

BAFFLE : Blockchain Based Aggregator Free Federated Learning

Paritosh Ramanan

*School of Industrial and Systems Engineering
Georgia Institute of Technology
Atlanta, Georgia
Email: paritoshpr@gatech.edu*

Kiyoshi Nakayama

*NEC Laboratories America
San Jose, California
Email: knakayama@nec-labs.com*

Abstract—A key aspect of Federated Learning (FL) is the requirement of a centralized aggregator to store and update the global model. However, in many cases orchestrating a centralized aggregator might be infeasible due to numerous operational constraints. In this paper, we introduce BAFFLE, an aggregator free, blockchain driven, FL environment that is inherently decentralized. BAFFLE leverages Smart Contracts (SC) to store the global copy of the model, delineate the FL mechanism into distinct rounds and aggregate local models and update the global copy after each round. BAFFLE boosts computational performance by first decomposing the global parameter space into distinct chunks followed by a novel score and bid strategy leading to significant reduction in computational costs on the blockchain. In order to validate our claims we conduct extensive experiments using a pertinent case study on a private Ethereum network and demonstrate the computational efficiency and scalability of BAFFLE. Further, our results also show that BAFFLE delivers similar performance as its centralized as well as classical FL counterparts in addition to minimizing the computational overhead of blockchain based decentralization.

Index Terms—Blockchain based decentralization, Aggregator Free Federated Learning, Ethereum driven Smart Contracts

1. Introduction

Federated Learning (FL) [1] is a distributed machine learning paradigm that accomplishes large scale learning tasks [2]. FL leverages data sets localized on end user devices in order to ensure privacy. A fundamental assumption of the FL paradigm is the presence of a centralized aggregator meant to coordinate the global computational progress. An aggregator discharges four main functions in the FL paradigm. First, it is responsible for delineating the global computational process into distinct rounds. Second, it maintains a global estimate of the machine learning model to be updated after every round. Third, the aggregator is responsible for selecting the devices and sending a copy of the global model estimate to each. Lastly, the aggregator is responsible for performing the critical step of updating the global model estimate with the aggregate of the selected local copies.

The requirement of central aggregator raises operational challenges, especially in FL applications wherein the global

model is of vital operational and diagnostic value to the end users [3], [4]. First, in many instances, implementing a central aggregator might not be a feasible option due to logistical challenges [3], [5]. Second, end users must have faith in the aggregator's selection and update mechanism for the user devices and their local models respectively. In case an aggregator holds a bias towards specific users, the final global model might not generalize well [6]. Third, the central aggregator results in a single point of failure for the FL task, thereby raising robustness concerns [7]. Lastly, the central aggregator is typically cloud based [1], [2]. Access to the cloud, might be out of reach for small organizations due to lack of technical skill and expertise [8]. As a result, a central aggregator might induce a high barrier of entry for small organizations which might be incapable of implementing large scale FL tasks.

Blockchain based decentralization can be effectively leveraged for alleviating operational issues concerning a centralized aggregator. However, careful consideration of the computational constraints imposed by the blockchain is required in order to realize an aggregator free FL scheme [9]. First, storage of data and computation on the blockchain incurs a significant cost. Second, pushing an entire machine learning model to the blockchain becomes computationally bulky potentially incurring heavy latency due to consensus. Lastly, there are limits on transaction size imposed by the blockchain protocols that restrict the amount of data that can be stored and updated on blockchain in a single transaction. These computational constraints place limitations on the model aggregation and update process in FL. Nevertheless, the need for a scalable framework that incurs low computational costs, retains the benefits of aggregator free FL and adheres to the computational constraints of the blockchain has so far not been addressed [9].

In this paper, we propose BAFFLE, a blockchain based aggregator free FL environment. BAFFLE leverages Smart-Contracts (SCs) to maintain the global model copy and the associated computational state of the users. By its very design, BAFFLE enables users to update the global model on the SC independently and in parallel, leading to significantly lower computational costs. On the operational front, for a particular round, selection of end users in BAFFLE is based on the worth of their local updates as assessed by

the SC. Further, BAFFLE ensures that rounds are delineated according to the reported computational state of all the users thereby avoiding bias. Lastly, owing to a fully decentralized, aggregator free approach, BAFFLE saves on cloud setup and operational costs and eliminates technical expertise requirements for maintaining centralized aggregators [10]. Therefore, BAFFLE is able to deliver high computational efficiency while successfully eliminating the operational limitations of an aggregator driven FL paradigm.

From a social standpoint, the computational benefits of BAFFLE coupled with elimination of cloud based costs and expertise requirements lowers the entry barrier for small organizations. BAFFLE can be used by micro scale organizations on public or private blockchains to self organize and leverage FL among their peers in a computationally friendly way. In doing so, each organization in a community can preserve their own data privacy but yet collaborate on building a global model that helps address challenges common to the entire community. As a result, BAFFLE can effectively be used to empower communities of users who would otherwise not have the capability to obtain robust machine learning models for their own internal challenges.

We show that BAFFLE is able to successfully eliminate the operational limitations of the aggregator driven FL paradigm in a computationally sound manner. We devise a budgeted approach to model update and aggregation steps and leverage SCs to delineate the rounds. We theoretically show that a classical FL scheme is equivalent to a BAFFLE driven approach with a linear relation between the respective learning rates. We provide a practical, production level implementation of BAFFLE on a private Ethereum network, with Solidity powered SC deployments. We demonstrate the merits of BAFFLE on a real world case study using a large Deep Neural Network(DNN) model.

Based on our case study, we perform exhaustive experiments to study the user benefits, robustness and scalability of BAFFLE compared to other benchmarks. Our results indicate that BAFFLE provides superior computational performance despite the highly restrictive constraints imposed by the blockchain.

Our paper is organized as follows. In Section 2 we provide an overview of related work pertaining to the fields of blockchain and decentralized ML. Section 3 discusses the novel strategies employed in BAFFLE to circumvent the restrictions imposed by the blockchain. Section 4 provides an overview of the local and global computational perspectives of BAFFLE. Section 5 introduces a real world case study of improving driver revenue where an aggregator free FL mechanism could be highly beneficial. Section 6 deals with the entire set of experiments and their analysis. We conclude the paper in Section 7 in addition to providing a quick overview of future work.

2. Related Work

Improving a global neural network model using distributed data with a privacy-preserving purpose was first studied in [11]. The authors provide a scheme of jointly learning an accurate model by multiple parties for a given

objective. More specifically, they consider a global shared memory model where parameters of the global model are held. Various agents participating in this framework can update a random subset of global parameters based on their local training.

Federated Learning was later proposed in [1], [12] with theoretical basis explored in [13]. The authors provide an effective method for building collective knowledge across a set of devices while preserving their individual autonomy and privacy. There are ongoing efforts to scale up the FL framework as presented in [2]. The framework considers multiple aggregators headed by a master in order to manage the entire FL process. Although the work proposes a distributed network of aggregators coordinated by a master, it is not inherently aggregator free.

[14], [15] propose a framework of fully decentralized FL in which users update their belief by aggregating information from neighbors. While the theoretical aspect of decentralized FL is explored in these works, numerous system and architectural issues persist in achieving true decentralization. As a result, such systemic issues need to be dealt with in order to obtain a FL framework that is feasible under practical settings.

Practical efforts to integrate AI onto the blockchain are largely confined to white paper proposals without any tangible real world implementations available. The framework proposed in [16] designs an SC based machine learning platform allowing users to upload tasks as well as contribute models to solve existing tasks. A distributed, AI computing platform has also been proposed in [17] where mining nodes earn their income from processing AI models.

There are also several projects that integrate federated learning into blockchain technologies. The work done in [18] supports implementing the FL framework into the mining mechanism of the underlying blockchain platform. However, owing to modification requirements to the underlying consensus protocols such approaches tend to be cumbersome to implement on off the shelf blockchain platforms. The work done in [9] proposes and implements a decentralized AI framework using the blockchain. However, a key requirement of this framework is that training data from devices needs to be published on the blockchain. As a result, the data privacy benefits of FL paradigm is eliminated. In fact, the authors note that a decentralized, blockchain based AI framework with full user data privacy is a key component of their future work.

Despite the above mentioned attempts, a concrete, practical framework for realizing decentralized aggregator free FL is so far lacking both in research and in industrial domains.

To the best of our knowledge, BAFFLE is the first production-level decentralized FL platform that could run over existing blockchain networks such as Ethereum.

3. Smart FL Contract Design: Decentralizing Role of Aggregator

As mentioned in Section 1, a number of technical aspects need to be considered in order to make the FL process

aggregator free. In this section, we examine the salient features of BAFFLE that allows us to circumvent blockchain based system constraints without compromising on solution quality. Even though BAFFLE has been implemented and evaluated on the Ethereum platform, the same technical principles would extend over to other blockchain based SC platforms as well.

3.1. Chunking

Most blockchain platforms have an upper limit pertaining to the data size of each transaction. For the Ethereum Virtual Machine (EVM) with the version we have used, this limit has been set to 24 kB by default. Such a limitation immediately results in a bottleneck for an aggregator free FL scheme since the underlying machine learning models are usually significantly larger than the transaction limit sizes. Such a system induced constraint necessitates the need for partitioning the machine learning model weight vector into numerous *chunks* such that each chunk size is less than the maximum transaction size.

However, chunking in turn introduces a few other notable aspects with regards to model sharing.

Serialization: Since storage on the SC is expensive, the machine learning model needs to be stored in a serialized format. However, partitioning the model after serialization leads could lead to inconsistencies. Therefore, for a specific FL task, it is important to first generate a partitioning scheme that must be used by all agents followed by individual serialization of the chunks. Such a *chunk-and-serialize* scheme has numerous benefits. First, the chunks can be read to and written from independently and seamlessly. Second, such an independence among chunks can be exploited for parallel updates from multiple devices at the same time. Lastly, a chunk independence scheme also leads to a potential scoring technique wherein parts of the model can be evaluated for their worth.

Budgets: A potential benefit of chunking is that user devices are empowered to decide their levels of contribution individually. Since, pushing chunks on the blockchain involves a computational cost as well as miner fees, users can independently evaluate their own cost to benefit ratio and decide the number of chunks that they wish to update in a round. The maximum limit on the number of chunks that a user device wishes to update is referred to as the budget for that device. As a result the set of budget values from all user devices can be heterogenous in nature.

3.2. Scoring and Bidding

Each chunk is assigned a score by the end user devices themselves based on a norm difference with respect to the latest available global copy. Depending on a random selection the user device submits bids on a set of chunks as allowed by the budget limit. The SC receives bids on a diverse set of chunks from different user devices in every round. For chunks on which multiple bids were submitted, the device submitting the maximum score is chosen as the sole updater of the chunk.

3.3. Delineation of Rounds

Owing to the decentralized nature of our approach, the onus of delineating the rounds rests with the end user devices themselves aided by the information maintained on the SC. Specifically, a Participation Level (PL) is chosen for every FL learning task which specifies the number of agents which must have submitted their bids in order for the round to start. Once the participation level criteria is met, the round begins and no new devices are allowed to participate. Devices upload the chunks on whom their bids were accepted and proceed to signal a close of their round.

4. Computational Perspectives of BAFFLE

The entire aggregator free blockchain based FL paradigm presented by BAFFLE can be viewed in terms of two important perspectives. The computational steps undertaken by the user devices in their interaction with the blockchain network forms the local perspective. The global perspective details the computational picture pertaining to the role of the SC in orchestrating the BAFFLE framework.

4.1. User Level Actions: The Local Perspective

Locally, model training and aggregation form the two important steps that every user device participating in BAFFLE must undertake.

4.1.1. Local Training

User devices continuously observe new data points from their environment which can be leveraged for the FL task at hand. The user device pulls the latest available model from the blockchain and performs an average with its latest available local copy. The resulting model is used to train on the locally available data to yield the new local copy.

4.1.2. Model Aggregation and Update

In order to aggregate with the other devices and push its update to the chain, every agent considers the local model copy obtained after local training. The steps taken by the user device for model aggregation and update can be traced with the help of the flowchart depicted in Figure 1 and summarized concisely in Algorithm 1. Each user device is initialized on the basis of the same given partition scheme. As soon as local training is complete, user devices average with the global model copy and check for the round status. In case a round is already underway and thus inactive, the user device returns to the task of collecting new data. If a round is active and accepting bids, devices choose randomly from their local chunks based on their budget size. A scalar score is assigned to each chunk based on the norm of the difference of the local weights with the global weights copy. These scores form the basis of the bid submitted to the SC which decides on which user device gets to update which chunk.

4.2. Global Perspective: The Overall Picture

Globally, the computational process employed by BAFFLE is divided into three distinct phases. We illustrate the global computational perspective with the help of an example shown in Figure 2. In our example we consider a BAFFLE system comprising of 5 asset devices A1, A2, A3, A4, A5 respectively. The model is divided into 5 chunks C1, C2,

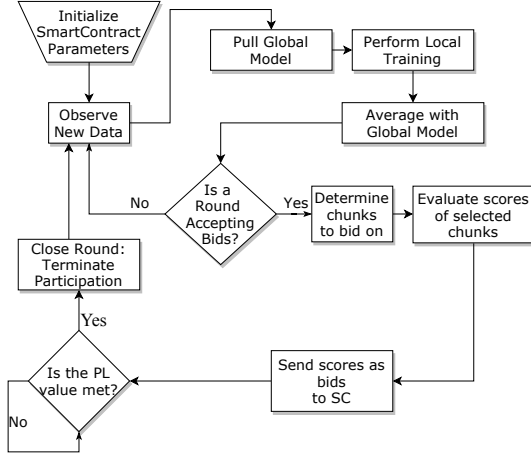


Figure 1. Flowchart depicting the sequence of events at the agent level

Algorithm 1 Agent based SC interaction

for $k = 0 \dots \text{do}$
 if round is open for participation **then**
 choose chunks $\tilde{C}^k \subseteq \mathcal{C}, |\tilde{C}^k| = B^j$ randomly
 calculate scores $\delta_c = \|Q^c - Q_{k+1}^{j,c}\|, \forall c \in \tilde{C}^k$
 submit bids $[c, \delta_c], \forall c \in \tilde{C}^k$ to SC
 determine accepted chunk set $C^k \subseteq \tilde{C}^k$
 push C^k to blockchain
 end if
end for

C3, C4 and C5. For this example, we consider a PL value of 4. In Phase 1, each device performs local training and aggregation to generate new bids. Next, every device attempts to submit bids for its randomly chosen chunks. The bids chronologically arrive in the order A1, A3, A5, A4 and A2. In Phase 2, owing to the PL value being met with the arrival of bids from A4, A2 is rejected from the current round. The accepted devices push the respective chunks for which their bids were accepted. In Phase 3, every device eventually signals the culmination of all its local steps to the SC to mark the end of Phase 3 as well as the current round.

From a theoretical perspective, using lemma 1, we show that the global computational process is equivalent to classical FL scheme with a learning rate that varies by a constant factor.

Lemma 1. *Given a BAFFLE framework and its aggregator driven counterpart with learning rates η_{BFL}, η_{FL} respectively, involving a total of C chunks with L participating agents each possessing a maximum budget potential of $B \leq C$, the following relation holds*

$$\eta_{BFL} = \frac{2 \cdot C \cdot (C - B + 1) \cdot \alpha_{FL}}{B \cdot L \cdot \alpha_{BFL}} \cdot \eta_{FL}$$

where $\alpha_{FL}, \alpha_{BFL}$ are the probability that an agent is selected for model aggregation for aggregator free decentralized FL and the classical FL respectively.

Proof. We know that for SGD the following relation holds:

$$\hat{w}_i^{k+1} = \hat{w}_i^0 - \eta \sum_{t=1}^k \nabla f(\hat{w}^t)_i \quad (1)$$

where \hat{w}_i^k is the estimate of the i^{th} component of the weight vector at round k , η is the learning rate. Further, $\nabla f(\hat{w}^k)_i$ is i^{th} component of the gradient estimated based on the globally available weight vector.

In case of BAFFLE, we can say that

$$\hat{w}_i^{k+1} = \frac{1}{2} [\hat{w}_i^k + \hat{w}_i^k - \eta_{BFL} \nabla f(\hat{w}^t)_i] \quad (2a)$$

$$\hat{w}_i^{k+1} = \hat{w}_i^k - \frac{\eta_{BFL}}{2} \nabla f(\hat{w}^t)_i \quad (2b)$$

Therefore, if at the t^{th} round, device j_t is active and the i^{th} component is chosen, it follows that the expected value of the weight vector is given by:

$$E[\hat{w}_i^{k+1}] = \hat{w}_i^0 - \frac{\eta_{BFL}}{2} E \left[\sum_{t=1}^k \nabla f_{j_t}(\hat{w}^t)_i \right] \quad (3)$$

At every round, we also assume that the probability of user device j_t being selected is denoted by α_{BFL} . Given a budget size B , the total number of chunks C , devices choosing their chunks randomly subject to budget B , the probability of picking the chunk containing the i^{th} weight element, is then determined as follows:

$$\frac{(C-1)}{(B-1)} / \binom{C}{B} = \frac{B}{(C-B+1)C} \quad (4)$$

Therefore, Equation 3 is equivalent to:

$$E[\hat{w}_i^{k+1}] = \hat{w}_i^0 - \frac{\eta_{BFL}}{2} \left[\sum_{t=1}^k \alpha_{BFL} \sum_{j=1}^n \frac{B \nabla f_{j_t}(\hat{w}^t)_i}{(C-B+1)C} \right] \quad (5)$$

which leads to:

$$E[\hat{w}_i^{k+1}] = \hat{w}_i^0 - \frac{B \cdot \eta_{BFL} \alpha_{BFL}}{2(C-B+1)C} E \left[\sum_{t=1}^k \sum_{j=1}^n \nabla f_{j_t}(\hat{w}^t)_i \right] \quad (6)$$

On the other hand, with aggregator driven FL, with L user devices aggregated in each round, we can similarly state:

$$E[\hat{w}_i^{k+1}] = \hat{w}_i^0 - \frac{\eta_{FL} \cdot \alpha_{FL}}{L} \left[\sum_{t=1}^k \sum_{j=1}^n \nabla f_{j_t}(\hat{w}^t)_i \right] \quad (7)$$

where α_{FL}, η_{FL} is the probability of choosing a device and the learning rate of aggregator driven FL respectively.

Therefore, equating 6 and 7, we can say that with a learning rate of

$$\eta_{BFL} = \frac{2 \cdot C \cdot (C - B + 1) \alpha_{FL}}{B \cdot L \cdot \alpha_{BFL}} \cdot \eta_{FL} \quad (8)$$

BAFFLE is equivalent to classical FL with learning rate η_{FL} . \square

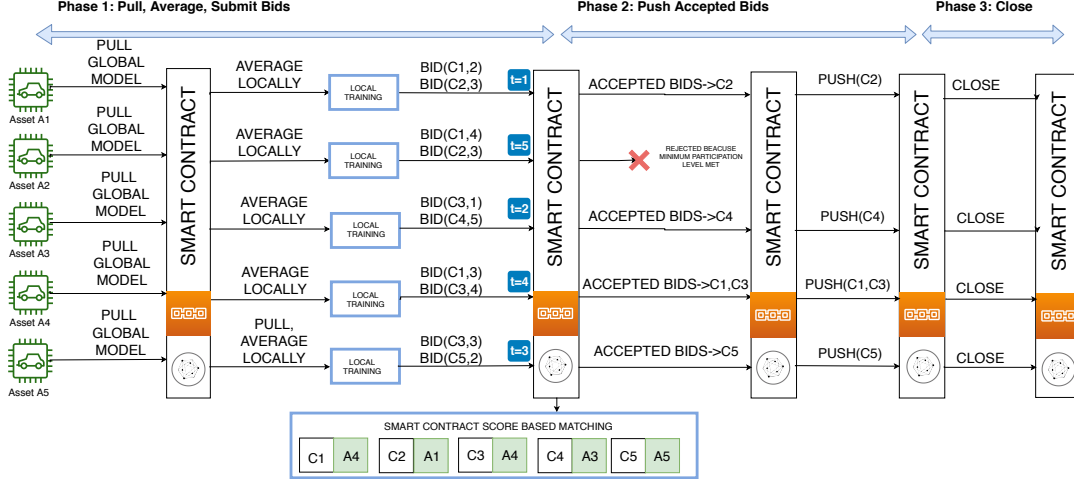


Figure 2. Global computational steps

4.3. Smart FL Contract Data Structure

The Smart FL Contract follows contract-oriented design principles that required to function on the blockchain network. Fields of significance contained in the Smart FL Contract are listed in Table 4.3.

For each attribute, we have set up an appropriate modifier that restricts the access to modify the value. In particular, the modifier for the AI Model data is designed only for the potential FL contributors that can benefit from the global model updated by the participants.

5. Case Study: Improving Taxi Driver Revenue with BAFFLE

A key problem in the taxi and ride sharing industry is to improve driver revenue by reducing idle time [19]. Drivers are often unable to find passengers at certain locations in the city at varying points of time during the day due to low demand [19]. As a result, they usually hover around the same location until they find a passenger. Idling time reduces vehicle utilization and leads to potential loss in revenue for the individual driver [20].

The application of machine learning to improve driver revenue by reducing idle time has been studied before [19], [20], [21]. Based on existing work, a Deep Reinforcement Learning (DRL) scheme is demonstrated to provide good quality improvement in driver revenues [21]. However, these approaches assume the presence of a centralized coordinator to steer the RL process. A central repository of ride information presents several privacy issues which have been successfully exploited to de-anonymize passenger information [22]. The work done in [23] as an extension of [21] introduces privacy preserving features and distributed computation as a means to improve driver revenue. However, [23] assumes a hierarchical computational setup that prevents all the benefits of decentralized computations from being realized in their entirety. The requirements of multiple control centres to perform the learning tasks leads to limited applicability of such approaches.

5.1. Benefits of Aggregator Free FL for Improving Driver Revenue

The taxi and ride sharing industry is a perfect example of micro scale enterprises that could benefit significantly from an aggregator free FL approach. The ride sharing and taxi industry remains largely an unorganized market where setting up a trusted coordinator remains a challenging proposition. Even in case of a central data repository, extracting intelligence from the anonymized data proves to be a futile exercise [22]. Moreover, drivers usually also do not have access to sophisticated computing platforms on which they could orchestrate learning tasks to improve their revenue. Therefore, a decentralized aggregator free FL environment allows drivers to leverage their collective ride experiences and improve their revenue without sharing their private ride data itself.

5.2. Deep Batch Reinforcement Learning for Taxis

We use a batch DRL paradigm to learn the Q function values and employ the Deep Neural Fitted Q [24] method to accomplish our learning task. Specifically, we define our states and actions as follows:

- Pickup State $s_i: \langle \text{pickup_location}, \text{pickup_time} \rangle$
- Dropoff State $s'_i: \langle \text{dropoff_location}, \text{dropoff_time} \rangle$
- Action $a: \text{action}(\text{dropoff_location})$
- Reward $r: \text{fare}$

State is defined by $S \times T$, where S is set of discrete cells that divide the city into distinct grids. T is set of 96 discrete intervals of 15 mins each for 24 hours. Therefore, given N rides, we denote the ride set $\mathcal{H} = \{(s_i, a_i, s'_i, r_i), \forall i \in \{1, \dots, N\}\}$.

$$\tilde{Q}_k(s_i, a_i) = r_i + \gamma \max_b Q_k(s'_i, b), \forall i \in \mathcal{H} \quad (9a)$$

$$Q_{k+1} \leftarrow \tilde{Q}_k - \eta \nabla \tilde{Q}(s_i, a_i) \quad (9b)$$

Equations 9a and 9b govern the functioning of the batch DRL framework at the k^{th} round. The Q function is updated

TABLE 1. SC ATTRIBUTES IN BAFFLE

Attribute	Description
Model ID	Unique identifier assigned for every FL task by the SC
Round Registration Details	List of users with submitted bids for the upcoming round
Participation Level	The minimum number of users with submitted bids required to begin a round
Chunk Core Array	A data structure for every chunk holding: last updated time; last user to update; set of submitted scores & their owners

based on Equation 9a before being trained on the DNN using Equation 9b.

Algorithm 2 BAFFLE for Improving Driver Revenues

```

for taxi:  $j = 1 \dots P$  do
  initialize model  $Q_0^j = Q^{init}$ , budget  $B$ 
  initialize chunk set  $\mathcal{C}$  based on given partition scheme.
  for  $k = 0 \dots$  do
    observe new ride set  $\mathcal{H}^k$ 
    pull latest available model  $Q$  from blockchain
    perform averaging  $Q_k \leftarrow \frac{Q_k^j + Q}{2}$ 
    update  $\tilde{Q}_k^j$  based on Equation 9a
    locally train  $Q_{k+1}^j$  via Equation 9b
    employ Algorithm 1 to push updates to SC
  end for
end for

```

In Algorithm 2, we consider P taxis and begin by initializing all user devices to the same initial state. Next the partition information and SC details is loaded on each device. The user devices utilize a new set of rides accumulated locally in every round. The local estimate of the Q function is updated and trained locally based on Equations 9 before being pushed onto the blockchain using Algorithm 1.

5.3. Data and Benchmarking Techniques

For our case study, we used the NYC taxi data set [25] for our experiments. Specifically, we randomly chose 2 million rides pertaining to May 2018 which was divided into two equal parts to denote the training and testing data sets. Restricting the rides specifically for the area of lower Manhattan resulted in approximately a little more than half million rides each in training and test data sets. The training set was used to assign rides to taxis participating in the FL process.

On the basis of the test set, we determine 50 taxi trajectories which form benchmark for FL tasks based on work done in [20]. Each trajectory comprises of 50 rides and assumes idling in case no ride is found. The sum total of fares accrued from the 50 benchmark trajectories is referred to as the Aggregated Simulation Revenue (ASR) which forms the *No Learning (NL)* baseline for our case study.

The benchmark trajectories and the accompanying simulation procedure are also used to calculate ASR values for various DRL models as well. However in this case, instead of hovering in the same location upon not finding a ride, the DRL model in question is used to determine a new location to transition into [20]. The sum total of fares from the ensuing trajectories denotes the ASR value for the DRL

model being considered. For robustness purposes, we perform this simulation multiple times for any DRL model and report the average ASR value.

We derive a RandomDFL mechanism that is inspired by the work done in [11] that can be directly applied for orchestrating a naive aggregator free FL approach. RandomDFL is described in detail in Section B

6. Experiments

In order to evaluate the efficacy of BAFFLE, we focus on four key experiments. We perform a benchmark study where we compare the potential benefits from BAFFLE with respect to classical FL as well as other non FL paradigms. Next, we examine the trends arising from varying number of chunks as well as budget sizes of user devices. We then move onto a scalability analysis that demonstrates the impact of varying the total number of active user devices on the model quality. Lastly, we demonstrate the robustness of BAFFLE to the participation level (PL) parameter of BAFFLE. Further, we also show superior computational performance of BAFFLE compared to the best possible aggregator free approach inspired by the current state-of-the-art.

6.1. Experimental Setup

BAFFLE was implemented and evaluated on a private Ethereum blockchain setup exclusively for our computational experiments. We employed `go-ethereum`, an official `go` based implementation of the Ethereum protocol [26] to orchestrate our private blockchain comprising of 16 Ethereum nodes. Proof-Of-Authority was used as the primary consensus protocol for all our experiments. The SC layer was developed using the Solidity programming language and deployed on the private blockchain using `go-ethereum`. The private blockchain was deployed on an Intel Xeon CPU with a clock rate of 2.40 GHz with 16 cores and 2 threads per core. We used `OpenMPI` [27] in conjunction with `mpi4py` [28] to spawn multiple distributed memory client processes intended to simulate the user devices on the field. Client processes were deployed on Intel Core i7 CPUs with 12 cores each. We used a 2 layer DNN with 500 perceptron in each layer for our experiments on Keras [29] with TensorFlow [30].

6.2. Benefits Study

In this experiment we compare the benefits accrued by drivers participating in BAFFLE with respect to two other types of learning paradigms. The first comprises of a *Local Learning (LL)* mechanism, wherein no model aggregation is involved. The second paradigms pertains to an aggregator driven Classical FL(CFL) scheme. We considered each taxi

TABLE 2. BENEFIT ANALYSIS

Category	ASR (USD)	Benefit (%)
No Learning (NL)	13387.31	-
Local Learning (LL)	16106.02	20.31
Classical FL (CFL)	18495.94	38.16
BAFFLE	18442.21	37.75

having accumulated approximately 700 rides in each round for a total of 50 rounds. For the FL cases we considered 16 taxis whereas for the LL case, we considered a single taxi. Table 2 presents the results with respect to LL, CFL and BAFFLE mechanisms in terms of their ASR value and benefit relative to the NL baseline.

The trends depicted in Table 2 provide numerous key insights into the performance of BAFFLE. Primarily, we observe that BAFFLE is able to provide a benefit of approximately 38% which rivals the CFL approach. Further, we observe that BAFFLE and CFL approaches improve driver benefit by close to 18% as compared to the LL case. Overall, the results demonstrate that blockchain driven FL paradigms are highly capable of delivering good quality machine learning models in an aggregator free, decentralized fashion.

6.3. Sensitivity Analysis

TABLE 3. FINAL BENEFIT (%) BASED ON AVERAGE ASR

Chunk Size (kB)	No. Of Chunks	Budget Size		
		16	24	32
2	738	38.32	38.18	36.51
4	356	36.37	36.87	39.17
8	181	40.23	34.79	38.11
16	88	39.07	38.82	38.02

TABLE 4. AVERAGE TOTAL TRAINING TIME(IN SECS) (STD DEV.)

Chunk Size (kB)	Budget Size		
	16	24	32
2	87.48(2.89)	85.97(3.26)	73.49(2.70)
4	79.92(3.18)	77.63(2.83)	73.22(3.87)
8	74.16(3.70)	71.90(2.53)	69.79(3.09)
16	73.44(4.01)	76.39(3.37)	71.79(3.51)

We perform a robustness study to analyze the impact of variation in chunk sizes as well as local budget sizes on the overall model quality. For this experiment, we considered a total of 64 taxis, with each taxi having accumulated approximately 70 rides in each round for 125 rounds overall. Table 3 shows the benefit percentage calculated for varying chunk and budget sizes. Figure 3 represents the overall trends with Figures 3(a), 3(b) depicting the boxplots pertaining to Gas Costs, Push Time respectively. Table 4 shows the mean and standard deviation with respect to the training time incurred by the individual agents.

The results for all the combinations in Table 3 depict benefits that closely mirror that of the CFL approach shown in Table 3 on the same training set. Therefore, on the basis

of data presented in Table 3 one can conclude that BAFFLE is significantly resilient to varying degrees of budget and chunk sizes.

On the basis of Table 4, we conclude that time incurred for training is marginal compared to the push time depicted in Figure 3(b) for all combinations of budget and chunk sizes. The relatively small training time implies that reducing the total push time is critical in ensuring a computationally efficient performance for a blockchain based FL mechanism.

We draw upon the trends shown in Figure 3 to reveal numerous key insights which elucidate the high computational efficiency of BAFFLE.

Primarily, in Figure 3(a) we observe a smaller variation in gas costs for the 2 kB chunk size irrespective of budget sizes. However, as the chunk size increases we see the variation in gas costs also increase substantially for all budget sizes. Second, despite the increased variation, the mean gas cost appears to saturate for higher chunk sizes. We also observe that for the budget size of 32 after the initial uptick there is a relatively more pronounced downward trend for higher chunk sizes. This trend can be clearly attributed to the scoring and bidding mechanism incorporated in BAFFLE. Since a higher chunk size implies lesser number of chunks, there is relatively more competition among user devices to update the same set of chunks. As a result for higher chunk sizes, only user devices which are able to consistently contribute higher scoring chunks will incur a higher gas cost. Therefore, owing to its underlying scoring and bidding mechanism, BAFFLE is able to achieve significant savings in gas costs for the users.

We observe that in Figure 3(b) despite the budget size increasing, the total push time increases only marginally owing to the scoring and bidding mechanisms. Therefore, we can safely say that BAFFLE is successfully able to circumvent the computational bottleneck posed by the push step of BAFFLE .

6.4. Scalability Analysis

TABLE 5. SCALABILITY ANALYSIS WITH VARYING NO. OF TAXIS

Taxis	Average ASR (USD)	Benefit (%)
16	14489.59	8.2
32	16547.20	23.6
64	18266.72	36.44
128	18414.48	37.55

We attempt to gauge the impact of the total number of active user devices on the performance of BAFFLE. For this experiment, we assumed each taxi having accumulated approximately 70 rides in each round for 62 rounds overall. Table 5 represents the ASR value and the ensuing benefit percentages for 16, 32, 64 and 128 taxis respectively. From the trends presented in Table 5 it is apparent that increasing number of user devices results in a sizeable improvement in the model quality. However, the trends in Table 5 also reveal that the improvement in model quality eventually saturates with increasing active devices potentially indicating a convergence to a globally superior model.




Figure 3. Performance analysis with respect to Chunk Size and Budget




Figure 4. Weak Scaling trends

Figure 4 depicts the trends pertaining to the gas costs as well as the push time with varying active user devices in Figures 4(b) and 4(a) respectively. Figure 3(a) shows a reduction in gas costs with increasing number of active devices. However, Figure 4(b) reveals little variation of push time with increase in active devices.

The reduction in gas costs in Figure 4(a) can be attributed to greater competition arising from an increase in total number of devices. Moreover, owing to a constant push time depicted in Figure 4(b) we infer that increase in number of participants leads to reduction in gas costs in BAFFLE.

6.5. Participation Level (PL) Analysis

In this experiment we study the impact of varying the PL on the performance of BAFFLE with 64 taxis, approximately 70 rides per round and a total of 62 rounds. Figure 5 presents results pertaining to PL values ranging from 5% to 75%. Further, we also compare the RandomDFL case in which devices update the global copy without any global coordination. Figures 5(a), 5(b) and 5(c) represent the trends pertaining to the growth in model quality, gas costs and the total push time pertaining to varying PL in every round.

From Figure 5(a), we observe that the fastest convergence of the model quality occurs in case of the RandomDFL case.

However, the convergence characteristics of BAFFLE with a 5% PL value closely mirrors the RandomDFL case. Overall, the trends in Figure 5(a) generally indicate that a lower PL value leads to a faster convergence. Figure 5(b) shows that a lower PL value in BAFFLE incurs a lower gas cost as well. Trends similar to Figure 5(b) are also exhibited in Figure 5(c) wherein a lower PL value in BAFFLE corresponds to a lower total push time as well. We observe that in general, BAFFLE incurs barely half the gas cost and push time as compared to the RandomDFL case. In fact, BAFFLE outperforms the RandomDFL case by a factor of more than 2 with a PL value of 5% both in terms of the gas cost as well as the push time.

Since fewer devices are pushing to the global model copy every round, the chances of multiple devices pulling the same global model is significantly higher in case of a lower PL value. This leads to greater stability in the decentralized process which ultimately leads to a faster convergence for a low PL value as shown in Figure 5(a).

BAFFLE incurs significantly lower gas costs compared to the RandomDFL case owing to minimization of redundant updates. Due to the decentralized round delineation and a robust scoring and bidding process, devices only push chunks

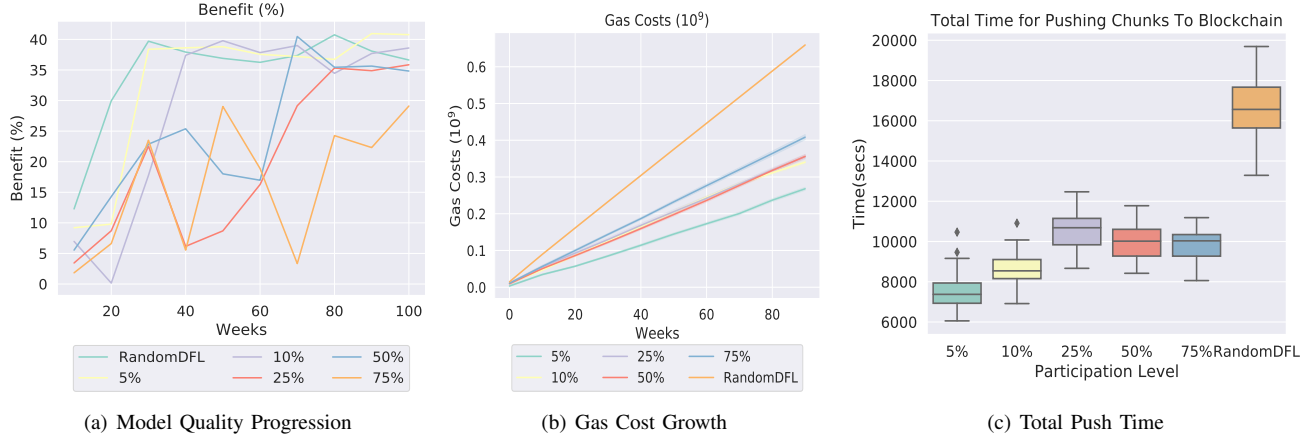


Figure 5. Performance analysis with respect to Participation Level (PL)

that are among the best in the round. As a result, collision among devices for the same chunk is completely eliminated leading to a much lower gas cost and push time.

We now move towards the concluding our work and discussing future directions.

7. Conclusion and Future Work

In this paper we investigate the use of the blockchain for realizing a decentralized aggregator free FL mechanism. We design and develop BAFFLE, a custom made blockchain based framework for aggregator free FL. In our framework, we successfully eliminate the role of a centralized aggregator by effectively decentralizing the concepts of round delineation, user device selection and model aggregation with the help of an SC. Further, in order to circumvent the computational restrictions imposed by the blockchain, we employ an effective model partitioning and serialization mechanism that enables independent and parallel model updates. We orchestrate BAFFLE on a private Ethereum blockchain network with a Solidity driven SC implementation.

We argue that the operational and computational benefits of aggregator free FL has significant potential for solving business problems for micro scale enterprises. We support our claims by applying BAFFLE to a case study pertaining to the ride sharing and taxi industry which serves as a perfect example of a micro scale enterprises. Our case study utilizes the BAFFLE framework to improve driver revenue based on a DRL model that is collectively augmented by all drivers using FL. We show that BAFFLE yields approximately a 40% improvement in driver revenues compared to non FL approaches. We further show that despite being aggregator free, BAFFLE's result quality matches that of classical FL schemes that require investment in an aggregator. Moreover, BAFFLE performs significantly better compared to other aggregator free approaches that are inspired by the current state of the art.

The issue of aggregator free FL opens up new avenues for research especially in the blockchain domain. Effective

aggregator free techniques for more complex models like CNNs and LSTM will go a long way to enable wider adoption of FL. Therefore, extending BAFFLE for handling such models forms our immediate future work. We also wish to investigate the use of differential privacy in an aggregator free setting.

Our work shows that an aggregator free approach to FL offers significant potential for revolutionizing small scale organizations and their businesses by delivering quality machine learning models at lower costs. Driven by a robust decentralized platform like the blockchain, the benefits of FL could impact a variety of domains leading to its widespread adoption.

References

- [1] J. Konecny, H. B. McMahan, F. X. Yu, P. Richtarik, A. T. Suresh, and D. Bacon, "Federated learning: Strategies for improving communication efficiency," in *NIPS Workshop on Private Multi-Party Machine Learning*, 2016.
- [2] K. Bonawitz, H. Eichner, W. Grieskamp, D. Huba, A. Ingerman, V. Ivanov, C. M. Kiddon, J. Konecny, S. Mazzocchi, B. McMahan, T. V. Overveldt, D. Petrou, D. Ramage, and J. Roselander, "Towards federated learning at scale: System design," in *SysML 2019*, 2019.
- [3] A. G. Roy, S. Siddiqui, S. Pölsterl, N. Navab, and C. Wachinger, "Braintorrent: A peer-to-peer environment for decentralized federated learning," *arXiv preprint arXiv:1905.06731*, 2019.
- [4] J. Passerat-Palmbach, T. Farnan, R. Miller, M. S. Gross, H. L. Flannery, and B. Gleim, "A blockchain-orchestrated federated learning architecture for healthcare consortia," *arXiv preprint arXiv:1910.12603*, 2019.
- [5] C. He, C. Tan, H. Tang, S. Qiu, and J. Liu, "Central server free federated learning over single-sided trust social networks," *arXiv preprint arXiv:1910.04956*, 2019.
- [6] M. Mohri, G. Sivek, and A. T. Suresh, "Agnostic federated learning," *arXiv preprint arXiv:1902.00146*, 2019.
- [7] I. Hegedüs, G. Danner, and M. Jelasity, "Gossip learning as a decentralized alternative to federated learning," in *IFIP International Conference on Distributed Applications and Interoperable Systems*, pp. 74–90, Springer, 2019.
- [8] M. Carcary, E. Doherty, and G. Conway, "The adoption of cloud computing by irish smes-an exploratory study," *Electronic Journal of Information Systems Evaluation*, vol. 17, no. 1, p. 3, 2014.

- [9] J. D. Harris and B. Waggoner, “Decentralized & collaborative AI on blockchain,” *arXiv preprint arXiv:1907.07247*, 2019.
- [10] A. Khajeh-Hosseini, I. Sommerville, and I. Sriram, “Research challenges for enterprise cloud computing,” *arXiv preprint arXiv:1001.3257*, 2010.
- [11] R. Shokri and V. Shmatikov, “Privacy-preserving deep learning,” in *Proceedings of the 22nd ACM SIGSAC conference on computer and communications security*, pp. 1310–1321, ACM, 2015.
- [12] B. McMahan and D. Ramage, “Federated learning: Collaborative machine learning without centralized training data,” 2017.
- [13] J. Konečný, B. McMahan, and D. Ramage, “Federated optimization: Distributed optimization beyond the datacenter,” *arXiv preprint arXiv:1511.03575*, 2015.
- [14] A. Lalitha, S. Shekhar, T. Javidi, and F. Koushanfar, “Fully decentralized federated learning,” *Proceedings of third workshop on Bayesian Deep Learning (NeurIPS)*, 2018.
- [15] A. Lalitha, O. C. Kilinc, T. Javidi, and F. Koushanfar, “Peer-to-peer federated learning on graphs,” *arXiv preprint arXiv:1901.11173*, 2019.
- [16] Z. Chen, W. Wang, X. Yan, and J. Tian, “Cortex - AI on blockchain: The decentralized AI autonomous system,” *Cortex White Paper*, 2018.
- [17] H. Yong, C. Lee, and D. Wang, “Artificial intelligence computing platform driven by blockchain,” tech. rep., DeepBrain Chain, Singapore, 2017.
- [18] H. Kim, J. Park, M. Bennis, and S.-L. Kim, “On-device federated learning via blockchain and its latency analysis,” *arXiv preprint arXiv:1808.03949*, 2018.
- [19] M. Han, P. Senellart, S. Bressan, and H. Wu, “Routing an autonomous taxi with reinforcement learning,” in *Proceedings of the 25th ACM International Conference on Information and Knowledge Management*, pp. 2421–2424, ACM, 2016.
- [20] T. Verma, P. Varakantham, S. Kraus, and H. C. Lau, “Augmenting decisions of taxi drivers through reinforcement learning for improving revenues,” 2017.
- [21] D. Shi, J. Ding, S. M. Errapotu, H. Yue, W. Xu, X. Zhou, and M. Pan, “Deep q-network based route scheduling for transportation network company vehicles,” in *2018 IEEE Global Communications Conference (GLOBECOM)*, pp. 1–7, Dec 2018.
- [22] M. Douriez, H. Doraiswamy, J. Freire, and C. T. Silva, “Anonymizing nyc taxi data: Does it matter?,” in *2016 IEEE International Conference on Data Science and Advanced Analytics (DSAA)*, pp. 140–148, Oct 2016.
- [23] D. Shi, J. Ding, S. M. Errapotu, H. Yue, W. Xu, X. Zhou, and M. Pan, “Deep q-network based route scheduling for tnc vehicles with passengers location differential privacy,” *IEEE Internet of Things Journal*, pp. 1–1, 2019.
- [24] M. Riedmiller, “Neural fitted q iteration—first experiences with a data efficient neural reinforcement learning method,” in *European Conference on Machine Learning*, pp. 317–328, Springer, 2005.
- [25] Travel and L. C. of New York, “Nyc 2018 yellow taxi data.”
- [26] GETH, “Go ethereum: Official go implementation of the ethereum protocol.”
- [27] E. Gabriel, G. E. Fagg, G. Bosilca, T. Angskun, J. J. Dongarra, J. M. Squyres, V. Sahay, P. Kambadur, B. Barrett, A. Lumsdaine, *et al.*, “Open mpi: Goals, concept, and design of a next generation mpi implementation,” in *European Parallel Virtual Machine/Message Passing Interface Users Group Meeting*, pp. 97–104, Springer, 2004.
- [28] L. Dalcin, “Mpi for python.”
- [29] F. Chollet *et al.*, “Keras,” 2015.
- [30] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng, “TensorFlow: Large-scale machine learning on heterogeneous systems,” 2015. Software available from tensorflow.org.

Appendix

1. Centralized Deep Batch Q Learning

Algorithm 3 details the centralized batch Deep Q Learning for improving driver revenue. It starts with observation of a new ride set every epoch. For every ride in the ride set, the existing Q-value estimate is updated with the fare collected for the ride and a discounted future reward. The discounted future reward is based on the action that gives highest Q value originating from the destination state. Based on the observed set of rides, a Deep Neural Network (DNN) is used to calculate the next Q function estimate.

Algorithm 3 Centralized Deep Neural Fitted Q

```

for  $k = 0 \dots \mathcal{d}$ 
  observe new ride set  $\mathcal{H}^k$ 
  pull latest available model  $Q$  from blockchain
  perform averaging  $Q_k \leftarrow \frac{Q_k + Q}{2}$ 
  update  $\tilde{Q}_k^j$  based on Equation 9a
  locally train  $Q_{k+1}^j$  via Equation 9b
end for

```

2. Random Decentralized FL (RandomDFL)

Algorithm 4 Randomized Decentralized Deep Neural Fitted Q

```

for taxi:  $j = 1 \dots P$  do
  for  $k = 0 \dots \mathcal{d}$ 
    observe new ride set  $\mathcal{H}^k$ 
    pull latest available model  $Q$  from blockchain
    perform averaging  $Q_k \leftarrow \frac{Q_k + Q}{2}$ 
    update  $\tilde{Q}_k^j$  based on Equation 9a
    locally train  $Q_{k+1}^j$  via Equation 9b
    push random set of chunks  $C^k \subseteq \mathcal{C}, |C^k| = B$ 
  end for
end for

```

In the randomized version represented in Algorithm 4, the SC is considered to be naive. User devices are free to update any chunks subject to their own budget values. In this naive randomized version, some chunk updates are bound to get wasted owing to the fact that they may be overwritten by another user device’s contribution before the previous update has had a chance to be read by the other agents.