

Sustainable Spectrum Crowdsensing

Yijing Zeng*, Bangya Liu*, Yilong Li, Domenico Giustiniano[†], Suman Banerjee
UW-Madison, [†]IMDEA Networks Institute
{yijingzeng, bangya, yilong, suman}@cs.wisc.edu, [†]domenico.giustiniano@imdea.org

Abstract—Spectrum crowdsensing is a paradigm where participants upload their collected spectrum data to the cloud for extracting analytics. First movers like Microsoft Spectrum Observatory and Electrosense, though with support from leading industry, research, and government, still suffer from sustainability challenges. In this paper, we present *Fiesta*, a sustainable framework for spectrum crowdsensing. On the technology side, we use federated learning and blockchain to decentralize the data analysis computations. For individual participants, minimal invasion of privacy suppresses concerns regarding large-scale adoption. From organizations’ perspectives, using blockchain avoids single point of failure and enhances the robustness of the entire system against malicious attacks. On the policy side, we propose a reward quantification mechanism to motivate engagement. Potential funding sources to ensure ongoing sustainability are also discussed. We have demonstrated *Fiesta* through simulation testbeds and real-world deployments with two demo tasks. Results show that *Fiesta*, as a decentralized framework, can preserve user privacy, enhance system robustness, maintain data fidelity compared with traditional methods, and fairly reward participants. We believe *Fiesta* is a stepping stone for the future spectrum crowdsensing paradigm.

I. INTRODUCTION

The scarcity of wireless spectrum and the exploding mobile data traffic compel regulatory authorities to consider opening previously allocated bands for secondary users, such as TV whitespace [25] and CBRS band [22], in addition to leveraging existing unlicensed bands [5]. The feasibility of such endeavors depends strongly on a comprehensive understanding of how these bands are used by primary users. This requires large-scale spectrum sensing measurements across frequency, temporal, and spatial domains.

To fulfill this requirement, crowdsensing [23] is a promising spectrum measurement paradigm. For example, Microsoft Spectrum Observatory [1] and Electrosense [20] make efforts to achieve national-wide or even continental-wise spectrum measurement using Software Defined Radio (SDR) platforms.

However, those prior projects were designed without considering sustainability, inevitably leading to the consequence that the cost of running and maintaining a network for the sole purpose of performing spectrum measurements is unjustified. In fact, both of them are shutting down.

With sustainability in mind, we envision that a large-scale spectrum crowdsensing system should have a three-tier architecture. At the bottom, mobile devices owned by individuals have spectrum sensing capability and collect spectrum measurements. In the middle, servers from organizations perform data analysis for spectrum-related applications. At the top,

the administrator coordinates the whole system, for example deciding the target frequency range and what kind of spectrum-related data analysis should be performed.

Despite the proposed three-tier architecture, there are still several challenges before realizing the vision at scale.

- Prior efforts, including TxMiner [26] and BigSpec [24], require contributors to upload raw data records with meta-data, e.g., GPS location and timestamps, to a centralized cloud for analysis. However, this information is private for participants, especially in the scenario of mobile sensing. The spectrum data itself could also be sensitive, and in some cases, information could be decoded from captured IQ data or PSD (Power Spectrum Density) data.
- Another technical challenge is that sending data to a centralized data analysis platform poses the risk of single point of failure and is more error-prone as one single entity owns the computing infrastructure entirely.
- On the policy side, without a proper reward mechanism to compensate, all these privacy concerns discourage mobile spectrum crowdsensing participants from contributing more data. Organizations may be unwilling to collaborate with little incentive.

To address these concerns, we present *Fiesta*, a sustainable framework for large-scale spectrum crowdsensing practicing the three-tier architecture.

First, we propose to move partial data analysis from the cloud server to mobile devices. We introduce and adapt federated learning [10] into the spectrum measurement context, where the user shares useful information in the format of a machine learning model instead of raw data. The scenario becomes challenging considering ordinary federated learning requires the participants to be synchronized in the sense that the model aggregation is performed after all of the participants have completed a local training round. However, synchronization is hard to achieve when scaling to a large scale where each device has its schedule of sampling and training, so tackling this asynchronous setup is a must. We propose a novel aggregation policy and empirically show its enhancement compared with vanilla federated learning.

Second, we eliminate the risk of single point of failure and enhance system robustness by using a distributed ledger, i.e., blockchain, to aggregate and archive the learned model. We design a block structure that contains the current global model so that sharing the data analysis result is naturally supported. To make it more efficient and pragmatic for large-scale deployment, we also solve several implementation issues,

*These authors contributed equally to this work.

including how to support oversized models and the evolution of the model structure.

Third, a reward quantification mechanism based on this blockchain is carefully designed to further promote the participation of both individuals and organizations. In the vision of Fiesta, individuals and organizations could gain rewards either by contributing useful information or joining the blockchain and helping with model aggregation. Administrators and authorities who could benefit from gathered spectrum data (e.g., FCC, 3GPP, operators, and research institutions) are expected to maintain this blockchain and provide rewards.

To demonstrate the efficacy of Fiesta, we have simulated, implemented, as well as deployed Fiesta using mobile devices connected to our spectrum sensor boards to learn the utilization of TV and LTE bands.

In summary, we make the following contributions:

- To the best of our knowledge, this is the first work that proposes and designs a sustainable spectrum crowdsensing system at scale. The three-tier architecture based on individuals who own mobile devices with sensing capabilities, data aggregation servers from multiple organizations, and the administrator coordinating the whole system is crucial for long-term operation.
- We solved multiple technical and policy-related challenges to realize this vision. By adapting federated learning and blockchain to the spectrum sensing context, we have mitigated the privacy concerns of the individual participants, addressed the risk of single point of failure, and encouraged larger engagement by designing a reward quantification mechanism based on the blockchain. Potential funding sources to ensure ongoing sustainability are also discussed.
- We have implemented and deployed Fiesta in a mid-size US city for 3 months and collected 220K lines of PSD data records for 600 hours using Android smartphones. We achieve 4dB accuracy for estimating TV channels' power across space and 5dB accuracy for reconstructing LTE band PSD data using machine learning models. The dataset is available at <https://www.kaggle.com/datasets/neutrino1iu/fiesta>, and the code is available at <https://fiesta4spectrum.github.io/>.

II. SYSTEM DESIGN

In this section, we present the design and the key ideas of Fiesta. We first present the architecture of Fiesta, next explain how the learning process works in Fiesta, then discuss the reward mechanisms to encourage participants, followed by a subsection justifying the necessity of blockchain.

A. Architecture

There are three different roles in Fiesta, as shown in Fig. 1, individuals with mobile devices, organizations with miners, and the administrator with seed node(s).

In Fiesta, *individuals* own the **mobile devices**, such as smartphones, with spectrum sensing capabilities. The frontend of these mobile devices has the capability of sensing the spectrum at fine granularity over a large set of frequency

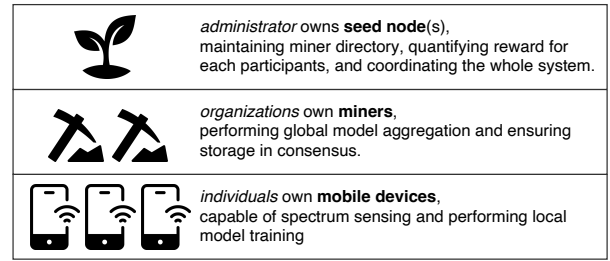


Fig. 1. Layered architecture of Fiesta.

bands. This sensing capability might probably be part of future generations of smartphones, depending on the application and the need of uploading IQ or PSD data. Mobile devices in Fiesta setup mainly operate in two modes, data collection mode and learning mode. When in the data collection mode, they record the spectrum readings and the corresponding meta-data, i.e. latitude, longitude, and timestamp. When in the learning mode, they compute the local updated model based on the current global model and the collected local spectrum data, then communicate with miners for model aggregation.

Organizations own the **miners** in Fiesta. Organizations could be cellular operators implementing an extension of Minimization of Drive Tests (MDT) measurements, or over-the-top organizations like Microsoft and Electrosense, who are interested in spectrum data. Miners will gather a bundle of updates from mobile devices and batch them into a block, then calculate the aggregated global model. They also feed the up-to-date global model with mobile devices. All the versions of the global model and the corresponding updates from mobile devices to get the global model are recorded in the form of a blockchain. The Proof-of-Work (PoW) serves as the consensus protocol among the miners.

The *administrator* of the community (e.g. the government or the regulatory authority like the FCC, 3GPP) owns the **seed node(s)** in Fiesta. These seed node(s) have two specific functions for the blockchain. First, they maintain a directory of working miners such that new mobile devices and miners joining the community know where to contact the existing miners. Second, they periodically scan the blockchain to quantify rewards for each participant. The administrator, e.g. FCC and 3GPP, also designs the machine learning models for spectrum-related applications, specifies which particular spectrum and spatio-temporal area is of interest, and offers monetary rewards for participants.

B. Learning process

Fig. 2 illustrates the model learning process in Fiesta, where red arrows indicate the cycle of federated learning and blue arrows indicate how a new block comes into birth.

Local training at mobile devices. The main idea is to use a learning model to describe the spectrum data in accordance to specific needs of the application. The mobile device starts training when the device is in charging as long as the user enables the learning mode in the app and will keep training whenever there are new global model downloadable from miner. For example, mobile device will train the pulled global

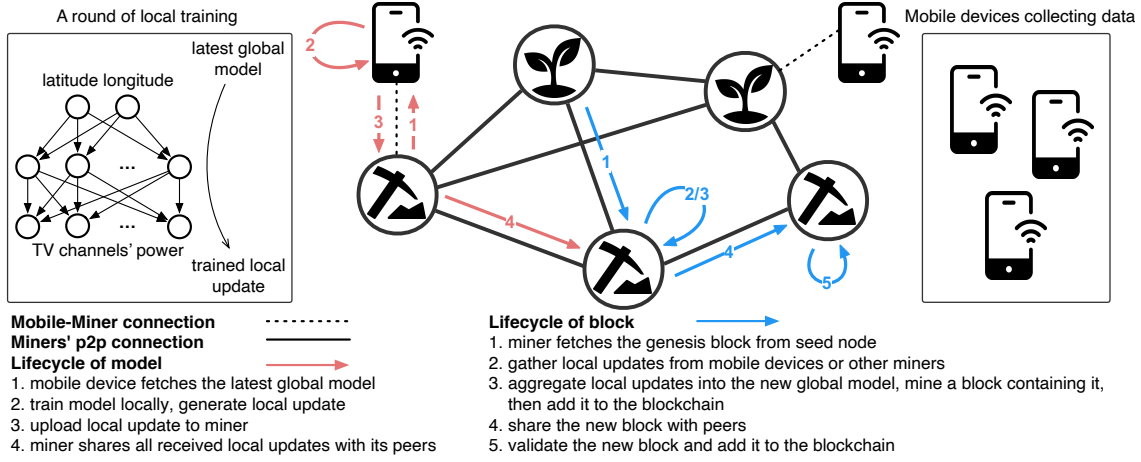


Fig. 2. Learning process of Fiesta.

model for one round (a fixed number of epochs) based on local PSD data (frequency information) and GPS readings (spatial information) to estimate TV channels' power at different locations, as shown in the left box of Fig. 2. Then, the mobile device uploads that updated model with other information including device identity, size of training set, and the training loss local training to any miner for performance monitoring and reward calculation.

Updates aggregation at miners. As shown in Fig. 2, miners receive updates, forward to neighbors, batch them into a block, and compute the new global model. We use a modified version of FedAvg [15] to deal with asynchronous model aggregation. FedAvg is the most popular model aggregation method in the synchronous setting. It computes the weighted average of the model updates with the weights equal to the number of data records, and it is guaranteed to converge with non-iid (independent and identically distributed) data for strongly convex and smooth problems [13], although slower than the iid case. Formally, suppose client $i \in \{1, \dots, C\}$ receives the same copy of the global model \mathbf{w}_t at the beginning of round $t + 1$ ($t \geq 0$). Then, after local training, all clients (or only a fraction of all clients selected by the central server) generate their local update \mathbf{w}_{t+1}^i using their local n_i number of data records. To get the new global model after round $t + 1$, the central server in FedAvg computes:

$$\mathbf{w}_{t+1} = \sum_{i=1}^C \frac{n_i}{N} \mathbf{w}_{t+1}^i, \quad (1)$$

where $N = \sum_{i=1}^C n_i$. Finally, this new global model \mathbf{w}_{t+1} is pushed to all clients for the next round of training.

Challenge of Asynchrony. Compared with ordinary federated learning, Fiesta needs to handle the challenge of asynchrony. Due to the heterogeneity of the capability of each mobile devices, the different amounts of data each mobile device collects, and the latency and throughput of the wireless connections during model learning, it is expected to encounter a variety of speeds per local round of training. Furthermore, due to the block propagation delay within the miners' network and the possible forking of the blockchain, mobile clients

might also download different versions of the global model at the same time. The challenge of asynchrony is further exacerbated by the freedom of joining and exiting Fiesta.

EWMA enhanced FedAvg. In the context of synchronous federated learning, it is expected that $block_{t+2}$ contains those model updates derived from only the global model w_{t+1} outputted in $block_{t+1}$. However, due to above-mentioned asynchronous nature of Fiesta, $block_{t+2}$ may also contain model updates based on the global model w_t outputted in $block_t$, uploaded by straggling clients. We define these updates in $block_{t+2}$ but based on global model w_t as "1 block delayed". In reality, the updates may be several blocks delayed instead of merely one. Due to this impurity of local updates, when we preform aggregation in $block_{t+2}$, we should also take the global model outputted in previous blocks into consideration. As a result, we use the exponential weighted moving average (EWMA) of the global model in Fiesta model aggregation. Formally, in $block_{t+1}$, the miner first aggregates the model updates within this block same as FedAvg, and gets the temporary model, say \mathbf{w}_{t+1} . Then, the global model will be:

$$\mathbf{w}_{t+1}^{global} = (1 - \alpha) \mathbf{w}_t^{global} + \alpha \mathbf{w}_{t+1}, \alpha \in (0, 1). \quad (2)$$

Note that the round notion in the synchronous setting does not apply here anymore, we define *async round* as each time a new block is generated with a new global model within. In addition, if miners receive multiple updates from the same mobile device in an async round (say \mathbf{w}_{t-1}^i and \mathbf{w}_t^i), they should only use the most recent one (\mathbf{w}_t^i) for aggregation so that too stale updates will be ignored.

A trivial fact is that in synchronous settings, using EWMA is equivalent to setting a lower learning rate. Yet in asynchronous settings, it is not equivalent considering each block contains model updates from different devices. There are also other ways to penalize staled local updates but EWMA introduces modification with the least complexity.

Block structure. Fig. 3 shows the structure of a block. The principle of block structure design is that all parameters that affect global model aggregation should be recorded in the block, besides standard blockchain required parameters. For example,

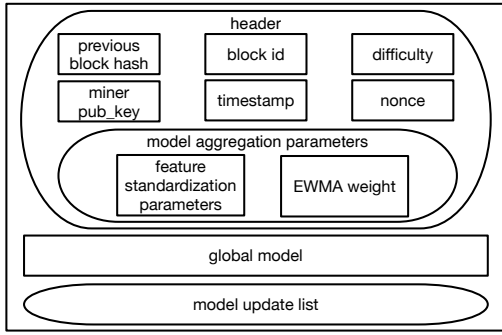


Fig. 3. Block structure of Fiesta.

we include the parameters for feature standardization. Local training related parameters are not included in the blockchain, e.g. optimizer type and the learning rate.

Mining difficulty and block generation speed. In a blockchain, blocks are chained by containing the hash of the previous block and it requires the hash of each block to contain at least d leading zeros, which is the mining difficulty (Difficulty here is slight different with the definition of those cryptos). The mining difficulty also decides the block generation speed accordingly. It balances a tradeoff between the first confirmation time (the time taken to generate the next valid block) and the amount of work wasted due to blockchain forking. A lower mining difficulty means a higher block generation speed, which leads to a higher possibility of blockchain forking and more work wastage but also a shorter time for first confirmation.

Block verification. Before adding the new block received from a neighbor to the blockchain and forwarding it to other neighbors, a miner needs to verify its validity. In addition to checking the PoW validity and block generation time and size, which are similar to bitcoin [18], the followings also need to be verified: a) compatible parameters compared with the previous block, e.g. difficulties and feature standardization parameters. b) the new global model is correctly computed based on the local updates within that block. c) there are no other updates from the same mobile device within a block.

C. Reward mechanisms

Rewards within a certain time period are calculated by the administrator based on the information recorded in the blockchain. Here we only propose a quantification policy. The exact reward could be monetary or using some other credit form. We distinguish two types of rewards, one for individuals and one for organizations.

Individual reward: Previous work on generalized blockchain and federated learning [8] utilizes a reward mechanism that is proportional only to the number of data records each device contributes. However, this approach is not suitable for spectrum sensing. If the reward is purely based on the number of data records, this device will tend to gain more reward without providing much insight into how this channel is utilized across the area, which is undesirable. Thus, we decide that the reward of an update for each mobile device/individual should be proportional to its contribution

in loss function minimization. Suppose a model update from client i is trained on local data set D_i with $|D_i| = n_i$, and it uses $block_t$'s global model \mathbf{w}_t^{global} as the initial parameters, and outputs \mathbf{w}_{t+1}^i . Denote the average loss on each data record as $l(\mathbf{w}; D_i)$. The individual reward is calculated as follows:

$$r_{t+1}^i = n_i \cdot \left[l(\mathbf{w}_t^{global}; D_i) - l(\mathbf{w}_{t+1}^i; D_i) \right]. \quad (3)$$

Since we assume the mobile clients report these values with the model update to the miners faithfully (this could be enforced by on device trusted execution environment, i.e. TEE), the individual rewards can be calculated reliably at seed node(s). In order to gain more rewards with the same n_i , one can enlarge the difference between $l(\mathbf{w}_t^{global}; D_i)$ and $l(\mathbf{w}_{t+1}^i; D_i)$ by measuring the locations that are not captured by the existing global model rather than visiting the locations that are already measured.

Organization reward: It is desirable that the reward to each organization is proportional to the contribution in the computing power that aggregates updates from mobile devices and generates the new global model. Therefore, the reward to each miner/organization is proportional to the number of blocks generated in the blockchain within the time period, similar to bitcoin.

D. Motivation of integrating BlockChain

There are actually abundant solutions to decentralizing a system and enhance its robustness, for example, Paxos [11], Raft [19], and PBFT [4]. But we still choose blockchain [18] (proof-of-work, to be exact) as the consensus algorithm for the following reasons:

Scalability. Conventional consensus algorithm usually has one or several strong leader nodes, either through designation or election. It is not a problem when there are only a few nodes. For example, Google's distributed lock service and file system, Chubby [3], is based on Paxos and consists of only five replicas. However, in the context of crowdsourcing, with more and more nodes joining into the community, communication and coordination cost will inflate dramatically and soon become the bottleneck of those leader nodes' performance [9]. PoW does not suffer from scalability problem that much since all that member node needs to preform is checking the hash value of each block it receives and follows the longest chain in the community [17].

Security. Node failure is not the only threat for robustness. Sometimes malicious nodes deteriorate the system in a more severe way, especially in the context of crowdsourcing where the community is open to everyone. One of the most common attack could be Sybil attack [6] where malicious miner injects overwhelming numbers of forged local updates to gain extra rewards for this dummy contribution. On one hand this could be partially prevented through a stricter local updates verification. On the other hand, the nature of PoW ensures that no matter how many malicious nodes one has instantiated (cloud virtual machine, for example, is a cheap way to do so) or how many dummy updates one has injected, it is the computation

power that has the impact on final consensus. This mechanism lower the volume of the attacker dramatically and increases their cost of collaring illegal rewards.

Alternative consensus mechanism for blockchain. Another well-known substitution to PoW is Proof-of-Stake (PoS) [7], where it is the earned token instead of the computation power that supports the volume of each node. But PoS suffers from the nothing-at-stake problem severely [12] and has not been widely tested. Hence we simply pick PoW in our implementation at current version.

III. IMPLEMENTATION & DEPLOYMENT

We implement the three roles of Fiesta’s layered architecture using different hardware and software stack.

Mobile device. A mobile device is connected with an external sensor board through a USB cable to gain spectrum sensing ability, which is shown in Fig. 4(a). The sensor board we use is from [14]. It is equipped with an AD9361 RF transceiver and a ZYNQ-7020 FPGA chip, featuring strong spectrum sensing capability ranging from 70MHz to 6GHz with a 56MHz sampling rate and high programmability. Users can check the status and interact through an Android app, as shown in Fig. 4(a). The app will enter the **sensing mode** when a sensor board is attached, and will start the **machine learning mode** when the mobile device is in charging and under stable WiFi access. This dual mode design (or lazy training design) dramatically reduces the impact of our spectrum sensing task onto the daily usage of host device. In terms of the software stack, `usb-serial-for-android` and `deeplearning4j` are integrated into the Android app to facilitate serial communication with the sensor board and on-device machine learning.

Miner. Miners are realized using Flask, a lightweight web application framework based on Python, and all the communications between mobile devices and miners use HTTP protocol for simplicity. An ordinary blockchain node using Proof-of-Work (PoW) consensus algorithm written in Python consists of the main body of each miner. Having received several local weight updates from mobile devices, the miner will integrate them into a new block and compute the new global model, for the next async round of training. For our experiments and deployment, several miner servers are deployed on different host machines including a 16-core 16 GB RAM personal desktop and an 8-core 8 GB RAM lab workstation.

Seed node. Seed node is also a Flask server and it uses HTTP to communicate with other roles as well. The seed node is backed up with two tables, one noting down the members registered for Fiesta and the other recording the contribution and reward for each member. In our design, there are a limited number of seed nodes and each of them should be reserved for only the administrator. For our experiments and deployment, a seed node is running on a 16-core 16 GB RAM desktop with a public domain name to allow other roles to query the seed node and register themselves through the Internet.

A. Implementation Challenges

We resolve several issues during implementation to make Fiesta’s implementation efficient and practical to be deployed.

Local update suppression. When the model converges and the pattern of the local data is already learned, the mobile device should not upload the local model to save the network bandwidth of both the mobile device and the miners as well as the computing resources at the miners. Given the reward mechanism we propose, this can be implemented as the mobile device stops uploading the update if the individual reward calculated based on local training loss reduction is smaller than a minimal threshold, i.e. when the local training loss does not decrease significantly, or the reward is non-positive.

Oversized models. Model size can be large, which makes block size large. This leads to two drawbacks: a heavy burden on the network transmission during miner communication, and a burden on the computation workload to the hash function when performing PoW. To address the former, we only keep full information of a block in its born miner and another archive server (which could be seed node), and share a streamlined block keeping only necessary components. To address the latter, we perform a first stage hash over those local updates from each different mobile devices, and use the generated fixed-length hash string to further generate the hash of the whole block when finding the nonce during PoW.

Model structure evolution. The structure of the model may need to grow more complex to better fit the increasing number of data records. At this time, seed node(s) will send out a broadcast message with their signature, which then will be verified by miners and other seed node(s), such that the miners start to work on a new blockchain upon receipt of this message. The genesis block of the new blockchain contains the new model structure. In order to fully utilize the data already obtained, there could be a knowledge transfer process before announcing the new structure, which means the global model in the genesis block of the new blockchain may not be randomly initialized.

B. Deployment and Applications

We deploy 10 sensor boards with mobile devices in a mid-sized US city for 3 months. 3 sensor boards are static (S1-S3), and the rest are deployed on personal vehicles of volunteers (M1-M7). Fig. 4(b) shows an antenna mounted on a vehicle. Due to the frequency constraint of the antennas, we configure the sensor boards to capture Power Spectrum Density (PSD) readings of 256 bins from 500MHz to 800MHz on a 50MHz basis, among which 500-700MHz is primarily for TV broadcast and 700-800MHz is mainly for LTE usage. PSD data is collected at the rate of one data record per minute for each 50MHz band. Fig. 4(c) shows the number of data records per sensor. Overall, we have gathered more than 220K data records of 600 hours with significant heterogeneity among different mobile devices. Fig. 4(d) illustrates the locations of measurements. It shows the static nodes and the GPS trace of one mobile node, only for illustration.

We have deployed two different machine learning applications in real world evaluation to demonstrate the functionality of our design. Structure of learning models are illustrated in Fig. 4(e), where digits on each block refers to how many

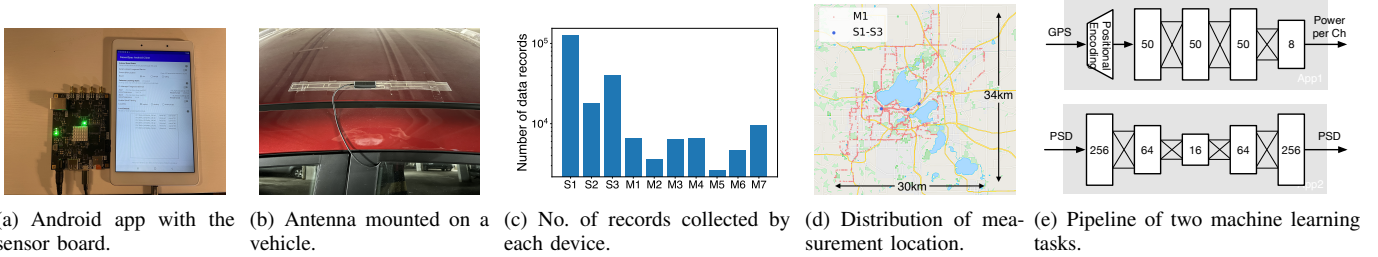


Fig. 4. HW/SW suite, real world deployment, and demo tasks of Fiesta.

neurons in a network layer and crossbar between blocks indicates that it is a fully connected layer.

App1: This application is to learn how the TV spectrum is used in the spatial domain through a neural network model predicting the signal power of TV channels at different locations. The model is a feed-forward neural network with latitude and longitude as inputs and signal power of TV channels as output, with mean square error as the loss function. In the actual implementation, a sin-cos positional encoding is adopted, inspired by NeRF [16]. In the future, environment information including weather condition could be fed into the model as an additional input to enhance model accuracy.

App2: This application is learning how the LTE band is used in the frequency domain by training an autoencoder [2] on the PSD data. The autoencoder is basically a neural network with symmetric structure design and hourglass shape. It takes the PSD data as input and the output of the autoencoder tries to reconstruct the input with high fidelity. The reconstruction result is expected to function as a precursor for detecting frequency domain anomalies, similar to [21].

IV. EVALUATION AND RESULTS

We first benchmark Fiesta’s performance by simulations using available spectrum datasets from [24], which contains TV band PSD data with spatial variations in a city. We also demonstrate Fiesta’s robustness through node failure and Sybil attack simulation on the miner side. Simulation framework and analysis could be found in our code repository under the directory `Fiesta/EdgeSim`. In this section we mainly report the evaluation results of our real-world deployment in a mid-size US city gathered by 10 mobile devices for 3 months.

A. Real-world evaluation results

We deploy Fiesta onto Android devices with sensor boards, along with two seed node and three miner nodes for each App. These devices present a wide variety in terms of their models (tablets and phones) and mobility status (stationary or portable), which cover the majority of real-world scenarios.

1) *Overhead on mobile devices:* Machine learning on mobile devices may be of low efficiency and time-consuming. Hence, profiling of its overhead is necessary.

CPU and memory. Fiesta Android client occupies less than 200MB RAM, and 50% CPU utilization on Google Pixel 2 when training two models. On the other hand, in sensing mode instead of machine learning mode, the CPU utilization

is negligible considering the only workload is receiving data from the sensing board. Considering Pixel 2 is a device released 7 years ago, powerful tensor accelerators like Google G2 in today’s smart phones have the potential to bear the training task and further reduce the CPU workload.

Network traffic. In order to fetch the global model from a miner and upload the trained local model to a miner, data packets sized 220KB for **App1** and 1.5MB for **App2** will be exchanged during one local round of training.

Training speed. A local training round (contains 10 epochs) with a training set with 4k data records takes 1.3 minutes for **App1** and 2 minutes for **App2** on Samsung Galaxy Tablet A 8.0, and slightly faster on Google Pixel 2. The difference between applications comes from the variance of their neural network structures. We also observe that one round of training time is proportional to the size of the local dataset and the number of epochs.

Energy consumption. Since local training is triggered only when the device gets plugged with a charger, energy consumption in machine learning mode is no big issue. For sensing mode, where the device is linked with a sensor board, tablets can last for 4 hours and only 2 hours for Google Pixel 2, because not only the Android device needs to power the board through OTG, GPS service also drains significant energy. In this case, battery life can be extended via external powering for the sensor board and using coarse network location to replace GPS with a compromise of accuracy.

2) *Convergence and learned models of the two applications:* We track the convergence during the whole federated learning procedure. Unlike an off-the-shelf training set with a fixed validation set in our microbenchmark simulation, for the real-world deployment, it is hard to evaluate the trained model directly, considering the training set itself is gathered in a streaming manner. As a compromise, we define *AppLoss* using the weighted average of training loss (MSE) among local updates within $block_t$.

$$AppLoss = \frac{1}{\sum_{i \in block_t} n_i} \sum_{i \in block_t} n_i \cdot l(w_i^i; D_i). \quad (4)$$

After calculation of the *AppLoss* for each block, “convergence curves” as functions of date for both apps are shown in Fig. 8. The convergence curves only visualize a subset of data points for clarity.

Before week 6, there are only 5 sensor nodes deployed, including 3 static and 2 mobile. We require each block




Fig. 5. App1: Fiesta distributedly learnt TV Ch. 46 power heatmap.




Fig. 6. App1: centralized learnt TV Ch. 46 power heatmap.




Fig. 7. App1: ground truth of TV Ch. 46 power heatmap.




Fig. 8. Model convergence over time.




Fig. 9. App2: example of reconstructed PSD data.




Fig. 10. App2: example outlier of LTE band.




Fig. 11. Compression and distribution of block size.




Fig. 12. Comparison of different reward policies.

contains 4~10 local updates. Reduction of both apps' losses is observed.

Starting from week 6, five additional sensor nodes are deployed and the losses got further reduced. With more devices joining, the asynchronous nature of this bigger system emerges and leads to inconsistent compositions between blocks, reflected as noticeable fluctuations afterward.

A distinct leap occurs for **App1** loss after week 8 because we online upgraded the machine learning model through transfer learning and the output dimension was increased from single one to eight causing a higher MSE loss. Detail is discussed in § IV-A3. Finally, the losses of **App1** and **App2** converge to 4dB and 5dB respectively.

We are also interested in what those models have learned. Fig. 5 visualizes the estimated power of TV channel 46 (662-668 MHz). The figure is generated through sampling **App1** model with a resolution of 100×100 within a specific latitude and longitude range. To intuitively demonstrate that the model has learnt the geographical feature of spectrum usage successfully, Fig. 7 visualize the raw data (available in kaggle) we collected from user devices directly. From either heat map and combining our prior knowledge with the city's building distribution, we notice that area with a higher density of high-rise buildings tends to have lower TV power and vice versa.

We also reconstruct the heat map through traditional centralized learning instead of federated learning way in Fig. 6, based on all of the gathered data. A general similarity among Fig. 5, Fig. 6, and Fig. 7 demonstrates the feasibility of using distributed federated learning in replacement of centralized learning or even raw data, while achieving good fidelity. A mismatch of regional details is also acceptable considering the raw data itself only covers main streets in the city (Fig. 4(d)) instead of every square feet. As a reference, test loss for **App1** in the centralized setup converges to 2.5 dB.

Fig. 9 shows the reconstruction result of a sample without

obvious signals using the autoencoder trained for **App2**. The non-flat noise floor of the sensor including a spike around 720MHz is successfully revealed after reconstruction. Fig. 10 shows a signal not fully reconstructed by the model, which might be a frequency domain outlier and could be used for further anomaly detection. Other spike due to non-flat noise floor of the sensor is still successfully reconstructed.

3) *Model structure evolution*: On week 9, we online upgrade the structure of the model for **App1** to output multiple TV channels' power instead of only one. We achieve knowledge transfer by duplicating the original single output neuron to initiate other output neurons corresponding to the multiple TV channels, and keep the hidden layers' parameter unchanged in the meantime. The MSE loss increases due to a larger output dimension. Furthermore, estimating multiple TV channels' power is intrinsically more challenging compared with a single one. Fig. 8 shows the ease of use of Fiesta when one needs to modify the machine learning model structure.

4) *Effectiveness of the mechanism to deal with oversized models*: Machine learning models with thousands or even millions of parameters will be a huge burden for the network and local storage. Fig. 11 shows our block compression policy with respect to the number of local updates within a block. Considering we replace the giant Python dictionary of each local weight with a much smaller unit (§ III-A), the more local updates a block contains, the more storage we could save.

We also show the CDF of the number of local updates within a block. We notice that due to the relatively small mining difficulty, most blocks contain only 4 local updates, the minimal number we set. The max number of local updates a block includes is 8. **App1** blocks contain more local updates than **App2** blocks. It is reasonable considering the machine learning model of **App2** contains much more parameters than that of **App1**, hence requiring a longer time for local training. More swiftly the local training completes, more local updates

can be included in one block given the mining difficulty.

5) *Reward mechanisms*: As mentioned in § II-C, the reward mechanism previous work [8] adopts fails to reflect the actual contribution of each member. To demonstrate the effectiveness of our proposed reward mechanism shown in formula 3, we compare it with the naive size-based reward mechanism as the baseline, in Fig. 12, using 4 representative devices, including two stationary and two mobile. All the reward values are normalized by the maximum.

Stationary device 1 performs local training on a large but roughly geographically homogeneous dataset considering it is deployed in a stationary environment. On the contrary, mobile device 1 and 2 have relatively smaller but more heterogeneous data because they are deployed on moving vehicles. From Fig. 12, we observe that the baseline mechanism grants higher rewards of the stationary device regardless to the little information within, while the proposed mechanism values the data diversity significantly. The difference is more distinct when it comes to **App1**, which tries to learn the spatial distribution of TV channels' power. Stationary device 2, though having more data points than mobile devices, but performs few local training and uploading, hence little reward earned.

V. CONCLUSION

A major pitfall of previous spectrum crowdsensing efforts is the lack of awareness of sustainability. We have presented Fiesta, a sustainable framework for spectrum crowdsensing based on a three-tier architecture. On the technical side, Fiesta addresses two deficiencies of prior spectrum sensing architectures: privacy concerns that require participants to upload sensitive information and a single point of failure due to centralized analysis, through a combination of federated learning and blockchain. On the policy side, we have proposed a fair reward quantification mechanism for participants and discussed potential funding sources to ensure long-term sustainability. We also release the code and dataset used in this paper to welcome further contributions and wider adoption from the community to realize Fiesta's long-term vision.

ACKNOWLEDGEMENT

We thank Hongyi Wang for discussions about federated learning in the early stage of the project. All UW-Madison authors were supported in part by the US National Science Foundation awards 2003129, 2112562, 2107060, 2212688, and through an award from the US Department of Commerce with award number 70NANB21H043. The work at IMDEA has been funded in part by the project MAP-6G, reference TSI-063000-2021-63, granted by the Ministry of Economic Affairs and Digital Transformation and the European Union-NextGenerationEU through the UNICO-5G R&D program of the Spanish Recovery, Transformation and Resilience Plan.

REFERENCES

- [1] Microsoft spectrum observatory. spectrum-observatory.cloudapp.net.
- [2] J. An and S. Cho. Variational autoencoder based anomaly detection using reconstruction probability. *Special Lecture on IE*, 2(1):1–18, 2015.
- [3] M. Burrows. The chubby lock service for loosely-coupled distributed systems. In *Usenix OSDI*, pages 335–350, 2006.
- [4] M. Castro, B. Liskov, et al. Practical byzantine fault tolerance. In *Usenix OSDI*, volume 99, pages 173–186, 1999.
- [5] E. Chai, K. Sundaresan, M. A. Khojastepour, and S. Rangarajan. Lte in unlicensed spectrum: Are we there yet? In *Proceedings of the 22nd Annual International Conference on Mobile Computing and Networking*, pages 135–148, 2016.
- [6] J. R. Douceur. The sybil attack. In *International workshop on peer-to-peer systems*, pages 251–260. Springer, 2002.
- [7] A. Kiayias, A. Russell, B. David, and R. Oliynykov. Ouroboros: A provably secure proof-of-stake blockchain protocol. In *Annual international cryptology conference*, pages 357–388. Springer, 2017.
- [8] H. Kim, J. Park, M. Bennis, and S.-L. Kim. Blockchain on-device federated learning. *IEEE Communications Letters*, 24(6):1279–1283, 2019.
- [9] M. Kogias and E. Bugnion. Hovercraft: Achieving scalability and fault-tolerance for microsecond-scale datacenter services. In *Proceedings of the Fifteenth European Conference on Computer Systems*, pages 1–17, 2020.
- [10] J. Konečný, B. McMahan, and D. Ramage. Federated optimization: Distributed optimization beyond the datacenter. In *NIPS Optimization for Machine Learning Workshop*, 2015.
- [11] L. Lamport. The part-time parliament. In *Concurrency: the Works of Leslie Lamport*, pages 277–317. 2019.
- [12] C. Lepore, M. Ceria, et al. A survey on blockchain consensus with a performance comparison of pow, pos and pure pos. *Mathematics*, 8(10):1782, 2020.
- [13] X. Li, K. Huang, W. Yang, S. Wang, and Z. Zhang. On the convergence of fedavg on non-iid data. In *ICLR*, 2020.
- [14] Y. Li, Y. Zeng, and S. Banerjee. Enabling wideband, mobile spectrum sensing through onboard heterogeneous computing. In *Proceedings of the 22nd International Workshop on Mobile Computing Systems and Applications*, pages 85–91, 2021.
- [15] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial Intelligence and Statistics*, pages 1273–1282. PMLR, 2017.
- [16] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021.
- [17] D. Mingxiao, R. Xiao Feng, Z. Zhe, W. Xiangwei, and C. Qijun. A review on consensus algorithm of blockchain. In *2017 IEEE international conference on systems, man, and cybernetics (SMC)*, pages 2567–2572. IEEE, 2017.
- [18] S. Nakamoto. Bitcoin whitepaper. URL: <https://bitcoin.org/bitcoin.pdf>, 2008.
- [19] D. Ongaro and J. Ousterhout. In search of an understandable consensus algorithm. In *Usenix ATC*, pages 305–319, 2014.
- [20] S. Rajendran, R. Calvo-Palomino, M. Fuchs, B. Van den Bergh, H. Corobés, D. Giustiniano, S. Pollin, and V. Lenders. Electrosense: Open and big spectrum data. *IEEE Communications Magazine*, 56(1):210–217, 2017.
- [21] S. Rajendran, W. Meert, V. Lenders, and S. Pollin. Saife: Unsupervised wireless spectrum anomaly detection with interpretable features. In *2018 IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN)*, pages 1–9, 2018.
- [22] S. Sarkar, M. Buddhikot, A. Baset, and S. K. Kasera. Deepadar: a deep-learning-based environmental sensing capability sensor design for cbrs. In *Proceedings of the 27th Annual International Conference on Mobile Computing and Networking*, pages 56–68, 2021.
- [23] Y. Zeng, R. Calvo-Palomino, D. Giustiniano, G. Bovet, S. Banerjee, et al. Adaptive uplink data compression in spectrum crowdsensing systems. In *IEEE International Symposium on Dynamic Spectrum Access Networks (DySPAN)*, 2021.
- [24] Y. Zeng, V. Chandrasekaran, S. Banerjee, and D. Giustiniano. A framework for analyzing spectrum characteristics in large spatio-temporal scales. In *The 25th Annual International Conference on Mobile Computing and Networking*, pages 1–16, 2019.
- [25] T. Zhang, N. Leng, and S. Banerjee. A vehicle-based measurement framework for enhancing whitespace spectrum databases. In *Proceedings of the 20th annual international conference on Mobile computing and networking*, pages 17–28. ACM, 2014.
- [26] M. Z. Zheleva, R. Chandra, A. Chowdhery, P. Garnett, A. Gupta, A. Kapoor, and M. Valerio. Enabling a nationwide radio frequency inventory using the spectrum observatory. *IEEE Transactions on Mobile Computing*, 17(2):362–375, 2017.