

# FedZero: Leveraging Renewable Excess Energy in Federated Learning

Philipp Wiesner  
wiesner@tu-berlin.de  
TU Berlin

Pratik Agrawal  
pratik.agrawal@tu-berlin.de  
TU Berlin

Ramin Khalili  
ramin.khalili@huawei.com  
Huawei

Lauritz Thamsen  
lauritz.thamsen@glasgow.ac.uk  
University of Glasgow

Dennis Grinwald  
dennis.grinwald@tu-berlin.de  
TU Berlin

Odej Kao  
odej.kao@tu-berlin.de  
TU Berlin

## ABSTRACT

Federated Learning (FL) is an emerging machine learning technique that enables distributed model training across data silos or edge devices without data sharing. Yet, FL inevitably introduces inefficiencies compared to centralized model training, which will further increase the already high energy usage and associated carbon emissions of machine learning in the future. One idea to reduce FL’s carbon footprint is to schedule training jobs based on the availability of renewable excess energy that can occur at certain times and places in the grid. However, in the presence of such volatile and unreliable resources, existing FL schedulers cannot always ensure fast, efficient, and fair training.

We propose FedZero, an FL system that operates exclusively on renewable excess energy and spare capacity of compute infrastructure to effectively reduce a training’s operational carbon emissions to zero. Using energy and load forecasts, FedZero leverages the spatio-temporal availability of excess resources by selecting clients for fast convergence and fair participation. Our evaluation, based on real solar and load traces, shows that FedZero converges significantly faster than existing approaches under the mentioned constraints while consuming less energy. Furthermore, it is robust to forecasting errors and scalable to tens of thousands of clients.

## CCS CONCEPTS

• **Social and professional topics** → Sustainability; • **Computing methodologies** → Distributed artificial intelligence.

## KEYWORDS

sustainable computing, carbon efficiency, electricity curtailment, federated learning, client selection, green AI

### ACM Reference Format:

Philipp Wiesner, Ramin Khalili, Dennis Grinwald, Pratik Agrawal, Lauritz Thamsen, and Odej Kao. 2024. FedZero: Leveraging Renewable Excess Energy in Federated Learning. In *The 15th ACM International Conference on Future and Sustainable Energy Systems (E-Energy '24)*, June 4–7, 2024, Singapore, Singapore. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3632775.3639589>

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

*E-Energy '24*, June 4–7, 2024, Singapore, Singapore

© 2024 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0480-2/24/06

<https://doi.org/10.1145/3632775.3639589>




Figure 1: Quarterly wind and solar curtailments by the California ISO [13]. Leveraging this renewable excess energy in FL can drastically reduce its operational carbon emissions.

## 1 INTRODUCTION

The majority of today’s machine learning (ML) solutions perform centralized learning, where all required training data are gathered in a single location, usually an energy-efficient data center with specialized hardware. Yet, in many practical use cases, it is not feasible to collect data across a distributed system due to security and privacy concerns or because large amounts of raw data cannot be migrated from the deep edge to the cloud. Federated Learning (FL) was introduced to address this issue by enabling distributed training of ML models without transmitting training data over the network [37]. In FL, we train a common ML model on clients that cannot or do not want to share their data, by iteratively distributing the model to a subset of them. Clients then train locally on their own data and send back the updated models to the server, which aggregates them before starting the next round.

Unfortunately, FL approaches require considerably more training rounds than traditional ML and are often executed on infrastructure that is less energy-efficient than centralized GPU clusters, resulting in a significant increase in overall energy usage and associated emissions [39, 43, 63]. Even without the application of FL, the training of large ML models is known to be an energy-hungry process and has increasingly raised concerns in recent years [17, 52, 61]. As models keep growing in size and complexity, this problem is expected to aggravate, which is why there are numerous efforts towards more energy-efficient algorithms and hardware to reduce the carbon footprint of AI. Yet, when focusing on reducing emissions, “using renewable energy grids for training neural networks is the single biggest change that can be made” [16, 42].

In this work, we study how the operational carbon emissions of synchronous FL training can be reduced to zero by operating under the hard constraint of only leveraging renewable excess energy

and spare computing capacity at cloud or edge resources. Excess energy, also called stranded energy, occurs in electric grids when more power is generated than demanded or when the grid does not have sufficient capacity for transmission. If the oversupply cannot be stored in batteries (which are expensive and only available in a limited capacity) or traded with neighboring grids (whose excess energy patterns often correlate) the last resort is curtailment, the deliberate reduction in production. Through curtailment, the California Independent System Operator wasted more than 27 million megawatt-hours of utility-scale solar energy in 2022, which is around 7% of their entire solar production [13]. Due to the increasing penetration of variable renewable energy sources, the amount of curtailed energy is only expected to grow, as shown in Figure 1. At the same time, many existing computing infrastructures are frequently underutilized or could be overclocked if the occurrence of excess energy justifies reduced energy efficiency [14]. To make better use of these resources, *carbon-aware computing*, i.e. considering the spatio-temporal availability of low-carbon energy during scheduling, has attracted much attention in recent years [6, 19, 35, 44, 59, 60, 66].

FL is a promising workload for carbon-aware computing, as it consists of energy-intensive batch jobs that are scheduled in geo-distributed environments (to leverage spatial resource and energy availability) without strict runtime requirements (to also leverage temporal variations). However, as excess energy and the availability of spare computing resources can be highly volatile, not explicitly taking them into account during client selection can lead to significantly longer training times due to *stragglers*: clients that perform less local training than expected, or become entirely unavailable during a training round. Furthermore, energy-agnostic selection strategies can introduce biases by disproportionately selecting clients that have a lot of excess resources available throughout the training. Yet, the idea of aligning FL scheduling with the availability of renewable energy has so far only been studied theoretically and under assumptions like independent and identically distributed (iid) data and fixed “energy arrival” patterns that are not realistic in practice and do not consider the above challenges [21].

To fill this gap, we propose *FedZero*, an FL system for heterogeneous and geo-distributed environments that utilizes forecasts for renewable excess energy and spare computing capacity to ensure fast, efficient, and fair training under energy and resource constraints. We summarize our **contributions** as follows:

- We propose a system design for executing FL trainings exclusively on renewable excess energy and spare computing capacity which allows clients to share common energy budgets at runtime.
- We introduce a scalable client selection strategy that results in fast convergence and fair client participation under variable energy and resource constraints.
- We evaluated our approach on different datasets, models, and scenarios, to show that FedZero enables fast and energy-efficient training while being robust to forecast errors.
- We implemented FedZero and all baselines using Flower [8] and Vessim [58] and made this code openly available<sup>1</sup>.

<sup>1</sup><https://github.com/dos-group/fedzero>

## 2 A CASE FOR FL ON EXCESS RESOURCES

FL was originally developed for use cases on mobile and edge devices, where individual clients usually consume little energy. However, in recent years a variety of new application domains have been explored to enable cross-device and cross-silo training, many of which include clients with significant computing capabilities and electricity demand. Examples of these novel FL settings include healthcare [47], the financial sector [62], remote sensing [49], autonomous driving [40], and smart cities [25], which all consist of complex models that require periodic re-training to adapt to changing environments.

We argue that in environments where individual clients require significant energy for participating in an FL training, it merits to explicitly consider the availability of renewable excess energy in client selection and during training to reduce carbon emissions.

### 2.1 Renewable Excess Energy

Due to the expanding deployment of variable renewable energy sources such as solar and wind, it is becoming increasingly challenging to match power supply and demand at all times. If locally occurring renewable excess energy cannot be passed on to neighboring grids due to limited grid capacity and cannot be buffered in some kind of energy storage, the only option left to operators is to throttle supply. In this section, we describe the two main scenarios in which renewable excess energy can occur.

The most direct way of operating IT infrastructure in a sustainable manner is through the use of on-site renewable energy, where the energy source is located close to the datacenter or powerful edge device. Within a microgrid, energy storage can buffer limited amounts of excess energy but it is expensive, entails losses, and frequent charge cycles accelerate battery aging [36]. Moreover, while more and more countries offer the possibility to sell energy to the public grid, feed-in tariffs are usually well below purchase prices [46]. Therefore, operators have a clear incentive to consume all generated electricity directly.

The more common practice to achieve datacenters “powered by 100% renewable energy”, as claimed by big cloud service providers [3, 20, 38], is through carbon accounting. Carbon accounting allows operators to offset their consumption of carbon-intensive grid energy through the purchase of renewable energy certificates. However, today’s certificates allow that power production and consumption can take place at vastly different times and locations, which is why their utility for achieving science-based targets is often questioned [9]. A prominent effort towards stricter carbon accounting, often called *24/7 matching*, are Google’s Time-based Energy Attribute Certificates (T-EACs) [54] that are issued hourly and location-specific [6]. Similar ideas have been brought forward by Microsoft [41] and Amazon [27]. During curtailment periods, these certificates are expected to be very cheap and can be used for the cost-effective execution of flexible workloads such as FL.

### 2.2 Computing on Spare Resources

As we do not want to promote the purchase of new hardware just to enable the flexible scheduling of training jobs, FedZero aims to schedule workloads exclusively on spare capacity of existing hardware. We believe this is realistic in a wide range of scenarios

as many IT infrastructures are over-dimensioned to accommodate peak loads. For example, public clouds, as well as emerging public edge solutions, must maintain a relatively high proportion of spare computing resources to sell their promise of infinite scalability. Similarly, on-site infrastructure is often designed for peak loads and can be severely underutilized outside these periods. A common measure in cluster managers to increase the utilization during off-peak hours is by deferring delay-tolerant workloads, for example in the form of *best-effort* jobs [55].

Lately, several ideas have been developed to not only shift load for increased resource utilization but for better aligning electricity usage with grid carbon-intensity or excess energy [35, 44, 59]. Moreover, even infrastructure that is already utilized close to capacity can be used to compute flexible workloads, if the use of otherwise curtailed excess energy justifies the reduced energy efficiency caused by overclocking and increased cooling [14].

### 3 PROBLEM STATEMENT

Our goal in this paper is to train a federated learning model with no operational carbon emissions in an efficient and fair manner. We aim to optimize for fast convergence and low overall energy usage in a setting, where clients are only allowed to train on excess resources. We do so by cherry-picking clients that are likely to have access to renewable excess energy and spare computing capacity (or potential for overclocking) and by operating within these constraints at runtime.


#### 3.1 Challenges

FL on excess resources poses a number of new challenges that have not yet been addressed by existing approaches.

**Convergence speed and efficiency.** The availability of excess resources can be highly variable. Thus, clients that have access to excess energy and spare computing resources during client selection can run out of resources over the duration of a training round. This leads to an increased number of stragglers that can severely harm training performance. A common way to alleviate the impact of stragglers in FL is to select more participants in each round than actually needed, but to only wait for a number of early responses before aggregating the results and starting a new round. The extent of over-selection can be adapted to the environment but is usually around 30% in the related work [10, 30, 31]. While this makes FL training more robust, it has the disadvantage of wasting computing capacity and energy, since the work of some clients is discarded in each round. Moreover, if multiple clients reside in the same power domain, over-selection can actively harm training progress as more clients share the same limited power source at runtime rather than only attributing power to the most useful clients.

**Common power budgets.** As clients can share the same source of excess energy, we need to treat energy as a shared and limited resource during client selection and at runtime. We use the term *power domain* to describe the clustering of FL clients into groups with access to the same source of renewable excess energy – either because they are physically connected within a datacenter’s microgrid, or because their operation is covered through a common budget of, for example, T-EACs (see Section 2.1). Power domains

are disjunct, meaning that one client can only be part of a single power domain. Excess energy occurring within a power domain must be shared by all clients within the domain. For example, in this work, we study the behavior of FL training under resource and energy constraints in two different solar-based scenarios: In the first scenario, clients are spread across ten globally distributed power domains (Figure 2a). In the second scenario, power domains are in close geographic proximity (Figure 2b).



(a) Distributed power domains. (b) Co-located power domains.

Figure 2: Excess power availability for different scenarios.

**Fairness of participation.** Simply optimizing for available resources in client selection will lead to a strong imbalance in favor of clients who have a lot of excess resources available throughout the training. In realistic settings, where the distribution of data varies between clients, this can lead to significant biases toward certain data in the training, which is unfair and ultimately harms model performance. Hence, even if the availability of excess resources is highly imbalanced, we want to ensure that all clients are able to participate in a similar number of rounds.

**Robustness and scalability.** To target the previous challenges, we try to spread the client selection over different power domains using forecasts of available excess energy and spare computing resources – while still considering system and statistical utility. However, forecasts usually come with a certain error, which is why we require a solution that is robust to inaccurate forecasts.

Moreover, we need to ensure that any underlying optimization comes with a low overhead and runtime complexity, as real-life FL scenarios can comprise large numbers of clients.

#### 3.2 Problem Formalization

We define  $C$  as the set of clients distributed over a disjunct set of power domains  $P$ . Clients are characterized by their energy efficiency  $\delta_c$  and maximum computing capacity  $m_c$ . For simplicity, we do not consider other resources like memory in this work, but they can be integrated the same way as  $m_c$ . We divide time into slots of duration  $t$ , so an estimated training round duration  $d$  is always a multiple of  $t$ . The duration of  $t$  depends on the problem setting but is usually in the order of one minute. We define the training of a mini-batch, from now on called batch, as an atomic operation to be performed within these time slots. Table 1 provides an overview of all introduced variables and constants.

As common for FL in heterogenous environments [10, 34], we allow clients to train a variable amount of batches, but require the configuration of a lower ( $m_c^{min}$ ) and upper ( $m_c^{max}$ ) bound per client (for example, 1 to 5 local epochs). Furthermore, the server should define the number of selected clients per round  $n$  as well as a maximum round duration  $d^{max}$  after which results get aggregated, even if not all clients did respond in time. We allow multiple clients

**Table 1: Overview of constants and variables**

System-related constants	
$C$	set of clients
$P$	set of power domains
$C_p$	set of clients in power domain $p$
$m_c$	maximum capacity of client $c$ (batches/timestep)
$\delta_c$	energy efficiency of client $c$ (energy/batch)
User-defined constants	
$n$	number of selected clients per round
$d^{\max}$	maximum round duration in multiples of $t$
$m_c^{\min}; m_c^{\max}$	minimum/maximum number of batches client $c$ must participate per round
Input variables (updated each round)	
$m_{c,t}^{\text{spare}} \in [0, m_c]$	spare capacity forecast for client $c$ at time $t$
$r_{p,t}$	excess energy forecast for power domain $p$ at time $t$
$\sigma_c$	fairness weighting of client $c$
Optimization variables (determined each round)	
$d$	expected round duration
$b_c \in \{0, 1\}$	whether or not client $c$ is selected
$m_{c,t}^{\text{exp}} \in [0, m_{c,t}^{\text{spare}}]$	expected number of batches client $c$ will compute at time $t$ considering energy and capacity constraints

to share a common excess energy budget by clustering them into power domains.  $C_p$  describes the clients of a power domain  $p \in P$ . Each power domain comprises a control plane, like an ecovisor [51], which is responsible for attributing power to clients. Section 4.3 describes our client selection algorithm and optimization problem.

### 3.3 Boundaries

In this work, we do not explicitly consider energy storage or feeding excess energy to the public grid, since these options are not always available and have drawbacks compared to consuming excess energy directly [60]. Moreover, as we target larger-scale infrastructure that is usually connected to the network via highly energy-efficient optical fiber, we do not model the energy usage for data transmission. Lastly, we require the availability of excess resources during training. Environments where relevant clients never have access to renewable excess energy or spare computing capacity need to default to a less radical approach and consider carbon-intensive grid energy consumption at times.

## 4 SYSTEM DESIGN

An overview of FedZero’s protocol is depicted in Figure 3. The training starts after a required amount of clients register themselves with the server (Section 4.1, ①). At the beginning of each round, the server requests forecasts on expected excess energy within power domains and spare capacity at clients (Section 4.2, ②). FedZero then selects  $n$  clients for training, for which it expects the shortest round duration under the given resource constraints (Section 4.3, ③). This selection is performed based on the forecasts and information on past participation or statistical utility of clients for ensuring

performance and fairness (Section 4.4). Next, the selected clients train locally on spare capacity and, via continuous exchange with the power domain controller, excess energy (Section 4.5, ④). Finally, all participating clients send their updated model back to the server which aggregates them and documents the participated batches and local loss for future decisions (⑤).

### 4.1 Client Registration

Before starting the training process, FedZero requires the following information for each client:

- (1) The *maximum computing capacity* of a client is denoted as  $m_c$  (batches/timestep) and can be derived from its FLOPS (floating point operations per second), the model’s MACs (multiply-accumulate operations), and the batch size. Alternatively, it can also be benchmarked before or during the training. For variable capacity datacenters [14],  $m_c$  should describe the actual maximum with overclocking.
- (2) The *energy efficiency* is denoted as  $\delta_c$  (energy/batch) and can be obtained through measurements or derived from the client’s system performance and power consumption characteristics. Linear power modeling is a meaningful simplification if we can assume power-proportional clients or sequential processing of workloads. If not,  $\delta_c$  can also change from round to round depending on the system utilization, which is especially relevant in variable-capacity datacenters [14].
- (3) The *control plane addresses* define a client’s power domain, as well as where to query load and excess energy forecasts from. Load forecasts can be provided by the client itself or its cluster manager/container orchestrator. Typical providers for energy forecasts are electricity providers, microgrid control systems, or ecovisors [51].

### 4.2 Forecasting Excess Energy and Load

To avoid picking clients with access to little or no resources during a round, FedZero relies on multistep-ahead forecasts of excess energy at power domains and spare capacity at clients.

Power production forecasts for variable renewable energy sources like solar [11, 28] and wind [2, 32] are usually based on weather models for mid- and long-term predictions as well as, in case of solar, satellite data for short-term predictions that enable up to 5-minute resolution. For on-site installations, there exist a large number of companies providing power production forecasts as a service. In the case of time-based power purchase agreements, it is the responsibility of the utility provider to inform their customers of future energy budgets. For determining future excess energy, the system furthermore needs to take load forecasts for IT infrastructure as well as co-located consumers into account. We define  $r_{p,t}$  to be the forecasted excess energy of power domain  $p$  at time  $t$ .

Load prediction is a widely researched field covering forecasts related to application metrics, such as requests per second, as well as the utilization of (virtualized) hardware resources like CPU, GPU, or RAM. They usually entail time series forecasting models trained on historical data but can also take additional context information into account. As  $m_c$  describes the maximum computing capacity of a client in batches/timestep, we define  $m_{c,t}^{\text{spare}} \in [0, m_c]$  to be its forecasted spare capacity at time  $t$ .

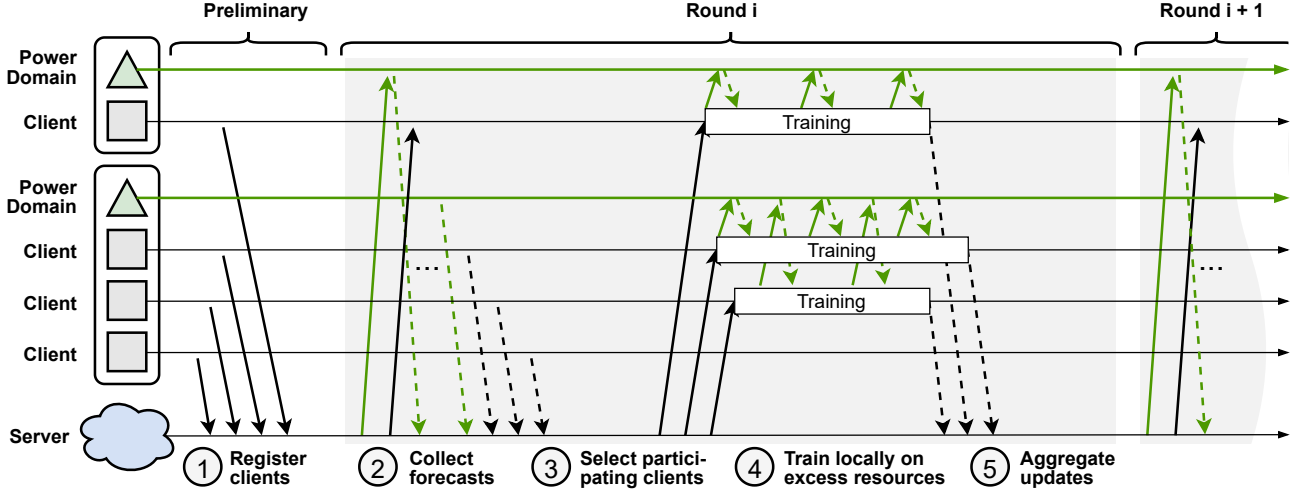


Figure 3: At each training round, FedZero queries excess energy forecasts of power domains and load forecasts of individual clients. Based on this information, it selects a set of clients for which it expects a short round duration at high statistical utility. At runtime, clients have to periodically adjust their training performance to align with the actual available excess energy.

### 4.3 Client Selection

FedZero selects clients based on their forecasted energy and capacity constraints as well as statistical utility.

**Iterative search.** FedZero optimizes for system utility by selecting clients that are expected to compute their  $m_c^{min}$  as fast as possible. We guarantee low computational overhead, by performing an iterative search over possible round duration  $d$ : For each round duration, we solve a simple mixed-integer program (MIP) which scales linearly with the number of clients and power domains

---

#### Algorithm 1 Determine clients and round duration

---

```

1:  $C \leftarrow$  set of clients
2:  $P \leftarrow$  set of power domains
3: # search for shortest possible round duration
4: for  $d \leftarrow 1$  to  $d^{max}$  do
5:   # filter out power domains without excess energy
6:    $\bar{P} \leftarrow \{\forall p \in P, \forall t = 1, \dots, d : r_{p,t} > 0\}$ 
7:   # filter out clients that over-participated in the past (see Section 4.4)
8:    $\bar{C} \leftarrow \{\forall c \in C : \sigma_c > 0\}$ 
9:   # filter out clients without sufficient computing capacity or energy
10:  for  $p \in \bar{P}$  do
11:     $\bar{C} \leftarrow \bar{C} \setminus \{\forall c \in C_p : \sum_{t=0}^d \min(m_{c,t}^{spare}, \frac{r_{p,t}}{\delta_c}) < m_c^{min}\}$ 
12:  # increase duration if there are not at least  $n$  valid clients
13:  if  $|\bar{C}| < n$  then
14:    continue
15:  # select optimal clients
16:   $b \leftarrow \text{findOptimalClients}(\bar{C}, \bar{P}, d)$ 
17:  if  $b$  is valid solution then
18:    return  $b, d$ 
19: # wait, if no solution is found for  $d = d^{max}$ 

```

---

(see Section 5.5). For simplicity, the iterative search is described as an incrementing for-loop in Algorithm 1. In practice, it can be implemented as a binary search with complexity  $O(\log n)$ .

On every iteration, Algorithm 1 heavily pre-filters entire power domains (Line 6) and individual clients (Lines 11 and 8) that cannot constitute valid solutions within the current  $d$ , to further reduce the runtime of the MIP. If no valid solution is found within the maximum round duration  $d^{max}$ , the algorithm waits for conditions to improve or it could resolve the situation by weakening constraints, e.g. by considering grid energy. In this work, we only operate under hard energy and capacity constraints.

**Optimization Problem.** For the MIP, we define two discrete optimization variables per eligible client:

- $b_c \in \{0, 1\}$  equals 1 iff client  $c$  participates in the round
- $m_{c,t}^{exp} \in [0, m_{c,t}^{spare}]$  denotes the expected number of batches client  $c$  will compute at time  $t$

The optimization problem is described as follows:

$$\max_{b_c, m_{c,t}^{exp}} \sum_{c \in C} b_c \cdot \sigma_c \sum_{t=0}^d m_{c,t}^{exp} \quad (1)$$

$$\text{s.t. } b_c = 1 \implies m_c^{min} \leq \sum_{t=0}^d m_{c,t}^{exp} \leq m_c^{max} \quad \forall c \in C \quad (1)$$

$$\sum_{c \in C_e} m_{c,t}^{exp} \cdot \delta_c \leq r_{e,t} \quad \forall e \in E, t = 0, \dots, d \quad (2)$$

$$\sum_{c \in C} b_c = n \quad (3)$$

We optimize for the maximum number of batches to be computed within the input duration  $d$ , weighted by each client's statistical utility  $\sigma_c$ . Equation (1) limits each selected client to compute between  $m_c^{min}$  and  $m_c^{max}$  batches. Equation (2) constrains all clients in a

power domain to not use more energy than available. Equation (3) ensures that exactly  $n$  clients are selected per round.

**Statistical utility.** We introduce a utility function  $f : C \rightarrow \{\sigma_c : \forall c \in C\}$  which is invoked in every round and returns a weighting that gives precedence to certain clients in the optimization problem. This function can be based on the previous participation of clients, an approximation of statistical client utility, or other user-defined metrics, for example, to respect fairness constraints like group parity.

The utility function applied in the remainder of this paper is based on the statistical utility function proposed in Oort [30]:

$$\sigma_c = \begin{cases} |B_c| \sqrt{\frac{1}{|B_c|} \sum_{k \in B_c} \text{loss}(k)^2}, & \text{if } p(c) \geq 1 \\ 1, & \text{otherwise} \end{cases}$$

Oort approximates statistical client utility based on the number of available training samples  $B_c$  and the local training loss, which is expected to correlate with the gradient norm.

#### 4.4 Ensuring Fair Participation

When performing FL in heterogenous environments under excess energy and capacity constraints, we actively have to take care of avoiding biases towards powerful clients with lots of spare capacity ( $m_c$ ) or clients within power domains with large amounts of excess energy ( $r_e$ ). This problem is exacerbated by FedZero, which prefers energy-efficient clients ( $\delta_c$ ) and – without further measures – tends to select similar sets of clients in consecutive rounds, which can harm the model’s generalization performance.

To mitigate the mentioned biases and reduce variance, we add clients to a blacklist after they participate in a training round. Blocked clients get assigned  $\sigma_c = 0$  and are hence excluded from future rounds. At the start of each round, clients can get released from the blacklist with probability  $P(c)$ :

$$P(c) = \begin{cases} (p(c) - \omega)^{-\alpha}, & \text{if } p(c) - \omega > 0 \\ 1, & \text{otherwise} \end{cases}$$

where  $p(c)$  describes the number of rounds a client previously participated and  $\alpha$  is a user-defined parameter that controls the speed at which clients get released. A high  $\alpha$  will cause overparticipating clients to remain longer on the blacklist, thereby reducing the set of clients that FedZero can pick from. This can extend training time but ensures fair participation. An  $\alpha$  close to 0 reduces the impact of the blacklist. We consider  $\alpha = 1$  for the remainder of this paper, which turned out to provide the best balance between training speed and performance in all evaluated experiments.

The parameter  $\omega$  avoids decreasing release probabilities over time and gets periodically updated to  $\omega = \text{mean}\{p(c) : \forall c \in C\}$ . Users can choose a different  $P(c)$  for their use case, for example, to improve group fairness or other custom metrics.

#### 4.5 Executing Training Rounds

This section describes the local control loop executed by clients during training rounds, see ④ in Figure 3. Using the actually available resources at runtime, each client tries to compute  $m_c^{min}$  batches as fast as possible. Upon completion, it notifies the server but continues computation until  $m_c^{max}$  is reached. The server signals the

end of a training round and gathers all updated models once all clients computed  $m_c^{min}$ , or once  $d^{max}$  has passed. If a client does not manage to compute at least  $m_c^{min}$  batches before  $d^{max}$ , its work is discarded to not impede the training progress, as commonly performed in the literature [10].

Below, we discuss the two main challenges of the local control loop: First, if multiple clients from the same power domain are participating in the same round, they have to share a common energy budget at runtime. Second, the actual available excess energy and spare capacity are subject to short-term fluctuations and usually differ from previously performed forecasts.

**Sharing power at runtime.** If only one client within a power domain is participating, it can make use of all available excess energy at runtime. In this case, the capacity available for training is defined as the minimum of the free capacity and the capacity that can be powered using excess energy [60].

However, if two or more clients of a power domain participate simultaneously and there is not enough energy for all of them, they have to share a common energy budget at runtime, which has to be coordinated by the power domain controller. To determine each client’s share of power, we propose a simple two-step approach: First, power is attributed to clients that have not yet reached their minimum round participation  $m_c^{min}$ , weighted by how much energy is still required to reach the threshold. If  $m_c^{comp}$  describes the number of batches a client has already computed in the active round, this can be written as:  $C_p \rightarrow \{\min(0, \delta_c m_c^{min} - \delta_c m_c^{comp}) : \forall c \in C_p\}$ . Second, if there is still power left, it is attributed to all clients below their maximum participation  $m_c^{max}$ , again weighted by the energy required to reach this limit:  $C_p \rightarrow \{\min(0, \delta_c m_c^{max} - \delta_c m_c^{comp}) : \forall c \in C_p\}$ . As clients oblige capacity constraints and may not be able to make use of their entire share of power, the actual distribution of power must be decided in constant consultation with clients.

**Short-term variations.** As FedZero aims to not interfere with other processes running on a client, the extent of local training must be adapted over time. For example, to determine excess energy, clients must periodically query their power domain controller, e.g. ecovisor [51] and use this information to perform power capping of tasks or entire containers [18, 33]. Recently, works have been proposed that do this in a more sophisticated manner than simply throttling the training process’ access to resources. For example, DISTREAL [45] handles the time-varying availability of resources in FL by dynamically adjusting the computational complexity of the trained neural network. This approach could be extended to also consider power consumption, which is ultimately composed of the utilization of computing resources. However, runtime behavior in the presence of short-term resource and energy fluctuations is currently not considered in our prototype.

## 5 EVALUATION

We implemented FedZero and all baselines using the FL framework Flower [8] and the energy system simulator Vessim [58]. We extended Flower to enable discrete-event simulation over time series datasets like excess energy availability and client load. It enables us to perform experiments faster than in real-time by, for example, skipping over time windows where the system is idle and waiting

for excess energy or spare capacity on clients. We use Gurobi<sup>2</sup> to solve the MIP. Our implementation and all datasets are openly available (see Section 1).

## 5.1 Experimental Setup

To evaluate FedZero, we simulate the power usage characteristics and performance of 100 FL clients using our Flower extension and perform the training on four NVIDIA V100 and two RTX 5000 GPUs. This allows us to evaluate our approach without training models over multiple weeks and consuming megawatt hours of energy.

**Table 2: Max energy consumption and training performance of the three types of clients.**

client type	max energy	performance (samples per minute)			
		DenseNet-121	EfficientNet-B1	LSTM	KWT-1
<b>small</b>	70 W	110	118	276	87
<b>mid</b>	300 W	384	411	956	303
<b>large</b>	700 W	742	795	1856	586

**Clients.** We model heterogeneity among clients by randomly assigning them to one of three types (*small, medium, large*) that are roughly based on the performance<sup>3</sup> and energy usage characteristics of T4, V100, and A100 GPUs, respectively. However, we downscaled their actual compute capabilities (samples per minute), as shown in Table 2. We use 100 randomly selected machines from the Alibaba GPU cluster trace dataset [57] to model client load (gpu\_wrk\_util) and load forecasts (gpu\_plan).

**Scenarios.** For modeling power domains, we focus on on-site solar energy generation in two scenarios based on real solar and solar forecast data provided by Solcast<sup>4</sup>: A *global scenario* (ten globally distributed cities from June 8-15, 2022) and a *co-located scenario* (ten largest cities in Germany from July 15-22, 2022), both displayed in Figure 2. The solar data is available in 5-minute resolution and we assume a constant power supply for steps within this period. Clients are randomly distributed over the ten power domains, which each have a maximum output of 800 W. If there is little sun, or multiple clients are selected within a domain, energy becomes a limiting resource. The power and client availability is depicted in Figure 4. The upper plot shows the energy availability within energy domains, where each domain is represented by one line. The lower plot depicts the availability of clients over time, color-coded by how much of their total computational capacity is available for training.

**Datasets, models, parameters.** We evaluate our approach on four datasets and models commonly used in FL evaluations.


- *CIFAR-100* [29] contains 60,000 32x32 color images across 100 classes. We model heterogeneous data by applying a Dirichlet distribution with  $\alpha = 0.5$ , similar to [22], which skews the number of samples as well as the number of samples per class and client. We train<sup>5</sup> the convolutional model DenseNet-121 [23] using FedProx [34] with  $\mu = 0.1$ .

<sup>2</sup><https://www.gurobi.com>

<sup>3</sup><https://developer.nvidia.com/deep-learning-performance-training-inference>

<sup>4</sup><https://solcast.com>

<sup>5</sup>SDG optimizer, learning rate = 0.001, weight decay = 5e-4, momentum = 0.8



**Figure 4: Power production and client availability over the course of both scenarios. While there are always some clients available in the global scenario, in the co-located scenario clients are always available around the same time.**

- *Tiny ImageNet* contains 100,000 64×64 color images of 200 classes. We distribute samples to clients using the same Dirichlet distribution as for CIFAR-100. We train<sup>6</sup> an EfficientNet-B1 [53] model using FedProx with  $\mu = 0.1$ .
- In the *Shakespeare* [12] dataset, each client represents one of 100 randomly selected speaking roles from a play. As in [34], we train a two-layer LSTM<sup>7</sup> using FedProx with  $\mu = 0.001$  to perform next character prediction.
- *Google Speech Commands* contains more than 100,000 audio samples of 30 different words. We randomly assigned speakers to the 100 clients and train<sup>8</sup> the keyword transformer model KWT-1 [7] for speech classification.

The number of samples computed per timestep was obtained through benchmarking runs and is stated in Table 2. All simulations use a timestep  $t = 1$  min and a max round duration  $d^{\max} = 60$  min. We select  $n = 10$  clients each round which have to compute 1 to 5 local epochs, so  $m^{\min}$  and  $m^{\max}$  depend on the locally available number of samples. Clients locally train on minibatches of size 10. We ran each experiment five times over the course of seven simulated days and report mean values.

**Baselines.** We compare FedZero with existing approaches by training six different baselines. First, we run all experiments using *Random* client selection as well as the guided selection strategy *Oort* [30]. We update each client’s system utility, an important factor in Oort’s scheduling, based on the available energy and capacity in every round. Both approaches can only select from clients, which currently have access to excess energy and spare resources.

Second, we train the above baselines again but this time allow them to select  $1.3n$  clients per round. Over-selection is commonly employed in the related work [10, 30] to counteract inefficiencies caused by stragglers in unreliable environments. Once  $n$  clients have returned their results a new round starts. The baselines are called *Random 1.3n* and *Oort 1.3n*.

Third, we want to demonstrate that access to forecasts alone is not the decisive advantage over existing approaches. For this, we train two baselines *Random fc* and *Oort fc* that only select 10 clients,

<sup>6</sup>Adam optimizer, learning rate = 0.001

<sup>7</sup>100 hidden units, 8D embedding layer, SDG optimizer, learning rate = 0.8, see [34]

<sup>8</sup>AdamW optimizer, learning rate = 0.001, weight decay = 0.1, see [7]

**Table 3: Best accuracy and time/energy to reach the target accuracy of FedZero and the best-performing baselines.**

Dataset & model	Approach	Global scenario				Co-located scenario			
		Target accuracy	Best accuracy	Time-to-accuracy	Energy-to-accuracy	Target accuracy	Best accuracy	Time-to-accuracy	Energy-to-accuracy
CIFAR-100 DenseNet-121	Random 1.3K		66.0 %	4.7 d	79.2 kWh		66.6 %	5.3 d	113.4 kWh
	Oort 1.3K		66.4 %	4.5 d	103.8 kWh		66.4 %	5.4 d	138.7 kWh
	Oort fc	64.7 %	65.8 %	5.3 d	102.4 kWh	65.5 %	66.1 %	6.4 d	126.7 kWh
	<b>FedZero</b>		<b>66.8 %</b>	<b>3.6 d</b>	<b>70.6 kWh</b>		<b>66.5 %</b>	<b>4.5 d</b>	<b>96.4 kWh</b>
Tiny ImageNet EfficientNet-B1	Random 1.3K		63.1 %	5.6 d	109.6 kWh		63.3 %	3.7 d	86.0 kWh
	Oort 1.3K		63.2 %	3.3 d	90.2 kWh		63.5 %	3.4 d	90.5 kWh
	Oort fc	62.4 %	63.1 %	3.9 d	89.0 kWh	62.8 %	62.7 %	-	-
	<b>FedZero</b>		<b>63.6 %</b>	<b>2.9 d</b>	<b>67.1 kWh</b>		<b>63.6 %</b>	<b>3.4 d</b>	<b>75.8 kWh</b>
Shakespeare LSTM	Random 1.3K		50.7 %	4.6 d	97.9 kWh		51.5 %	4.5 d	90.0 kWh
	Oort 1.3K		50.2 %	-	-		51.7 %	4.5 d	95.4 kWh
	Oort fc	50.4 %	50.5 %	6.7 d	157.4 kWh	50.9 %	50.5 %	-	-
	<b>FedZero</b>		<b>53.1 %</b>	<b>1.8 d</b>	<b>40.0 kWh</b>		<b>53.1 %</b>	<b>2.3 d</b>	<b>42.8 kWh</b>
Google Speech KWT-1	Random 1.3K		85.2 %	4.8 d	103.5 kWh		85.1 %	4.3 d	80.8 kWh
	Oort 1.3K		86.9 %	3.6 d	99.0 kWh		86.4 %	3.4 d	85.0 kWh
	Oort fc	83.6 %	87.0 %	3.7 d	86.2 kWh	82.8 %	84.9 %	3.7 d	76.6 kWh
	<b>FedZero</b>		<b>87.2 %</b>	<b>3.6 d</b>	<b>79.0 kWh</b>		<b>87.7 %</b>	<b>2.6 d</b>	<b>65.8 kWh</b>

but have access to load and energy forecasts for filtering out clients that are not expected to reach their  $m_c^{\min}$  within  $d^{\max}$ .

Lastly, for each experiment, we define an *Upper bound* in convergence speed and performance, by training a model that uses random client selection but is not subject to any energy constraints or existing load on clients (clients are still heterogeneous). This baseline is not limited to renewable excess energy.

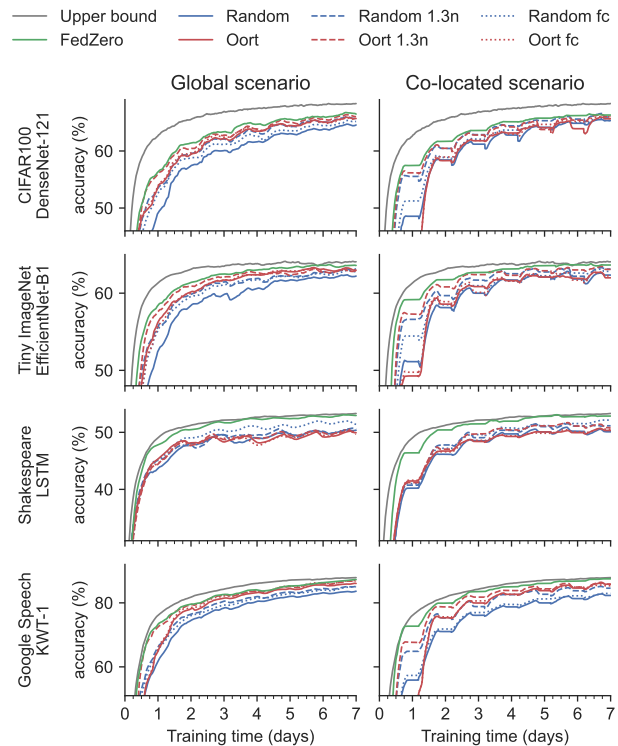
## 5.2 Performance Overview

**Training performance.** Figure 5 displays the training progress of FedZero and the baselines over the different experiments. We can observe that FedZero consistently outperforms all baselines in terms of top accuracy. While in some cases the performance of *Oort/Oort 1.3n/Oort fc* is comparable (see Appendix A for details), the gap is considerable in scenarios with heavy sample imbalance like Shakespeare (2365 ± 4674 samples per client; min = 730; max = 27950). This is due to the fact that none of the baselines considers common power budgets during client selection. We found that this problem is exacerbated for strategies that select clients based on statistical utility, such as Oort: If a power domain does not have access to excess energy for an extended period of time, the statistical utility of its clients is usually high as they have not participated for many rounds. Once excess energy is available, Oort heavily targets these clients which leads to increased competition for energy at runtime and, therefore, slower training progress.

For all datasets but CIFAR-100, FedZero reaches, or almost reaches, the top accuracy of the *Upper bound*, suggesting that varying resource availability does not necessarily harm training performance, but only increases the total training time.

**Time-to-accuracy and energy-to-accuracy.** To further demonstrate how FedZero improves FL under energy and capacity constraints, we define the top accuracy of the *Random* baseline as our

target accuracy for a specific experiment. Table 3 reports the time and energy required for FedZero and the baselines to reach this target accuracy. The table only contains the best-performing baselines; a full table with all results can be found in Appendix A.

**Figure 5: Training progress of all experiments.**



FedZero has the lowest time-to-accuracy and energy-to-accuracy across all experiments. On average, it reached the target accuracy around 35 % faster in the global scenario and around 26 % faster in the co-located scenario than *Random 1.3n* and *Oort 1.3n*, which were among the fastest and most energy-efficient baselines. At the same time, FedZero was using 36 % less energy on average in the global scenario and 30 % less in the co-located scenario. Oort-based baselines generally outperform the Random-based ones in terms of top accuracy and convergence speed but at the cost of higher energy usage. However, some Oort-based baselines do not reach the target accuracy at all, due to the previously described inefficiencies caused by selecting clients from the same power domain.

**Round durations.** As FedZero knows about the system utility and resource availability, it avoids combining clients with vastly different expected round durations. For example, in the global scenario on CIFAR-100 it required  $15.1 \pm 8.5$  min per round. For comparison, the *Random* baseline had an average round duration of  $33.7 \pm 19.6$  min, which was lowered to  $22.7 \pm 17.7$  min and  $27.8 \pm 17.3$  min by *Random 1.3n* and *Random fc*, respectively. The round duration of Oort-based baselines was  $18.6 \pm 14.5$  min on average.

The same applies to the co-located scenario, where all clients are available around the same time. Here, Random-based baselines take  $15.5 \pm 12.7$  min and Oort-based baselines  $12.0 \pm 13.0$  min on average. FedZero only requires  $9.7 \pm 7.6$  min, allowing it to perform considerably more training rounds within the same time. This observation is consistent across all experiments.

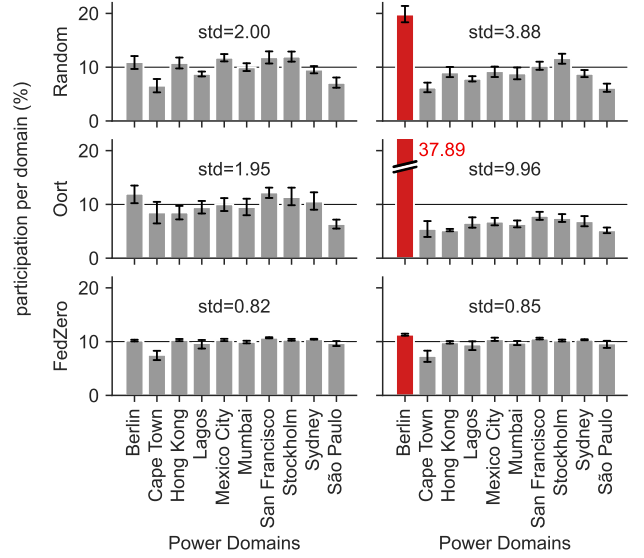
### 5.3 Fairness of Participation

When training under energy and capacity constraints, we inevitably introduce biases towards clients that have lots of excess resources available. To illustrate this, Figure 6a displays the average percentage of rounds in which clients have participated in the training for the CIFAR-100 global scenario, grouped by power domain. As we select 10 out of 100 clients per round, we ideally expect an average client participation of 10 %. However, some power domains have access to more excess energy than others, resulting in the *Random* and *Oort* strategies to favor clients in these domains. We can observe that FedZero exhibits a much more balanced participation within (marked by the error bar) as well as between power domains (*std* as stated on each figure).

We conducted an additional set of experiments on the same scenario, where the Berlin power domain has access to unlimited excess energy and all clients within Berlin have unlimited computing resources available. The results of this experiment are displayed in Figure 6b, where Berlin is colored in red. As clients in this power domain are now always available for training, the *Random* baseline almost doubles their participation from  $11.0 \pm 2.1$  to  $19.8 \pm 2.6$  %.

**Table 4: CIFAR-100 performance on the global scenario under imbalanced conditions (Berlin has unlimited resources).**

	Best accuracy	Time-to-acc.	Energy-to-acc.
Random	64.6 %	6.7 d	95.7 kWh
Oort	65.6 %	4.5 d	189.4 kWh
<b>FedZero</b>	<b>66.9 %</b>	<b>3.5 d</b>	<b>83.4 kWh</b>



**(a) Client participation per power domain. (b) Client participation with unlimited resources for Berlin.**

**Figure 6: FedZero ensures fair participation of clients, even under highly imbalanced conditions.**

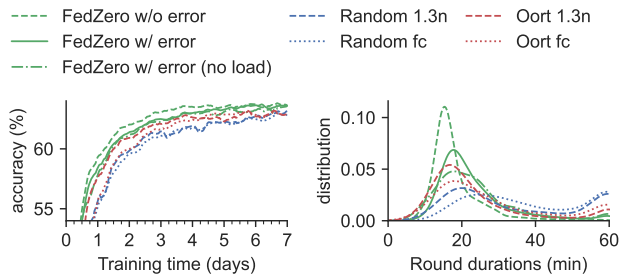
Even worse, Oort, which actively targets clients with high system utility, more than triples the participation of Berlin clients from  $12.0 \pm 2.9$  to  $37.9 \pm 1.3$  %. Oort describes a mechanism for combining its selection with a user-defined fairness metric. However, we found that Oort must rely almost entirely on our fairness metric, hence, disregard system and statistical utility, when attempting to achieve fairness comparable to FedZero. Other than all baselines, which introduce significant biases, FedZero only slightly leverages the additional resources by increasing the mean participation of clients in the domain from  $10.2 \pm 0.3$  to  $11.3 \pm 0.3$  %.

Table 4 displays the training performance of the three approaches in the scenario where Berlin has unlimited resources. We observe that *Random* used 19 % more energy and *Oort* even twice the energy in the imbalanced scenario (Figure 6b) for reaching a comparable accuracy as in the base scenario (Figure 6a). FedZero used only 4 % more energy and reduced its time-to-accuracy.

Based on our results on fairness of participation, we expect FedZero to also improve other fairness metrics such as accuracy parity between clients. However, these metrics depend on various additional factors like non-iid data distribution among power domains, which is why we leave this analysis to future work.

### 5.4 Robustness Against Forecasting Errors

To investigate the impact of forecast quality on FedZero’s performance, we performed further experiments based on the global scenario on the Tiny ImageNet and Google Speech datasets. Figure 7 shows the training progress and distribution of round durations for Tiny ImageNet. *FedZero w/ error* uses forecasts with realistic errors as in all previous experiments, *FedZero w/o error* uses perfect forecasts, and *FedZero w/ error (no load)* uses realistic errors for



**Figure 7: Analysis of convergence behavior and round durations of FedZero under forecasts of different quality.**

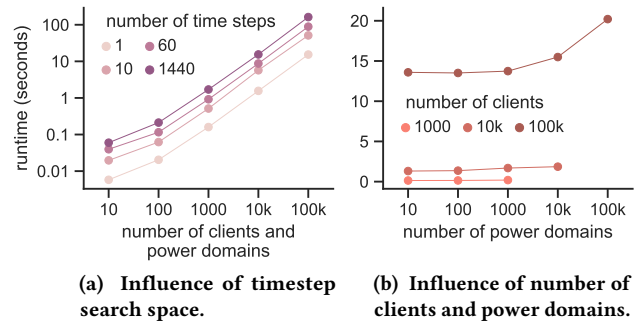
excess energy but has no forecasts for spare capacity available, as short term load might not always be predictable in every setting. Note, that FedZero is not able to operate if there are no predictions of excess energy at all, for example, due to communication loss.

The three experiments based on FedZero show small differences in convergence speed and energy usage. While *FedZero w/ error* takes 2.8 d and 65.2 kWh to reach the target accuracy, using perfect forecasts it requires 15.4 % less time and 15.2 % less energy. This is because FedZero becomes better at avoiding stragglers through the use of accurate predictions, resulting in shorter, hence, more efficient rounds, as shown in the right of Fig 7. FedZero without load forecasts takes 8.2 % more time to reach the target accuracy, using 10.0 % more energy. This result is of course specific to how we modeled load in our evaluation – the effect in other contexts may be bigger or smaller. Still, we can see that all three experiments converge to the same accuracy of 63.8%, while consistently exhibiting better time-to-accuracy and energy-to-accuracy, showing that FedZero can perform well even with suboptimal forecast quality.

We additionally performed this analysis on the Google Speech experiment and got comparable results. For example, FedZero without errors converged 5.2 % faster using 6.7 % less energy than FedZero with realistic forecasts.

## 5.5 Overhead and Scalability

We analyze the overhead of FedZero’s client selection by profiling the runtime of Algorithm 1, including the MIP, on an Apple M1 processor. Each experiment was repeated 5 times and we report mean values. Figure 8a shows the linear growth of runtime in regards to the number of clients: Even at the biggest evaluated setting, 100k clients distributed over 100k power domains searching over 1440 timesteps (24 hours in 1-minute resolution), the algorithm returns within two minutes. For scenarios in the scale of the previous evaluation (100 clients, 10 power domains, 60 timesteps) it only takes around 0.1 seconds to decide on a set of clients. We observe that due to the  $O(\log n)$  runtime of the binary search, increasing the timestep search space from 60 to 1440 (factor 24) only increases the runtime by factor 1.8. Figure 8b shows the runtime of a single MIP for different numbers of clients and power domains (note, that the y-axis is linear in this figure). We observe, that the number of power domains has little to no impact on the runtime for up to 10k domains. Increasing the number of power domains from 10k to 100k only increases the runtime from 15.4 to 20.1 seconds.



**Figure 8: FedZero overhead analysis.**

In terms of communication overhead, FedZero requires excess energy and load forecasts at the beginning of each round. Furthermore, clients periodically (in our evaluation minutely) sync with their power domain to align their performance with the actual available excess energy at training time. As both payloads are in the order of kilobytes and we are assuming clients are connected via fiber, this overhead is negligible.

## 6 RELATED WORK

**Carbon-aware computing.** As carbon pricing mechanisms, such as emission trading systems or carbon taxes, are starting to be implemented around the globe [5], the IT industry is pushing to increase the usage of low-carbon energy in datacenters. Carbon-aware computing tries to reduce the emissions associated with computing by shifting flexible workloads towards times [19, 44, 59] and locations [66, 67] with clean energy. For example, Google defers delay-tolerant workloads when power is associated with high carbon intensity as a measure to reach their 24/7 carbon-free target by 2030 [44].

While most research in carbon-aware computing aims at consuming cleaner energy from the public grid [35, 44, 59, 67], recent works also try to better exploit excess energy, similar to FedZero. Cucumber [60] is an admission control policy that accepts low-priority workloads on underutilized infrastructure, only if they can be computed using excess energy. Similarly, Zheng et al. [66] explore workload migration on underutilized data centers as a measure to reduce curtailment. The Zero-Carbon Cloud [15] already targets the problem of curtailment at the level of infrastructure planning, by placing data centers close to renewable energy sources.

**Carbon footprint of ML.** The training of large ML models is a highly relevant domain for carbon-awareness, due to the often excessive energy requirements on the one hand, and certain flexibility in scheduling on the other. Other than inference, which is usually expected to happen at low latency, ML training jobs can often be stopped and resumed, scaled up or down, or even migrated between locations. Because of this, many papers have previously addressed the carbon emissions of centralized ML [16, 17, 42, 52].

Qiu et al. [43] were the first to broadly study the energy consumption and carbon footprint of FL and state that, "depending on the configuration, FL can emit up to two orders of magnitude more carbon than centralized machine learning." Further studies

investigate the carbon impact of hyperparameters such as concurrency rate [64] or the cost of differential privacy [48]. Carbon awareness in the context of FL has so far only been explored by Güler and Yener [21] who define a model for intermittent energy arrivals and propose a scheduler with provable convergence guarantees. However, their assumptions regarding energy arrivals are highly simplified as they neither consider spare capacity on clients nor non-iid data distributions. Moreover, other than FedZero, their model does not allow multiple clients to share the same power domain.

**FL client selection.** Active (or guided) client selection in FL has received significant attention in recent years, as researchers try to improve the final accuracy, convergence speed, reliability, fairness, or reduce communication overhead compared to random client selection [50]. For example, Oort [30] exploits heterogeneous device capabilities and data characteristics by cherry-picking clients with high statistical model efficiency as well as high system utility, resulting in faster convergence and better final accuracy. Similarly, other novel approaches like FedMarl [65] or Power-of-Choice [24] utilize the local training loss of clients to bias the selection. FedZero does not try to compete with but can be adapted to integrate with other client selection strategies.

Few energy-aware client selection strategies exist, like EAFL [4], that extend Oort's utility function to additionally consider the battery level of clients. However, like most research addressing energy usage in FL [56, 63], it aims to increase the operating time of battery-constrained end devices and is not concerned with the overall emissions associated with the training. FedZero describes the first approach for FL training solely on excess energy and spare computing capacity.

## 7 CONCLUSION

This paper proposes FedZero, a system design for fast, efficient, and fair training of FL models using only renewable excess energy and spare computing capacity. Our results show, that FedZero's client selection strategy converges significantly faster than all baselines under the mentioned resource constraints while ensuring fair participation of clients, even under highly imbalanced conditions. Moreover, our approach is robust against forecasting errors and scalable to large-scale, globally distributed scenarios.

In future work, we want to investigate the impact of periodic patterns in excess energy availability on training performance [68]. Furthermore, we plan to integrate FedZero with novel asynchronous [26] or semi-synchronous [1] strategies, while explicitly taking energy storage and grid carbon intensity into account.

## ACKNOWLEDGMENTS

We sincerely thank Solcast for providing us with free access to their solar forecast APIs. We furthermore want to thank the anonymous reviewers of ICDCS '23 and e-Energy '24 for their helpful comments. This research was supported by the German Academic Exchange Service (DAAD) as ide3a and IFI as well as the German Ministry for Education and Research (BMBF) as BIFOLD (grant 01IS18025A) and Software Campus (grant 01IS17050).

## REFERENCES

- [1] Ahmed M. Abdelmoniem, Atal Narayan Sahu, Marco Canini, and Suhaib A. Fahmy. 2023. REFL: Resource-Efficient Federated Learning. In *EuroSys*. ACM. <https://doi.org/10.1145/3552326.3567485>
- [2] David B. Alencar, Carolina de Mattos Affonso, Roberto C. L. Oliveira, Jorge Laureano Moya Rodriguez, Jandecy Cabral Leite, and Jose Carlos R. Filho. 2017. Different Models for Forecasting Wind Power Generation: Case Study. *Energies* 10 (2017). <https://doi.org/10.3390/en10121976>
- [3] Amazon. 2022. Amazon's 2022 Sustainability Report. (2022).
- [4] Amna Arouj and Ahmed M. Abdelmoniem. 2022. Towards Energy-Aware Federated Learning on Battery-Powered Clients. In *Workshop on Data Privacy and Federated Learning Technologies for Mobile Edge Network at ACM MobiCom*.
- [5] World Bank. 2022. *State and Trends of Carbon Pricing 2022*. Technical Report. Washington, DC: World Bank.
- [6] Noman Bashir, David Irwin, Prashant Shenoy, and Abel Souza. 2022. Sustainable Computing - Without the Hot Air. In *HotCarbon*.
- [7] Axel Berg, Mark O'Connor, and Miguel Tairum Cruz. 2021. Keyword Transformer: A Self-Attention Model for Keyword Spotting. In *Proc. Interspeech 2021*. 4249–4253. <https://doi.org/10.21437/Interspeech.2021-1286>
- [8] Daniel J Beutel, Taner Topal, Akhil Mathur, Xinchu Qiu, Titouan Parcollet, and Nicholas D Lane. 2020. Flower: A Friendly Federated Learning Research Framework. *arXiv preprint arXiv:2007.14390* (2020).
- [9] Anders Bjørn, Shannon M. Lloyd, Matthew Brander, and H. Damon Matthews. 2022. Renewable energy certificates threaten the integrity of corporate science-based targets. *Nature Climate Change* 12, 6 (2022).
- [10] Keith Bonawitz, Hubert Eichner, Wolfgang Grieskamp, Dzmitry Huba, Alex Ingemar, Vladimir Ivanov, Chloé Kiddon, Jakub Konečný, Stefano Mazzocchi, Brendan McMahan, Timon Van Oortveldt, David Petrou, Daniel Ramage, and Jason Roselander. 2019. Towards Federated Learning at Scale: System Design. In *MLSys*. <https://proceedings.mlsys.org/paper/2019/file/bd686fd640be98efaae0091fa301e613-Paper.pdf>
- [11] Jamie M. Bright, Sven Killinger, David Lingfors, and Nicholas A. Engerer. 2018. Improved satellite-derived PV power nowcasting using real-time power data from reference PV systems. *Solar Energy* 168 (2018). <https://doi.org/10.1016/j.solener.2017.10.091>
- [12] Sebastian Caldas, Sai Meher Karthik Duddu, Peter Wu, Tian Li, Jakub Konečný, H Brendan McMahan, Virginia Smith, and Ameet Talwalkar. 2019. LEAF: A Benchmark for Federated Settings. In *Workshop on Federated Learning for Data Privacy and Confidentiality at NeurIPS*.
- [13] California ISO. 2024. Managing oversupply. <http://www.caiso.com/informed/Pages/ManagingOversupply.aspx>. accessed Jan. 2024.
- [14] Andrew Chien, Chaojie Zhang, Liuzixuan Lin, and Varsha Rao. 2022. Beyond PUE: Flexible Datacenters Empowering the Cloud to Decarbonize. In *HotCarbon*.
- [15] Andrew A Chien, Chaojie Zhang, and Hai Duc Nguyen. 2019. Zero-carbon Cloud: Research Challenges for Datacenters as Supply-following Loads. *University of Chicago, Tech. Rep. CS-TR-2019-08* (2019).
- [16] Payal Dhar. 2020. The carbon impact of artificial intelligence. *Nature Machine Intelligence* 2 (2020), 423–425.
- [17] Jesse Dodge, Taylor Prewitt, Remi Tachet des Combes, Erika Odmark, Roy Schwartz, Emma Strubell, Alexandra Sasha Luccioni, Noah A. Smith, Nicole DeCario, and Will Buchanan. 2022. Measuring the Carbon Intensity of AI in Cloud Instances. In *ACM FAccT*. <https://doi.org/10.1145/3531146.3533234>
- [18] Jonatan Enes, Guillaume Fieni, Roberto R. Expósito, Romain Rouvoy, and Juan Touriño. 2020. Power Budgeting of Big Data Applications in Container-based Clusters. In *IEEE CLUSTER*.
- [19] Gilbert Fridgen, Marc-Fabian Körner, Steffen Walters, and Martin Weibelzahl. 2021. Not All Doom and Gloom: How Energy-Intensive and Temporally Flexible Data Center Applications May Actually Promote Renewable Energy Sources. *Business & Information Systems Engineering* 63, 3 (2021).
- [20] Google. 2022. 2022 Environmental Report. (2022).
- [21] Başak Güler and Aylin Yener. 2021. A Framework for Sustainable Federated Learning. In *2021 19th International Symposium on Modeling and Optimization in Mobile, Ad hoc, and Wireless Networks (WiOpt)*. <https://doi.org/10.23919/WiOpt52861.2021.9589930>
- [22] Harry Hsu, Hang Qi, and Matthew Brown. 2019. Measuring the Effects of Non-Identical Data Distribution for Federated Visual Classification. *arXiv preprint arXiv:1909.06335* (2019).
- [23] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q. Weinberger. 2017. Densely Connected Convolutional Networks. In *CVPR*.
- [24] Yae Jee Cho, Jianyu Wang, and Gauri Joshi. 2022. Towards Understanding Biased Client Selection in Federated Learning. In *AISTATS*.
- [25] Ji Chu Jiang, Burak Kantarci, Sema Oktug, and Tolga Soyata. 2020. Federated Learning in Smart City Sensing: Challenges and Opportunities. *Sensors* 20, 21 (2020).
- [26] Zhifeng Jiang, Wei Wang, Baochun Li, and Bo Li. 2022. Pisces: Efficient Federated Learning via Guided Asynchronous Training. In *ACM Symposium on Cloud Computing (SoCC)*. <https://doi.org/10.1145/3542929.3563463>

- [27] Lucas Joppa. 2021. Made to measure: Sustainability commitment progress and updates. Microsoft. Retrieved Sept. 2023 from <https://blogs.microsoft.com/blog/2021/07/14/made-to-measure-sustainability-commitment-progress-and-updates>
- [28] Alexandra I. Khalyasmaa, Stanislav A. Eroshenko, T. Chakravarthy, Venu Gopal Gasi, Sandeep Kumar Yadav Bollu, Raphael Caire, Sai Kumar Reddy Atluri, and Suresh Karrolla. 2019. Prediction of Solar Power Generation Based on Random Forest Regressor Model. In *IEEE SIBIRCON*. <https://doi.org/10.1109/SIBIRCON48586.2019.8958063>
- [29] Alex Krizhevsky. 2009. *Learning multiple layers of features from tiny images*. Technical Report.
- [30] Fan Lai, Xiangfeng Zhu, Harsha V. Madhyastha, and Mosharaf Chowdhury. 2021. Oort: Efficient Federated Learning via Guided Participant Selection. In *USENIX OSDI*. <https://www.usenix.org/conference/osdi21/presentation/lai>
- [31] Chenning Li, Xiao Zeng, Mi Zhang, and Zhichao Cao. 2022. PyramidFL: A Fine-Grained Client Selection Framework for Efficient Federated Learning. In *ACM MobiCom*. <https://doi.org/10.1145/3495243.3517017>
- [32] Qing'an Li, Chang Cai, Yasunari Kamada, Takao Maeda, Yuto Hiromori, Shuni Zhou, and Jianzhong Xu. 2021. Prediction of power generation of two 30 kW Horizontal Axis Wind Turbines with Gaussian model. *Energy* 231 (2021). <https://doi.org/10.1016/j.energy.2021.121075>
- [33] Shaohong Li, Xi Wang, Xiao Zhang, Vasileios Kontorinis, Sreekumar Kodakara, David Lo, and Parthasarathy Ranganathan. 2020. Thunderbolt: Throughput-Optimized, Quality-of-Service-Aware Power Capping at Scale. In *USENIX OSDI*. <https://www.usenix.org/conference/osdi20/presentation/li-shaohong>
- [34] Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. 2020. Federated Optimization in Heterogeneous Networks. In *MLSys*.
- [35] Liuzixuan Lin, Victor M. Zavala, and Andrew Chien. 2021. Evaluating Coupling Models for Cloud Datacenters and Power Grids. In *ACM e-Energy*. <https://doi.org/10.1145/3447555.3464868>
- [36] Longjun Liu, Hongbin Sun, Chao Li, Tao Li, Jingmin Xin, and Nanning Zheng. 2017. Managing Battery Aging for High Energy Availability in Green Datacenters. *IEEE Transactions on Parallel and Distributed Systems* 28, 12 (2017). <https://doi.org/10.1109/TPDS.2017.2712778>
- [37] H. B. McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. 2016. Communication-Efficient Learning of Deep Networks from Decentralized Data. In *AISTATS*.
- [38] Microsoft. 2022. 2022 Environmental Sustainability Report. (2022).
- [39] Rakshit Naidu, Harshita Diddee, Ajinkya K Mulay, Aleti Vardhan, Krithika Ramesh, and Ahmed Zamzam. 2021. Towards Quantifying the Carbon Emissions of Differentially Private Machine Learning. In *Workshop on Socially Responsible Machine Learning at ICML*.
- [40] Anh Nguyen, Tuong Do, Minh Tran, Binh X. Nguyen, Chien Duong, Tu Phan, Erman Tjiputra, and Quang D. Tran. 2022. Deep Federated Learning for Autonomous Driving. In *2022 IEEE Intelligent Vehicles Symposium (IV)*.
- [41] Jake Oster. 2022. How we count carbon emissions from electricity matters. Amazon. Retrieved Sept. 2023 from <https://www.amazon.science/blog/how-we-count-carbon-emissions-from-electricity-matters>
- [42] David Patterson, Joseph Gonzalez, Urs Hölzle, Quoc Le, Chen Liang, Lluís-Miquel Munguia, Daniel Rothchild, David R. So, Maud Texier, and Jeff Dean. 2022. The Carbon Footprint of Machine Learning Training Will Plateau, Then Shrink. *Computer* 55, 7 (2022). <https://doi.org/10.1109/MC.2022.3148714>
- [43] Xinchu Qiu, Titouan Parcollet, Javier Fernandez-Marques, Pedro Porto Buarque de Gusmao, Daniel J. Beutel, Taner Topal, Akhil Mathur, and Nicholas D. Lane. 2021. A first look into the carbon footprint of federated learning. *arXiv preprint arXiv:2102.07627* (2021).
- [44] Ana Radovanovic, Ross Koningstein, Ian Schneider, Bokan Chen, Alexandre Duarte, Binz Roy, Diyu Xiao, Maya Haridasan, Patrick Hung, Nick Care, Saurav Talukdar, Eric Mullen, Kendal Smith, Mariellen Cottman, and Walfredo Cirne. 2022. Carbon-Aware Computing for Datacenters. *IEEE Transactions on Power Systems* (2022).
- [45] Martin Rapp, Ramin Khalili, Kilian Pfeiffer, and Jörg Henkel. 2022. DISTREAL: Distributed Resource-Aware Learning in Heterogeneous Systems. In *AAAI*.
- [46] REN21. 2022. Renewables 2022 Global Status Report. (2022).
- [47] Nicola Rieke, Jonny Hancox, Wenqi Li, Fausto Milletari, Holger R. Roth, Shadi Albarqouni, Spyridon Bakas, Mathieu N. Galtier, Bennett A. Landman, Klaus Maier-Hein, Sébastien Ourselin, Micah Sheller, Ronald M. Summers, Andrew Trask, Daguang Xu, Maximilian Baust, and M. Jorge Cardoso. 2020. The future of digital health with federated learning. *npj Digital Medicine* 3, 1 (2020).
- [48] René Schwermer, Ruben Mayer, and Hans-Arno Jacobsen. 2023. Energy vs Privacy: Estimating the Ecological Impact of Federated Learning. In *ACM e-Energy*.
- [49] Jinhyun So, Kevin Hsieh, Behnaz Arzani, Shadi Noghahi, Salman Avestimehr, and Ranveer Chandra. 2022. FedSpace: An Efficient Federated Learning Framework at Satellites and Ground Stations. *arXiv preprint arXiv:2202.01267* (2022).
- [50] Behnaz Soltani, Venus Haghighi, Adnan Mahmood, Quan Z. Sheng, and Lina Yao. 2022. A Survey on Participant Selection for Federated Learning in Mobile Networks. In *Workshop on Mobility in the Evolving Internet Architecture (MobiArch)* at *MobiCom*.
- [51] Abel Souza, Noman Bashir, Jorge Murillo, Walid Hanafy, Qianlin Liang, David Irwin, and Prashant Shenoy. 2023. Ecovisor: A Virtual Energy System for Carbon-Efficient Applications. In *ASPLOS*.
- [52] Emma Strubell, Ananya Ganesh, and Andrew McCallum. 2020. Energy and Policy Considerations for Modern Deep Learning Research. In *AAAI*.
- [53] Mingxing Tan and Quoc Le. 2019. EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks. In *ICML*.
- [54] Maud Texier. 2021. A timely new approach to certifying clean energy. Google. Retrieved Sept. 2023 from <https://cloud.google.com/blog/topics/sustainability/t-eacs-offer-new-approach-to-certifying-clean-energy>
- [55] Muhammad Tirmazi, Adam Barker, Nan Deng, Md E. Haque, Zhijing Gene Qin, Steven Hand, Mor Harchol-Balder, and John Wilkes. 2020. Borg: The next Generation. In *EuroSys*.
- [56] Cong Wang, Bin Hu, and Hongyi Wu. 2022. Energy Minimization for Federated Asynchronous Learning on Battery-Powered Mobile Devices via Application Co-running. In *ICDCS*.
- [57] Qizhen Weng, Wencong Xiao, Yinghao Yu, Wei Wang, Cheng Wang, Jian He, Yong Li, Liping Zhang, Wei Lin, and Yu Ding. 2022. MLaaS in the Wild: Workload Analysis and Scheduling in Large-Scale Heterogeneous GPU Clusters. In *USENIX NSDI*.
- [58] Philipp Wiesner, Ilja Behnke, and Odej Kao. 2023. A Testbed for Carbon-Aware Applications and Systems. [arXiv:2306.09774](https://arxiv.org/abs/2306.09774) [cs.DC]
- [59] Philipp Wiesner, Ilja Behnke, Dominik Scheinert, Kordian Gontarska, and Lauritz Thamsen. 2021. Let's Wait Awhile: How Temporal Workload Shifting Can Reduce Carbon Emissions in the Cloud. In *ACM Middleware*.
- [60] Philipp Wiesner, Dominik Scheinert, Thorsten Wittkopp, Lauritz Thamsen, and Odej Kao. 2022. Cucumber: Renewable-Aware Admission Control for Delay-Tolerant Cloud and Edge Workloads. In *International European Conference on Parallel and Distributed Computing (Euro-Par)*.
- [61] Carole-Jean Wu, Ramya Raghavendra, Udit Gupta, Bilge Acun, Newsha Ardalani, Kiwan Maeng, Gloria Chang, Fiona Aga Behram, Jinshi Huang, Charles Bai, Michael Gschwind, Anurag Gupta, Myle Ott, Anastasia Melnikov, Salvatore Candido, David Brooks, Geeta Chauhan, Benjamin Lee, Hsien-Hsin S. Lee, Bugra Akyildiz, Maximilian Balandat, Joe Spisak, Ravi Jain, Mike Rabbat, and Kim M. Hazelwood. 2022. Sustainable AI: Environmental Implications, Challenges and Opportunities. In *MLSys*.
- [62] Qiang Yang, Yang Liu, Yong Cheng, Yan Kang, Tianjian Chen, and Han Yu. 2019. *Federated learning*. Morgan & Claypool Publishers.
- [63] Zhaohui Yang, Mingzhe Chen, Walid Saad, Choong Seon Hong, and Mohammad Shikh-Bahaei. 2021. Energy Efficient Federated Learning Over Wireless Communication Networks. *IEEE Transactions on Wireless Communications* 20, 3 (2021).
- [64] Ashkan Yousefpour, Shen Guo, Ashish Shenoy, Sayan Ghosh, Pierre Stock, Kiwan Maeng, Schalk-Willem Krüger, Michael Rabbat, Carole-Jean Wu, and Ilya Mironov. 2023. Green Federated Learning. [arXiv:2303.14604](https://arxiv.org/abs/2303.14604) [cs.LG]
- [65] Sai Qian Zhang, Jieyu Lin, and Qi Zhang. 2022. A Multi-Agent Reinforcement Learning Approach for Efficient Client Selection in Federated Learning. In *AAAI*. <https://doi.org/10.1609/aaai.v36i8.20894>
- [66] Jiajia Zheng, Andrew A. Chien, and Sangwon Suh. 2020. Mitigating Curtailment and Carbon Emissions through Load Migration between Data Centers. *Joule* 4, 10 (2020).
- [67] Zhi Zhou, Fangming Liu, Yong Xu, Ruolan Zou, Hong Xu, John C.S. Lui, and Hai Jin. 2013. Carbon-Aware Load Balancing for Geo-distributed Cloud Services. In *21st Int. Symposium on Modelling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS)*.
- [68] Chen Zhu, Zheng Xu, Mingqing Chen, Jakub Konečný, Andrew Hard, and Tom Goldstein. 2022. Diurnal or Nocturnal? Federated Learning of Multi-branch Networks from Periodically Shifting Distributions. In *ICLR*.

## A TABLE WITH ALL RESULTS

Note, that the *Upper Bound* baseline is not constrained by capacity or energy availability and therefore also uses grid energy.

Dataset & model	Approach	Global scenario				Co-located scenario			
		Target accuracy	Best accuracy	Time-to-accuracy	Energy-to-accuracy	Target accuracy	Best accuracy	Time-to-accuracy	Energy-to-accuracy
CIFAR-100 DenseNet-121	Upper bound		68.3 %	1.6 d	91.1 kWh		68.3 %	2.0 d	117.5 kWh
	Random		64.7 %	6.7 d	80.6 kWh		65.5 %	6.7 d	101.0 kWh
	Random 1.3n		66.0 %	4.7 d	79.2 kWh		66.6 %	5.3 d	113.4 kWh
	Random fc		65.4 %	6.4 d	89.8 kWh		65.7 %	6.5 d	97.8 kWh
	Oort	64.7 %	65.9 %	4.7 d	96.1 kWh	65.5 %	65.9 %	6.5 d	130.9 kWh
	Oort 1.3n		66.4 %	4.5 d	103.8 kWh		66.4 %	5.4 d	138.7 kWh
	Oort fc		65.8 %	5.3 d	102.4 kWh		66.1 %	6.4 d	126.7 kWh
	<b>FedZero</b>		<b>66.8 %</b>	<b>3.6 d</b>	<b>70.4 kWh</b>		<b>66.5 %</b>	<b>4.5 d</b>	<b>96.4 kWh</b>
Tiny ImageNet EfficientNet-B1	Upper bound		64.1 %	1.4 d	81.3 kWh		64.1 %	1.8 d	105.7 kWh
	Random		62.4 %	6.7 d	99.8 kWh		62.8 %	5.7 d	92.0 kWh
	Random 1.3n		63.1 %	5.6 d	109.6 kWh		63.3 %	3.7 d	86.0 kWh
	Random fc		63.0 %	5.6 d	96.4 kWh		63.1 %	6.5 d	102.1 kWh
	Oort	62.4 %	63.3 %	3.8 d	88.1 kWh	62.8 %	62.7 %	-	-
	Oort 1.3n		63.2 %	3.3 d	90.2 kWh		63.5 %	3.7 d	112.4 kWh
	Oort fc		63.1 %	3.9 d	89.0 kWh		59.7 %	-	-
	<b>FedZero</b>		<b>63.7 %</b>	<b>2.8 d</b>	<b>64.8 kWh</b>		<b>63.8 %</b>	<b>3.4 d</b>	<b>76.6 kWh</b>
Shakespeare LSTM	Upper bound		53.3 %	1.4 d	82.5 kWh		53.3 %	1.8 d	104.2 kWh
	Random		50.4 %	6.7 d	131.5 kWh		50.9 %	5.7 d	93.2 kWh
	Random 1.3n		50.7 %	4.6 d	97.9 kWh		51.5 %	4.5 d	90.0 kWh
	Random fc		52.0 %	2.8 d	57.3 kWh		52.1 %	3.7 d	59.2 kWh
	Oort	50.4 %	50.2 %	-	-	50.9 %	50.7 %	-	-
	Oort 1.3n		50.4 %	-	-		51.7 %	4.5 d	95.4 kWh
	Oort fc		50.5 %	6.7 d	157.4 kWh		50.6 %	-	-
	<b>FedZero</b>		<b>53.1 %</b>	<b>1.8 d</b>	<b>40.0 kWh</b>		<b>53.1 %</b>	<b>2.3 d</b>	<b>42.8 kWh</b>
Google Speech KWT-1	Upper bound		87.9 %	2.7 d	105.6 kWh		87.9 %	2.3 d	91.0 kWh
	Random		83.6 %	7.0 d	102.4 kWh		82.8 %	6.7 d	84.2 kWh
	Random 1.3n		85.2 %	5.5 d	103.5 kWh		85.1 %	4.3 d	80.8 kWh
	Random fc		85.0 %	5.7 d	95.6 kWh		83.0 %	6.7 d	83.8 kWh
	Oort	83.6 %	86.1 %	4.5 d	99.0 kWh	82.8 %	86.2 %	3.7 d	78.2 kWh
	Oort 1.3n		86.8 %	3.9 d	103.5 kWh		86.4 %	3.4 d	85.0 kWh
	Oort fc		87.0 %	3.8 d	86.2 kWh		85.6 %	3.7 d	78.1 kWh
	<b>FedZero</b>		<b>87.3 %</b>	<b>3.6 d</b>	<b>77.4 kWh</b>		<b>87.5 %</b>	<b>2.6 d</b>	<b>65.6 kWh</b>