

DynamicFL: Balancing Communication Dynamics and Client Manipulation for Federated Learning

Bocheng Chen, Nikolay Ivanov, Guangjing Wang, Qiben Yan
Computer Science & Engineering, Michigan State University, East Lansing, USA.
Email:{chenboc1, ivanovn1, wanggu22, qyan}@msu.edu

Abstract—Federated Learning (FL) is a distributed machine learning (ML) paradigm, aiming to train a global model by exploiting the decentralized data across millions of edge devices. Compared with centralized learning, FL preserves the clients’ privacy by refraining from explicitly downloading their data. However, given the geo-distributed edge devices (e.g., mobile, car, train, or subway) with highly dynamic networks in the wild, aggregating all the model updates from those participating devices will result in inevitable long-tail delays in FL. This will significantly degrade the efficiency of the training process. To resolve the high system heterogeneity in time-sensitive FL scenarios, we propose a novel FL framework, DynamicFL, by considering the communication dynamics and data quality across massive edge devices with a specially designed client manipulation strategy. DynamicFL actively selects clients for model updating based on the network prediction from its dynamic network conditions and the quality of its training data. Additionally, our long-term greedy strategy in client selection tackles the problem of system performance degradation caused by short-term scheduling in a dynamic network. Lastly, to balance the trade-off between client performance evaluation and client manipulation granularity, we dynamically adjust the length of the observation window in the training process to optimize the long-term system efficiency. Compared with the state-of-the-art client selection scheme in FL, DynamicFL can achieve a better model accuracy while consuming only 18.9% – 84.0% of the wall-clock time. Our component-wise and sensitivity studies further demonstrate the robustness of DynamicFL under various real-life scenarios.

I. INTRODUCTION

Federated Learning (FL) enables globally distributed devices to participate in the model training without leaking raw data [1], [2]. FL has been applied across various fields where millions of edge devices participate in the training with their local data, including connected autonomous vehicle’s motion planning [3]–[5], activity prediction [6]–[8], text prediction [9], [10], and smart home automation [11]–[14].

Status Quo and Limitations. In the FL training process, a round of training comprises two steps: first, the server dispatches the model to each client; second, the server aggregates the uploaded model updates once all clients have completed their training. It can take hundreds of rounds for the model to achieve its final accuracy, which can take hours. One key problem in FL is achieving the desired model accuracy within a short time, commonly referred to as **time-to-accuracy** [15], [16]. On the server side, reducing the time required to achieve the final model accuracy can minimize the overall energy cost and bandwidth consumption for each client. On the client side,

each client can receive the best model in less time. Time-to-accuracy is particularly important in time-sensitive FL tasks, such as collision avoidance for autonomous driving [3], [4], where model updates from vehicles should be gathered in real-time. In recent years, as the FL system relies on the participation of a large number of clients, there has been increasing attention on improving the efficiency of FL systems. The client selection method [15], [16] is the state-of-the-art (SOTA) method for improving the efficiency of the FL system. This method has shown significant improvement in the time-to-accuracy of FL training.

However, in a realistic network scenario, clients often experience varying and fluctuating network conditions. This presents a significant challenge for FL systems. For example, when a client selected by the server happens to be in a tunnel, the downgraded network connectivity creates a bottleneck in the training round and results in a significant delay during the model upload stage. Moreover, for clients prioritized by the selection algorithm due to their fast communication speed in previous rounds, an unstable network can cause adverse outcomes when they are faced with unstable networks.

We conduct an experiment on how real-world bandwidth data [17], [18] affects the SOTA FL system [15], with results shown in Figure 1(a). Following the experimental settings in Oort [15], we select 100 clients in each round for the image classification task using the FEMNIST dataset. Our findings indicate that the time-to-accuracy can be extended by 20%-30% compared to scenarios without realistic network fluctuations. This degradation leads to prolonged delays during training rounds, which occur frequently throughout the training process. Such delays could hinder the parameter updates on devices that hold essential data, further exacerbating the difficulties in achieving efficient FL training.


Overview of Our Approach. To improve the FL system efficiency, we propose DynamicFL, an FL framework designed to actively manipulate the selection of well-performed clients in a real-world *dynamic* network environment. Our approach first leverages a bandwidth prediction model to help with client selection, enabling the server to prioritize clients with fast network connections and reducing the selection of bottleneck clients. In contrast, existing greedy selection algorithms [15], [16] are susceptible to dynamic networks in the wild. These algorithms rely on data from a single previous round, making accurate bandwidth prediction impossible, as illustrated in Figure 1(b). Additionally, the performance of existing greedy

algorithms in client selection is compromised by dynamic network conditions, as a client can perform worse due to sudden changes in the network in the next training round. To address these challenges, we introduce a long-term observation window to gather data to enable accurate bandwidth prediction in DynamicFL. Within this window, we pause the client selection, allowing us to monitor the overall performance of clients and prioritize those with stable network connections and high-quality data. To strike a balance between long-term scheduling and fine-grain control of client selection, we adaptively adjust the length of the observation window throughout the training process, optimizing the overall efficiency of the FL system from a long-term perspective. In the end, DynamicFL can work seamlessly with other optimization methods for client selection. We overcome the following challenges in designing DynamicFL:

- **Challenge #1 [Accounting for bandwidth]:** It is hard to achieve accurate bandwidth prediction relying solely on data from the previous round in the SOTA solutions.
- **Challenge #2 [Client selection accuracy in a long-term]:** Current greedy strategy in client selection is more susceptible to a dynamic bandwidth environment.
- **Challenge #3 [Trade-off between long-term schedule and time efficiency]:** There is a trade-off between the confidence of bandwidth prediction and observation time used to make client selection.

We implement DynamicFL and conduct experiments on four datasets with real-world user data. Compared with the SOTA methods, DynamicFL significantly improves the time-to-accuracy in the FL training process. DynamicFL achieves a higher final test accuracy while preserving user privacy. We implement DynamicFL on top of two FL paradigms with different optimization methods [15], [19], highlighting its compatibility with existing FL approaches to enhance their time-to-accuracy performance in real-world scenarios with dynamic network conditions. We summarize our contributions as follows:

- To overcome Challenge #1, we propose a framework for FL in real-world networks with a bandwidth prediction module, which prioritizes clients with better networking conditions.
- To overcome Challenge #2, we propose a long-term greedy strategy in each observation window, which temporally freezes client selection, improves the system robustness in a dynamic network, and optimizes the system efficiency with client selection.
- To overcome Challenge #3, we propose an algorithm to determine the size of the bandwidth observation and data accumulation window. This algorithm effectively handles the trade-off between client performance evaluation and client manipulation granularity.
- We implement DynamicFL and integrate it with other client selection algorithms, allowing the existing FL paradigms to perform better in real-world dynamic networks.



(a) Dynamic bandwidth affects the time-to-accuracy ('w' indicates with dynamic bandwidth, 'w/o' indicates without it).

(b) Prediction results with different window sizes in LSTM model: larger window size yields better prediction.

Figure 1. Illustration of the limitation of SOTA work and challenge of designing DynamicFL.

II. BACKGROUND AND MOTIVATION

There are various approaches for optimizing time-to-accuracy while achieving reasonable model accuracy. A recent SOTA FL system, called Oort [15], proposes a client selection strategy that considers both data and system heterogeneity in its system design. The main idea of client selection is to choose clients who have training data that can contribute the most to the model training and finish local training in a short time period. Although those methods can improve system efficiency and achieve better time-to-accuracy performance, they assume all the clients upload updates in a stable network. *In this work, we show that dynamic bandwidth can downgrade the performance and undermine the optimization methods on client selection.*

We conduct preliminary experiments on Oort to show how the dynamic network undermines its client selection optimization methods, as shown in Figure 1(a). We use FEMNIST [20] dataset to check the time-to-accuracy and final accuracy on the real-world bandwidth data [21], [22]. Our findings are summarized in the following two aspects.

Impractical Client Selection Setting in Wireless Network. From the result, we find that the time-to-accuracy in a dynamic bandwidth environment is much longer compared with the previous static-bandwidth data [15]. This problem motivates us to utilize a bandwidth prediction model for selecting clients. Without relying on any additional data, collected bandwidth data can be used in guiding client selection. The bandwidth prediction model helps to select clients with high-quality network conditions and reduces possible bottleneck clients.

Inherent Limitation of Short-term Greedy Algorithm. The short-term greedy algorithm in previous work is more susceptible to bandwidth dynamics in a dynamic network environment. The major limitation of the greedy client selection algorithm is the fact that the clients' model update uploading time in the last round does not indicate the communication quality in future rounds. With dynamic bandwidth data, the client with the highest utility in the last round might become the bottleneck in the future. Here, we use a long-term greedy strategy to collect more data and make a longer observation of client performance.




Figure 2. DynamicFL Framework.

III. DYNAMICFL FRAMEWORK

A. DynamicFL Overview

We first provide an overview of how DynamicFL improves the efficiency between chunks of the window and among the training rounds inside the window. Then, we discuss each step and component of DynamicFL in details.

Figure 2 shows the DynamicFL architecture with the proposed observation window. DynamicFL aims at gathering data within the observation window and using this data to control client selection based on client utility and bandwidth prediction. Inside the window, DynamicFL accumulates the bandwidth data and the utility score of clients' performance in each round and temporarily freezes the client selection in ②. After several rounds of observation, DynamicFL predicts the bandwidth for each client based on historical bandwidth. The accumulated utility (i.e., the performance of a client) is also averaged to show whether the client has stable and high performance. The average utility and bandwidth prediction results in ① are used for modifying feedback and selecting clients for the next window. Next, we describe each step of DynamicFL workflow in detail.

The server first initializes the model and dispatches it to each selected client in the pool, where clients have different configurations [15], [16], [19]. Then, the server starts watching the performance of each client in the observation window. On the client side, each client starts training the model with local data, followed by transmitting the model updates to the server. All model updates collected by the server are averaged and contributed to the shared model. Based on the training time duration and importance of the update, a utility score (i.e., the ranking-based info) will be calculated and assigned to each device client. The utility score and training duration are accumulated in the observation window. Clients continue to train and update the model and share the updated model in each round but without further client selection. When the number of rounds of training reaches the window size, the server makes predictions over the following rounds based on the bandwidth of each device. The server will select the client

based on the history of utility performance and bandwidth prediction results.

B. Prediction of Dynamic Network Conditions

To improve time-to-accuracy, DynamicFL should be able to give rewards to clients who will have a fast network connection speed and penalize clients who might suffer from bottlenecks. This is not trivial since the server can know a client has become the bottleneck only after a delay has occurred, *how can a server impose a penalty to a client that may become a bottleneck before it appears in the training round?*

To address the challenge, we propose to establish a communication protocol between the server and clients considering the bandwidth prediction method.

1) *Bandwidth prediction module*: Our first challenge is to obtain the bandwidth data without requesting any additional information in order to preserve privacy.

Bandwidth Calculation. As the server has the knowledge of the time cost of each client, we can calculate the bandwidth information as follows:

$$T(C_i, R_i) = \Delta \text{Comp}(C_i, R_i) + \Delta \text{Comm}(C_i, R_i)$$

$$\Delta \text{Comm}(C_i, R_i) = \frac{\mathbb{U}(\text{pull}) + \mathbb{U}(\text{push})}{b_t}. \quad (1)$$

Here, the time $T(C_i, R_i)$ that each client C_i consumes in round R_i comprised of two components: server-client communication ΔComm and local client computation ΔComp . Previous work [16] showed that the communication step usually takes much more time (about 90%) than device computation (about 10%). For each client, communication time is proportional to the bandwidth, because the model update size $\mathbb{U}$ that a client needs to upload changes slightly during the training. Thus, we directly use the time cost of each client in the bandwidth prediction instead of collecting more data.

Bandwidth Penalty. In DynamicFL, after collecting clients' data for several rounds of training, the server can use the historical data for drawing predictions. The prediction will be used in modifying feedback, the utility score for client i , for client selection in the following round. Based on the existing feedback function in Oort, we include an additional factor in the feedback as follows:

$$\text{Util}(i) = \underbrace{|B_i * \mathbb{F}| \sqrt{\frac{1}{|B_i|} \sum_{k \in B_i} L(k)^2}}_{\text{Statistical utility } U(i)} \times \underbrace{\left(\frac{T * \mathbb{F}}{t_i} \right)^{\mathbb{I}(T < t_i) \times \alpha}}_{\text{System utility}} \quad (2)$$

$$\mathbb{F} = \text{Norm}(P(b_H)) \quad (3)$$

In the statistical utility part, B_i is the local data samples set, and $L(k)^2$ is the training loss of data sample k . We add the bandwidth prediction result $\mathbb{F}$ as a factor to adjust statistical utility in feedback. In the system utility part, T is the developer-preferred duration, t_i is the wall clock time duration for client i and α is the penalty factor. Oort uses $\mathbb{I}(x)$ as an indicator function that takes value one if x is true and 0 otherwise. We use $\mathbb{F}$ to adjust the system utility, where

$\mathbb{F}$ is the normalized result of bandwidth prediction model P on data history b_H .

In Eq. (2), the bandwidth is important in selecting clients for the next round. The result of bandwidth prediction is normalized due to the different ranges of bandwidth on different devices. We prefer to select clients in a fast network with high utility. Our reward and penalty factor allows us to exclude a client from training when its network quality is inadequate.

Offline Bandwidth Prediction Model. DynamicFL uses an LSTM model to predict the bandwidth of clients in the next round. To ensure that the offline model meets the memory and latency requirements, we use the lightweight three-layers LSTM model. Given hundreds of bandwidth traces in the FL training process, to make DynamicFL privacy-preserving, we only use one bandwidth trace to train the offline model and leave the other hundreds of client bandwidth traces available for clients in the system evaluation. More details on the mechanism of assigning bandwidth traces to devices can be found in Section IV.

2) *Bandwidth prediction in DynamicFL:* Algorithm 1 shows how the bandwidth prediction module cooperates with the current client selection strategy in making decisions in a dynamic network environment. Since Oort [15] selects clients based on client utility and time duration, our goal is to modify the client duration time with our bandwidth prediction module (Lines 4-26). Inside the prediction function, DynamicFL uses an observation window to keep the bandwidth data for every client at the current time, which comes from the communication data used in the current round (Lines 9-10). After collecting all the bandwidth data of the last round, the model continues to update the parameters and train for the next round but stops selecting new clients (Lines 8-12). The server starts to make client selections until the number of observed rounds is equal to the window size. It uses the bandwidth history to predict the communication of each client for the next round (Lines 16-25). After predicting the bandwidth, we compute a reward or penalty that is dependent on the ratio between the predicted bandwidth and a predefined threshold. This reward or penalty is then incorporated into the training duration time and the client utility calculation.

When applying the reward and penalty scores, α increases as a reward when the bandwidth prediction result is higher than the reward threshold. Thus, the server increases the priority of those clients (Line 20). The penalty score reduces the occurrence of bottleneck clients when the prediction result is close to zero (Line 23). When the prediction result locates between the award/penalty thresholds, DynamicFL does not modify the client feedback (Line 25). While the SOTA solutions use exploration and exploitation strategy in client selection to keep clients with high utility scores and explore for new clients, DynamicFL utilizes varying award and penalty scores, allowing some devices to remain in the training due to the high client utility and future communication ability.

C. Long-Term Scheduling Strategy

In this section, we analyze the long-term scheduling strategy in dynamic network environments and design an algorithm

Algorithm 1: Bandwidth Prediction in FL.

Input: Client set $\mathbb{C}$, local training iteration base i , bandwidth penalty threshold TH_L , bandwidth reward threshold TH_H , current round R_i , bandwidth prediction model P , client utility set $\mathbb{U}$, client time duration set $\mathbb{D}$, observation window size W , adjustment coefficient c .

Output: Modified Feedback $(\mathbb{U}, \mathbb{D})$.

```

/* Initialize global variables */
1  $B_H \leftarrow \text{Function Observation}(\mathbb{C}, R_i, W)$ 
2  $\mathbb{U}, \mathbb{D} \leftarrow \text{PerformanceFeedback}(R_i)$ 
3  $\mathbb{U}, \mathbb{D} \leftarrow \text{Function Prediction}(B_H, \mathbb{U}, \mathbb{D})$ 
4  $\mathbb{C}_U \leftarrow \text{ClientSelection}(\mathbb{U}, \mathbb{D}, \mathbb{C})$ 
5 Function Observation  $(\mathbb{C}, R_i, W)$ 
/* Initialize bandwidth historical set */
6  $B_H \leftarrow \emptyset$ 
7 While  $R_i \bmod W$ 
/* Record bandwidth for each client */
8 Loop client  $j \in \mathbb{C}$  :
9    $t \leftarrow \text{CurrentTime}()$ 
10   $b_t^j = \text{GetCommData}(j, t)$ 
11   $B_H^j.append(b_t^j)$ 
12   $i \leftarrow i + 1$ 
13  Freeze Client Selection
14  ▷ Continue Model Training and Model Update
15 Return  $B_H$ 
16 Function Prediction  $(\mathbb{C}, TH_H, TH_L, R_i, P, \mathbb{U}, \mathbb{D}, c)$ 
/* Bandwidth prediction for each client */
17 Loop client  $j \in \mathbb{C}$  :
18    $\alpha = P(B_H^j)$ 
19   if  $\alpha \geq TH_H$  then
20      $\alpha \leftarrow (-\log(1 - \alpha) + c)$ 
21   else
22     if  $\alpha \leq TH_L$  then
23        $\alpha \leftarrow \exp(\alpha + c);$ 
24     else
25        $\alpha \leftarrow 1;$ 
26   end
27 end
/* Update client selection feedback */
28  $\mathbb{U}(j) \leftarrow \mathbb{U}(j) \times \alpha$ 
29  $\mathbb{D}(j) \leftarrow \frac{\mathbb{D}(j)}{\alpha}$ 
30 Return  $(\mathbb{U}, \mathbb{D})$ 

```

to offer a better client selection method. The clients with the current best performance found by the greedy algorithm cannot guarantee the best selection in the long run. The short-term client selection suffers from the instability caused by the dynamic bandwidth data. A device that underperforms due to a temporary network fluctuation is less likely to participate in future training. The goal of DynamicFL is to provide a longer observation window to slow down the decision-making process. The long-term greedy strategy in DynamicFL makes the system more robust and efficient.

In Algorithm 2, we show the long-term greedy strategy in DynamicFL. The goal is to evaluate the overall performance of each client who participates in the past several rounds. We use an accumulation window, to sum up the time duration for each client at every round. This average value represents the overall utility of a client. We freeze the client selection when accumulating the data. The server resumes client selection until data is collected enough to make decisions based on the overall understanding of each client. In this way, we select clients with stable performance and high utility scores.

Algorithm 2: Long-term Greedy Strategy.

Input: Client set $\mathbb{C}$, local training iteration base i , current round R_i , observation window size W .
Output: General time duration set of client performance $\mathbb{D}$.

```

1 Function LongTermGreedy ( $\mathbb{C}, R_i, W$ )
2   While  $R_i \bmod W$ 
3     /* Keep time cost of every client */
4     Loop client  $j \in \mathbb{C}$  :
5        $d_t^j \leftarrow \text{ClientTimeDuration}(j)$ 
6        $D^j \leftarrow D^j + d_t^j$ 
7       Freeze Client Selection;
8        $i \leftarrow i + 1$ 
9     Loop client  $j \in \mathbb{C}$  :
10       $D^j \leftarrow \frac{D^j}{W}$ 
11   Return  $\mathbb{D}$ 

```

D. Trade-off Between Bandwidth Prediction and Long-Term Schedule

The efficiency of DynamicFL depends on two factors: the prediction accuracy of bandwidth and the window size of data accumulation in the long-term greedy strategy. However, a trade-off exists on the window size of the server watching data for prediction and accumulating data before making client selection. Intuitively, we can improve bandwidth prediction by using a larger window with more bandwidth data. However, the long observation window causes a longer delay for the server to make necessary adjustments in client selection. On the contrary, a short window can force the server to make quick updates on removing lagged clients, but less data collected to predict bandwidth lowers the accuracy of client selection.

Our goal is to optimize the window size for monitoring bandwidth history data and data buffering. We apply an adaptable approach in adjusting the window size, which takes the system performance with the previous window size into account. We formulate the problem as follows:

$$W_{\text{opt}} = \underset{W}{\operatorname{argmax}} \left(P([b_1, \dots, b_W]) \times \frac{c}{T_W} \right) \quad (4)$$

Our goal is to find the optimal window size W_{opt} that can well predict the historical bandwidth data $b_1, \dots, b_W$ with a short observation time T_W .

Next, we present Algorithm 3 to optimize the window size in DynamicFL. We aim to adjust the size of the frozen window

based on the overall client performance, which is evaluated by the time duration of the round. We increase the control strength in selecting clients at a faster pace when the global training time is longer than our pre-defined threshold. So the server should make a quick update to remove the slow clients and try to find some new clients. When the general network speed becomes faster, we use a longer window in bandwidth observation and data accumulation to obtain better bandwidth prediction and system performance.

Algorithm 3: Trade-off on Window Size.

Input: Client set $\mathbb{C}$, local training iteration base i , current round R_i , observation window size W .
Output: Adjusted window size W .

```

1 Function WindowAdjustment( $\mathbb{C}, R_i, W$ )
2    $d_i \leftarrow \text{GlobalTimeDuration}(i)$ 
3   if  $d_i \geq D_H$  then
4      $W \leftarrow W \times \frac{D_H}{d_i}$ 
5   else
6     if  $d_i \leq D_S$  then
7        $W \leftarrow W \times \frac{D_S}{d_i}$ 
8     end
9   end
10  Return  $W$ 

```

By adaptively adjusting the window used for bandwidth prediction and long-term greedy strategy, DynamicFL achieves a local optimum for different network environments in the training process and improves the overall time-to-accuracy.

IV. EVALUATION

A. Experimental Setup

Implementation. We develop DynamicFL on FedScale [15], a benchmark framework for FL based on Oort [15]. FedScale can simulate the computation and communication over different types of devices. Although FedScale uses real computation capacities and network connection data, they are fixed through the training process without considering the network dynamics. We introduce the bandwidth dynamics into FedScale. We use Yogi [23] as the FL aggregation method and build our model with PyTorch on top of FedScale. We use the lightweight LSTM model to help predict bandwidth using hidden size as 2 and a learning rate of 0.01.

Datasets. To simulate the real-world network activity, we extract bandwidth data from two classes of traces: the HSDPA dataset [17] and the NYC dataset [18]. These two datasets are large real-time bandwidth datasets, including transportation traces on the train, ferry, car, bus, and metro. For each data trace, it records the bandwidth between an LTE mobile phone and a remote server every second. We show the raw data distribution in Figure 3(a). There are multiple experiments at different times of the day to improve the credibility of the dataset. In the evaluation, we map each client in the client pool with a device bandwidth trace from the two bandwidth




Figure 3. Raw trace data CDF for selected transportation types and LSTM prediction loss.

datasets using the division method of hashing. Therefore each client can upload their model updates with real-time high-fluctuated bandwidth.

To evaluate DynamicFL, we use four datasets with various real-world trace data and client configurations:

- *Image Classification* [20], [24]. For the two computer vision tasks, we evaluate DynamicFL on two image datasets: OpenImage [24] and FEMNIST [20]. OpenImage dataset has 8,000 clients collaborating on 1.1 million images. We use the MobileNet [24] as the training model with Top-1 accuracy for image classification. The FEMNIST dataset has 3,400 clients and contains 640K images, in which we also use the Top-1 accuracy.
- *Speech Recognition* [25]. In this task, We evaluate DynamicFL on the Google Speech recognition task [25]. We conduct a training experiment with 2,618 clients and 105K audio commands. We use ResNet-18 [26] for our 20-class speech recognition task.
- *Human Activity Recognition* [6]. We use the HARBox dataset [6], containing 34,115 data samples collected from 121 smartphones for human activity recognition, with a two-layer DNN model in the task.

Metrics and Baselines. We follow the same evaluation metrics as in the SOTA FL solutions [15], [16], i.e., time-to-accuracy and final model accuracy. The time-to-accuracy evaluates the wall-clock time for the central server model to achieve a certain accuracy. The model accuracy is the accuracy that the model achieves when it converges. We use the top 1 accuracy by default. We compare DynamicFL with two baselines: random client selection methods and the Oort [15] client selection method. We use the SOTA optimizer Yogi [23] in the training process.

LSTM Model. We train a three-layer LSTM model with only one airline passengers bandwidth data trace and test the model with other trace data. From Figure 3(b), we find the test loss with a window size much higher than that with a window size of 5. We utilize this bandwidth prediction model in DynamicFL to help to better select clients.

Parameter Configurations. We follow the same parameter setting in previous work [15], [16], [19]. We use the batch size 20 for all the datasets, for both training and testing. DynamicFL selects 100 clients from 130 candidates in each round by default for OpenImage, FEMNIST, and Speech tasks.

Due to the limitation of the size of clients in HAR, we select 5 clients for the task [30]. Each client needs to train the local model for 20 epochs on each round, and we test the model every 10 rounds. The learning rate for the FEMNIST dataset is 0.01, and for the remaining dataset is 0.005.

B. Speedup Performance

DynamicFL Reduces Wall-clock Time to Achieve Final Accuracy. We show a clear wall-clock time reduction of DynamicFL with the two baseline models on all four experiments, as shown in Table I. On the OpenImage dataset with real-world device traces and client data, DynamicFL is $6.13\times$ faster than the Oort approach, consumes 16.3% of the wall-time clock time, and improves the final model accuracy. This largest improvement is caused by the diverse data distribution in the OpenImage dataset. The high non-independent and identically distributed data are better exploited with our long-term greedy strategy and bandwidth prediction module. The non-i.i.d. feature of the OpenImage dataset also makes the Oort model susceptible to the dynamic network environment. On the contrary, the time efficiency of DynamicFL on the HAR dataset, with a lower non-i.i.d. level, is the least improved one with only $1.19\times$ speed-up. When each client shares a small number of samples, the long-term observation data has little difference from the short-term observation data. Figure 4 shows more details of the accuracy change along with training time, corroborating the outstanding performance of DynamicFL with dynamic bandwidth data. Compared with DynamicFL, both baseline client selection methods suffer from the greedy selection method with the fast-changing dynamic network, where the client prioritized in the last round can cause a huge lag in the next round of training.

DynamicFL reduces the time wasted in waiting for the client with bottleneck network connectivity in the training round and reduces the overall training time. Without the prediction ability on the future bandwidth, the client with a high utility score in the last round can be selected due to the short-term greedy strategy, while causing a delay in the current network environment. From Figure 5, we can find that DynamicFL does not reduce the number of training rounds significantly to achieve its final accuracy. This can be explained by reducing the time wasted on waiting for the bottleneck clients in each round, as communication time is the major time cost in FL model training.

DynamicFL Improves the Final Model Accuracy in Limited Training Time. From Table I, DynamicFL slightly improves the system performance in terms of the final accuracy of the training model. The main reason is previous design gives enough time on waiting for the bottleneck client, so they can finally upload the updates to the server. Therefore, the difference in the final accuracy is limited. Considering the setting in the real-world scenario, the model accuracy can be improved by 200% when we stop model training in 24 hours on the FEMNIST dataset. Though the bandwidth prediction model may not select the clients with high data utility temporarily due to current network connectivity, the long-term greedy

Table I
SUMMARY OF DYNAMICFL'S IMPROVEMENT ON TIME-TO-ACCURACY OVER OORT [15]. WE COMPARE THE OVERALL IMPROVEMENT IN FINAL ACCURACY AND WALL CLOCK TIME TO REACH THE FINAL ACCURACY.

Federated Applications	Dataset	Model	Oort [15]+Yogi [23]		DynamicFL +Yogi	
			Acc	Time	Δ Metric	TimeCost
Image Classification	#1 [24]	MobileNet [27]	35.84%	33.77h	2%$\uparrow$	76.9%
Speech Recognition	#2 [25]	ResNet-34 [26]	58.71%	116.48h	5%$\uparrow$	73.5%
Image Recognition	#3 [20]	Shufflenet [28]	74.91%	145.28h	3%$\uparrow$	16.3%
Activity Recognition	#4 [29]	Customized	68.9%	1.76h	2%$\uparrow$	84.1%




Figure 4. Time-to-Accuracy for DynamicFL and Oort on FEMNIST and OpenImage dataset.




Figure 5. Round-to-Accuracy for DynamicFL and Oort on FEMNIST and OpenImage dataset.

strategy can explore and exploit those clients by freezing client selection and getting a general view of client performance in the long run. With our improved client selection method, our model is likely to converge into the sub-optimal. Our long-term greedy strategy allows the server to make selections more carefully and deliver a faster convergence towards final accuracy.

C. Ablation Study

We evaluate the bandwidth prediction module and long-term schedule module respectively. In the ablation study, we deliver our understanding of the importance of each part in DynamicFL.

- **DynamicFL w/o Bandwidth Prediction.** We first conduct experiments on DynamicFL without using the bandwidth prediction module. DynamicFL freezes client selection inside the window and makes decisions based on the historical performance of clients. But the server does not consider the possible bandwidth changes when selecting clients for future rounds. As such, clients with network connectivity downgrading can become the bottleneck in the future.
- **DynamicFL w/o Long Term Greedy strategy.** we disable the long-term greedy method and conduct the bandwidth prediction based on data only from the last round. We aim to

Table II
SUMMARY OF DYNAMICFL'S IMPROVEMENT ON TIME-TO-ACCURACY OVER OORT [15]: THE ABLATION STUDY ON THE IMPROVEMENT IN FINAL ACCURACY AND WALL CLOCK TIME.

Dataset+Model with Yogi	w/o Long-term		w/o Prediction	
	Δ Metric	Speedup	Δ Metric	Speedup
#1+MobileNet	1%$\uparrow$	1.04$\times$	1%$\uparrow$	1.25$\times$
#2+Shufflenet	2%$\uparrow$	1.01$\times$	1%$\uparrow$	4.72$\times$

select clients based on the client utility and training duration from the last round. Without a general consideration of the client performance along with the quick change of network connectivity, the server is likely to select clients causing delays in future training.

DynamicFL Optimizes Client Selection via Bandwidth Prediction in a Dynamic Network. Table II shows the evaluation result on the image classification tasks with only the bandwidth prediction module compared to Oort with a dynamic network. We find the improvements in speedup is limited for these two tasks. The bandwidth prediction part can improve the client selection with the selection of clients with fast network connection speed after observation. However, relying solely on data from the last round makes the prediction imprecise and insufficient for effectively guiding client selection, where the performance falls short of achieving the level of DynamicFL.

DynamicFL Improves the Time Efficiency Using a Long-term Greedy Strategy. The long-term greedy strategy provides the main contribution to the improvement of time-to-accuracy, as shown in Table II. We explain this result through the fact that the long-term greedy strategy can get a better understanding of the client performance. This can better evaluate the clients in a dynamic environment and guide selection in the next round. We find the long-term greedy strategy can bring a 4 $\times$ improvement on wall-clock time.

By combining the bandwidth prediction and long-term greedy method in DynamicFL, we achieve a more precise prediction result and get a more accurate evaluation score on clients in the dynamic network environment. These two strategies contribute to better client selection outcomes.

D. Robustness and Sensitivity Analysis

Impact of Different Optimization Functions. To examine the impact of different optimizers in our design, we conduct experiments on the FEMNIST dataset [20] with three types of optimizers: Prox [31], Yogi [23] and FedAvg [1]. From Figure 6, we find that DynamicFL has a better performance on all the three optimizers, compared with Oort [15] with




Figure 6. DynamicFL improves time-to-accuracy with different optimization functions in comparison with Oort.




Figure 7. DynamicFL achieves better performance under the various number of participants.

dynamic bandwidth dataset. The Yogi optimizer is the best one among the three optimizers. We find that the Oort baseline performance is unstable and fluctuates along with the changed network connectivity. With the long-term greedy strategy and bandwidth prediction, DynamicFL removes the bottleneck clients in advance and achieves high convergence with a much shorter wall-clock time.

Impact of Number of Participants. We evaluate DynamicFL with the different numbers of clients participating in the training in each round, viz., 50, 100, and 150. We compare the time-to-accuracy and round-to-accuracy with the Oort model simulated with dynamic bandwidth on the FEMNIST dataset. From Figure 7, we observe that our design can achieve a better model accuracy in far less time while requiring fewer rounds in training with all three settings. Considering the performance with different participant numbers, using more clients leads to greater performance degradation for the SOTA work with real-time bandwidth data. We explain this result through the fact that the increasing number of participants also increases the chance of bottleneck clients appearing in the training round. On the contrary, DynamicFL can remove the slow clients in an effective way and achieve a similar time-to-accuracy for all three settings.

Impact of Penalty Score. In Algorithm 1, DynamicFL uses the reward and penalty factor to update the feedback of clients. We use four settings to check the impact of the coefficient that multiplies the factor of reward and penalty respectively in Figure 8. From setting 1 (s1) to setting 4 (s4), the reward and penalty coefficient varies from (1.5,5), (2,6), (2,3), and (1.5,10), where greater reward and penalty coefficients indicate a stronger strength in client manipulation. DynamicFL achieves better performance over Oort on time-to-accuracy with different settings. With a selected factor, it can improve performance on the dynamic bandwidth data.




Figure 8. DynamicFL improves the time-to-accuracy across different penalty/reward factors.

Although our model does not finish training in fewer rounds, it significantly reduces the time cost in each round.

V. RELATED WORK

Bandwidth Prediction in Wireless Networks. Bandwidth prediction is an important problem in computer network research, which is first studied as a problem of throughput estimation using historical data [32], [33]. Most recent studies consider using machine learning models in bandwidth prediction, including the DNN model [34]–[37], reinforcement learning model [38], and LSTM model [39]. In our work, we use an LSTM model in bandwidth prediction to guide client selection in the FL model training. Moreover, we integrate bandwidth prediction with long-term scheduling in client selection to optimize the training efficiency in a dynamic network condition.

System Performance Optimization in FL. A number of studies have proposed different methods for improving the efficiency of FL systems. Some studies [15], [19], [31], [40]–[44] dispatch different models to clients based on their different capabilities to improve the model performance and training efficiency. Client selection [15] is an emerging approach to optimize both the model accuracy and the time efficiency in FL. The follow-up work uses a fine-grained method that considers the importance of different data samples [30] and adaptive training method [16] in the synchronized FL system. Although they use real-world capacities for different client devices, none of these studies consider applying FL in real-world dynamic networks. In this research, we are the first to address the problem of client selection with dynamic networks in FL using a real-world bandwidth dataset.

VI. CONCLUSION

In this paper, we focus on managing the unstable network dynamics across massive edge devices in FL by designing a guided client selection strategy. Its core idea is to combine bandwidth prediction and client selection to improve the training efficiency of FL systems in real-world networks. DynamicFL predicts the network bandwidth of participating clients by observing their training time duration, with which it can coordinate the client selection wisely. To optimize the long-term scheduling policy in client selection, we further balance the trade-off between the reliability of network prediction and the client manipulation granularity. Our evaluation demonstrates DynamicFL’s effectiveness under various

application scenarios. Compared with the state-of-the-art FL systems, DynamicFL achieves a better model accuracy while only consuming 18.9%-84.0% of the wall-clock time.

REFERENCES

- [1] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Artificial intelligence and statistics*. PMLR, 2017, pp. 1273–1282.
- [2] G. Wang, H. Guo, A. Li, X. Liu, and Q. Yan, "Federated iot interaction vulnerability analysis," in *2023 IEEE 39th International Conference on Data Engineering (ICDE)*. IEEE, 2023.
- [3] Y. Fu, C. Li, F. R. Yu, T. H. Luan, and Y. Zhang, "A selective federated reinforcement learning strategy for autonomous driving," *IEEE Transactions on Intelligent Transportation Systems*, 2022.
- [4] X. Liang, Y. Liu, T. Chen, M. Liu, and Q. Yang, "Federated transfer reinforcement learning for autonomous driving," in *Federated and Transfer Learning*. Springer, 2022, pp. 357–371.
- [5] I. Donevski, J. J. Nielsen, and P. Popovski, "On addressing heterogeneity in federated learning for autonomous vehicles connected to a drone orchestrator," *Frontiers in Communications and Networks*, vol. 2, p. 709946, 2021.
- [6] X. Ouyang, Z. Xie, J. Zhou, J. Huang, and G. Xing, "Clusterfl: a similarity-aware federated learning system for human activity recognition," in *Proceedings of the 19th Annual International Conference on Mobile Systems, Applications, and Services*, 2021, pp. 54–66.
- [7] L. Tu, X. Ouyang, J. Zhou, Y. He, and G. Xing, "Feddl: Federated learning via dynamic layer sharing for human activity recognition," in *Proceedings of the 19th ACM Conference on Embedded Networked Sensor Systems*, 2021, pp. 15–28.
- [8] C. Li, Z. Liu, Y. Yao, Z. Cao, M. Zhang, and Y. Liu, "Wi-fi see it all: Generative adversarial network-augmented versatile wi-fi imaging," in *Proceedings of ACM SenSys*, 2020.
- [9] D. Leroy, A. Coucke, T. Lavril, T. Gisselbrecht, and J. Dureau, "Federated learning for keyword spotting," in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 6341–6345.
- [10] B. Chen, N. Ivanov, G. Wang, and Q. Yan, "Understanding multi-turn toxic behaviors in open-domain chatbots," in *Proceedings of the 26th International Symposium on Research in Attacks, Intrusions and Defenses*, 2023.
- [11] Y. Zhao, J. Zhao, L. Jiang, R. Tan, D. Niyato, Z. Li, L. Lyu, and Y. Liu, "Privacy-preserving blockchain-based federated learning for iot devices," *IEEE Internet of Things Journal*, vol. 8, no. 3, pp. 1817–1829, 2020.
- [12] G. Wang and Q. Yan, "Federated threat detection for smart home iot rules," in *International Workshop on Federated Learning for Distributed Data Mining*, 2023.
- [13] G. Wang, N. Ivanov, B. Chen, Q. Wang, T. Nguyen, and Q. Yan, "Graph learning for interactive threat detection in heterogeneous smart home rule data," *Proceedings of the ACM on Management of Data*, vol. 1, no. 1, pp. 1–27, 2023.
- [14] M. Alhanahnah, C. Stevens, B. Chen, Q. Yan, and H. Bagheri, "Iotcom: Dissecting interaction threats in iot systems," *IEEE Transactions on Software Engineering*, 2022.
- [15] F. Lai, X. Zhu, H. V. Madhyastha, and M. Chowdhury, "Oort: Efficient Federated Learning via Guided Participant Selection," in *Proceedings of USENIX OSDI*, 2021.
- [16] C. Li, X. Zeng, M. Zhang, and Z. Cao, "Pyramidfl: Fine-grained data and system heterogeneity-aware client selection for efficient federated learning," in *Proceedings of ACM MobiCom*, 2022.
- [17] H. Riiser, P. Vigmostad, C. Griwodz, and P. Halvorsen, "Dataset: Hsdpa-bandwidth logs for mobile http streaming scenarios," 2012.
- [18] L. Mei, R. Hu, H. Cao, Y. Liu, Z. Han, F. Li, and J. Li, "Realtime mobile bandwidth prediction using lstm neural network and bayesian fusion," *Computer Networks*, vol. 182, p. 107515, 2020.
- [19] F. Lai, Y. Dai, X. Zhu, H. V. Madhyastha, and M. Chowdhury, "Fedscale: Benchmarking model and system performance of federated learning," in *Proceedings of the First Workshop on Systems Challenges in Reliable and Secure Federated Learning*, 2021, pp. 1–3.
- [20] G. Cohen, S. Afshar, J. Tapson, and A. Van Schaik, "Emnist: Extending mnist to handwritten letters," in *2017 international joint conference on neural networks (IJCNN)*. IEEE, 2017, pp. 2921–2926.
- [21] H. Riiser, P. Vigmostad, C. Griwodz, and P. Halvorsen, "Commute path bandwidth traces from 3g networks: analysis and applications," in *Proceedings of the 4th ACM Multimedia Systems Conference*, 2013, pp. 114–118.
- [22] L. Mei, R. Hu, H. Cao, Y. Liu, Z. Han, F. Li, and J. Li, "Realtime mobile bandwidth prediction using lstm neural network and bayesian fusion," *Computer Networks*, vol. 182, p. 107515, 2020.
- [23] S. Reddi, Z. Charles, M. Zaheer, Z. Garrett, K. Rush, J. Konečný, S. Kumar, and H. B. McMahan, "Adaptive federated optimization," in *Proceedings of ICLR*, 2021.
- [24] A. Kuznetsova, H. Rom, N. Alldrin, J. Uijlings, I. Krasin, J. Pont-Tuset, S. Kamali, S. Popov, M. Mallocci, A. Kolesnikov *et al.*, "The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale," *International Journal of Computer Vision*, vol. 128, no. 7, pp. 1956–1981, 2020.
- [25] P. Warden, "Speech commands: A dataset for limited-vocabulary speech recognition," *arXiv preprint arXiv:1804.03209*, 2018.
- [26] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of IEEE/CVF CVPR*, 2016.
- [27] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen, "Mobilenetv2: Inverted residuals and linear bottlenecks," in *Proceedings of the IEEE/CVF CVPR*, 2018.
- [28] X. Zhang, X. Zhou, M. Lin, and J. Sun, "Shufflenet: An extremely efficient convolutional neural network for mobile devices," in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2018, pp. 6848–6856.
- [29] X. Ouyang, Z. Xie, J. Zhou, J. Huang, and G. Xing, "ClusterFL: A Similarity-Aware Federated Learning System for Human Activity Recognition," in *Proceedings of ACM MobiSys*, 2021.
- [30] J. Shin, Y. Li, Y. Liu, and S.-J. Lee, "Fedbalancer: Data and pace control for efficient federated learning on heterogeneous clients," 2022.
- [31] T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, A. Talwalkar, and V. Smith, "Federated optimization in heterogeneous networks," *Proceedings of Machine learning and systems*, vol. 2, pp. 429–450, 2020.
- [32] H. Jiang and C. Dovrolis, "Passive estimation of tcp round-trip times," *ACM SIGCOMM Computer Communication Review*, vol. 32, no. 3, pp. 75–88, 2002.
- [33] Q. He, C. Dovrolis, and M. Ammar, "On the predictability of large transfer tcp throughput," *ACM SIGCOMM Computer Communication Review*, vol. 35, no. 4, pp. 145–156, 2005.
- [34] Z. Ali, L. Jiao, T. Baker, G. Abbas, Z. H. Abbas, and S. Khaf, "A deep learning approach for energy efficient computational offloading in mobile edge computing," *IEEE Access*, vol. 7, pp. 149 623–149 633, 2019.
- [35] H. Guo, G. Wang, Y. Wang, B. Chen, and Q. Yan, "Black-box query-efficient audio adversarial attack via split-second phoneme injection," in *Proceedings of the 26th International Symposium on Research in Attacks, Intrusions and Defenses*, 2023.
- [36] N. Ivanov, C. Li, Z. Sun, Z. Cao, X. Luo, and Q. Yan, "Security threat mitigation for smart contracts: A survey," *arXiv preprint arXiv:2302.07347*, 2023.
- [37] C. Li, Z. Cao, and Y. Liu, "Deep ai enabled ubiquitous wireless sensing: A survey," *ACM Computing Surveys (CSUR)*, vol. 54, no. 2, pp. 1–35, 2021.
- [38] H. Mao, R. Netravali, and M. Alizadeh, "Neural adaptive video streaming with pensieve," in *Proceedings of the conference of the ACM special interest group on data communication*, 2017, pp. 197–210.
- [39] P. Li, X. Jiang, G. Jin, Y. Yu, and Z. Xie, "Alstm: An attention-based lstm model for multi-scenario bandwidth prediction," in *2021 IEEE 27th International Conference on Parallel and Distributed Systems (ICPADS)*. IEEE, 2021, pp. 98–105.
- [40] E. Diao, J. Ding, and V. Tarokh, "Heterofl: Computation and communication efficient federated learning for heterogeneous clients," *arXiv preprint arXiv:2010.01264*, 2020.
- [41] A. Li, J. Sun, P. Li, Y. Pu, H. Li, and Y. Chen, "Hermes: an efficient federated learning framework for heterogeneous mobile clients," in *Proceedings of ACM MobiCom*, 2021.
- [42] Y. Wang, H. Guo, G. Wang, B. Chen, and Q. Yan, "Vsmask: Defending against voice synthesis attack via real-time predictive perturbation," *arXiv preprint arXiv:2305.05736*, 2023.
- [43] G. Li, Z. Cao, and T. Li, "Echoattack: Practical inaudible attacks to smart earbuds," in *Proceedings of the 21st Annual International Conference on Mobile Systems, Applications and Services*, 2023, pp. 383–396.
- [44] C. Li, H. Guo, S. Tong, X. Zeng, Z. Cao, M. Zhang, Q. Yan, L. Xiao, J. Wang, and Y. Liu, "Nelora: Towards ultra-low snr lora communication with neural-enhanced demodulation," in *Proceedings of ACM SenSys*, 2021.