IEEE Access
Multidisciplinary | Rapid Review | Open Access Journal

Received 11 June 2024, accepted 24 June 2024, date of publication 28 June 2024, date of current version 15 July 2024.

Digital Object Identifier 10.1109/ACCESS.2024.3420814

RESEARCH ARTICLE

Federated Learning Game in IoT Edge Computing

STEPHANE DURAND[1], KINDA KHAWAM[1,2], DOMINIQUE QUADRI[1], SAMER LAHOUD[3], AND STEVEN MARTIN[1]
[1]ROCS, LISN, Université Paris Saclay, 91190 Gif-sur-Yvette, France
[2]DAVID Lab, Université de Versailles Saint-Quentin-en-Yvelines, 78000 Versailles, France
[3]Faculty of Computer Science, Dalhousie University, Halifax, NS B3H 4R2, Canada

Corresponding author: Kinda Khawam (kinda.khawam@uvsq.fr)

This work was supported in part by the IE6 project.

ABSTRACT Edge Computing provides an effective solution for relieving IoT devices from the burden of handling Machine Learning (ML) tasks. Further, given the limited storage capacity of these devices, they can only accommodate a restricted amount of data for training, resulting in higher error rates for ML predictions. To address this limitation, IoT devices can leverage Edge Computing and collaborate in the learning process through a designated peer acting as an Edge device. However, the transmission of offloaded tasks over a wireless access network poses challenges in terms of time and energy consumption. Consequently, although collaborative learning can diminish the variance of the learned model, it introduces a communication cost, dependent on the chosen Edge device. In light of these considerations, this paper introduces a coalition formation game that proposes a distributed Federated Learning approach, where devices autonomously and efficiently select the most suitable Edge device, aiming to minimize both their learning error and communication cost.

INDEX TERMS Edge computing, game theory, federated learning, linear regression, IoT.

I. INTRODUCTION

As the Internet of Things (IoT) emerges, substantial data is gathered by resource-constrained devices. However, such massive data cannot feed efficiently ML models, locally on the IoT devices because of their limited capabilities. Consequently, these devices typically transmit the locally collected data to central servers to efficiently train and infer Machine Learning (ML) models. However, this process imposes undesired costs in terms of computation, storage, and communication on IoT devices, exacerbating the risk of privacy breaches.

However, for FL to give satisfying accuracy, it is essential to adhere to certain assumptions, with the two most crucial ones being the following:

- The initial assumption asserts that the data samples captured by various IoT devices represent independent and identically distributed (i.i.d.) random variables. However, in practice, since the local dataset of a single IoT device may not be representative of the overall population distribution. Hence, in this work, we relax this unrealistic assumption and take into account the Mean Square Error of the ML prediction model, that no longer amounts to zero.
- The second assumption stipulates that the local datasets generated by federated learners are of approximately equal size, ensuring balanced distributions. We maintain this assumption, as dissimilarity in dataset sizes predominantly arise from the diverse types of application scenarios. However, in the context of this study, we focus on a scenario where IoT devices are employed for the same application.

IoT devices collaborate by combining their learned parameters with a set of devices forming a learning cluster to enhance their predictive learning capabilities. However, while model aggregation through Federated Learning (FL) diminishes error variance by tapping into a broader dataset, it concurrently elevates error bias due to the diversity in the considered data distributions. Hence, the efficacy of learning is contingent on the number of IoT devices within the same cluster. Furthermore, increasing the number of federated learners poses an additional challenge by heightening the communication cost among these learning devices.

The associate editor coordinating the review of this manuscript and approving it for publication was Miguel López-Benítez.

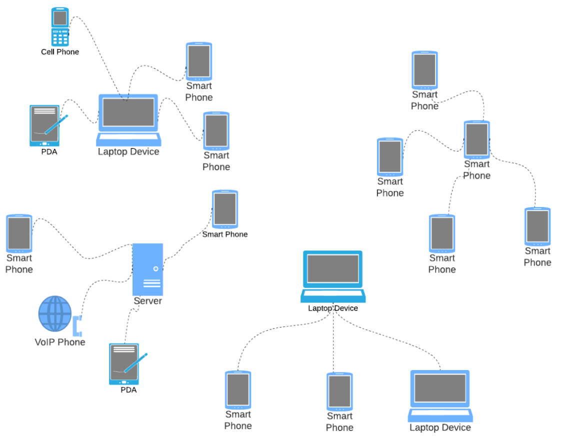

FIGURE 1. **Federated learning game.**

We build upon a prior work in [1] where IoT devices form Federated Learning coalitions through cooperative game theory, to strike a good compromise between communication efficiency and learning accuracy. The Edge device is one of the devices of the formed coalition, also denoted by *coalition leader*. The latter is chosen dynamically to avoid the risk of single-point-of-failure encountered with fixed Edge servers. Clustering devices around the coalition leader/Edge device is a coalition formation game [2] that we define and solve by considering Nash stable coalitions that coincide with optimal cluster formation. However, in our previous work, IoT devices were considered to be identical with regard to their computing power and ability to federate learning. In this work, we tackle a realistic heterogeneous setting where IoT devices are dissimilar which is particularly challenging as we loose key properties in the portrayed game. In fact, the communication cost in a coalition takes into account the disparity in the computation capacity of devices, which in turn impacts the selection of the coalition leader.

Fig. 2 gives the general framework of the paper. Initially, we have a set of devices that need to learn some model (subfig. (2a)). As they cannot send their dataset to a central server nor train the learning model on their own, a shared learning model is trained by each device and combined at the central server. However, to reduce the signaling cost with the central server, we propose to partition the devices into coalitions (subfig. (2b)), choose an aggregating node in each coalition and have each coalition train its learning model through the selected aggregating node (subfig. (2c)).

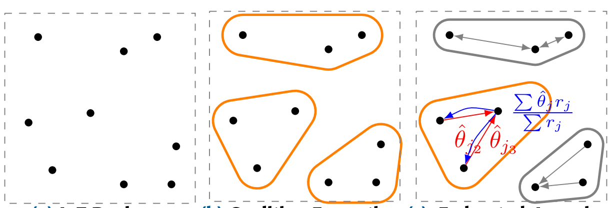

(a) **IoT Devices** **(b)** **Coalition Formation** **(c)** **Federated Learning per Cluster**

FIGURE 2. **General framework.**

The rest of the paper is organized as follows. We provide in section III the system model. In section IV, we portray the Federated Learning game by devising astutely the coalition cost function. In Section V, we identify the optimal coalition configuration by devising astute dynamics. We conclude in section VI.

II. RELATED WORK

Federated Learning is a privacy-preserving approach to machine learning that can efficiently leverage edge computing to train the learning models of edge clients without the need for hefty centralizing of data. This solution circumvents the need for central servers, ensuring swift computation and training of learning models.

Our work addresses the challenge of Federated Learning in Edge Computing through the lens of game theory. Numerous studies in the existing literature have explored FL within the context of non-cooperative game theory. Nevertheless, the majority of articles primarily serves the purpose of motivating end-users to engage in the FL process. In works such as [3] and [4], the framework of Stackelberg games is applied. In this context, end devices act as sellers, determining the pricing for joining a Federated Learning service platform, which acts as the buyer, seeking prediction services. The primary objective of these games is to enhance the accuracy of the global model for the leader while generating revenue for the followers. Another approach presented in the literature is the utilization of a non-cooperative Bayesian game, as proposed in [5]. This game aims to motivate participating end-users to provide precise results for prediction services. In the work detailed in [6], the aggregation process within the FL system is also depicted as a non-cooperative game. In the latter, both end-devices, capable of transmitting either accurate or erroneous updates, and the Edge server, with the ability to accept or disregard these updates, act as players. When an end-device submits accurate updates that are accepted by the Edge server, both parties experience an increase in their respective payoffs. A recent work in [7] propose a two-stage distributed collaborative architecture that leverages Edge Computing and FL to allow multiple Edge servers to collaboratively build an efficient learning model. However, the non-cooperative game is related to the provisioning of CPU resources and not the FL process. Similarly to our work, the authors in [8] explore the relation between the number of participating devices in FL and the accuracy of the global learning model, stressing on the energy consumption aspect. However, the study is done out of the scope of game theory.

An important work that bridged Federated Learning and cooperative game theory is the work in [9]. The authors showed that, depending on the characteristics of the dataset, devices should all participate in the same learning coalition (where aggregation is made in the Cloud) or should apply solely local learning (without having recourse to FL). However, considering only these two extreme cases is insufficient. Another relevant work is found in [10] which addresses the intersection of edge computing and FL using Hedonic games. Nonetheless, it focuses on finding Nash stable coalitions in the presence of communication cost without characterizing the utility function. In a prior work [1], we detailed the cost function and proved that when communication cost comes into play, devices often federate the learning in smaller coalitions. This highlights the need for a more nuanced understanding of Federated Learning scenarios.

The objective of our research is to address the challenges of distributed Federated Learning in Edge computing through game theory. In this context, constrained IoT devices must autonomously and judiciously form resilient coalitions for enhanced learning with energy efficient cost.

III. THE SYSTEM MODEL

We detail in this section the FL model which uses Linear Regression amounting to a specific error prediction model. Then, we give the communication model and the entailed cost. Finally, we give the cost function that will be used by the devised coalition formation game.

N, n	Set of the devices and their number
r, t	Number and indexes of repetitions to obtain the training data-set
X_t^i, Y_t^i	Measure and prediction of efficiency
$\theta, \hat{\theta}_K$	Prediction model and its estimation through a set of device $K \subset N$
η, μ_e	random bias of the regression model and its mean
E	Partition of N as a set of coalitions
m	Number of coalition, $m = \lvert E \rvert$
E_k	k^{th} part in configuration E, a subset of N
$E(i)$	Part containing the device i in configuration E
K, L	Partitions of E
w_{E_k}, w_k	Computation cost of the k^{th} coalition (4)
$\delta_i(E)$,δ_i	Communication cost of a device i under E
α	Normalizing factor
$c_i(E)$	Full cost of a device i under E, encoding both learning noise and communication cost. (6) .
$C(E)$	Total cost of the configuration, $C(E) = \sum_{i \in N} c_i(E)$
β	Folded constant $\beta = \frac{\mu_e}{\alpha(r-2)}$
$u_i(E)$	Marginal cost as utility 9 Difference in cost of a coalition from device i
m_{opt}	optimal number of coalition
$\mathcal{H}^w$	Ordered set of values of successive marginal costs given a set of weight w 1
$\mathcal{H}^w(E)$	Subset of $\mathcal{H}^w$ of the sizes lower than those in E ie the intermediary sizes when building E
$\mathcal{H}_i^w$	First coordinate (cost) of the i^{th} element of $\mathcal{H}$
$d(m)$	Number of extra device adding a coalition without adjusting the target marginal cost
$O(f)$	Asymptotically less than a multiple of f
$\Theta(f)$	Asymptotically more than a non zero multiple of f

FIGURE 3. Notations.

A. DATASET MODEL AND ERROR PREDICTION

We consider a set N of n constrained IoT devices that aim to perform learning through linear regression on a limited dataset. Due to the small dataset size and the limited storage capacity of such devices, the learning performed has relatively high error. Thus, an IoT device can federate the learning process with other devices to reduce collectively the learning error.

The Federated Learning process typically unfolds through the following sequential steps:

1) Initialization: An initial global model is established, often derived from an existing model or provided by the aggregating Edge device (the coalition leader).
2) Local Training: Each IoT device, denoted as i, undergoes training using its local dataset. This dataset consists of r_i independent and identically distributed (i.i.d.) samples for training, denoted by X_i^t for $t = 1, \ldots, r_i$, while the predicted variable is denoted by Y_i^t. The relationship between Y_i^t and X_i^t is linear, given by $Y_i^t = \theta_i^t \cdot X_i^t + \eta_i^t$, where θ_i^t denotes the slope of the linear regression, and η_i^t represents the random bias of the regression model with a mean of μ_e. Each device i aims to estimate the mean $\hat{\theta}_i$ based on its dataset of r_i samples. The training typically involves an iterative process, such as stochastic gradient descent, where the model is updated in multiple iterations based on the local data of each participating device.
3) Model Aggregation: Following local training, the updated model $\hat{\theta}_i$ from any device i is transmitted back to the coalition leader for aggregation. The latter combines the models using weighted averaging to create a new global model as follows:
$$\hat{\theta}_{E_k} = \frac{1}{\sum_{j=1}^{|E_k|} r_j} \sum_{j=1}^{|E_k|} \hat{\theta}_j \cdot r_j, \tag{1}$$
where $|E_k| \leq n$ is the number of devices that are federating the learning in the same coalition of device i, denoted E_k.
4) Model Distribution: The updated global model is sent back to the participating devices, and the process iterates with new rounds of local training.

This iterative cycle of local training, model aggregation, and model distribution continues until the total model converges or goes below an acceptable level of error.

Following the model in [9], the Mean Square Error (MSE) for linear regression learning by device i is given by what follows:

$$MSE_i = \frac{\mu_e}{(\sum_{j=1}^{|E_i|} r_j)^2} \sum_{j=1}^{|E_i|} \frac{r_j^2}{(r_j - 2)}, \tag{2}$$

We consider that devices have the same dataset size denoted by r, accordingly the MSE simplifies to:

$$MSE_i = \frac{\mu_e}{(r-2) \cdot |E_i|} \tag{3}$$

Note that the prediction error in (3) decreases as the learning coalition size increases. Hence, devices are interested in grouping altogether to federate the learning through a Cloud server. However, we will see that when communication cost comes into play, devices will prefer to federate the learning in smaller coalitions, through an astutely chosen Edge device.

B. COMMUNICATION MODEL

When the learning model is computed in a given coalition, each device i joining the latter inflicts an additional communication cost among devices in coalition E_i. We adopt a Time Division Multiple Access type of channel access like with NB-IoT (Narrow Band IoT) [11] or RedCap (Reduced Capability) [12]. In such a setting, devices endure a delay related to the coalition size as devices send in turn (in a given time slot) their estimated parameters to the coalition leader that will operate the aggregation according to (1), hence, roughly after $|E_i| - 1$ time slots. Further, the coalition energy consumption is also proportional to its size $|E_i|$.

Moreover, we consider that one packet is sufficient to send the learned parameter and that it is sent in one time slot. Further, we consider that orthogonal channels are re-allocated among coalitions to cancel out interference. It is reasonable to assume that a coalition will form among devices close together geographically to reduce the communication cost. However, not all aggregating nodes are equivalent and can have different computation capacities, which constitutes another limiting factor. Thus, the communication cost of device i in coalition E_i is denoted by δ_i and given by:

$$\delta_i = w_{E_k} \cdot (|E_k| - 1) \tag{4}$$

where the weight w_{E_k} characterizes the computation cost of coalition E_k, which is dependent on the members of the coalition.

The communication cost in a coalition takes into account the disparity in the computation capacity and geographical position of IoT devices which directly affects the selection of the coalition leader. Hence, the weight w_{E_k} of a coalition E_k is the weight of its leader, which is the device with the fastest computational speed or best centrality (the closest geographically from all devices in the coalition), corresponding to the smallest weight:

$$w_{E_k} = \min\{w_k, k \in E_j\} \tag{5}$$

C. COST FUNCTION

After defining the prediction error and communication cost, we devise the cost of device i in coalition E_k as follows:

$$c_i = MSE_i + \alpha\delta_i \tag{6}$$

where α is a normalizing factor.

We define the state $E = \{E_1, \cdots, E_m\}$, a partition of devices into coalitions, specifically noting E_k the k^{th} coalition and m the total number of coalitions.

The total cost that we seek to minimize is the sum of the cost over all devices $C(E) = \sum_{i \in N} c_i(E)$. We re-write the cost function of device i in coalition E_k:

$$c_i(E) = \frac{\beta}{|E_k|} + w_{E_k} \cdot (|E_k| - 1), \tag{7}$$

where $\beta = \left(\frac{\mu_e}{\alpha(r-2)}\right)$. Note that when device i acts as a singleton and learns on its own, the communication cost vanishes.

The total cost becomes:

$$C(E) = \beta|E| + \sum_{E_k \in E} w_{E_k} \cdot |E_k|(|E_k| - 1) \tag{8}$$

which is only dependent on the number, size and weight of coalitions, thus necessitates minimal signaling among players.

IV. FEDERATED LEARNING (GFL) GAME

We define the Federated Learning game as a game with IoT devices as the players, whose possible strategies are: to join another single device to form a new coalition, to join an existing coalition or to separate back out from the current coalition. We set the utility of a player as the marginal cost, defined as the difference between the current total cost (including the player) and the total cost without the said player:

$$u_i(E) := C(E) - C\left(\{K \setminus \{i\}, K \in E\} \setminus \{\emptyset\}\right) \tag{9}$$

Accordingly,

$$u_i(E) := \begin{cases} E_j = \{i\} & \Rightarrow u_i(E) = \beta \\ E_j \neq \{i\} & \Rightarrow u_i(E) = 2w_{E_j}(|E_j| - 1) \end{cases} \tag{10}$$

In the homogeneous case when $w_{E_k} = 1, \forall k$, the game is a potential game. However, in the heterogeneous case addressed in the present work, the game is no longer a potential game as the cost of a coalition after a device defection can depend on factors not known by the device. Therefore, in such a complex setting, we need to have recourse to a practical way to compute optimal solutions, which will be explained hereafter.

A. FORMULATED ABSTRACTION

We will introduce as an abstraction, denoted by $\mathcal{H}$ a set of successive breakpoints in the effective cost of filling the coalitions. This will allow us to express the structure and cost of a solution in a parametric way using said equivalent cost, which also facilitates comparing solutions among each others. Then, we will devise a few algorithms using this abstraction to build such solutions.

1) DEFINITION OF $\mathcal{H}$

The validity and total cost of a solution only depend on the identity of the leaders of coalitions and on the number of devices joining each coalition. Thus, it is not affected by the detail of which device joined which coalition.This remains true when building a solution by having the devices join gradually, only keeping the index of the joined coalition still keeps enough information to characterize the solution.

We then associate to those indexes the number of devices in the coalition (other than the leader) prior to their current joining of the coalition. This removes the need to track the order of events, and instead treat the coalitions list as a set. We also associate a second information to each element, the marginal cost of joining that coalition at that size. Namely the difference in the coalition cost incurred by the additional device: a device joining coalition j which already contains k non leaders devices will have a marginal cost of $2 \cdot w_j \cdot k$. (writing $w_j = w_{E_j}$ for ease of reading). This second label does not add any information but instead will serve as an order on the elements of the set.

We obtain a finite set of triplets with enough information to characterize the desired solution. Formally, we obtain what follows:

Definition 1: The set $\mathcal{H}^{\mathbf{w}}$ and subset $\mathcal{H}^{\mathbf{w}}(E)$:

For a given set or subset of weights $\mathbf{w} = w_1 \cdots w_m$, we consider:

- For each coalition j and positive integer k, the triplet of the marginal cost of coalition j when going from size k to $k+1$, the coalition identity j and coalition size k are given by: $(2 \cdot w_j \cdot k, j, k)$, also noted $(2w_j \cdot k)_j^k$.
- $\mathcal{H}^{\mathbf{w}}$: the set of all such possible triplets.
- For a given configuration E, $\mathcal{H}^{\mathbf{w}}(E)$ is the subset of $\mathcal{H}^{\mathbf{w}}$ verifying $0 \leq k < |E_j|$.
  It is also a list of events consisting of devices joining coalitions and allowing to build a solution equivalent to the configuration E.
- When sorting $\mathcal{H}^{\mathbf{w}}$ by marginal cost first then by coalition index, we denote by $\mathcal{H}_i^{\mathbf{w}}$ the first coordinate of the i^{th} element of the set $\mathcal{H}^{\mathbf{w}}$, corresponding to the marginal cost,
- Likewise, for E a solution on n devices, $\mathcal{H}_i^{w}(E)$ is the first coordinate of the i^{th} element of the set $\mathcal{H}(E)$.

We can note that $\mathcal{H}^{\mathbf{w}}$ is such that the number of elements whose first coordinate is below a given value b can be bounded by a finite number, specifically $\sum_{j \leq |\mathbf{w}|} \frac{b}{w_j}$. This means that the number of elements whose first coordinate is lower or equal to the first coordinate of a specific element is finite, and both the i^{th} element of the set and its first coordinate $\mathcal{H}_i^{w}$ are well defined.

Remark 1: Remarks on the construction of the coalition set

1) For a given finite subset $H \subset \mathcal{H}^{\mathbf{w}}$, there exists a solution E such that $H = \mathcal{H}^{\mathbf{w}}(E)$ if and only if H verifies:

$$\forall j,\ \forall 0 \leq \ell \leq k,\ \left((2 \cdot k \cdot w_j)_j^k \in H\right) \Rightarrow \left((2 \cdot \ell \cdot w_j)_j^\ell \in H\right) \tag{11}$$

2) For a given solution E, j is a leader in E if and only if $0_j^0 \in \mathcal{H}^{\mathbf{w}}(E)$, which corresponds to the addition of the leader to its own coalition.

Proof: For the first remark, the first conditional statement is easy: $\mathcal{H}^{\mathbf{w}}(E)$ is of the form $\bigcup_j \{(2w_j k)_j^k, k < |E_j|\}$, so if we take j, k and ℓ such that $\ell \leq k$, then $(2w_j k)_j^k \in \mathcal{H}^{w}(E)$ means that $k < |E_j|$, thus $\ell \leq k < |E_j|$ and in turn $(2w_j \ell)_j^\ell \in \mathcal{H}^{w}(E)$.

In the other direction, since H is finite, for any j, $\{k \in \mathbb{N}, (2w_j \cdot k)_j^k \in H\} = H \cup \{x_j^k, x \in \mathbb{R}, k \in \mathbb{N}\}$ is finite and its complementary to $\mathbb{N}$, $\{k \in \mathbb{N}, (2w_j \cdot k)_j^k \notin H\}$ is non empty and admits a smallest element. Consider a solution E using those values as the sizes of the coalitions: for any device j, noting x_j the smallest positive integer such that $(2w_j x_j)_j^{x_j}$, if $x_j \neq 0$ then E admits a coalition whose leader is j, of size x_j; else, j is not a leader in E. Then, computing the set associated with this solution E gives us exactly H, from which we started, proving the first assertion of the remark that any H verifying property (11) is associated with a valid solution.

The second remark states what follows: in a solution, the leaders become the coalition indexes such that $H^{\mathbf{w}}(E)$ restricted to those values of indexes is nonempty, with at least one element. In turn, this subset contains the elements corresponding to all integers lower than the one of that first element, including 0, whose corresponding element is $(2 \cdot w_j \cdot 0)_j^0 = 0_j^0$. Conversely, if the device is not a leader, then the subset is empty and $\mathcal{H}^{w}(E)$ does not contain that element. ■

Lemma 1: Properties of the defined order

1) *The set $\mathcal{H}^{\mathbf{w}}$ is increasing in $\mathbf{w}$, using the inclusion as the order on both sides.*
2) *For any potential leader $j \in N$, and size $k \in \mathcal{N}$, for any vector of weights $\mathbf{w}$, there is a unique corresponding element t, j, k in $\mathcal{H}^{\mathbf{w}}$. Furthermore, the first coordinate, which is the marginal cost t, is increasing in the values of the coordinates of $\mathbf{w}$.*
3) *The value of $\mathcal{H}_i^{\mathbf{w}}$ for any given index i is increasing in the values w_i of the weights, using the order on the reals.*
4) *The values of $\mathcal{H}_i^{\mathbf{w}}$ for any given i is decreasing in $\mathbf{w}$ using the inclusion.*
5) *$\mathcal{H}^{\mathbf{w}}(E)$ verifies a form of convexity: $\forall j, k \in \mathbb{N}$, $\left(\left((2w_j \cdot k)_j^k \in \mathcal{H}^{\mathbf{w}}(E)\right) \Rightarrow \left(\forall \ell \leq k, (2w_j \cdot \ell)_j^\ell \in \mathcal{H}^{\mathbf{w}}(E)\right)\right)$*

Proof: For the first property, we can separate $\mathcal{H}^{w}$ by coalitions, and rewrite it as the disjoint union $\mathcal{H}^{w} = \bigcup_j \{(2 \cdot w_j \cdot k), k \in \mathbb{N}\} = \bigcup_j \mathcal{H}^{\{w_j\}}$ making the assertion obvious.

For property 2, The first coordinate of any triplet is a function of the two other coordinates and $\mathbf{w}$, and $\mathcal{H}^{\mathbf{w}}$ is the set of all possible such triplets, so for any pair of indexes $j \in [1 \cdots m]$, $k \in \mathbb{N}$, there is exactly one triplet in $\mathcal{H}^{\mathbf{w}}$ whose second and third coordinates are j and k. This expression is $2w_j k$, with 2 and k both positive integers, so the first coordinate of this triplet is increasing in one of the coordinate of $\mathbf{w}$ and untouched by the others.

For property 3, we also need to show that this property is also true when identifying the elements by the order of the first coordinate, we can consider a first weight vector $\mathbf{w}^{(1)}$ and take note of the i smallest elements and values, by definition all lower or equal to $\mathcal{H}_i^{\mathbf{w}^{(1)}}$. We then choose $\mathbf{w}^{(2)} \leq \mathbf{w}^{(1)}$ (i.e. $\forall j$, $\mathbf{w}_j^{(2)} \leq \mathbf{w}_j^{(1)}$) and apply the previous observation that the triplets identified by their second and third coordinates are increasing in $\mathbf{w}$ to see that the triplet we are considering have decreased or preserved their values, so remained lower or equal to $\mathcal{H}_i^{\mathbf{w}^{(1)}}$. Since we can show i elements of $\mathcal{H}^{\mathbf{w}^{(2)}}$ whose values are each lower than $\mathcal{H}_i^{\mathbf{w}^{(1)}}$, then $\mathcal{H}_i^{\mathbf{w}^{(2)}} \leq \mathcal{H}_i^{\mathbf{w}^{(1)}}$. This is true for all $\mathbf{w}^{(1)}$ and all $\mathbf{w}^{(2)} \leq \mathbf{w}^{(1)}$, so $\mathcal{H}_i^{\mathbf{w}}$ is increasing in $\mathbf{w}$.

For property 4, we can see the assertion as a consequence of property 1: increasing $\mathbf{w}$ for the inclusion means inserting new elements into $\mathcal{H}^{\mathbf{w}}$. Either none of the coordinates of those new elements is lower than $\mathcal{H}_i^{w}$, in which case, the identity and first coordinate of the i^{th} element is unchanged, or some of the new elements are lower than $\mathcal{H}_i^{w}$, meaning the i^{th} element of $\mathcal{H}$ is replaced either by one of those new elements, or by an element of $\mathcal{H}$ that had a lower ranking, in both cases reducing or preserving the first coordinate.

Finally, property 5 is a direct result of the definition of $\mathcal{H}^{\mathbf{w}}(E)$. ∎

Without loss of generality, we can consider that the list of weights is sorted in increasing order.

An instance of GFL on n devices of weights $\mathbf{w}$ translates into choosing a finite subset $H(E)$ of $\mathcal{H}^{\mathbf{w}}$ of size n verifying Remark 1. We can separate it into choosing a set of leaders $\{j_1, \cdot j_m\} \subset \{1 \cdots n\}$ and choosing $\{0^0_{j_1}, 0^0_{j_2} \cdots 0^0_{j_m}\} \subset H(E) \subset \mathcal{H}^{\{w_{j_1}, \cdot w_{j_m}\}} \subset \mathcal{H}^{\mathbf{w}}$ as attributing the devices to coalitions headed by those leaders.

We give the following examples for clarity:
With three devices of weights 2, 3 and 5, the set is:

$$\mathcal{H}^{(2,3,5)} = \left\{0^0_1, 0^0_2, 0^0_3, 4^1_1, 6^1_2, 8^2_1, 10^1_3, 12^3_1, 12^2_2, 16^4_1, \cdots\right\}$$

A solution E on this set with two coalitions, one coalition with leader indexed by 1 of size 2 and the other with leader indexed by 2 of size 1, which translates into:

$$\mathcal{H}^{(2,3,5)}(E) = \left\{0^0_1, 0^0_2, 4^0_1\right\}$$

As another example, we consider two devices with weight 1 and $\sqrt{2}$ which gives:

$$\mathcal{H}^{1,\sqrt{2}} = \left\{0^0_1, 0^0_2, 2^1_1, (2\sqrt{2})^1_2, 4^2_1, (4\sqrt{2})^2_2, 6^3_1, 8^4_1, (6\sqrt{2})^3_2 \cdots\right\}$$

with no other equality of the first coordinate other than in 0.

2) COST FUNCTION OF THE GFL GAME

Recall that according to (5), the cost of a device in this game, with the assumption that the leader is indeed the least costly, is as follows:

$$c_i(E) = \frac{\beta}{|E_j|} + w_j\left(|E_j| - 1\right) \tag{12}$$

and the total cost is:

$$C(E) = \beta|E| + \sum_j w_j|E_j|\left(|E_j| - 1\right) \tag{13}$$

Theorem 1: Total cost of $\mathcal{H}^{w}(E)$: The cost of any solution E can be expressed as $\sum_{i \leq |\mathcal{H}^{w_1 \cdots w_m}(E)|} \mathcal{H}_i^{w_1 \cdots w_m}(E) + \beta|E|$.

Proof: Intuitively, we defined the first coordinate of each triplet as the marginal cost of adding the device to that coalition on a solution which already had the coalitions and leaders fixed, with the empty solution having only the cost $\beta \cdot |E|$.

We can use the identity $w_j|E_j|\left(|E_j| - 1\right) = \sum_{k \leq |E_j|} 2\, w_j \cdot k$ on every coalition in the formula from (13), and obtain $C(E) - \beta|E| = \sum_j \sum_{k \leq |E_j|} 2\, w_j \cdot k = \sum_{x_j^k \in \mathcal{H}(E)} x = \sum_{i \leq n} \mathcal{H}_i(E)$. ∎

Corollary 1: Optimal cost from $\mathcal{H}^{\mathbf{w}}$: The total cost of an optimal solution of $GFL(w_1 \cdots w_m)$ is equal to $\sum_{i \leq n} \mathcal{H}_i^{w_1 \cdots w_m} + \beta m$, where $m = |E|$.

Proof: This is a corollary of Theorem 1: the cost of any solution is the sum of the n first coordinates of the elements of $\mathcal{H}^{\mathbf{w}}$, so the smallest such sum is a natural lower bound for it. Since choosing the smallest elements verifies the pseudo convexity property from Lemma 1 item 1, it corresponds to a feasible solution and thus the actual optimal. ∎

B. OPTIMAL NUMBER OF COALITIONS

The main advantages of the result from Corollary 1 are to give a way to obtain the optimal cost for a specific set of coalition leaders without needing to build the solution explicitly, and to express the effect of altering $\mathbf{w}$ on the cost. In particular, it can be used to show the following result:

Theorem 2: Convexity in the number of coalitions m: The function associating m to the cost of the best solution of GFL game on exactly m coalitions is convex.

Proof: We use the expression from equation (1). In order to express the effect of varying the number of coalitions, we consider the best solution using $m+1$ coalitions and remove the one whose leader has of largest weight w_{m+1} from the game to start building the best solution for m coalitions. The new set of triplet is a subset of the previous one: $\mathcal{H}^{w_1 \cdots w_m} \subset \mathcal{H}^{w_1 \cdots w_{m+1}}$ and by construction, the optimal solution for $m+1$ coalition takes the smallest elements of this set: $!\exists (t, u) \in \mathcal{H}^{w_1 \cdots w_{m+1}} \setminus \mathcal{H}^{w_1 \cdots w_{m+1}}(E_opt(m)) \times \mathcal{H}^{w_1 \cdots w_{m+1}}(E_opt(m)), t < u$, thus the restriction of the solution to the reduced set also consist in the smallest elements of the new set: $!\exists (t, u) \in \mathcal{H}^{w_1 \cdots w_m} \setminus (\mathcal{H}^{w_1 \cdots w_{m+1}}(E_opt(m)) \cap \mathcal{H}^{w_1 \cdots w_m}) \times \mathcal{H}^{w_1 \cdots w_m}(E_opt(m)), t < u$. With the rest of the elements of the new best solution being the elements of $\mathcal{H}^{w_1 \cdots w_m}$ of indexes $n - \left|(x)^k_{m+1} \in \mathcal{H}^{w_1 \cdots w_{m+1}}\left(E_{opt}(m+1)\right)\right| + 1$ to n. We write $d(m+1)$ the number of those removed elements.

We can write the difference in cost when going from $m+1$ to m as the sum over all moved devices of the difference in cost between the two positions:

$$C(m) - C(m+1) = \sum_{i=1}^{d(m+1)} \left(\mathcal{H}_{n-d(m+1)+i}(m) - i \cdot w_{m+1}\right) - \beta$$

$$= \sum_{i=0}^{d(m+1)-1} \left(\mathcal{H}_{n-i}(m) - \sum_{i=1}^{d(m+1)} i \cdot w_{m+1} \right) - \beta$$
$$= \sum_{i=1}^{d(m+1)} (\mathcal{H}_{n-i+1}(m) - i \cdot w_{m+1}) - \beta. \quad (14)$$

In this expression, we can see that:

- By nature of $d(m)$,
$$i \cdot w_{m+1} \leq d(m+1) w_{m+1} \leq \mathcal{H}_n^{w_1 \cdots w_{m+1}} = \mathcal{H}_{n-d(m+1)}^{w_1 \cdots w_m}.$$
As $\mathcal{H}_{n-d(m+1)}^{w_1 \cdots w_m} \leq \mathcal{H}_{n-d(m+1)+i}^{w_1 \cdots w_m}$, all terms of the sum are positive.
- Using Lemma 1 item 1, $H_n^{w_1 \cdots w_m}$ is decreasing in the set $w_1 \cdots w_m$ for the inclusion which itself is increasing in m, while the value w_m is increasing in m. Thus $d(m)$ is decreasing in m.
- Also using Lemma 1 item 4, for a fixed i, $\mathcal{H}_{n-i+1}(m)$ is decreasing in m.
- Since w_m is increasing in m, for a fixed i, $i \cdot w_{m+1}$ is increasing in m.
- β is constant.

Thus, the difference between the optimal total cost for m minus the one for $m+1$ is decreasing in m and in turn the optimal cost on m coalitions, for which this difference is the opposite of the discrete derivative, is convex in m. ■

This gives a way to check whether a candidate value of m is above or below the optimal, and find the said optimal either by testing values successively or by dichotomy.

V. COALITION FORMATION ALGORITHM

The formulated abstraction devised in the previous section gives us a first method to obtain the best solution for a specific number of coalitions (m is fixed) by building the first n elements of the set $\mathcal{H}(m)$ and by using the number of each class of elements as the size of the corresponding coalition.

We can then obtain the optimal solution by iterating on m. Being a small optimization problem, we use the fact that it is easier to compute $\mathcal{H}^{w_1 \cdots w_{m+1}}$ from $\mathcal{H}^{w_1 \cdots w_m}$ and w_{m+1} than to rebuild the set fully, and simply keep and update the number of elements of the different coalitions instead of having to recount them.This gives us Algorithm 1, which builds and updates $\mathcal{H}(i)$ and makes use of them to compute the successive costs, until the corresponding total cost increases again, at which point the next to last value is the optimal number of coalitions, and the associated set gives the globally optimal solution.

Algorithm 1 Finding E_{opt} Through H

Variables : Ordered triplet set H; real C and Ct; arrays of integers X and Y
sort w
/* initialization, $H := \mathcal{H}^{\{w_1\}}$ */
$H \leftarrow \emptyset$
for $0 \leq k < n$ **do**
  H**.insert(** $(2k \cdot w_1, 1, k)$**,end)**
end

$X \leftarrow [n]$ // temporary list of sizes
$C \leftarrow w_1 \cdot n \cdot (n-1) + \beta$ // temporary total cost
/* loop */
while $1 \leq m \leq n$ *and flag* **do**
  /* update H */
  for $0 \leq 2 \cdot k \cdot w_m < (H.last).first$ // ((H.last).first is $\mathcal{H}_n^{w_1 \cdots w_m}$)
  do
    H**.insert(** $(2k \cdot w_m, m, k)$**)**
    H**.delete(last)** // preserves size= n
  end
  /* list of sizes, to avoid going through H fully */
  $Y \leftarrow$ **empty array of size** m
  for $1 \leq j \leq m$ **do**
    $Y[j] \leftarrow \left\lfloor \frac{(H.last).first}{2 \cdot w_j} \right\rfloor + 1$ // the +1 is the leader
    if $\frac{(H.last).first}{2 \cdot w_j}$ *integer and* $(H.last).second > j$
    then $Y[j] \leftarrow Y[j] + 1$
    // second term of the lexicographical order value then coalition
  end
  /* total cost */
  $Ct \leftarrow m \cdot \beta$
  for $1 \leq j \leq m$ **do**
    $Ct \leftarrow Ct + w_j \cdot Y[j] \cdot Y[j-1]$
  end
  if $Ct < C$ // before the convex minimum
  then
    $C \leftarrow Ct$; $X \leftarrow Y$
  else
    Break and return C, X
  end
  return C, X
end

A. COMPLEXITY OF THE ALGORITHM

Sorting **w** (line 1) takes a time $O(nlogn)$ if the list is unsorted.

Initializing the loop with the solution for $m = 1$ i(n lines 2 to 7) takes n insertions at the end of the set and is in $O(n)$.

Line 8 to 30 is the main loop, repeated until the cost increases again, meaning from $m = 2$ to $m = m_{opt} + 1$.

- In the loop, the first step (lines 9 to 13) is to update $\mathcal{H}$ by the addition of a new coalition. This means inserting $d(m) \leq \frac{\frac{1}{w_m}}{\sum_{i \leq m} \frac{1}{w_i}} \leq \frac{n}{m}$ pre-sorted elements in a set maintained at size n. It can be done in $O(d(m) \log n) \leq O\left(\frac{n}{m} \log n\right)$ operations.
- The second step uses $\mathcal{H}^{w_1 \cdots w_m}$ to build the optimal solution on m coalitions (in lines 14 to 18) in $O(m)$

operations. Line 17 comes from the fact that we are using a lexicographical order: in case of equality of the first coordinate, the arbitration is done by the second (since the three coordinates are linked linearly, the second being also equal means the third will be as well, so it does not need to be checked).

- Then, the algorithm computes the total cost of that solution (in lines 19 to 22), using the original equation (13). This takes $O(m)$ operations.
- The end of the loop consists in just comparing the new obtained value to the previous one.

The complexity of the loop is the sum of the complexity of its steps, for m going from 2 to $m_{opt}+1$, which takes $O(m_{opt}^2 + n\log n \log m_{opt})$. This is also the global complexity of the algorithm.

This expression depends on the number of devices n, and a characteristic of the final solution which is the optimal number of coalitions, m_{opt}, which we know is at most n.

We show in appendix that if we take the reasonable assumption that the ratio of β over the average weight is neither vanishing nor unbounded, we can show that m_{opt} is not vanishing compared to n, which means the complexity bound is the actual complexity, and will allow to compare with the other variants.

As a remark, with this algorithm, if instead of testing the number of coalitions in increasing order, we try to use a dichotomy search or an oracle, we would not gain much, as the cost of building $\mathcal{H}$ directly for the final set of leaders is close to the cumulative sum of the computation cost of each incremental update. In order to use a more precise search, we will need a more direct method to compute the optimal total cost for a specific set of leaders.

B. SUB-MODULARITY AND GREEDY ALGORITHM

The GFL game with a fixed set of coalition leaders is submodular. This allows to use a greedy algorithm in order to add or remove devices from an optimal solution, changing the number of devices n on the same set of coalitions (same m and $\mathbf{w}$) while preserving the optimality.

Theorem 3: Sub-modularity: The GFL game constrained on the number of coalitions is sub-modular as any optimal solution for $n+1$ players can be constructed from an optimal solution for n players with the best coalition that the additional player should join.

The proof is in appendix .

We can use this method to obtain an optimal solution for a specific n in three ways:

The empty solution is an optimal solution for $n = m$. Starting from this solution and using a greedy algorithm is simply a new perspective on building and using $\mathcal{H}$ directly: at each step, the possible choices for the next element to be added greedily is exactly the element or the set of elements of $\mathcal{H}$ with the smallest first coordinate.

An optimal solution for a set of leaders, restricted to a strict subset, is still an optimal solution, except for a reduced number of devices. This allows us to link the solutions for consecutive values of m, and complete the solutions greedily. Again, this behaves similarly as the update of $\mathcal{H}$ from Algorithm 1 reversing the order of the loop.

However, using the sub-modularity to correct an imperfect oracle giving an optimal solution for a slightly incorrect number of devices is distinct from the previously depicted algorithms, and will give us a faster method.

1) THRESHOLD METHOD

Previously, when modeling the solutions as a subset of $\mathcal{H}^{w_1\cdots w_m}$, we have characterized in Corollary 1 the optimal solutions as the ones taking the n elements of smallest marginal costs, which translates directly to choosing for each coalition the size that would put their marginal cost immediately below or equal to the real $\mathcal{H}_n^{w_1\cdots w_m}$, with some care taken in case of equality. This approach is not limited to using $\mathcal{H}_n^{w_1\cdots w_m}$ as we will show that any value in $\left[\mathcal{H}_n^{w_1\cdots w_m}, \mathcal{H}_{n+1}^{w_1\cdots w_m}\right]$ works as well.

Definition 2: Fitted Sizes: For a fixed set of leaders, a set of coalition sizes is fitted to a real τ if for each leader j, the corresponding size is the closest integer to $\frac{\tau}{w_j}+\frac{1}{2}$, or one of the two closest integers in case of equality.

We replace the control on the number of devices n by a parameter that can be seen as an incentive for the leaders to have bigger coalitions. Accordingly, Theorem 1 becomes:

Corollary 2: Optimal Fitting: For any coalition configuration and set of weights, a set of sizes $\mathbf{X} = X_1 \cdots X_m$ fitted to a positive real is an optimal solution for the number of devices corresponding to the total size $m + \sum X_j$.

Proof: If we build $\mathcal{H}^{\mathbf{w}}$, and fit the sizes to $\mathcal{H}_\ell^{\mathbf{w}}$, Theorem 1 and Corollary 1 tell us that the total cost will be the sum of the ℓ smallest elements of $\mathcal{H}^{\mathbf{w}}$ and that it is optimal for $m+\ell$ devices. ■

Another consequence of the same theorem is choosing the right value that allows us to get the optimal solution for n devices:

Corollary 3: Existence of a fitted solution for any number of devices: For any number of coalitions and set of weights and for any number of devices n, there exists a threshold $\tau \in \mathbb{R}$ and a set of sizes fitted to τ that is an optimal solution for GFL on n devices.

Proof: We can simply take $\mathcal{H}_n^w$, which allows to build the same solution as using $\mathcal{H}$. ■

However, finding that real, denoted by τ, is as difficult as finding the optimal solution. Considering which reals will give us the correct number of devices, we can use $\mathcal{H}$ to see that these reals are the interval $\left[\mathcal{H}_n^{w_1\cdots w_m}, \mathcal{H}_{n+1}^{w_1\cdots w_m}\right]$. If we consider any optimal solution E, this interval is the intersection of the intervals $\left[2\cdot w_j\cdot |E_j|, 2\cdot w_j\cdot \left(|E_j|+1\right)\right]$. Finally, we can compute the average of the bound of those intervals and use it to approximate an element of the intersection. Then, we can correct the error from using an approximation instead of an element of the interval by adjusting the potential difference in the number of devices via a greedy algorithm owing to the sub-modularity property.

Lemma 2: Approached value of the Threshold τ: For a given positive integer n, a solution of GFL fitted to $\tilde{\tau}(n, m) := \frac{n-\frac{m}{2}}{\sum \frac{1}{w_i}}$ is at most $\frac{m}{2}$ distant from n. In addition, for uniformly and independently chosen weights, the absolute value of the error is $\frac{\sqrt{m}}{12}$.

The proof is in the appendix, in section .

Consequently, we have obtained a method to compute a solution and its total cost corresponding to a specific number of coalitions without needing an initial partial result, and thus allowing the use of dichotomy or guesses based on properties of $\mathbf{w}$. In fact, since we have a fast method to compute the cost of the best solution for a given m and do not rely on computations for neighboring values, we search for the best m by dichotomy, using the convexity, to test whether a value is too high or too low. For each value of m, we guess the corresponding pairs of marginal costs of coalitions in an optimal solution as being close to $\tilde{\tau}$. We build a solution of potentially an incorrect size by taking the sizes of the coalitions so that $\tilde{\tau}$ is in each of the intervals of the pairs of marginal costs, then use a greedy algorithm to adjust the number of players without losing optimality until we obtain a solution both optimal and of the correct size. This method is encoded in Algorithm 2.

Algorithm 2 Approached Threshold Method

Variables : real $\tilde{\tau}$, real A, array of m integers X, sorted list of m pairs Y, integer o

`/* approximated value */`

$A \leftarrow 0$

for $0 < j \leq m$ **do**

  $A \leftarrow A + \frac{1}{w_j}$

end

$\tilde{\tau} \leftarrow \frac{n-\frac{m}{2}}{A}$

`/* approximated solution */`

for $0 < j \leq m$ **do** $X_j \leftarrow \left\lfloor \frac{\tilde{\tau}}{w_j} + \frac{1}{2} \right\rfloor$

`/* check */`

$o \leftarrow 0$

for $0 < j \leq m$ **do** $o \leftarrow o + X_j$

if $o = n$ **then** return X

`/* adjustment */`

if $o < n$ **then**

  $Y \leftarrow \emptyset$

  for $0 < j \leq m$ **do** Y.insert($(2 \cdot w_j \cdot (X_j + 1)), j$)

  `// Y is the list of coalition and of the marginal costs of increasing its size by one, sorted by increasing cost`

  while $o < n$ **do**

    $j_0 \leftarrow (Y$.first).second

    Y.insert(((Y.first).first $+ w_{j_0}, j_0$))

    Y.remove(first)

    $X_{j_0} \leftarrow X_{j_0} + 1$

    $o \leftarrow o + 1$

  end

else

  $Y \leftarrow \emptyset$

  for $0 < j \leq m$ **do** Y.insert($(2 \cdot w_j \cdot X_j), j$)

  `// Y is the list of coalition and of the marginal gain of decreasing its size by one, sorted by decreasing gain`

  while $o > n$ **do**

    $j_0 \leftarrow (Y$.last).second

    Y.insert(((Y.last).first $- w_{j_0}, j_0$))

    Y.remove(last)

    $X_{j_0} \leftarrow X_{j_0} - 1$

    $o \leftarrow o - 1$

  end

end

return X

2) COMPLEXITY OF THE ALGORITHM

Computing $\tilde{\tau}$ (in lines 1 to 5) takes $O(m)$ operations. Fitting the sizes, in line 6, also takes $O(m)$ operations. The rest of the loop consists in either adding (lines 11 to 19) or removing (lines 21 to 29) devices until reaching the correct number of devices n. Each greedy insertion or removal consists in a replacement in an ordered list of size m and a direct access, and takes a time $O(\log m)$. The number of those steps is bounded by $\frac{m}{2}$ in the worst case, and $\frac{\sqrt{m}}{12}$ in average, for a complexity in $O(m \log m)$ and $O(\sqrt{m} \log m)$.

Without a heuristic to obtain or approach m, assuming m to be in the same order of n and using a (weighted) dichotomy, the complexity cost becomes $O(n \log^2 n)$ in the worst case, and $O(n \log n)$ in average (building $\tilde{\tau}$ and the initial solution become the most costly operation).

An example to show the algorithm general behavior is portrayed in Figure 4 where $\tau(n) = \mathcal{H}_n^{\mathbf{w}}$ and $\tilde{\tau}(n)$ are depicted for $n = 2000$ and for $\mathbf{w}$ chosen uniformly between 1 and 2. The two values are almost indistinguishable, the highest value of their difference on this experiment is 8 while the average is slightly higher than 1. In Figure 5, the total cost is displayed as a function of m for the same simulation as in Figure 4 for $\beta = 400$. Furthermore, we added the total cost with no cooperation among devices as the latter perform on device learning (fully distributed and denoted as $C(n)$) and that corresponding to the grand coalition (centralized and denoted as $C(n)$). Those two schemes are representative of the state-of-the-art approaches, and clearly render worst performances than our devised scheme.

C. SEMI-DISTRIBUTED CONTEXT

In a semi-distributed context, we assume that a device can obtain, through signaling messages with a central entity, the information about the number of devices in play as well as its own weight, and can identify or inquire about the list of leaders, their weights and the size of their corresponding

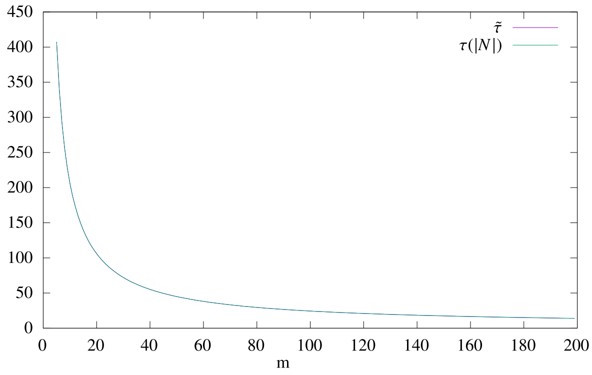

FIGURE 4. Discrepancy between $\tau(n) = \mathcal{H}_n^{\mathbf{w}}$ and $\tilde{\tau}(n)$ as a function of m.

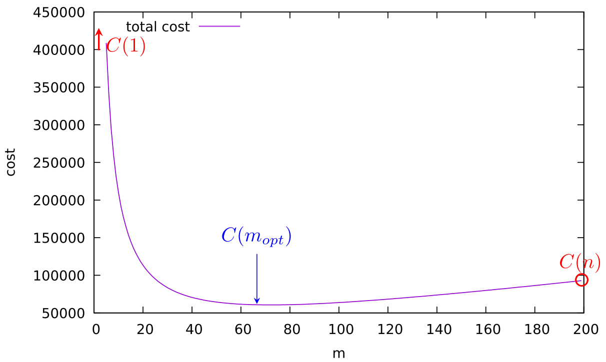

FIGURE 5. Total cost of the optimal solutions as a function of m for $\beta = 400$.

coalition, but not the weight of the other (non-leader) devices nor their specific repartition.

With those conditions, a player currently serving as the leader of a non trivial coalition and deciding to quit being a leader, does not know which member of the coalition has the second lowest weight and thus cannot predict the cost for the coalition members. In other words, we cannot build a set of utilities as in (5) such that the total cost in (13) is a potential function.

However, we can still solve the problem as a dynamic by separating it into two parts: the choice of the leaders (and hence the forming of coalitions headed by those leaders) and the attribution of non-leader devices to the formed coalitions.

We have from item 3 of Lemma 1 the monotony of the total cost in the value of the weights $\mathbf{w}$. This implies that the coalition leaders of an optimal solution are the devices of smallest weights, giving us a first criterion to identify if a device D is entitled to be a leader or not. Accordingly, we define what follows:

Criterion 1: in an optimal solution, if there is a leader whose weight is greater than the weight of device D, then D is a leader as well.

From Theorem 2, equation (14), we can compute the difference in total cost of the best solution when adding a new leader, giving us a second criterion:

Criterion 2: In an optimal solution minimizing first the total cost and then the number of leaders, device D is a leader if and only if $\sum_{i=1}^{d(m+1)} \left(\mathcal{H}_{n-i+1}^{\mathbf{w}\cup D}(m) - i \cdot w_D\right) > \beta$ with $\mathbf{w}$ the list of weights of leaders other than device D.

These two criterions are sufficient to characterize the set of leaders minimizing the total cost and can be computed autonomously by the devices despite their access to limited information.

Then, in parallel of that evolution, the devices can also select the coalition they choose to join according to a best response dynamics using the marginal cost as utility, behaving as with fixed leaders.

After convergence of the set of leaders, this second dynamics is a potential game which will converge in turn.

We obtain the following algorithms: Algorithm 3 simply emulates the semi-distributed dynamics where devices have access to limited information and call for Algorithm 4 and Algorithm 5.

Algorithm 5 uses the criterions to constitute the list of leaders, and Algorithm 4 attributes the devices to those leaders' coalitions using best response dynamics.

1) CORRECTNESS OF ALGORITHM 3

We provide the correctness of Algorithm 3 by showing the correctness of the two embedded Algorithms 4 and 5.

a: CORRECTNESS OF THE DYNAMICS ON THE LIST OF LEADERS

By construction, the list $1 \cdots m_{opt}$ of the smallest weights is stable by both criterions: trivially for Criterion 1, and through the monotonicity of the total cost in the values of the weights for Criterion 2.

Let's consider another list of leaders noted ℓ. We can build a path from ℓ to $1 \cdots m_{opt}$: by calling all elements that are in $1 \cdots m_{opt}$ but not in ℓ in increasing order, then by calling all elements that are in ℓ but not in $1 \cdots m_{opt}$ in decreasing order.

The devices called will apply Criterion 1, until either m_{opt} or the maximum of ℓ is reached, and add themselves to the list. Then, the list of leaders will be either a super-set or a subset of the optimal list, with elements called from the farthest from the equilibrium point to the closest. Using the convexity of the total cost on $\mathbf{w}$ as a set, we can see that Criterion 2 will add (if subset) or remove (if super-set) those elements, until the list of leaders is the optimal list $1 \cdots m_{opt}$: from any state, there is a path of nonzero probability (specifically $n^{-|\Delta(\ell, 1\cdots m)|} \geq n^{-n} > 0$) to the only stable state, thus the dynamics on the list of leaders converges to that stable state.

b: CORRECTNESS OF THE DYNAMICS OF DEVICES ATTRIBUTION TO COALITIONS

Once the list of leaders has converged, it remains fixed. The next step is to attribute devices to those leaders. When

Algorithm 3 GFL as a Dynamics

```
input      : List of weights w
/* This part emulates the context of
   a dynamics where devices have
   limited access and communications
   */
Variables : sorted list of leader ℓ, state E
repeat
    Construct w_t := w_{ℓ_1}, w_{ℓ_2} ··· w_{ℓ_{|ℓ|}}
    Device D is activated
    j_0 is the current leader of D
    bool isDLeader = D.call(Algo 4 Leader Selection
      (w_t, w_D, n ))
    if j_0 = D // currently a leader
     then
        if not isDLeader then
            /* D is removed from the
               leaders ℓ                        */
            Remove D from ℓ
            w_t is the weight of current leaders ℓ
            X is the sizes of the coalitions of current
              leaders ℓ
            for i ∈ E_D /* including D          */
             do
                k ← i.call(Algo 4 Best
                  Response(w_t,X))
                E_{ℓ_k} ← E_{ℓ_k} ∪ {i}
            end
            E_D ← ∅
        end
    else
        E_{j_0} ← E_{j_0} \ D
        if isDLeader then
            /* D is added to the leaders
               ℓ                                */
            E_D ← {D}
            insert D in ℓ
        else
            /* D moves between
               coalitions                       */
            w_t is the weight of the current leaders ℓ
            X is the sizes of the coalitions of the
              current leaders ℓ
            k ← D.call( Algo 4 Best Response(w_t,X))
            E_{ℓ_k} ← E_{ℓ_k} ∪ {D}
        end
    end
until convergence
```

the leaders are fixed, the game of the attribution of devices where devices use the marginal costs as utilities becomes a potential game admitting the total cost as a potential. A best response dynamics on that game converges to the unique Nash equilibrium, which is the actual optimum for the list of leaders, and in turn, for the global problem.

Algorithm 4 Best Response, Device Attribution for an Unassociated Non Leader Device, Used in Algo 3

```
Input     : Current ordered list of the weight w of the
            leaders, current list of size X
Output    : Index of joined coalition
for j leaders do
    Compute marginal cost u_j ← 2X_j w_j
end
return argmin_j(u_j)
```

Algorithm 5 Leader Selection, Used in Algo 3

```
Input      : current ordered list of leaders' weight w,
             own weight v, number of devices n
Variables : list of leader weight w_p, previous set H^p
             from the last call on this device
Output     : boolean isLeader
if v < w.last // criterion 1
 then
    return true
else
    if w_p ≠ w then
        Build H^{w∪v}_{1...n}
        H^p ← H^{w∪v}_{1...n}
        w_p ← w
    end
    Compute d ← ⌊ H^p_n / w_D ⌋
    Compute the cost differential
      ΔC ← Σ_{i=1}^{d} (H^p_{n-i+1}(m)) − (d(d−1)/2) w_D − β
    return ΔC > 0 // criterion 2
end
```

2) COMPLEXITY OF ALGORITHM 3

We provide the complexity analysis of Algorithm 3 by analyzing that of the two embedded algorithms 4 and 5.

a: CONVERGENCE OF THE LIST OF LEADERS

Starting from all devices being leaders (or from the list of leaders being all devices with weights lower than a sufficiently high value), the only device that can change its status is the one with the greatest weight, removing itself from the list and reducing the size of the leaders set by one. This preserves the property that the set of leaders being the ones of lowest values and Criterion 2 ensures that the size of the list remains greater or equal to m_{opt}.

This repeats until reaching m_{opt}, in at most an average of $(n - m_{opt}) \cdot n$ calls.

Conversely, starting from an arbitrary list will trigger calls to devices having weights lower than the greatest weight among leaders and adding themselves to the list of leaders until having the previous structure, in a variant of the coupon collector [13] on the missing elements, so in $O(n \log n)$ calls.

This bounds the expectancy of the number of calls needed for the list of leaders to converge to $O(n^2)$, so a total complexity of $O(n^2)$ in terms of communication and $O(n^3 \log n)$ in terms of computation, mostly from building $\mathcal{H}$ in the Leader Selection algorithm (Algorithm 3).

The computation cost could be made lower by keeping $\mathbf{w}$ and $\mathcal{H}$ in memory for each device and simply updating $\mathcal{H}$ when needed.

b: ATTRIBUTION OF DEVICES TO FORMED COALITIONS

This part is characterized by the fact that the branches in line 8 and line 21 are not taken, but the Leader Selection procedure is still called.

To show that this dynamics converges, we can compare the coalitions' sizes with those characterizing the closest optimal solution and use the sum of the positive part of the difference $d(X, X^{opt}) = \sum_{i\in\ell} \max\left(X_i - X_i^{opt}, 0\right) = \frac{||X-X^{opt}||_1}{2}$ as a distance. Specifically, we consider the smallest distance between such configuration and the set of optimal solutions, which is always positive and only reaches 0 on one of the optimal configurations.

We consider a step of the dynamics, and the intermediary configuration obtained after separating the device but before reinserting it to the coalition corresponding to its best response. This intermediary configuration only has $n-1$ devices, one less than the closest optimal solution which has n devices. Thus, it has at least one coalition with a size strictly lower than its size in the closest optimal solution.

If we recall Lemma 1, we can see that the marginal cost of the coalitions is monotonic in their sizes and that the optimal solutions have consecutive values of $\mathcal{H}$ as marginal costs, thus a marginal cost of a coalition of the current solution is strictly lower than the marginal costs of all coalitions in any optimal solution if and only if the corresponding size is strictly lower than that of the current solution.

Moreover, the best response dynamics chooses the coalition of smallest marginal cost. Since there is a coalition whose size is strictly lower than its size in the optimal solution, then there is a coalition of marginal cost below the smallest marginal cost of the optimal solution. In turn, the coalition chosen by the best response has a marginal cost strictly lower than this threshold, thus a size strictly lower than the corresponding size in the closest optimal solution. This means that if the device was chosen from a coalition with too many elements, the number of devices in excess decreases by 1, else, it remains unchanged.

We can thus bound the number of devices in excess both by the number of devices in coalitions that have a too large size, and by the total number of devices. The distance to the closest optimal solution during the dynamics can be treated as a Markov chain. The transitions between states are either to keep the same value or to decrease it by 1. The probability of the latter is lower bounded by the proportion of devices in coalitions whose size is greater in the current configuration than in the closest optimum $\frac{\sum_{X_i \geq X_i^{(opt)}} X_i}{n}$. This is itself greater than the proportion of devices in excess, which is the distance divided by the number of devices $\frac{d}{n}$.

As a state of the chain is a number of devices, so at most equal to n. The solution is optimal when the Markov chain state reaches 0. Thus, we can use the Markov chain to bound the number of expected calls by the chain convergence time. The expected convergence time can be obtained through the Poisson equations, writing $T(i)$ the expected time to convergence from state i, with $T(0) = 0$, we have what follows: $T(i) = 1 + P(i \to i-1)T(i-1) + (1 - P(i \to i-1))T(i)$, $T(i) = \frac{1}{(1-P(i\to i-1))} + T(i-1) = \sum \frac{1}{(1-P(i\to i-1))}$. Substituting the probability above, we obtain $\sum\limits_{0\leq i\leq n-1} \frac{n}{n-i} = \sum \frac{n}{i} = H_n = n \log n + \gamma + o(1) = O(n \log n)$.

c: TOTAL COMPLEXITY

This gives us a total of $O(n^2)$ calls during the first phase, which corresponds to the time between the start of Algorithm 3 and the convergence of the list of leaders, and $O(n \log n)$ calls during the second phase, from the convergence of the list of leaders to the termination of Algorithm 3. Each call requires signaling to inform the device about the actual configuration and the computation cost of building or updating $\mathcal{H}$, so the complexity of the communications would be $O(n^2)$, most of which comes from the first phase. In fact, during the first phase, the devices need to rebuild $\mathcal{H}$ fully in most steps, taking a time $O(n \log m)$ in each step. During the second phase, the list of leaders remains unchanged, allowing devices to avoid those computations, and reducing the cost to $O(m)$ per step.

This gives us a global complexity of $O(n^3 \log n)$.

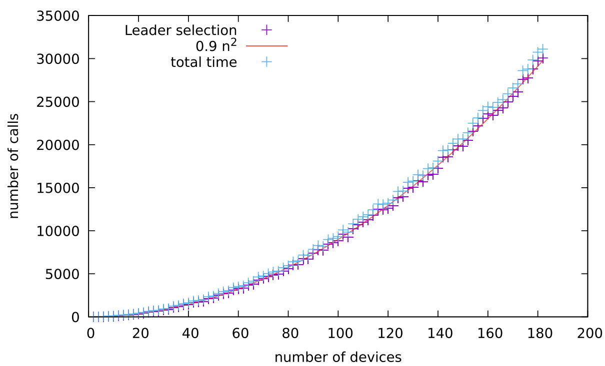

FIGURE 6. Number of calls to convergence of the list of leaders algorithm and of the attribution of devices algorithm through Algorithm 3.

In Figure 6, we depict the number of calls necessary to Algorithms 4 and 5 before the global Algorithm 3 converges. The dynamics was emulated by sequential and uniform activation of devices. The weight were chosen uniformly between 1 and 3. The additional line (in red) is a regression of the convergence time of the list of leaders, confirming that the proposed bound is very tight.

VI. CONCLUSION

Federated learning enables devices to learn collaboratively from information collected from all devices without sharing the original data, addressing privacy concerns and reducing the need for hefty data transfer to a central server. Learning on larger datasets reduces the variance in the learned model, and in turn its error. However, federating the learning process inflects a communication cost among learning devices that must be taken into account. Therefore, autonomous IoT devices engage in a non cooperative game to form learning coalitions to reduce both the learning error and the communication cost. In this paper, the optimal size of learning coalitions is computed in a realistic setting where devices are supposed to be heterogeneous with regard to their ability to federate the learning process. Various algorithms are designed to form the optimal sized learning coalitions in a semi-distributed fashion and their complexity duly assessed. In future work, we intend to apply the proposed framework to more complex learning models.

We provide in the appendix the proofs for both the comparison parenthesis, the proof for Theorem 3 and the proof for Lemma 2.

APPENDIX. PROOF OF THE COMPLEXITY LOWER BOUND PARENTHESIS

For that, recall the proof of Theorem 2, mainly, the expression of the difference in total cost (14):

$$C(m) - C(m+1) = \sum_{i=1}^{d(m+1)} \left(\mathcal{H}_{n-i+1}(m) - i \cdot w_{m+1}\right) - \beta$$

Replacing d by its definition, and bounding the elements $\mathcal{H}_{n-d}$ to $\mathcal{H}_n$ by $\mathcal{H}_n$, we get:

$$\begin{aligned} &C(m) - C(m+1) \\ &\leq d(m)\mathcal{H}_n - w_{m+1}\frac{d(m)\cdot(d(m)+1)}{2} - \beta \\ &\leq \frac{\mathcal{H}_n}{w_{m+1}}\mathcal{H}_n - \left(\frac{\mathcal{H}_n^2}{2w_{m+1}}\right) - \beta = \frac{\mathcal{H}_n^2}{2w_{m+1}} - \beta \\ &\leq \frac{\left(n-\frac{m}{2}\right)^2}{2w_{m+1}\left(\sum_{j\leq m+1}\frac{1}{w_j}\right)^2} - \beta \end{aligned}$$

As this expression is a (reversed) difference of successive values in a function we know to be convex, we can express a lower bound of m_{opt} as the 0 of that expression. Using the increasing order on the w_j to lower bound the denominator $\sum_{j\leq m+1}\frac{1}{w_j}$ by $\frac{m}{w_{m+1}}$ and $\frac{m}{\bar{w}}$, upper bounding the expression and in turn lower bounding its zero again, this gives us:

$$C(m) - C(m+1) \leq \frac{\left(n-\frac{m}{2}\right)^2}{m^2}\bar{w} - \beta \quad (15)$$

$$\Rightarrow m_{opt} \geq n\left(\sqrt{\frac{\beta}{\bar{w}}} - \frac{1}{2}\right) = \Theta(n) \quad (16)$$

APPENDIX. PROOF OF THEOREM 2

Proof: Let's take the number of coalitions m fixed ($w_1 \cdots w_m$ are also fixed). We define E_n an optimal solution for n devices and F^{n+1} an optimal solution for $n+1$ devices. We will show that there exists $j \in [1, m]$ such that adding the $n+1^{\text{th}}$ device to E^n gives a solution at least as good as F^{n+1}

First, since F is a partition of a greater set than E, there exists j such that $|E_j^n| < |F_j^{n+1}|$. Since E_j^n is a non trivial set, then F_j^{n+1} contains at least 2 elements, so removing one does not modify the number of coalitions in the solution. Hence, removing one element from this coalition form a valid partition F^n of n devices in exactly m coalitions with weights $w_1 \cdots w_m$. Conversely, adding the $n+1^{\text{th}}$ device to E_j creates a valid partition E^{n+1} of the $n+1$ devices over the same weights $w_1 \cdots w_m$.

E^n is optimal for n, $(w_i)_i$, so in particular, it is at least as good as F^n: $C(F^n) - C(E^n) \geq 0$.

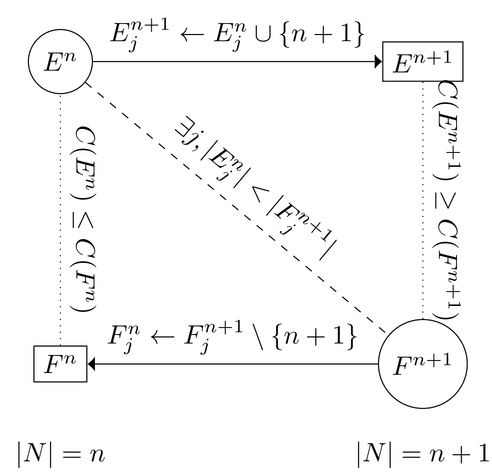

FIGURE 7. Relation between E^n, F^n, E^{n+1} and F^{n+1}: E^n and F^{n+1} are optimal solutions, E^{n+1} and F^n are built from them.

Recall that $E_{k\neq j}^{n+1} = E_k^n$, $E_j^{n+1} = E_j^n \cup \{n+1\}$, $F_{k\neq j}^{n+1} = F_k^n$ and $F_j^{n+1} = F_j^n \cup \{n+1\}$, accordingly, we have what follows:

$$\begin{aligned} C(F^n) - C(E^n) &= \sum_{i=1}^{m} w_i|F_i^n|\left(|F_i^n| - 1\right) \\ &\quad - \sum_{i=1}^{m} w_i|E_i^n|\left(|E_i^n| - 1\right) \\ C(F^n) - C(E^n) &= \sum_{i\leq m, i\neq j} w_i|F_i^{n+1}|\left(|F_i^{n+1}| - 1\right) \\ &\quad + w_j\left(|F_j^{n+1}| - 1\right)\left(|F_j^{n+1}| - 2\right) \\ &\quad - \sum_{i=1}^{m} w_i|E_i^n|\left(|E_i^n| - 1\right) \\ C(F^n) - C(E^n) &= \sum_{i=1}^{m} w_i|F_i^{n+1}|\left(|F_i^{n+1}| - 1\right) \end{aligned}$$

$$- 2w_j \left(|F_j^{n+1}| - 1\right)$$
$$- \sum_{i=1}^{m} w_i|E_i^n| \left(|E_i^n| - 1\right)$$

Now, we write the same difference between the solutions for $n+1$, since F^{n+1} is the optimal. This difference is negative or zero $C(F^{n+1}) - C(E^{n+1}) \leq 0$ and is equal to what follows:

$$C(F^{n+1}) - C(E^{n+1}) = \sum_{i=1}^{m} w_i \left(|F_i^{n+1}| - 1\right)^2 - \sum_{i=1}^{m} w_i \left(|E_i^{n+1}| - 1\right)^2$$

$$C(F^{n+1}) - C(E^{n+1}) = \sum_{i=1}^{m} w_i \left(|F_i^{n+1}| - 1\right)^2 - \sum_{i\leq m, i\neq j}^{m} w_i \left(|E_i^n| - 1\right)^2 + w_j \left(|E_j^n|\right)^2$$

$$C(F^{n+1}) - C(E^{n+1}) = \sum_{i=1}^{m} w_i \left(|F_i^{n+1}| - 1\right)^2 - \sum_{i=1}^{m} w_i \left(|E_i^n| - 1\right)^2 + 2|E_j^n| - 1$$

We can notice that the difference is the same in both expressions, so

$$\left(C(F^{n+1}) - C(E^{n+1})\right) - \left(C(F^n) - C(E^n)\right) = 2\left(|F_j^{n+1}| - |E_j^n| - 1\right)$$

Since we choose j such that $E_j^n < F_j^{n+1}$, then the expression $2\left(|F_j^{n+1}| - |E_j^n| - 1\right)$ is positive or 0. We have then that the sum of two negative or null terms is positive or zero, thus both are null. We deduce that E^{n+1} is as good as the optimal solution F^{n+1}, and is thus optimal itself. This proves that extending an optimal solution optimally gives an optimal solution.

Similarly, F^n is as good as the optimal solution E^n. We deduce that for any optimal solution for $n+1$ devices, there exists an optimal solution for n devices that can be extended into the former. ■

APPENDIX. PROOF OF LEMMA 1

Proof: Recall the definition of the fitted size in Definition 2: each size X_i takes the value (or pair of values) $\left[\frac{\tilde{\tau}}{w_i} + \frac{1}{2}\right]$, which can be rewritten as $\frac{\tilde{\tau}}{w_i} + \frac{1}{2} + r_i$ with r_i being the rounding error, bounded in absolute value by $\frac{1}{2}$.

Summing the coalition sizes and separating the r_i gives what follows: $\sum x_i = \sum \frac{\tilde{\tau}}{w_i} + \sum \frac{1}{2} + \sum r_i = \left(\sum \frac{1}{w_i}\right)\tilde{\tau} + \frac{m}{2} + \sum r_i = \left(\sum \frac{1}{w_i}\right)\frac{n - \frac{m}{2}}{\sum \frac{1}{w_i}} + \frac{m}{2} + \sum r_i = n + \sum r_i$. Thus, summing the actual value and comparing to n gives the sum of the rounding errors r_i, each bounded by $\pm\frac{1}{2}$. The sum of the bounds over the coalitions bound the sum of the r_i, so the difference between the sum of the sizes and n is in the interval $\left[-\frac{m}{2}, \frac{m}{2}\right]$. ■

Lemma 3: Average Error: The average error of the effective number of devices in a solution with fixed m built using $\tilde{\tau}$ is $\frac{\sqrt{m}}{12}$.

Proof: In order to define the average error, we replace $\tilde{\tau}$ by a list of independent random variables chosen uniformly in an interval whose size is a multiple of the weight of the corresponding coalition: It is mostly equivalent to choosing the inverse of the weights uniformly while keeping their actual values for the final expression, and it allows us to see the rounding error as a uniformly chosen variable in $\left]-\frac{1}{2}, \frac{1}{2}\right[$.

Thus, the error caused by the rounding is a sum of independent uniform laws of same amplitude, an Irvin Hall distribution of parameter m with an offset of $\frac{m}{2}$. Its variance is $\frac{m}{12}$, and thus the average error is $\frac{\sqrt{m}}{2\sqrt{3}}$. ■

REFERENCES

[1] K. Khawam, S. Lahoud, C. Adjih, S. Makhoul, R. Al Tawil, and S. Martin, "Edge learning as a hedonic game in LoRaWAN," in *Proc. IEEE Int. Conf. Commun.*, Jun. 2023, pp. 3163–3168.
[2] J. Hajduková, "Coalition formation games: A survey," *Int. Game Theory Rev.*, vol. 8, no. 4, pp. 613–641, Dec. 2006.
[3] L. U. Khan, S. R. Pandey, N. H. Tran, W. Saad, Z. Han, M. N. H. Nguyen, and C. S. Hong, "Federated learning for edge networks: Resource optimization and incentive mechanism," *IEEE Commun. Mag.*, vol. 58, no. 10, pp. 88–93, Oct. 2020.
[4] R. Gupta and J. Gupta, "Federated learning using game strategies: State-of-the-art and future trends," *Comput. Netw.*, vol. 225, Apr. 2023, Art. no. 109650.
[5] J. Weng, J. Weng, H. Huang, C. Cai, and C. Wang, "FedServing: A federated prediction serving framework based on incentive mechanism," in *Proc. IEEE Conf. Comput. Commun.*, May 2021, pp. 1–10.
[6] E. Tahanian, M. Amouei, H. Fateh, and M. Rezvani, "A game-theoretic approach for robust federated learning," *Int. J. Eng.*, vol. 34, no. 4, pp. 1–12, Oct. 2021.
[7] Z. A. E. Houda, B. Brik, A. Ksentini, L. Khoukhi, and M. Guizani, "When federated learning meets game theory: A cooperative framework to secure IIoT applications on edge computing," *IEEE Trans. Ind. Informat.*, vol. 18, no. 11, pp. 7988–7997, Nov. 2022.
[8] Z. Xu, D. Li, W. Liang, W. Xu, Q. Xia, P. Zhou, O. F. Rana, and H. Li, "Energy or accuracy? Near-optimal user selection and aggregator placement for federated learning in MEC," *IEEE Trans. Mobile Comput.*, vol. 23, no. 3, pp. 1–15, Sep. 2023.
[9] K. Donahue and J. Kleinberg, "Model-sharing games: Analyzing federated learning under voluntary participation," in *Proc. AAAI Conf. Artif. Intell.*, vol. 35, 2021, pp. 1–17.
[10] C. Hasan, "Incentive mechanism design for federated learning: Hedonic game approach," 2021, *arXiv:2101.09673*.
[11] E. M. Migabo, K. D. Djouani, and A. M. Kurien, "The narrowband Internet of Things (NB-IoT) resources management performance state of art, challenges, and opportunities," *IEEE Access*, vol. 8, pp. 97658–97675, 2020.
[12] S. Moloudi, M. Mozaffari, S. N. K. Veedu, K. Kittichokechai, Y.-P. E. Wang, J. Bergman, and A. Höglund, "Coverage evaluation for 5G reduced capability new radio (NR-RedCap)," *IEEE Access*, vol. 9, pp. 45055–45067, 2021.
[13] L. E. Baum and P. Billingsley, "Asymptotic distributions for the coupon collector's problem," *Ann. Math. Statist.*, vol. 36, no. 6, pp. 1835–1839, Dec. 1965.

STEPHANE DURAND received the master's degree from ENS Lyon, in 2015. He defended his thesis with the University of Grenoble in the domain of game theory, in 2018. He is currently a Research Assistant with LISN. His research interests include modelizations of systems to predict behavior, convergence, or cinetic.

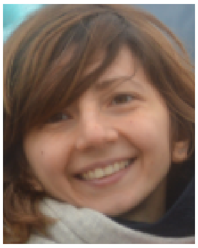

KINDA KHAWAM received the B.E. degree in telecommunications and networking engineering from the Saint Joseph University of Beirut, in 2001, and the M.E. and Ph.D. degrees in computer science from Telecom ParisTech, in 2003 and 2006, respectively. She is currently an Associate Professor with the Université de Versailles Saint-Quentin-en-Yvelines. Her research interests include distributed and centralized radio resource management, the green networking and IoT, 5G wireless networks, game theory, machine learning, and optimization.

DOMINIQUE QUADRI is currently a Full Professor with the LISN Laboratory, Paris Saclay University. Her research interest includes the alliance between non-linear mathematical programming and networks.

SAMER LAHOUD received the Ph.D. degree in computer science and networks from IMT Atlantique, Rennes, and the Habilitation degree from the University of Paris-Saclay, France. He is currently an Associate Professor with the Faculty of Computer Science, Dalhousie University. Prior to joining Dalhousie, in 2023, he was an Associate Professor with the University of Rennes and has been a Researcher with the IRISA Laboratory, France, since 2007. He was also a Visiting Scholar with Saint Joseph University, Beirut, Lebanon, from 2016 to 2022. Additionally, he gained experience as a Research Engineer with the Nokia Bell Laboratories Europe, in 2006 and 2007. His research interest includes optimizing radio resource allocation in wireless networks.

STEVEN MARTIN received the Engineering degree, in 2000, and the Ph.D. degree in France, in 2004. He is currently a Professor with Université Paris Saclay. His research interests include wireless access technologies for the Internet of Things, user association and radio access technology selection, resource allocation and interference mitigation, and energy efficiency. He has carried out projects involving both academia and industry at national and international level.