

PARALLEL FEDERATED LEARNING OVER HETEROGENEOUS DEVICES

Anonymous authors

Paper under double-blind review

ABSTRACT

Federated Learning (FL) is a cutting-edge distributed machine learning framework that enables multiple devices to collaboratively train a shared model without exposing their own data. In the scenario of device heterogeneity, the synchronous FL suffers from a latency bottleneck induced by network stragglers, which hampers the training efficiency significantly. In addition, due to the diverse structures and sizes of local models, simple and fast averaging aggregation is not feasible anymore. Instead, complicated aggregation operations, such as knowledge distillation, are required. The time cost for complicated aggregation becomes a new bottleneck that limits the computational efficiency of FL. In this work, we claim that the **root cause** of training latency actually lies in the aggregation-then-broadcasting workflow of the server. By swapping the computational order of aggregation and broadcasting, we propose a new parallel federated learning (PFL) framework that unlocks the edge nodes during global computation and the central server during local computation. This fully asynchronous and parallel pipeline enables handling device heterogeneity and network stragglers, allowing flexible device participation as well as achieving scalability in computation. We theoretically prove that synchronous and asynchronous PFL can achieve a similar convergence rate as FL. Extensive experiments empirically show that our framework can tolerate both network stragglers and complicated aggregation, which brings up to $5.56\times$ speedup compared with traditional FL.

1 INTRODUCTION

Federated Learning (FL) is an emerging distributed machine learning framework that enables numerous devices to collaboratively train a shared model without exposing their private data (Shokri & Shmatikov, 2015; Konečný et al., 2016; Kairouz et al., 2019; Li et al., 2020). The main advantage of FL is the decoupling of model training from the necessity of directly accessing to the training data. Thus, in many data-sensitive scenarios, such as applications in biomedicine (Xu et al., 2021; Courtiol et al., 2019; Xu et al., 2021) and finance (Long et al., 2021; Pfitzner et al., 2021), FL can significantly reduce the concern of privacy and security.

The most popular FL system FedAvg (McMahan et al., 2017) works in synchronous training, in which each round includes the following steps (Kairouz et al., 2019): 1) the central server broadcasts the latest global model to edge devices; 2) each edge device then updates its local model with the private data, and uploads the local model update to the central server; 3) once *all* local updates are received, the central server conducts the aggregation operation to produce the next global model. FedAvg works in a serial computation manner, in which no local model optimization is allowed during communication and aggregation while no global computation is allowed during local training and local updates transmission, as shown in Fig.1-(a).

Synchronous FL (SyncFL) faces two main challenges that **hinder the training efficiency**. First, device heterogeneity inevitably introduces network stragglers, which take much longer time to complete local training than common ones. According to the rule that aggregation is conducted after *all* local updates are received, the slowest-responding client, due to either the slowest training speed or dropping out mid-round, becomes the bottleneck of SyncFL system. Second, the weighted average manner is widely applied for global aggregation, which is simple and fast. As a result, in most existing FL schemes, the aggregation time spent in the central server is ignored. However, for

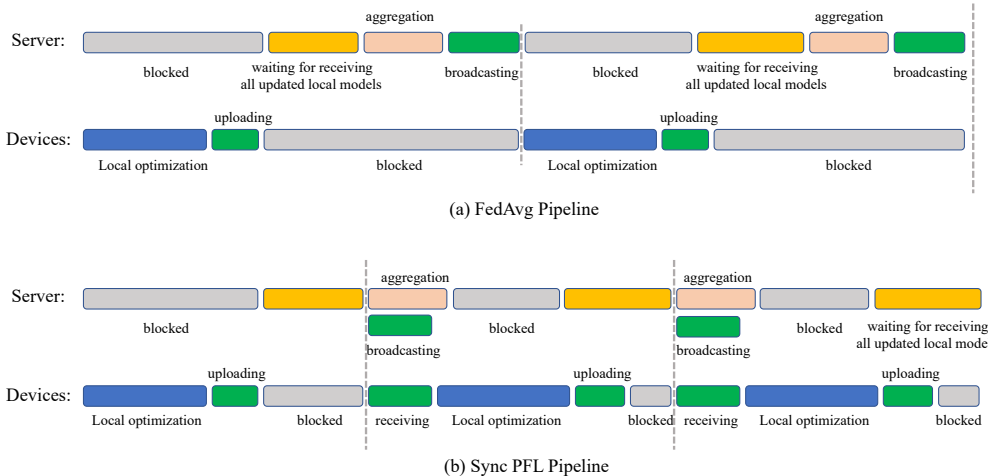


Figure 1: The pipeline of FedAvg and DGA. (a) FedAvg works in a synchronous manner, in which local and global computation are conducted serially. (b) Sync PFL exchange the execution order of global aggregation and broadcasting. As can be seen, Sync PFL achieves a significant speedup under the same conditions.

the scenario of FL over heterogeneous devices, directly averaging of local models is not feasible anymore, since local models may have diverse structures and sizes. Recently, some methods attempt to address the model heterogeneity by performing knowledge distillation in the server side (Zhu et al., 2021b; Lin et al., 2020), leading to complicated aggregation operation. When the time cost for aggregation is significant, in SyncFL, as shown in Fig.1-(a), devices are blocked to wait for the central server, which limits the utilization (*i.e.*, the fraction of processors actively computing at any time) of the system. Accordingly, the time cost for global aggregation becomes another key factor influencing the training efficiency of FL.

The network latency induced by straggler has been well-studied, while the training latency induced by complicated aggregation is overlooked by the literature. The state-of-the-art synchronous and asynchronous FL training strategies attempt to remedy the network latency issue from the edge side. A solution called DGA is offered in (Zhu et al., 2021a), which allows extra local optimization during the process of communication, and thus saves the overall time cost to some extent. However, as a synchronous approach, DGA still suffers from the high network latency induced by network stragglers. An alternative strategy is to allow clients to update the global model asynchronously by dividing fast and slow clients into different subsets. Avdiukhin *et al.* (Avdiukhin & Kasiviswanathan, 2021b) propose an asynchronous local SGD model, where all the client iterates evolve in synchrony with respect to the global clock, but communicate with the server in asynchrony at arbitrary time intervals. Instead, we claim that the bottleneck of training latency actually lies in the workflow of the server. In the literature, this research line has not been explored yet.

In this paper, we propose a simple yet efficient scheme called *parallel federated learning* (PFL) to achieve training efficiency over heterogeneous devices, which can solve the above challenges simultaneously. We claim that the bottleneck of training efficiency actually lies in the workflow of the server. The popular FedAvg conducts serial computation between edge nodes and the central server: uploading \rightarrow global aggregation \rightarrow broadcasting \rightarrow local optimization. In this pipeline, as shown in Fig.1-(a), the server would be blocked to wait for receiving all local updates in order to perform the global aggregation; the edge nodes would be blocked to wait for receiving the broadcasted global model and then can begin the next local optimization. By carefully examining the above process, it can be found that the cause of training inefficiency actually lies in the dependency of server on all local updates uploading and the dependency of clients on global aggregation. An intuitive idea is to decouple these dependencies such that the edge clients and server can run in parallel.

Without redesigning the whole FL framework, we propose a very simple operation—exchange the execution order of global aggregation and broadcasting. The new workflow of the server becomes: uploading (server receives only the completed local updates) \rightarrow broadcasting (server broadcasts the stale global model in the buffer) \rightarrow global aggregation (server averages the currently received local updates to get the next global model). Under this setup, for an edge device, once the local optimization is completed, it uploads the local update to the server; right after the server receives the new local update, it immediately broadcasts the current global model (stale one) in the buffer to the edge device; the edge device no longer needs to wait for the aggregation result but starts

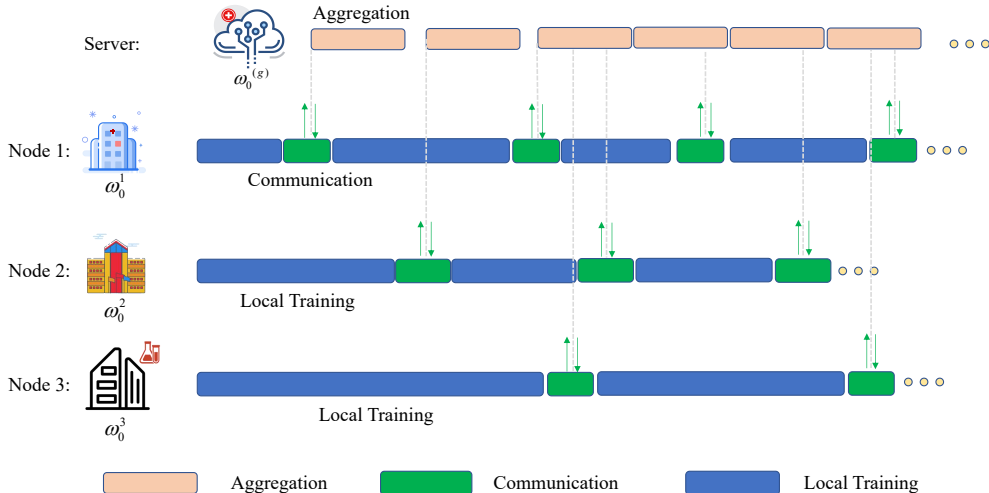


Figure 2: The pipeline of PFL. The local training process of the edge device and the aggregation process of the central server can be carried out simultaneously, thus greatly improving the training efficiency of FL.

the next local optimization upon the received global model; the global aggregation on new updates are then carried out. In this way, the global aggregation is decoupled from the communication and can be performed in parallel with local training. PFL includes both synchronous and asynchronous versions, as shown in Fig. 1-(b) and Fig. 2. We theoretically prove that PFL can achieve the similar convergence rate as FedAvg, and empirically show that our framework can tolerate both stragglers and complicated aggregation tasks, which brings $1.77\times$ to $5.56\times$ speedup. In summary, PFL enjoys the following merits:

- **Handling Device Heterogeneity and Network Stragglers.** The server and edge clients conduct model updates in a fully asynchronous fashion, which is favorable in handling device heterogeneity. Moreover, the server no longer needs to wait for the straggling devices, which contribute their local updates to aggregation whenever they are available to the server.
- **Flexible Device Participation.** In cross-device FL, it is expected that each client would only participate in an arbitrary number of update rounds. Our PFL can achieve flexible device participation naturally. **Once a local client gets disconnected, it can download the global model into the buffer, upon which it can conduct a local model update and rejoin the collaborative training circle.**
- **Scalability in Computation.** In PFL, the global aggregation is executed in parallel with communication and local updates, thus it can be generalized to conduct more complicated operations beyond simple averaging. For instance, we can employ knowledge distillation to deal with the model heterogeneity; use homomorphic encryption to boost security.

2 METHODOLOGY

We consider federated learning with N edge devices. Each device i has access to a private dataset $\mathcal{S}_i = \{(\mathbf{x}_j^i, y_j^i)\}_{j=1}^{m_i}$ including m_i training samples from an unknown and fixed distribution \mathcal{P}_i over $\mathcal{X} \times \mathcal{Y}$. The sample number m_i and the data distribution \mathcal{P}_i can be diverse across edge devices. The overall goal is to train a global model $f(\omega)$ using $m = \sum_{i=1}^N m_i$ from all the edge nodes, where ω is denoted as the model parameters to be optimized. Formally, for each device i , an objective function $F^{(i)}(\omega)$ is defined for optimization over \mathcal{S}_i :

$$F^{(i)}(\omega) = \mathbb{E}_{(\mathbf{x}, y) \in \mathcal{S}_i} [L(f(\mathbf{x}, \omega), y)], \quad (1)$$

2.1 PARALLEL FEDERATED LEARNING

The proposed PFL can be conducted in both synchronous and asynchronous modes. For the synchronous mode, in each round, the server needs to wait until receiving all updates of edge devices

Algorithm 1: Parallel Federated Learning**Input:** N edge devices with private datasets $\{\mathcal{S}_i\}_{i=1}^N$, the maximum global clock number T .**Output:** The final global model $\omega_T^{(g)}$

```

1 Initialization:  $\omega_0^{(g)}$ ;
2 for global clock  $t = 1, \dots, T$  in parallel do
3   for edge device  $i$  corresponding to global clock  $t$ 
4     – Received global models: get  $\omega_{t-1}^{(g)}$ 
5     – Local Update:  $\omega_t^{(i)} \leftarrow \omega_{t-1}^{(g)} - \gamma G_{t-1}^{(i)}$ ;
6     – Uploading: Uploading the gradient updates  $\delta^{(i)}$  to the central server;
7     Global Computation on Server
8     – Communication:
9     • Receive the gradient update  $\delta^{(i)}$  of arriving device  $i$  at global clock  $(t + n + 1)$ ;
10    • Send back the stale global model  $\omega_{t+n-r}^{(g)}$  to device  $i$ ;
11    – Aggregation:
12    • Update the global model:  $\omega_{t+n+1}^{(g)} \leftarrow \omega_{t+n+1-p}^{(g)} + \frac{1}{N} \sum_{i \in \mathcal{C}_t} \delta^{(i)}$ ,

```

prior to conducting the global aggregation; while for the asynchronous mode, the server allows asynchronous updates.

2.1.1 SYNCHRONOUS PFL

Due to the limitation of synchronized aggregation, our synchronous PFL (SPFL) cannot handle the bottleneck induced by network stragglers, which instead is tailored to alleviate the training latency induced by complicated aggregation. Specifically, for the t -th round, SPFL works as follows:

- **Local Update:** Based on the received global model $\omega_{t-1}^{(g)}$, each edge node i updates the local model by stochastic gradient descent using its private data:

$$\omega_t^{(i)} \leftarrow \omega_{t-1}^{(g)} - \gamma G_{t-1}^{(i)}, \quad t = 1, \dots, T, \quad (2)$$

where $G_{t-1}^{(i)}$ represents the gradient updates computed by the i -th node based on $\omega_{t-1}^{(g)}$.

- **Uploading:** Once the edge node i finishes the local training, it uploads the gradient update $\delta^{(i)}$ to the central server.
- **Broadcasting:** Right after receiving the uploaded signals $\{\delta^{(i)}\}_{i=1}^N$ from all edge devices, the central server sends back to them *the stale global model* in the buffer. Note that it is the aggregation result of the last round but not the up-to-date one. If the aggregation procedure of the last round has not yet completed, the server needs to wait until its completion prior to broadcasting the global model.
- **Global Aggregation:** Along with broadcasting, the central server performs the aggregation operation based on the received local updates:

$$\omega_{t+1}^{(g)} = \omega_t^{(g)} + \frac{1}{N} \sum_{i=1}^N \delta^{(i)}, \quad t = 1, \dots, T, \quad (3)$$

When the central server is performing aggregation, the edge nodes are performing local training. Accordingly, by exchanging the execution order of broadcasting and aggregation, we decouple local training and global aggregation and achieve parallel speedup.

2.1.2 ASYNCHRONOUS PFL

Compared with the synchronous mode, the asynchronous mode enjoys more degree of freedom. In the proposed asynchronous PFL (APFL), we attempt to simultaneously remedy the training latency induced by network stragglers on the edge side and complicated aggregation operation on the server side. In APFL, each edge device and the central server have different clocks. To handle this issue, we

define the global clock: whenever any device communicates with the central server, the global clock will increase by 1. In other words, each global clock corresponds to only one device. Specifically, for the t -th global clock, the proposed PFL works as follows:

- **Local Update:** Edge device i completes the communication with the central server, which receives the global model denoted as $\omega_{t-1}^{(g)}$. Based on $\omega_{t-1}^{(g)}$, edge device i updates the local model by stochastic gradient descent using its private data:

$$\omega_t \leftarrow \omega_{t-1}^{(g)} - \gamma G_{t-1}^{(i)}, \quad t = 1, \dots, T. \quad (4)$$

- **Uploading:** After completing the local training, edge device i then uploads the gradient update $\delta^{(i)}$ to the central server. During the local training and the uploading process of device i , there would be n other devices communicating with the central server. Therefore, when the central server receives the local update of the device i , the global clock becomes $t + n + 1$.
- **Broadcasting:** Right after receiving the uploaded signal $\delta^{(i)}$ from edge device i , **the central server sends back to it the stale global model in the buffer, which is the most recent aggregation result. It would be $\omega_{t+n}^{(g)}$ if the central server is idle, or $\omega_{t+n-r}^{(g)}$ if the central server is performing aggregation, where r denotes the number of local devices that previously arrived at the server but did not participate in aggregation. The former is a special case of the latter with $r = 0$. For simplicity, we denote the stale global model as $\omega_{t+n-r}^{(g)}$. If the local model uploaded in the previous clock by that device has not completed aggregation, then the device must wait until it has.** Note that the local signal $\delta^{(i)}$ received this time does not **participate** in the current aggregation process. Therefore, the edge device i no longer needs to wait for the global aggregation, which immediately starts the next local training relying on $\omega_{t+n-r}^{(g)}$. The central server then performs the next aggregation in parallel with the broadcasting and the new local training on device i .
- **Global Aggregation:** The new broadcasting-then-aggregation mechanism, which sends to clients the stale global models but not the latest ones, makes the global aggregation decoupled from the communication. However, this fully asynchronous manner leads to a scenario that, when the global aggregation is being executed, there are some local updates arriving. This situation is easy to handle: the arrivals just wait for the current aggregation to be done and then launch a new aggregation. The averaging operation is thus performed on the updates of the new received edge devices:

$$\omega_{t+n+1}^{(g)} \leftarrow \omega_{t+n+1-p}^{(g)} + \frac{1}{N} \sum_{i \in \mathcal{C}_t} \delta^{(i)}, \quad t = 1, \dots, T, \quad (5)$$

where \mathcal{C}_t represents the devices that arrived during the last aggregation, p is the number of devices in \mathcal{C}_t . Normally, \mathcal{C}_t contains only one device.

2.1.3 COMPARISON OF SYNCHRONOUS PFL AND ASYNCHRONOUS PFL

The main difference between SPFL and APFL is whether the central server needs to wait for all devices. SPFL is subject to the requirement of synchronized aggregation, while APFL is not. This results in SPFL and APFL having different speed-up ratios compared to standard federated learning. Specifically, we divide the training time of federated learning into four parts: the time cost T_l for local training, the time cost T_c for communication, the time cost T_w for waiting, and the time cost T_g for global aggregation. Standard federated learning works in a serial manner, for which the training time of each round is $T = T_l + T_c + T_w + T_g$. In SPFL, the aggregation of the central server is performed in parallel with the local training of the edge nodes and the communication process, the training time for each round is $T = \max(T_l + T_c + T_w, T_g)$. In APFL, there is no waiting time, we have $T = \max(T_l + T_c, T_g)$.

2.2 CONVERGENCE ANALYSIS

In this subsection, we present the convergence analysis to demonstrate in theory that the proposed synchronous/asynchronous PFL both can get a similar convergence as FedAvg. It is worth noting




Figure 3: The speedup ratio comparison under different setting. T_w/T_l represents the ratio of the straggler device’s dropout time to the local training time. Ratio of T_g/T_l is associated with the number of devices and the computing power of the central server. Experimental results show that PFL yields the highest speedup ratios in both situations.

that there is a standard routine for FL convergence analysis, thus we just follow it without the necessity of creating a new one. We take the following standard assumptions (Stich, 2018; Avdiukhin & Kasiviswanathan, 2021a; Nguyen et al., 2022):

- **Smoothness:** All local functions $f^i (i \in [N])$ are L -smooth.
- **Bounded second moment:** There exists a constant $G_{max} > 0$ such that:

$$\mathbb{E} \left[\|\nabla F^{(i)}(x)\|^2 \right] \leq G_{max}^2, \quad \forall i \in [N], \forall \mathbf{x} \in \mathbb{R}^d, \quad (6)$$

where $\nabla F^{(i)}(x)$ is an unbiased stochastic gradient of $f^{(i)}$ at x .

- **Bounded variance:** There exists a constant $\sigma \geq 0$ such that:

$$\mathbb{E}_{\xi \sim \mathcal{S}_i} \|\nabla F_i(\mathbf{x}) - \nabla f_i(\mathbf{x})\| \leq \sigma, \quad \forall i \in [N], \forall \mathbf{x} \in \mathbb{R}^d. \quad (7)$$

Analysis for Synchronous PFL. In synchronous PFL, we only exchange the execution order of global aggregation and broadcasting, which makes each edge device receive the stale global model. At first glance, it looks a bit like asynchronous FL. The difference is that in asynchronous FL, only part of the devices (stragglers) upload the stale gradient updates. In synchronous PFL, the global model received by the devices is only from the last round. Due to the above similarities, we can analyze the convergence of synchronous PFL with the help of existing studies on asynchronous FL. Here, we follow (Avdiukhin & Kasiviswanathan, 2021a), but give a sharper bound.

Definition 1 (Synchronous PFL Sequence) Let $\omega_t^{(g)}$ represents the global model of synchronous PFL, $G_t^{(i)}$ represents the uploaded gradient updates computed by the i -th client based on $\omega_t^{(g)}$, the global model $\omega_t^{(g)}$ can be expressed as:

$$\omega_t^{(g)} = \omega_0^{(g)} - \sum_{\tau=2}^t \gamma \text{avg}_i \left(G_{\tau-2}^{(i)} \right), \quad t = 2, \dots, T, \quad (8)$$

where $\omega_0^{(g)}$ and $\omega_1^{(g)}$ are two different initialized models.

Definition 2 (Virtual Sequence) We construct a virtual sequence in which each device at round t uses the $\omega_{t-1}^{(g)}$ to calculate the uploaded gradient updates, this virtual sequence’s global model can be expressed as:

$$\zeta_t = \omega_0^{(g)} - \sum_{\tau=1}^t \gamma \text{avg}_i \left(G_{\tau-1}^{(i)} \right), \quad t = 1, \dots, T. \quad (9)$$

Similar virtual sequences have been used before in decentralized optimization in various contexts (Lian et al., 2017; Yuan et al., 2016; Nedić et al., 2018; Stich, 2018; Avdiukhin & Kasiviswanathan, 2021a).

Proposition 3 (Distance Bound) For ζ_t of virtual sequence and $\omega_t^{(g)}, \omega_t^{(i)}$ in synchronous PFL we have:

$$\max \left(\mathbb{E} \left[\|\zeta_t - \omega_t^{(i)}\|^2 \right], \mathbb{E} \left[\|\zeta_t - \omega_t^{(g)}\|^2 \right] \right) \leq 4\gamma^2 G_{max}^2 \quad (10)$$

Table 1: Comparison of synchronous FedAvg and our PFL’s accuracy (%) on 3 datasets with both i.i.d and non-i.i.d data partitions. Speed-up is measured for each global clock, normalized by the run-time of FedAvg. The best result is **boldfaced**.

Method	MNIST		CIFAR-10		Tiny-ImageNet		Time for per global clock
	i.i.d	non-i.i.d	i.i.d	non-i.i.d	i.i.d	non-i.i.d	
FedAvg	93.21	85.31	85.41	75.26	59.52	48.14	$T = T_l + T_c + T_w + T_g$
	1×		1×		1×		
DGA	92.12	83.72	83.72	73.24	58.11	46.26	$T = T_l + (T_c + T_w + T_g)$ $= T_l + T_E$
	1×		1×		1×		
MIFA	92.88	84.73	85.14	74.78	59.24	47.12	$T = T_l + T_c + T_g$
	1.57×		1.86×		1.98×		
SPFL	93.14	85.26	85.57	75.31	60.21	48.31	$T = \max(T_l + T_c, T_g) + T_w$
	1.22×		1.30×		1.33×		
PFL	93.08	85.13	85.34	75.46	60.18	48.07	$T = \max(T_l + T_c, T_g)$
	2.20×		3.28×		3.91×		

This Proposition shows that ζ_t are close to $\omega_t^{(g)}$ and $\omega_t^{(i)}$ for all devices. Its proof can be seen in Appendix.

Theorem 4 Let $f_{\max} = f(\omega_0^{(g)}) - f(\omega^*)$, where ω^* is the minimizer for f , $\gamma = \sqrt{N}/\sqrt{T}$ we have:

$$\frac{1}{T} \sum_{t=0}^T \mathbb{E} \left[\left\| \nabla f(\omega_t^{(g)}) \right\|^2 \right] = O \left(\frac{1}{\sqrt{NT}} + \frac{N}{T} \right). \quad (11)$$

Compared to the corresponding result for Basu et al. (2019) and Avdiukhin & Kasiviswanathan (2021a), $O\left(\frac{1}{\sqrt{NT}}\right)$, when N is fixed, our results get a sharper convergence rate. And illustrates a clearer relationship between T and N , i.e. as the number of nodes N increases, the number of communication rounds T should be larger as well.

Analysis for Asynchronous PFL. Convergence of asynchronous PFL can be demonstrated in a similar way as above, with the difference that in asynchronous PFL, the global model received by the device will be more than one clock staler than the current global model when aggregated. For asynchronous PFL, we can use almost the same method of Avdiukhin & Kasiviswanathan (2021a) to prove the convergence of our method. Our contribution does not lie in the novelty of the method of proof, but rather in pointing out that our approach can be fully accommodated in existing proofs of asynchronous FL and achieve greater speed-up ratios. For convenience, we use a notion similar to that used in synchronous PFL to describe the global model and the local model. Except for local models, we use w_t to represent the local model, since each global clock only corresponds to one device.

Definition 5 (Asynchronous PFL Sequence) Let $\omega_t^{(g)}$ represents the global model of asynchronous PFL, $G_t^{(i)}$ represents the uploaded gradient updates computed by the i -th client based on $\omega_t^{(g)}$, the global model $\omega_t^{(g)}$ can be expressed as:

$$\omega_t^{(g)} = \omega_0^{(g)} - \sum_{\tau=1}^t \gamma G_{\tau-\delta_t}^{(i)}, \quad t = 1, \dots, T, \quad (12)$$

where $\tau - \delta_t$ denotes the clock of the global model received by the device participating in the t -th round of aggregation, and δ_t can represent the degree of asynchrony.

Proposition 6 (Distance Bound) For ζ_t of virtual sequence and $\omega_t^{(g)}, \omega_t$, we have:

$$\max \left(\mathbb{E} \left[\left\| \zeta_t - \omega_t \right\|^2 \right], \mathbb{E} \left[\left\| \zeta_t - \omega_t^{(g)} \right\|^2 \right] \right) \leq 4(\delta - 1)^2 \gamma^2 G_{\max}^2, \quad (13)$$

Theorem 7 Let $f_{\max} = f(\omega_0^{(g)}) - f(\omega^*)$, where ω^* is the minimizer for f , if $T > N^3$ and $\delta \leq T^{1/4}/N^{3/4}$ we have:

$$\frac{1}{T} \sum_{t=0}^T \mathbb{E} \left[\left\| \nabla f(\omega_t^{(g)}) \right\|^2 \right] = O \left(\frac{L f_{\max}}{\sqrt{NT}} + \frac{G_{\max}^2}{\sqrt{NT}} + \frac{L \sigma^2}{\sqrt{NT}} \right), \quad (14)$$

where $\delta = \max(\delta_t)$.

Table 2: Comparison of FedAvg and our PFL’s training efficiency (%) on 3 datasets with different local epochs under distillation-based aggregation. The communication time T_c and the stragglers are ignored, the times in the table are the average times for each training between edge device and central server.

Dataset	Local epoch	T_l (s)	T_g (s)	Fedavg		PFL		Speedup
				Time(s)	Acc.	Time(s)	Acc.	
MNIST	5	4	5	9	99.31	5	99.18	$1.80 \times$
	10	7	5	12	98.17	7	98.09	$1.71 \times$
	15	11	5	16	96.13	11	96.24	$1.45 \times$
CIFAR-10	5	19	20	39	89.22	20	89.13	$1.95 \times$
	10	38	20	58	87.81	38	87.76	$1.53 \times$
	15	60	20	80	86.48	60	86.37	$1.33 \times$
Tiny-ImageNet	5	185	180	365	62.88	185	62.79	$1.97 \times$
	10	375	180	555	60.13	375	60.06	$1.48 \times$
	15	563	180	743	57.25	563	57.22	$1.32 \times$

3 EXPERIMENTS

In this section, we provide a thorough evaluation of the proposed PFL over heterogeneous devices. We first compare the convergence speed and accuracy of the PFL under average-based aggregation. Next, to investigate the speedup ratio of the PFL under different scenarios, experiments are carried out with different T_w/T_l and T_g/T_l , where T_w represents the reconnection time of the straggler devices, T_l represents the local training time, and T_g represents the aggregation time. We then show that our method is also feasible for distillation-based aggregation. Experimental results show that our method is resilient to stragglers and can achieve a parallel pipeline to improve computational efficiency without sacrificing accuracy.

3.1 EXPERIMENTAL SETUP

Straggler Devices and Sample Rate. Following the previous study (So et al., 2022), we randomly select $p \times N$ devices where p is the dropout rate. To constrain the asynchronous level between devices, we set $T_w/T_l = 3$, which represents the ratio of waiting time to local training time. We consider the worst-case scenario (Bonawitz et al., 2017), where the straggler devices artificially drop after receiving the global model. According to the realistic FL system (Bonawitz et al., 2019), we set $p = 0.2$.

Datasets and Models. We conduct experiments on three datasets, including MNIST, CIFAR-10, and Tiny-Imagenet (100,000 images with 200 classes). For all three datasets, we use the same ResNet-18 architecture for a fair and clear comparison.

Number of devices and Communication bandwidth. In our experiments, we train up to $N = 100$ devices. We are particularly focused on mobile devices in FL, the real measured bandwidth of a mobile phone is 4MB/s. The cost of time to upload and broadcast updated local and global models, denoted as T_c .

Baselines. We analyze and compare the performance of PFL with three baseline schemes: FedAvg (McMahan et al., 2017), DGA (Zhu et al., 2021a), and MIFA (Gu et al., 2021). DGA allows extra local optimization during the communication process to implicitly reduce wasted computing power during communication. MIFA is an asynchronous communication FL method to solve straggler problems. These two methods address the challenge of training efficiency in FL from two different perspectives.

Data Partition: Following Avdiukhin & Kasiviswanathan (2021a), for IID setup, we randomly split the data equally into N sub-datasets; for Non-IID setup, each device has a corresponding class, where μ fraction of local data is from the corresponding class, while the rest are randomly selected from other classes. In our experiments, $\mu = 0.3$.

3.2 PERFORMANCE EVALUATION

To fully explore the communication efficiency of FL, we simulate all the three scenarios mentioned in the above section, concretely, 1) devices will randomly drop out, 2) the upload and download time of the models is measured by bandwidth 4Mb/s, and 3) the central server needs to perform additional complex computations such as knowledge distillation or toxic device identification.

Average-based Aggregation. Although various novel aggregation methods have been proposed, FedAvg is still the most widely used method, so we first investigate the performance of our method

under the average-based approach. To simulate complex operations on the central server in the real environment without introducing additional effects, we perform additional complex tasks on the central server to demonstrate the superiority of the training efficiency of our approach. Specifically, an additional stage of toxic device identification is added before the central server performs the aggregation. The local models uploaded from each device need to be validated by the public dataset on the central server before participating in aggregation. This is a very efficient but time-consuming detection method and, therefore, is rarely used in real-world scenarios. As can be seen in Table 1, our method can achieve almost the same accuracy as FedAvg, while achieving a speed increase of approximately $3.91\times$. Speed-up is measured in run time per global clock between device and central server, normalized by the run time of FedAvg. DGA does not reduce the computation time for each communication but rather speeds up convergence by allowing the nodes to perform local training while communicating. Therefore, it has an acceleration ratio of only $1\times$. MIFA can get good speedups in the presence of straggler, e.g. $1.86\times$ faster on the CIFAR-10 dataset. Our approach can improve the computational efficiency of federated learning in three dimensions simultaneously, the time for per global clock is $T = \max(T_l + T_c, T_g)$, for example, when on the CIFAR-10 dataset, our method can obtain a $3.28\times$ speedup.

Impact of T_w/T_l on speedup ratio. T_w/T_l represents the ratio of reconnection time of straggler devices to local training time. The longer the dropout time, the greater its impact on the overall computational efficiency of SyncFL and the greater its impact on the accuracy of AsyncFL. Therefore, it is necessary to investigate the effect of different T_w/T_l on computational efficiency, here we set $T_w/T_l = 1$ and $T_w/T_l = 1, 2, 3, 4, 5$, and the experiment is carried out on the CIFAR-10 with 100 devices. As can be seen in Fig. 3, as T_w/T_l increases, the computational speed-up ratio between our method and MIFA’s method gets higher and higher, and when $T_w/T_l = 5$, the computational efficiency of our PFL can be improved by up to $5.56\times$, while MIFA can only improve the speed-up ratio by $3.16\times$.

Impact of T_g/T_l on speedup ratio. T_g/T_l represents the ratio of the central server’s computation time to the local device’s training time, which is related to the number of devices and the computation capacity of the central server, and is further complicated when the computation power of the edge devices is different from each other. To simplify the computational model, we assume that the local training time is the same for each edge device. Here, we set $T_g/T_l = 1$ and $T_w/T_l = 0.25, 0.5, 1, 2, 4$, and the experiment is carried out on the CIFAR-10 with 100 devices. As can be seen in Fig. 3, the speed-up ratio of our method reaches its maximum $2.52\times$ when $T_g/T_l = 1$.

Distillation-based Aggregation. To verify the effectiveness of our method for complex aggregations, we have chosen the first distillation-based aggregation method (Lin et al., 2020). The distillation-based approach makes better use of information from edge devices, but introduces serious computational reductions in efficiency that make it difficult to apply to real-world applications. Following (Lin et al., 2020), we added additional public datasets to the central server and then fine-tuned the global model by distillation after coarse aggregation was completed by averaging operations. In this experiment, the straggler node is ignored. As can be seen from Table 2, our method is still feasible under distillation-based aggregation and can improve the computational efficiency of federated learning without sacrificing accuracy, achieving $2\times$ speedup when the training time for edge devices is the same as the aggregation time for the central server.

4 CONCLUSION

The traditional FL works in a serial computation manner, in which communication is deemed a major bottleneck, leading to poor training efficiency. Most existing methods attempt to improve training efficiency either by performing additional computations or by reducing the communication frequency or data amount. In this paper, alternatively, we propose a simple yet efficient scheme called parallel federated learning (PFL). Instead of remedying the efficiency issue from the edge side as done in existing methods, we claim that the bottleneck of computational efficiency actually lies in the workflow of the server. Through exchanging the execution order of global aggregation and broadcasting, global aggregation is decoupled from the communication and can be performed in parallel with local training. We theoretically prove that PFL can achieve the same convergence rate as FedAvg. The proposed scheme has the potential to serve as a new baseline FL framework enabling a wide range of applications.

REFERENCES

- Dmitrii Avdiukhin and Shiva Kasiviswanathan. Federated learning under arbitrary communication patterns. In *International Conference on Machine Learning*, pp. 425–435. PMLR, 2021a.
- Dmitrii Avdiukhin and Shiva Kasiviswanathan. Federated learning under arbitrary communication patterns. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 425–435. PMLR, 18–24 Jul 2021b. URL <https://proceedings.mlr.press/v139/avdiukhin21a.html>.
- Debraj Basu, Deepesh Data, Can Karakus, and Suhas Diggavi. Qsparse-local-sgd: Distributed sgd with quantization, sparsification and local computations. *Advances in Neural Information Processing Systems*, 32, 2019.
- Keith Bonawitz, Vladimir Ivanov, Ben Kreuter, Antonio Marcedone, H Brendan McMahan, Sarvar Patel, Daniel Ramage, Aaron Segal, and Karn Seth. Practical secure aggregation for privacy-preserving machine learning. In *proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, pp. 1175–1191, 2017.
- Keith Bonawitz, Hubert Eichner, Wolfgang Grieskamp, Dzmitry Huba, Alex Ingerman, Vladimir Ivanov, Chloe Kiddon, Jakub Konečný, Stefano Mazzocchi, Brendan McMahan, et al. Towards federated learning at scale: System design. *Proceedings of Machine Learning and Systems*, 1: 374–388, 2019.
- Pierre Courtiol, Charles Maussion, Matahi Moarii, Elodie Pronier, Samuel Pilcer, Meriem Sefta, Pierre Manceron, Sylvain Toldo, Mikhail Zaslavskiy, Nolwenn Le Stang, et al. Deep learning-based classification of mesothelioma improves prediction of patient outcome. *Nature medicine*, 25(10):1519–1525, 2019.
- Xinran Gu, Kaixuan Huang, Jingzhao Zhang, and Longbo Huang. Fast federated learning in the presence of arbitrary device unavailability. *Advances in Neural Information Processing Systems*, 34, 2021.
- Peter Kairouz, H. Brendan McMahan, and et al. Advances and open problems in federated learning. *CoRR*, abs/1912.04977, 2019. URL <http://arxiv.org/abs/1912.04977>.
- Jakub Konečný, H Brendan McMahan, Felix X Yu, Peter Richtárik, Ananda Theertha Suresh, and Dave Bacon. Federated learning: Strategies for improving communication efficiency. *arXiv preprint arXiv:1610.05492*, 2016.
- Tian Li, Anit Kumar Sahu, Ameet Talwalkar, and Virginia Smith. Federated learning: Challenges, methods, and future directions. *IEEE Signal Processing Magazine*, 37(3):50–60, 2020.
- Xiangru Lian, Ce Zhang, Huan Zhang, Cho-Jui Hsieh, Wei Zhang, and Ji Liu. Can decentralized algorithms outperform centralized algorithms? a case study for decentralized parallel stochastic gradient descent. *arXiv preprint arXiv:1705.09056*, 2017.
- Tao Lin, Lingjing Kong, Sebastian U Stich, and Martin Jaggi. Ensemble distillation for robust model fusion in federated learning. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 2351–2363. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/file/18df51b97ccd68128e994804f3eccc87-Paper.pdf>.
- Guodong Long, Yue Tan, Jing Jiang, and Chengqi Zhang. Federated learning for open banking. *CoRR*, abs/2108.10749, 2021. URL <https://arxiv.org/abs/2108.10749>.
- 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.
- Angelia Nedić, Alex Olshevsky, and Michael G Rabbat. Network topology and communication-computation tradeoffs in decentralized optimization. *Proceedings of the IEEE*, 106(5):953–976, 2018.

John Nguyen, Kshitiz Malik, Hongyuan Zhan, Ashkan Yousefpour, Mike Rabbat, Mani Malek, and Dzmityr Huba. Federated learning with buffered asynchronous aggregation. In *International Conference on Artificial Intelligence and Statistics*, pp. 3581–3607. PMLR, 2022.

Bjarne Pfitzner, Nico Steckhan, and Bert Arnrich. Federated learning in a medical context: A systematic literature review. *ACM Trans. Internet Technol.*, 21(2), jun 2021. ISSN 1533-5399. doi: 10.1145/3412357. URL <https://doi.org/10.1145/3412357>.

Reza Shokri and Vitaly Shmatikov. Privacy-preserving deep learning. In *Proceedings of the 22nd ACM SIGSAC conference on computer and communications security*, pp. 1310–1321, 2015.

Jinhyun So, Corey J Nolet, Chien-Sheng Yang, Songze Li, Qian Yu, Ramy E Ali, Basak Guler, and Salman Avestimehr. Lightsecagg: a lightweight and versatile design for secure aggregation in federated learning. *Proceedings of Machine Learning and Systems*, 4:694–720, 2022.

Sebastian U Stich. Local sgd converges fast and communicates little. *arXiv preprint arXiv:1805.09767*, 2018.

Jie Xu, Benjamin S Glicksberg, Chang Su, Peter Walker, Jiang Bian, and Fei Wang. Federated learning for healthcare informatics. *Journal of Healthcare Informatics Research*, 5(1):1–19, 2021.

Kun Yuan, Qing Ling, and Wotao Yin. On the convergence of decentralized gradient descent. *SIAM Journal on Optimization*, 26(3):1835–1854, 2016.

Ligeng Zhu, Hongzhou Lin, Yao Lu, Yujun Lin, and Song Han. Delayed gradient averaging: Tolerate the communication latency for federated learning. *Advances in Neural Information Processing Systems*, 34, 2021a.

Zhuangdi Zhu, Junyuan Hong, and Jiayu Zhou. Data-free knowledge distillation for heterogeneous federated learning. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 12878–12889. PMLR, 18–24 Jul 2021b. URL <https://proceedings.mlr.press/v139/zhu21b.html>.

A APPENDIX

In the Appendix, we first give the details of the convergence analysis that are missing in the text Section 2.2. And conduct an experiment to investigate the influence of stale global model.

B CONVERGENCE ANALYSIS OF PFL

Proof for Synchronous PFL.

Proposition 3 (Distance Bound) For ζ_t of full synchronous FL and $\omega_t^{(g)}, \omega_t^{(i)}$ of Algorithm 1, we have:

$$\max \left(\mathbb{E} \left[\left\| \zeta_t - \omega_t^{(i)} \right\|^2 \right], \mathbb{E} \left[\left\| \zeta_t - \omega_t^{(g)} \right\|^2 \right] \right) \leq 4\gamma^2 G_{\max}^2. \quad (15)$$

Proof:

$$\begin{aligned} \mathbb{E} \left[\left\| \zeta_t - \omega_t^{(g)} \right\|^2 \right] &= \mathbb{E} \left[\left\| \sum_{\tau=2}^t \gamma \text{avg}_i \left(G_{\tau-2}^{(i)} \right) - \sum_{\tau=1}^t \gamma \text{avg}_i \left(G_{\tau-1}^{(i)} \right) \right\|^2 \right] \\ &\leq \gamma^2 \mathbb{E} \left[\left\| \sum_{\tau=2}^t \text{avg}_i \left(\mathbf{G}_{\tau-2}^{(i)} - \mathbf{G}_{\tau-1}^{(i)} \right) - \gamma \text{avg}_i \mathbf{G}_0^{(i)} \right\|^2 \right] \\ &\leq \gamma^2 \mathbf{G}_{\max}^2. \end{aligned} \quad (16)$$

Similarly, $\mathbb{E} \left[\left\| \omega_t^{(g)} - \omega_t^{(i)} \right\|^2 \right] \leq \gamma^2 \mathbf{G}_{\max}^2$. Combining these bounds, we have the following.

$$\begin{aligned} \mathbb{E} \left[\left\| \zeta_t - \omega_t^{(i)} \right\|^2 \right] &= \mathbb{E} \left[\left\| \zeta_t - \omega_t^{(g)} + \omega_t^{(g)} - \omega_t^{(i)} \right\|^2 \right] \\ &\leq 2 \left(\mathbb{E} \left[\left\| \zeta_t - \omega_t^{(g)} \right\|^2 \right] + \mathbb{E} \left[\left\| \omega_t^{(g)} - \omega_t^{(i)} \right\|^2 \right] \right) \\ &\leq 4\gamma^2 \mathbf{G}_{\max}^2. \end{aligned} \quad (17)$$

Theorem 4 Let $f_{\max} = f(\omega_0^{(g)}) - f(\omega^*)$, where ω^* is the minimizer for f , we have:

$$\frac{1}{T} \sum_{t=0}^T \mathbb{E} \left[\left\| \nabla f(\omega_t^{(g)}) \right\|^2 \right] = O \left(\frac{1}{\sqrt{NT}} + \frac{N}{T} \right). \quad (18)$$

Proof: Similar to Avdiukhin & Kasiviswanathan (2021a) Theorem 2.4.

From smoothness Lipschitz condition on the gradients:

$$\begin{aligned} \mathbb{E} \left[\left\| \nabla f(\omega_t^{(g)}) - \nabla f(\zeta_t) \right\|^2 \right] &\leq L^2 \mathbb{E} \left[\left\| \omega_t^{(g)} - \zeta_t \right\|^2 \right] \leq L^2 \gamma^2 \mathbf{G}_{\max}^2, \text{ and} \\ \mathbb{E} \left[\left\| \nabla f^{(i)}(\omega_t^{(i)}) - \nabla f^{(i)}(\zeta_t) \right\|^2 \right] &\leq L^2 \mathbb{E} \left[\left\| \omega_t^{(i)} - \zeta_t \right\|^2 \right] \leq 4L^2 \gamma^2 \mathbf{G}_{\max}^2. \end{aligned} \quad (19)$$

First, we bound ζ_t , for ζ_t , we have:

$$\zeta_{t+1} = \zeta_t - \gamma \text{avg}_i \left(\mathbf{G}_t^{(i)} \right). \quad (20)$$

By the smoothness property:

$$\mathbb{E} [f(\zeta_{t+1})] \leq \mathbb{E} [f(\zeta_t)] - \mathbb{E} \left[\left\langle \nabla f(\zeta_t), \gamma \text{avg}_i \left(\mathbf{G}_t^{(i)} \right) \right\rangle \right] + \frac{L}{2} \mathbb{E} \left[\left\| \gamma \text{avg}_i \left(\mathbf{G}_t^{(i)} \right) \right\|^2 \right]. \quad (21)$$

The last term in Eq. 21 can be rewritten as:

$$\begin{aligned} &\frac{L}{2} \mathbb{E} \left[\left\| \gamma \text{avg}_i \left(\mathbf{G}_t^{(i)} \right) \right\|^2 \right] \\ &= \frac{\gamma^2 L}{2} \mathbb{E} \left[\left\| \text{avg}_i \left(\mathbf{G}_t^{(i)} + \nabla f^{(i)}(\omega_t^{(i)}) - \nabla f^{(i)}(\omega_t^{(i)}) \right) \right\|^2 \right] \\ &= \frac{\gamma^2 L}{2} \mathbb{E} \left[\left\| \text{avg}_i \left(\nabla f^{(i)}(\omega_t^{(i)}) + \left(\mathbf{G}_t^{(i)} - \nabla f^{(i)}(\omega_t^{(i)}) \right) \right) \right\|^2 \right] \\ &= \frac{\gamma^2 L}{2} \mathbb{E} \left[\left\| \text{avg}_i \left(\nabla f^{(i)}(\omega_t^{(i)}) \right) \right\|^2 \right] + \frac{\gamma^2 L}{2} \mathbb{E} \left[\left\| \text{avg}_i \left(\mathbf{G}_t^{(i)} - \nabla f^{(i)}(\omega_t^{(i)}) \right) \right\|^2 \right] \\ &= \frac{\gamma^2 L}{2} \mathbb{E} \left[\left\| \text{avg}_i \left(\nabla f^{(i)}(\omega_t^{(i)}) \right) \right\|^2 \right] + \gamma^2 \frac{L\sigma^2}{2}. \end{aligned} \quad (22)$$

Substituting this into the Eq. 21, get:

$$\begin{aligned} \mathbb{E} [f(\zeta_{t+1})] &\leq \mathbb{E} [f(\zeta_t)] - \mathbb{E} \left[\left\langle \nabla f(\zeta_t), \gamma \text{avg}_i \left(\nabla f^{(i)}(\omega_t^{(i)}) \right) \right\rangle \right] \\ &\quad + \frac{\gamma^2 L}{2} \mathbb{E} \left[\left\| \text{avg}_i \left(\nabla f^{(i)}(\omega_t^{(i)}) \right) \right\|^2 \right] + \gamma^2 \frac{L\sigma^2}{2} \\ &\leq \mathbb{E} [f(\zeta_t)] - \gamma \mathbb{E} \left[\left\langle \nabla f(\zeta_t), \text{avg}_i \left(\nabla f^{(i)}(\zeta_t) \right) \right\rangle \right] \\ &\quad - \gamma \mathbb{E} \left[\left\langle \nabla f(\zeta_t), \text{avg}_i \left(\nabla f^{(i)}(\omega_t^{(i)}) - \nabla f^{(i)}(\zeta_t) \right) \right\rangle \right] \\ &\quad + \frac{\gamma^2 L}{2} \mathbb{E} \left[\left\| \text{avg}_i \left(\nabla f^{(i)}(\omega_t^{(i)}) \right) \right\|^2 \right] + \gamma^2 \frac{L\sigma^2}{2}. \end{aligned} \quad (23)$$

The second term in Eq. 23 can be simplified by $\text{avg}_i (f^{(i)}(\zeta_t)) = f(\zeta_t)$:

$$\gamma \mathbb{E} \left[\left\langle \nabla f(\zeta_t), \text{avg}_i \left(\nabla f^{(i)}(\zeta_t) \right) \right\rangle \right] = \gamma \mathbb{E} \left[\|\nabla f(\zeta_t)\|^2 \right]. \quad (24)$$

For the third term in Eq. 23 we have:

$$\begin{aligned} & \gamma \mathbb{E} \left[\left\langle \nabla f(\zeta_t), \text{avg}_i \left(\nabla f^{(i)}(\omega_t^{(i)}) - \nabla f^{(i)}(\zeta_t) \right) \right\rangle \right] \\ & \leq \frac{\gamma}{2} \left(\mathbb{E} \left[\|\nabla f(\zeta_t)\|^2 \right] + \mathbb{E} \left[\left\| \text{avg}_i \left(\nabla f^{(i)}(\omega_t^{(i)}) - \nabla f^{(i)}(\zeta_t) \right) \right\|^2 \right] \right) \\ & \leq \frac{\gamma}{2} \left(\mathbb{E} \left[\|\nabla f(\zeta_t)\|^2 \right] + \text{avg}_i \left(\mathbb{E} \left[\left\| \nabla f^{(i)}(\omega_t^{(i)}) - \nabla f^{(i)}(\zeta_t) \right\|^2 \right] \right) \right) \\ & \leq \frac{\gamma}{2} \left(\mathbb{E} \left[\|\nabla f(\zeta_t)\|^2 \right] + 4L^2\gamma^2 G_{\max}^2 \right). \quad (\text{According to Eq. 19}) \end{aligned} \quad (25)$$

For the fourth term in Eq. 23 we have:

$$\begin{aligned} & \frac{\gamma^2 L}{2} \mathbb{E} \left[\left\| \text{avg}_i \left(\nabla f^{(i)}(\omega_t^{(i)}) \right) \right\|^2 \right] \\ & = \frac{\gamma^2 L}{2} \mathbb{E} \left[\left\| \text{avg}_i \left(\nabla f^{(i)}(\zeta_t) + \left(\nabla f^{(i)}(\omega_t^{(i)}) - \nabla f^{(i)}(\zeta_t) \right) \right) \right\|^2 \right] \\ & \leq \gamma^2 L \left(\mathbb{E} \left[\left\| \text{avg}_i \left(\nabla f^{(i)}(\zeta_t) \right) \right\|^2 \right] + \mathbb{E} \left[\left\| \text{avg}_i \left(\nabla f^{(i)}(\omega_t^{(i)}) - \nabla f^{(i)}(\zeta_t) \right) \right\|^2 \right] \right) \\ & \leq \gamma^2 L \left(\mathbb{E} \left[\|\nabla f(\zeta_t)\|^2 \right] + 4L^2\gamma^2 G_{\max}^2 \right). \quad (\text{According to Eq. 19}) \end{aligned} \quad (26)$$

Substituting Eq. 24, 25, 26 into the Eq. 23, get:

$$\begin{aligned} \mathbb{E} [f(\zeta_{t+1})] & \leq \mathbb{E} [f(\zeta_t)] - \gamma \left(1 + \frac{1}{2} - \gamma L \right) \mathbb{E} \left[\|\nabla f(\zeta_t)\|^2 \right] - 2\gamma^3 L^2 G_{\max}^2 + 4\gamma^4 L^3 G_{\max}^2 \\ & \quad + \gamma^2 \frac{L\sigma^2}{2N} \quad (\text{Assume } \gamma \leq 1/(2L)) \\ & \leq \mathbb{E} [f(\zeta_t)] - \gamma \mathbb{E} \left[\|\nabla f(\zeta_t)\|^2 \right] + \gamma^2 \frac{L\sigma^2}{2N}. \end{aligned} \quad (27)$$

Move $\mathbb{E} [\|\nabla f(\zeta_t)\|^2]$ to the left of the inequality;

$$\mathbb{E} \left[\|\nabla f(\zeta_t)\|^2 \right] \leq \frac{(\mathbb{E} [f(\zeta_t)] - \mathbb{E} [f(\zeta_{t+1})])}{\gamma} + \gamma \frac{L\sigma^2}{2N}. \quad (28)$$

Taking the sum over all iterations:

$$\frac{1}{T} \sum_{t=0}^T \mathbb{E} \left[\|\nabla f(\zeta_t)\|^2 \right] \leq \frac{(\mathbb{E} [f(\zeta_0)] - \mathbb{E} [f(\zeta_{T+1})])}{\gamma T} + \gamma \frac{L\sigma^2}{2N}. \quad (29)$$

Finally, we can bound $\|\nabla f(\omega_t^{(g)})\|$ in terms of $\|\nabla f(\zeta_t)\|$ as:

$$\begin{aligned} \mathbb{E} \left[\left\| \nabla f(\omega_t^{(g)}) \right\|^2 \right] & \leq 2 \left(\mathbb{E} \left[\left\| \nabla f(\omega_t^{(g)}) - \nabla f(\zeta_t) \right\|^2 \right] + \mathbb{E} \left[\|\nabla f(\zeta_t)\|^2 \right] \right) \\ & \leq 2\mathbb{E} \left[\|\nabla f(\zeta_t)\|^2 \right] + 2L^2 \mathbb{E} \left[\left\| \omega_t^{(g)} - \zeta_t \right\|^2 \right] \\ & \leq 2\mathbb{E} \left[\|\nabla f(\zeta_t)\|^2 \right] + \gamma^2 L^2 G_{\max}^2. \end{aligned} \quad (30)$$

Substituting this into the inequality above on $\frac{1}{T} \sum_{t=0}^T \mathbb{E} \left[\|\nabla f(\zeta_t)\|^2 \right]$ gives the claimed bound:

$$\frac{1}{T} \sum_{t=0}^T \mathbb{E} \left[\left\| \nabla f(\omega_t^{(g)}) \right\|^2 \right] = O \left(\frac{f_{\max}}{\gamma T} + \gamma^2 L^2 G_{\max}^2 + \gamma \frac{L\sigma^2}{N} \right). \quad (31)$$

Using the step size $\gamma = \sqrt{N}/\sqrt{T}$, we get:

$$\begin{aligned} & \frac{1}{T} \sum_{t=0}^T \mathbb{E} \left[\left\| \nabla f \left(\omega_t^{(g)} \right) \right\|^2 \right] \\ &= O \left(\frac{f_{\max}}{\sqrt{NT}} + \frac{N}{T} L^2 \mathbf{G}_{\max}^2 + \frac{L\sigma^2}{\sqrt{NT}} \right). \end{aligned} \quad (32)$$

Proof for Asynchronous PFL. We use the same notation in synchronous PFL for convenience. In asynchronous PFL, we can use almost the same method of Avdiukhin & Kasiviswanathan (2021a) to prove the convergence of our method. Our contribution does not lie in the novelty of the method of proof, but rather in pointing out that our approach can be fully accommodated in existing proofs of asynchronous FL and achieve greater speed-up ratios.

Definition 5 (Asynchronous PFL Sequence) Let $\omega_t^{(g)}$ represents the global model of asynchronous PFL, $G_t^{(i)}$ represents the uploaded gradient updates computed by the i -th client based on $\omega_t^{(g)}$, the global model $\omega_t^{(g)}$ can be expressed as:

$$\omega_t^{(g)} = \omega_0^{(g)} - \sum_{\tau=1}^t \gamma G_{\tau-\delta}^{(i)}. \quad (33)$$

Proposition 6 (Distance Bound) For ζ_t of virtual sequence and $\omega_t^{(g)}, \omega_t$ of Algorithm 1, we have:

$$\max \left(\mathbb{E} \left[\left\| \zeta_t - \omega_t \right\|^2 \right], \mathbb{E} \left[\left\| \zeta_t - \omega_t^{(g)} \right\|^2 \right] \right) \leq 4(\delta - 1)^2 \gamma^2 \mathbf{G}_{\max}^2 \quad (34)$$

Proof:

$$\begin{aligned} \mathbb{E} \left[\left\| \zeta_t - \omega_t^{(g)} \right\|^2 \right] &= \mathbb{E} \left[\left\| \sum_{\tau=1}^t \gamma \left(\mathbf{G}_{\tau-\delta}^{(j)} \right) - \left(\sum_{\tau=1}^t \gamma G_{\tau-1}^{(i)} \right) \right\|^2 \right] \\ &\leq \gamma^2 \mathbb{E} \left[\left\| \sum_{\tau=1}^t \left(\mathbf{G}_{\tau-\delta}^{(i)} - \mathbf{G}_{\tau-1}^{(i)} \right) \right\|^2 \right] \\ &\leq \gamma^2 (\delta - 1)^2 \mathbf{G}_{\max}^2. \end{aligned} \quad (35)$$

Similarly, $\mathbb{E} \left[\left\| \omega_t^{(g)} - \omega_t \right\|^2 \right] \leq \gamma^2 (\delta - 1)^2 \mathbf{G}_{\max}^2$. Combining these bounds, we have the following.

$$\begin{aligned} \mathbb{E} \left[\left\| \zeta_t - \omega_t \right\|^2 \right] &= \mathbb{E} \left[\left\| \zeta_t - \omega_t^{(g)} + \omega_t^{(g)} - \omega_t \right\|^2 \right] \\ &\leq 2 \left(\mathbb{E} \left[\left\| \zeta_t - \omega_t^{(g)} \right\|^2 \right] + \mathbb{E} \left[\left\| \omega_t^{(g)} - \omega_t \right\|^2 \right] \right) \\ &\leq 4\gamma^2 (\delta - 1)^2 \mathbf{G}_{\max}^2. \end{aligned} \quad (36)$$

Theorem 7 Let $f_{\max} = f \left(\omega_0^{(g)} \right) - f \left(\omega^* \right)$, where ω^* is the minimizer for f , we have:

$$\frac{1}{T} \sum_{t=0}^T \mathbb{E} \left[\left\| \nabla f \left(\omega_t^{(g)} \right) \right\|^2 \right] = O \left(\frac{f_{\max}}{\gamma T} + \gamma^2 L^2 \mathbf{G}_{\max}^2 \delta^2 + \gamma \frac{L\sigma^2}{N} \right). \quad (37)$$

Proof: Similar to Theorem 4 B.

If $T > N^3$ and $\delta \leq T^{1/4}/N^{3/4}$, we get:

$$\frac{1}{T} \sum_{t=0}^T \mathbb{E} \left[\left\| \nabla f \left(\omega_t^{(g)} \right) \right\|^2 \right] = O \left(\frac{L f_{\max}}{\sqrt{NT}} + \frac{\mathbf{G}_{\max}^2}{\sqrt{NT}} + \frac{L\sigma^2}{\sqrt{NT}} \right). \quad (38)$$

C STUDY ON THE INFLUENCE OF THE STALE GLOBAL MODEL

In our scheme, the key to achieving parallel federated learning is to send the stale global model to the edge devices. It would be interesting to investigate the influence of the stale global model on the final performance. In this section, we conduct a convergence analysis through empirical experiments with respect to both accuracy and communication rounds on the CIFAR-10 dataset with 20 edge devices. The convergence curve is provided in Fig. 4. It can be found that, in the early training stage, the proposed PFL suffers from lower accuracy and stronger oscillations compared with FedAvg; while with the communication round increases, the PFL finally achieves similar performance to FedAvg. This is consistent with the intuition that, when the algorithm converges, there is no significant difference between using the stale global model or using the up-to-date global model.

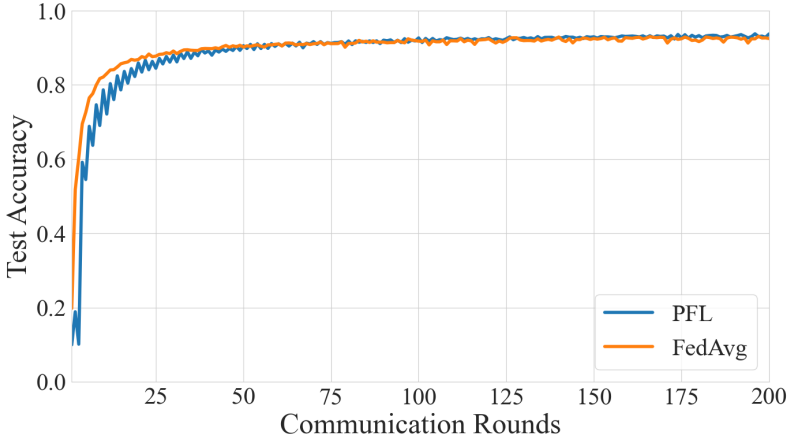


Figure 4: The convergence curve of PFL and FedAvg on CIFAR-10.

D SPECIAL CASE IN SYNCHRONOUS PFL FOR $T_g = T_l + T_c + T_w$

Here we give a timing diagram for a special case in synchronous PFL when $T_g = T_l + T_c + T_w$. As can be seen in Fig. 5, in this setting, the time spent on each round is T_g .

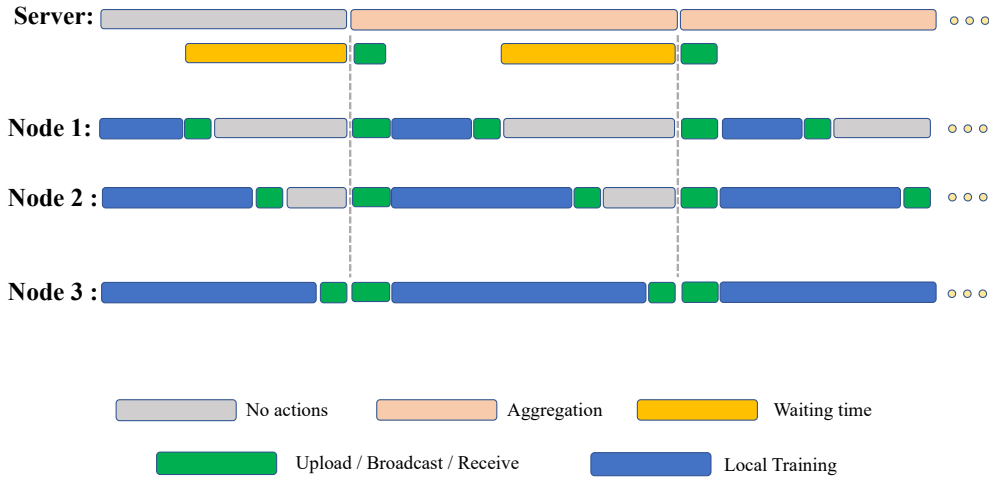


Figure 5: The timing diagram when the aggregation time is $T_g = T_l + T_c + T_w$.