n$ , $agg\_cost$ , $comm\_cost$ )(Send global model after all updatesServer( $W_i$ , $n$ , $agg\_cost$ , $comm\_cost$ )(Send global model)Event 1: Send the global model to client $c_i$ .end forfor $t = 1$ to $n \times T$ doif local model is received from any client $c_i$ thentime+ = CommTime( $W_{t+1}^{(i)}$ , server, $comm\_cost$ ) (Receive local model)Event 4: Receive and aggregate the local model from client $c_i$ .time+ = AggregateTime( $agg\_cost$ ) (Aggregation)Event 1: Send updated global model back to client $c_i$ .time+ = CommTime( $W_{t+1}^{(i)}$ , $c_i$ , $comm\_cost$ ) (Send updated model)end forfor $c_i$ from remaining unterminated clients dotime+ = CommTime(Termination Signal, $c_i$ , $comm\_cost$ )Event 5: Send termination signal to each client.end forClient( $c_i$ , train\_cost, comm\_cost)(Receive global model)etime+ = CommTime( $W_{t+1}^{(i)}$ , $c_i$ , $comm\_cost$ )(Send updated dlients dotime+ = CommTime( $W_t$ , $c_i$ , $comm\_cost$ )(Receive global model)Event 5: Se
	Exportation 2: FEEDES-ASYNC (EXEMPTIMES ASYNCTIONOUS FL.)Imput: Global model $W_t$ , Number of clients $n$ , Communication cost $comm\_cost$ , Training cost $train\_cost$ , Aggregation cost $agg\_cost$ Output: Final global model after all updatesServer( $W_t$ , $n$ , $agg\_cost$ , $comm\_cost$ )(Send flow 0 to $n - 1$ dotime+ = CommTime( $W_t$ , $c_i$ , $comm\_cost$ ) (Send global model)Event 1: Send the global model to client $c_i$ .end forfor $t = 1$ to $n \times T$ doif local model is received from any client $c_i$ thentime+ = CommTime( $W_{t+1}^{(i)}$ , server, $comm\_cost$ ) (Receive local model)Event 1: Send the global model back to client $c_i$ .time+ = AggregateTime( $agg\_cost$ ) (Aggregation)Event 1: Send updated global model back to client $c_i$ .time+ = CommTime( $W_{t+1}^{(i)}$ , $c_i$ , $comm\_cost$ ) (Send updated model)end forfor $c_i$ from remaining unterminated clients dotime+ = CommTime(W_{t+1}^{(i)}, $c_i$ , $comm\_cost$ ) (Send updated model)end forColspan="2">Colspan="2">CommTime(W_{t+1}^{(i)}, $c_i$ , $comm\_cost$ ) (Send updated model)end forCommTime(W_{t+1}^{(i)}, $c_i$ , $comm\_cost$ ) (Send updated model)end forCommTime(W_{t+1}^{(i)}, $c_i$ , $comm\_cost$ )Event 5: Send termination Signal, $c_i$ , $comm\_cost$ )Event 5: Send termination signal to each client. <t< td=""></t<>
	Argorithm 2: FCDES-Asylic (EXEMPTINES ASYNCTIONOUS FL.)Imput: Global model $W_t$ , Number of clients $n$ , Communication cost $comm\_cost$ , Training cost $train\_cost$ , Aggregation cost $agg\_cost$ Output: Final global model after all updatesServer( $W_t$ , $n$ , $agg\_cost$ , $comm\_cost$ )(Send flow $n - 1$ do(Send global model)Event 1: Send the global model to client $c_i$ .end foror $t = 1$ to $n \times T$ doif local model is received from any client $c_i$ thentime+ = CommTime( $W_{t+1}^{(i)}$ , server, $comm\_cost$ ) (Receive local model)Event 4: Receive and aggregate the local model from client $c_i$ .time+ = CommTime( $W_{t+1}^{(i)}$ , server, $comm\_cost$ ) (Send updated model)Event 1: Send updated global model back to client $c_i$ .time+ = CommTime( $W_{t+1}^{(i)}$ , $c_i$ , $comm\_cost$ ) (Send updated model)end ifend foror $c_i$ , from remaining unterminated clients dotime+ = CommTime(Termination Signal, $c_i$ , $comm\_cost$ )Event 5: Send termination signal to each client.end forOf $c_i$ from remaining unterminated clients dotime+ = CommTime( $W_{t+1}$ , $c_i$ , $comm\_cost$ )(Receive global model)Event 5: Send termination Signal, $c_i$ , $comm\_cost$ )Event 5: Send termination signal not received dotime+ = CommTime( $W_t$ , $c_i$ , $comm\_cost$ ) (Receive

Algorithm 3: FedDES-Compass (Exemplifies Semi-Asynchronous FL with Client Grouping)
<b>Input:</b> Global model $W_t$ , Number of clients $n$ , Number of epochs $T$ , Communication cost comm_cost,
Training cost $train\_cost$ , Aggregation cost $agg\_cost$ , Maximum local steps $max\_local\_steps$ ,
Group ratio $q_ratio$ , Group delay factor $\lambda$
Output: Final global model after all epochs
1 Server $(W_t, n, T, agg_cost, comm_cost)$
$2 time + = \text{Comm Fime}(W_t, \text{clients}, comm\_cost)$ (Simulate communication)
3 Event 1: Broadcast global model to all clients.
4 while $t < n \times 1$ do
5 If local model from client $c_i$ received then $c_i = \frac{1}{1000} \frac{1}{100$
<b>b</b> Event 4. Receive local model model from then $t_i$ .
<b>if</b> group undate is determined by compass algorithm or time limit reached <b>then</b>
9 Perform group aggregation for the buffered models
$time + = AggregateTime(agq\_cost)$ (Simulate aggregation)
$11 \qquad   \qquad t+= \text{group} $
12 <b>Event 1:</b> Send updated global model back to all clients in the group.
13 $time + = \text{CommTime}(W_{t+1}, \text{clients}, comm\_cost)$ (Simulate communication)
14 end if
15 end if
16 end while
17 for $c_i$ from unterminated clients do
18 $time+ = \text{CommTime}(\text{Termination Signal}, c_i, comm\_cost)$
19 <b>Event 5:</b> Terminate the remaining clients.
20 end for
21 Client( $c_i$ , $n$ , $train\_cost$ , $comm\_cost$ , $max\_local\_steps$ )
22 (Same as in Algorithm 2)

#### 3.3 HIGH-USABILITY: REAL-WORLD SIMULATION METHODOLOGY

As shown in Figure 2 FedDES enhances the usability and accuracy of FL simulations by integrating real-world considerations through a systematic process of logging, tracing, and profiling. The simulation framework incorporates real-world variability, such as system jitters, network delays, and stragglers, into its event generation process.

To achieve this, FedDES first logs and traces small-scale real-world FL workloads, capturing detailed event patterns such as local training, communication, and aggregation. These logs help abstract state machine transition rules for different FL paradigms. By understanding the timing and sequencing of these real-world events, FedDES accurately represents how the system transitions between states, including computation and communication phases.

FedDES also profiles computation and communication workloads. Profiling measures local training
 times, aggregation costs, and network delays for different system configurations. This data is then
 scaled to simulate large-scale deployments, with users only needing to input high-level parameters,
 such as the number of clients, system jitter distribution (like Gaussian), and network heterogeneity.





320

311

312

313 314

315

316

295 296

297



This high-level input allows developers to easily generate accurate large-scale simulations without needing detailed knowledge of the system's internals. For example, users can simulate scenarios with varying system jitters or straggler behaviors by simply adjusting the parameters, while FedDES generates the corresponding detailed events and execution traces. Using profiling and tracing ensures that FedDES produces accurate, real-world-like performance metrics, offering insights into how FL
 systems behave under different conditions.

Furthermore, all events in FedDES, including computation, communication, and aggregation, are logged with precise timestamps. This comprehensive event logging enables a deep analysis of system behavior, bottleneck identification, and performance evaluation across large-scale FL deployments. Integrating real-world profiling and tracing makes FedDES a highly usable and accurate tool for simulating and analyzing federated learning at scale.

#### 4 CORRECTNESS OF FEDDES

332

333 334

335

336

337

338

339

340 341

342

343

344

345

346

347

348 349 350

351

352 353

356

357 358

359 360

361

362

363 364

377

This section proves FedDES's correctness for modeling and FL across three aggregation paradigms: synchronous, asynchronous, and semi-asynchronous. The correctness of FedDES is demonstrated by proving that the state transitions follow the definitions of each aggregation paradigm.

**Theorem 1.** FedDES-Avg correctly simulates synchronous FL by ensuring that the system maintains the invariant that no global model update occurs until all clients have completed their local training and sent their updates.

*Proof.* Section 3.1 formalizes FedDES's state machine for FL event management. The state machine is represented by the tuple:

$$\mathcal{M} = \langle \mathcal{S}, \mathcal{E}, \mathcal{T}, \mathcal{C}, \mathcal{A} \rangle$$

where S denotes system states (clients, server, network),  $\mathcal{E}$  represents discrete events (communication, training, aggregation),  $\mathcal{T} : S \times \mathcal{E} \to S$  is the state transition function,  $\mathcal{C}$  are the conditions for state transitions, and  $\mathcal{A}$  are actions triggered by events.

We define the state transitions for FedDES-Avg as:

$$\mathcal{T}_{\text{FedDES-Avg}}: S_{\text{server}}(t), S_{\text{client}_i}(t) \longrightarrow S_{\text{server}}(t+1), S_{\text{client}_i}(t+1),$$

where  $S_{\text{server}}(t)$  is the server's state at time t, and  $S_{\text{client}_i}(t)$  is the state of client  $c_i$  at time t.

The correctness of FedDES-Avg is based on the following invariant:

 $\forall t, \quad C_{\text{all\_clients\_done}}(t) \implies \mathcal{A}_{\text{aggregate}}(W_{t+1}^{(1)}, \dots, W_{t+1}^{(n)}),$ 

indicating that aggregation occurs only when all client updates are received.

1. Initial Condition: At t = 0, the server distributes the global model  $W_0$  to all clients, triggering the state transition:

$$S_{\text{server}}(0) \to S_{\text{server}}(1), \quad S_{\text{client}_i}(0) \to S_{\text{client}_i}(1),$$

where  $S_{\text{server}}(1)$  reflects the server waiting for client updates.

2. State Transition: For t > 0, each client  $c_i$  completes local training and sends its updated model  $W_{t+1}^{(i)}$  to the server. The condition  $C_{\text{all\_clients\_done}}(t)$  is satisfied only when all client updates are received, triggering the aggregation action:

$$\mathcal{A}_{\text{aggregate}}(W_{t+1}^{(1)},\ldots,W_{t+1}^{(n)}).$$

365 366 367 368 3. Invariant Preservation: The transition function guarantees that the invariant is maintained at every time step. If any client  $c_i$  has not completed its update, the server continues waiting, preventing premature aggregation and ensuring synchronous behavior.

Thus, FedDES-Avg correctly models synchronous FL by ensuring that no aggregation occurs until all client updates are received, preserving the synchronous nature of the system.  $\Box$ 

Theorem 2. FedDES-Async correctly simulates asynchronous FL by ensuring that each client up date is aggregated as soon as it is received, without waiting for updates from other clients.

The proof of Theorem 2 is provided in Appendix A.1

Theorem 3. FedDES-Compass correctly simulates semi-asynchronous FL by ensuring that client
 updates within each group are synchronized while groups themselves operate asynchronously.

The proof of Theorem 3 is provided in Appendix A.2

### <sup>378</sup> 5 EVALUATION

379 380 381

This section evaluates FedDES' accuracy by comparing the simulation's event distributions against real-world experimental results. We aim to determine how well the simulated FL events reflect the system dynamics compared to a real FL environment.

#### 384 5.1 EXPERIMENTAL SETUP

**Testbed and Software Settings:** We use a Vision Transformer (ViT) Alexey (2020) on the parti-386 tioned CIFAR-10 Krizhevsky et al. (2009) dataset for both real-world and simulated FL experiments. 387 The real-world FL experiments involve more than 1,000 clients and three aggregation strategies: Fe-388 dAvg, FedAsync, and FedCompass. Our FedDES simulations mirror these strategies using the *class* 389 partitioning method. The real experiments are performed on the NCSA Delta High-Performance Computing (HPC) system, while the simulations run on a server equipped with two NVIDIA A40 390 GPUs and an Intel(R) Xeon(R) Gold 6336Y CPU. The real-world experiment captures communica-391 tion events across multiple clients during an FL task, while the same FL task is simulated through 392 FedDES with identical client configurations, communication steps, and task settings. In both the 393 simulation and experiment, events are logged with timestamps, indicating when clients send or re-394 ceive models and when the server performs aggregation. 395

Metric Settings We assess the similarity between the event distributions of real-world FL exper iments and FedDES simulations across three dimensions: time, communication steps, and client ID. Evaluating the similarity between actual and simulated event distributions in only one or two dimensions (e.g., time alone or step-wise comparisons) would fail to capture the full complexity of the system. By analyzing the distribution of events across all three dimensions—time, step, and client ID—we ensure a more holistic and accurate comparison of the behavior between real-world experiments and simulations. The metrics we employ for this 3D evaluation are:

- Kullback-Leibler (KL) Divergence Kullback & Leibler (1951): This metric measures how much the
   simulated distribution deviates from the real distribution across the three dimensions. A lower value
   indicates a higher similarity between the real and simulated data.
- Jensen-Shannon (JS) Divergence Lin (1991): This symmetric measure compares the real and simulated distributions over time, step, and client ID. It captures the overall similarity of the distributions in three dimensions.

Bhattacharyya Distance Bhattacharyya (1943): The Bhattacharyya distance measures the overlap between the real and simulated distributions in the three-dimensional space of time, communication steps, and client ID. It is particularly sensitive to differences in distribution shape and spread.

By evaluating these metrics in three dimensions, we ensure that the simulation's accuracy is assessed not only with respect to time but also to the sequence of communication steps and the behavior of individual clients. This approach allows us to capture the full complexity of the system dynamics and ensures that the simulation faithfully reproduces the interaction patterns observed in real-world federated learning environments.

417 418

## 410 5.2 RESULTS AND ANALYSIS

Table 12 and 3 compare the event distribution similarities between FedDES-Avg, FedDES-Async, and FedDES-Compass and their respective real workloads across three different simulation settings: simulation with no noise ( $\mathcal{N}(0,0)$ ), simulation with Gaussian noise ( $\mathcal{N}(0,0.12)$ ), and simulation with system heterogeneity observed from real workloads. We employ the discussed metrics to measure the divergence between the simulated and real workload distributions.

As shown in Table 1 across all three metrics (*KL*, *Jensen-Shannon*, and Bhattacharyya), the simulation with no noise ( $\mathcal{N}(0,0)$ ) has the highest divergence from the real workload. Introducing Gaussian noise ( $\mathcal{N}(0,0.12)$ ) slightly improves the similarity, as reflected by marginal decreases in all three metrics, suggesting that noise approximates the behavior of real systems better. The simulation with system heterogeneity further improves the similarity, showing the lowest values across all metrics. This demonstrates that accounting for system jitters and heterogeneity aligns the simulation more closely with real-world workloads. Figure 3 visualize the event distribution comparison among simulation settings (FedDES-Avg) and real-world experiments (FedAvg), showcasing visually perceivable similarities and correct communication patterns in event distributions. When accounting for system jitter and heterogeneity, the simulation yields highly accurate execution time results in just 1.47s, with a simulated execution time error of just 0.69% (2336s simulated vs. 2320s real). 

Table 1: Comparison of 3D event distribution evaluation in FedDES-Avg and real workloads across KL, Jensen-Shannon, and Bhattacharyya distances under different simulation settings.

Simulation Settings / Metrics	Kullback-Leibler	Jensen-Shannon	Bhattacharyya
Simu. w/ $\mathcal{N}(0,0)$	9.125	0.495	1.25
Simu. w/ $\mathcal{N}(0, 0.12)$	9.121	0.494	1.249
Simu. w/ system heterogeneity	9.106	0.491	1.24



Figure 3: Event Distribution of FedDES for Synchronous FL training compared to Real Runs. The S1-4 in Communication Steps denote for: S1, Server sends global model; S2, Clients receive global model; S3, Clients send local model; S4, Server receives local model. Note that there's an explicit synchronization block after clients receive global model.

As shown in Table 2, similar to FedDES-Avg, the no-noise simulation shows the highest divergence (KL: 9.687, JS: 0.546, Bhattacharyya: 1.532). Introducing Gaussian noise improves the distribu-tion similarity slightly, as indicated by minor reductions in KL and Jensen-Shannon Divergence, although *Bhattacharyya* Distance slightly increases. The simulation with system heterogeneity per-forms best, showing small reductions across all three metrics, indicating that heterogeneity captures real-world conditions better than noise alone. Figure 4 visualize the event distribution comparison among simulation settings (using FedDES-Async) and real-world experiments (FedAsync), which also showcased visually perceivable similarities and correct communication patterns in event distri-butions. When accounting for system jitter and heterogeneity, the simulation yields near-identical execution time results in just 2.03s, with a simulated execution time error of just 0.04% (2596s simulated vs. 2595s real). 

Table 2: Comparison of 3D event distribution evaluation in FedDES-Async and real workloads across KL, Jensen-Shannon, and Bhattacharyya distances under different simulation settings. 

Simulation Settings / Metrics	Kullback-Leibler	Jensen-Shannon	Bhattacharyya
Simu. w/ $\mathcal{N}(0,0)$	9.687	0.546	1.532
Simu. w/ $\mathcal{N}(0, 0.12)$	9.721	0.547	1.544
Simu. w/ system heterogeneity	9.658	0.544	1.533

As shown in Table 3, for FedDES-Compass, we observe a larger difference in divergence. The sys-tem heterogeneity simulation shows the lowest divergence across all metrics, with KL divergence dropping to 8.709 from 9.591 in the no-noise setting. Gaussian noise also reduces divergence com-pared to the no-noise scenario. Still, the reduction is more significant in FedDES-Compass than in FedDES-Avg or FedDES-Async, indicating that this semi-asynchronous approach might be more sensitive to noise and system variations. The Bhattacharyya Distance improves significantly in the



Figure 4: Event Distribution of FedDES for Asynchronous FL training compared to Real runs. Additionally, *S5* in Communication Steps denote for Server signals the finalization.

system heterogeneity setting, showing that this simulation closely approximates real-world event distributions. Figure <sup>5</sup>/<sub>5</sub> visualizes the event distribution comparison between simulation settings (using FedDES-Compass) and real-world experiments (FedCompass). Due to the dynamic grouping in FedCompass, which dynamically adjusts client training step size based on client speed, the event distribution exhibits higher randomness and sensitivity to system jitter and heterogeneity. Nonetheless, accounting for system jitter and heterogeneity in the simulation yields accurate results in only **2.77s**, with a simulated execution time error of just **1.95%** (2309s simulated vs. 2355s real).

Table 3: Comparison of 3D event distribution evaluation in FedDES-Compass and real workloads across *KL*, *Jensen-Shannon*, and *Bhattacharyya* distances under different simulation settings.

Simulation Settings / Metrics	Kullback-Leibler	Jensen-Shannon	Bhattacharyya
Simu. w/ $\mathcal{N}(0,0)$	9.591	0.539	1.49
Simu. w/ $\mathcal{N}(0, 0.12)$	9.392	0.528	1.425
Simu. w/ system heterogeneity	8.709	0.489	1.214



Figure 5: Event Distribution of FedDES for Semi-Async FL training compared to Real Runs. Note that clients have to explicitly block for synchronization before the amount of waiting clients reaches the preset threshold.

#### 

#### 6 CONCLUSION

FedDES provides an efficient, scalable solution for simulating large-scale FL systems using Discrete
Event Simulation (DES). By modeling client selection, training, communication, and aggregation
as discrete events, it enables precise analysis of various FL strategies, including synchronous, asynchronous, and semi-asynchronous paradigms. FedDES is framework-agnostic, allowing researchers
to integrate custom datasets and models, while accurately simulating heterogeneous client behaviors
and network conditions. Evaluations on over 1,000 clients show FedDES delivers high accuracy,
with less than 2% error compared to real-world experiments. This makes FedDES a powerful tool
for debugging, performance analysis, and prototyping new FL algorithms.

# 540 REFERENCES 541

542 543	fl-mobile simulator. https://github.com/PaddlePaddle/PaddleFL/tree/ master/python/paddle fl/mobile.
:11	
	Federated Learning Simulator (FLSim). https://github.com/facebookresearch/
:45	FLSim?tab=readme-ov-file.
.47	
.40	Dosovitskiy Alexey. An image is worth $16x16$ words: Transformers for image recognition at scale.
48	arxiv preprint arxiv: 2010.11929, 2020.
49	Jerry Banks Discrete event system simulation Pearson Education India 2005
50	Serry Danks, Discrete even system simulation, reason Davedion mela, 2005.
51	Jerry Banks, John S. Carson, Barry L. Nelson, and David M. Nicol. Discrete-event system simula-
52 53	tion. 2010.
54	Daniel J Beutel, Taner Topal, Akhil Mathur, Xinchi Qiu, Javier Fernandez-Marques, Yan Gao,
5	Lorenzo Sani, Kwing Hei Li, Titouan Parcollet, Pedro Porto Buarque de Gusmão, et al. Flower:
6	A friendly federated learning research framework. arXiv preprint arXiv:2007.14390, 2020.
7	
58	Anil Bhattacharyya. On a measure of divergence between two statistical populations defined by
9	their probability distributions. Bulletin of the Calcutta Mathematical Society, 35:99–109, 1943.
0	Keith Bonawitz Hubert Fichner Wolfgang Grieskamp Dóra Huba Alex Ingerman Vladimir
1	Ivanov Chloé Kiddon Jakub Konecny Steven Mazzocchi H Brendan McMahan et al Towards
2	federated learning at scale: System design. In <i>Proceedings of Machine Learning and Systems</i>
3	2019 (MLSvs), pp. 374–388, 2019.
<u>л</u>	
5	Henri Casanova. Simgrid: A toolkit for the simulation of application scheduling. In Proceedings
6	First IEEE/ACM International Symposium on Cluster Computing and the Grid, pp. 430–437.
7	IEEE, 2001.
r 5	
2	Cheng Chen, Ziyi Chen, Yi Zhou, and Bhavya Kalikhura. Fedcluster: Boosting the convergence of federated learning ris cluster and in a 2020 IEEE later stimul Conformation on Ris Data (Ris
9	Data) nr 5017 5026 IEEE 2020
4	<i>Duiu)</i> , pp. 5017–5020. IEEE, 2020.
1	Emily Ekaireb, Xiaofan Yu, Kazim Ergun, Quanling Zhao, Kai Lee, Muhammad Huzaifa, and Ta-
2	jana Rosing. ns3-fl: Simulating federated learning with ns-3. In Proceedings of the 2022 Work-
3	<i>shop on ns-3</i> , pp. 97–104, 2022.
+	
	Chaoyang He, Songze Li, Jinhyun So, Xiao Zeng, Mi Zhang, Hongyi Wang, Xiaoyang Wang, Pra-
	neeth Vepakomma, Abhishek Singh, Hang Qiu, et al. Fedml: A research library and benchmark
	tor federated machine learning. arXiv preprint arXiv:2007.13518, 2020.
	Peter Kairouz H Brendan McMahan Brendan Avent Aurálian Ballat Mahdi Bannis Ariun Nitin
)	Bhagoii Kallista Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, et al. Ad-
	vances and open problems in federated learning Foundations and trends® in machine learning
	14(1-2):1-210, 2021.
	Sai Praneeth Karimireddy, Satyen Kale, Mehryar Mohri, Sashank Reddi, Sebastian Stich, and
	Ananda Theertha Suresh. Scaffold: Stochastic controlled averaging for federated learning. In
	International conference on machine learning, pp. 5132–5143. PMLR, 2020.
;	
7	Alex Kriznevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images.
3	2009.
1	Solomon Kullback and Richard A Leibler. On information and sufficiency. The annals of mathe-
)	matical statistics, 22(1):79–86, 1951.
	Baochun Li, Ningxin Su, Chen Ying, and Fei Wang. Plato: An open-source research framework for
	production federated learning. In Proceedings of the ACM Turing Award Celebration Conference- China 2023, pp. 1–2, 2023a.

606

- Qinbin Li, Zeyi Wen, Zhaomin Wu, Sixu Hu, Naibo Wang, Yuan Li, Xu Liu, and Bingsheng He. A survey on federated learning systems: Vision, hype and reality for data privacy and protection. *IEEE Transactions on Knowledge and Data Engineering*, 35(4):3347–3366, 2021a.
- Tian Li, Maziar Sanjabi, Ahmad Beirami, and Virginia Smith. Fair resource allocation in federated learning. *arXiv preprint arXiv:1905.10497*, 2019.
- Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith.
   Federated optimization in heterogeneous networks. *Proceedings of Machine learning and systems*, 2:429–450, 2020.
- Tian Li, Shengyuan Hu, Ahmad Beirami, and Virginia Smith. Ditto: Fair and robust federated
   learning through personalization. In *International conference on machine learning*, pp. 6357–6368. PMLR, 2021b.
- Zilinghan Li, Pranshu Chaturvedi, Shilan He, Han Chen, Gagandeep Singh, Volodymyr Kindratenko, Eliu A Huerta, Kibaek Kim, and Ravi Madduri. Fedcompass: efficient cross-silo federated learning on heterogeneous client devices using a computing power aware scheduler. *arXiv* preprint arXiv:2309.14675, 2023b.
- Jianhua Lin. Divergence measures based on the shannon entropy. *IEEE Transactions on Information theory*, 37(1):145–151, 1991.
- Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas.
   Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, pp. 1273–1282. PMLR, 2017.
- Philipp Moritz, Robert Nishihara, Stephanie Wang, Alexey Tumanov, Richard Liaw, Eric Liang, Melih Elibol, Zongheng Yang, William Paul, Michael I Jordan, et al. Ray: A distributed framework for emerging {AI} applications. In *13th USENIX symposium on operating systems design and implementation (OSDI 18)*, pp. 561–577, 2018.
- Vaikkunth Mugunthan, Anton Peraire-Bueno, and Lalana Kagal. Privacyfl: A simulator for privacy preserving and secure federated learning. In *Proceedings of the 29th ACM International Confer- ence on Information & Knowledge Management*, pp. 3085–3092, 2020.
- John Nguyen, Kshitiz Malik, Hongyuan Zhan, Ashkan Yousefpour, Mike Rabbat, Mani Malek, and
   Dzmitry Huba. Federated learning with buffered asynchronous aggregation. In *International Conference on Artificial Intelligence and Statistics*, pp. 3581–3607. PMLR, 2022.
- 627
   628
   629
   629 Shawn Pearce. Ross: Rensselaer's optimistic simulation system user's guide christopher d. carothers david bauer. 2002.
- Jae Hun Ro, Ananda Theertha Suresh, and Ke Wu. Fedjax: Federated learning simulation with jax.
   *arXiv preprint arXiv:2108.02117*, 2021.
- Minseok Ryu, Youngdae Kim, Kibaek Kim, and Ravi K Madduri. Appfl: open-source software
   framework for privacy-preserving federated learning. In 2022 IEEE International Parallel and
   Distributed Processing Symposium Workshops (IPDPSW), pp. 1074–1083. IEEE, 2022.
- Ziteng Sun, Peter Kairouz, Ananda Theertha Suresh, and H Brendan McMahan. Can you really backdoor federated learning? *arXiv preprint arXiv:1911.07963*, 2019.
- Cong Xie, Sanmi Koyejo, and Indranil Gupta. Asynchronous federated optimization. *arXiv preprint arXiv:1903.03934*, 2019.
- Haotian Xu, Zhaorui Zhang, Sheng Di, Benben Liu, Alharthi Khalid, and Jiannong Cao. Fedfa: A
  fully asynchronous training paradigm for federated learning. *arXiv preprint arXiv:2404.11015*, 2024.
- Qiang Yang, Yang Liu, Tianjian Chen, and Yongxin Tong. Federated machine learning: Concept and applications. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 10(2):1–19, 2019.
- 647 Yue Zhao, Meng Li, Liangzhen Lai, Naveen Suda, Damon Civin, and Vikas Chandra. Federated learning with non-iid data. *arXiv preprint arXiv:1806.00582*, 2018.