Uplink Scheduling in Federated Learning: an Importance-Aware Approach via Graph Representation Learning

Marco Skocaj[1], Pedro Enrique Iturria Rivera[2], Roberto Verdone[1] and Melike Erol-Kantarci[2]

[1]University of Bologna, *DEI* & WiLab - National Laboratory for Wireless Communications, *CNIT*, Italy
[2]University of Ottawa, *School of Electrical Engineering and Computer Science*, Canada

Abstract—Federated Learning (FL) has emerged as a promising framework for distributed training of AI-based services, applications, and network procedures in 6G. One of the major challenges affecting the performance and efficiency of 6G wireless FL systems is the massive scheduling of user devices over resource-constrained channels. In this work, we argue that the uplink scheduling of FL client devices is a problem with a rich relational structure. To address this challenge, we propose a novel, energy-efficient, and importance-aware metric for client scheduling in FL applications by leveraging Unsupervised Graph Representation Learning (UGRL). Our proposed approach introduces a relational inductive bias in the scheduling process and does not require the collection of training feedback information from client devices, unlike state-of-the-art importance-aware mechanisms. We evaluate our proposed solution against baseline scheduling algorithms based on recently proposed metrics in the literature. Results show that, when considering scenarios of nodes exhibiting spatial relations, our approach can achieve an average gain of up to 10% in model accuracy and up to 17 times in energy efficiency compared to state-of-the-art importance-aware policies.

Index Terms—Federated Learning, Graph Representation Learning, Scheduling, Communication-efficient FL, Energy-efficient FL, 6G, Spatial Correlation

I. INTRODUCTION & MOTIVATION

Federated Learning (FL) [1] recently emerged as a new privacy-preserving paradigm for distributed training of Machine Learning (ML) algorithms without the need for explicit data sharing between users and a centralized computational unity. This framework is particularly appealing for next-generation 6G systems, which are foreseen to support ubiquitous Artificial Intelligence (AI) services and AI-native design of users' network procedures [2]. Indeed, users of a 6G network will naturally benefit from decentralized training that bypasses sharing and storing data in a centralized location. This will lead to the support of a new kind of AI-related traffic over wireless networks: frequent exchange of ML models introduces significant communication overhead, which raises a series of interesting novel challenges. This paved the way to a recent research area on communication-efficient FL for 6G [3]. Wireless FL is an example of goal-oriented communication [4], for which traditional Radio Resource Management (RRM) methods are typically inadequate, and customized protocols must be developed [5].

Arguably, one of the major challenges towards scalable and efficient 6G wireless FL systems is the massive scheduling of user devices, from now on referred to as *clients*. In fact, the central aggregation unity, namely the *Parameter Server (PS)*, generally needs to link a vast number of User Equipment (UE)s through a resource-constrained spectrum and thus can allow only a limited number of UEs to send their trained weights via unreliable channels for global aggregation [6]. To this end, the concept of *data importance*, or *importance-aware communications* in FL has taken hold in the recent literature [7]–[13]. The main idea is that by prioritizing users with high data importance, the distributed ML training is accelerated [13]. Because explicit information on clients' data is infeasible due to privacy concerns, state-of-the-art importance-aware approaches rely on feedback information from the training of local clients' models. This approach, even though proven to be effective, has the major disadvantage of being energy and computationally inefficient. In fact, feedback-based importance-aware methods lead to the training of local models on all clients, regardless of the number of scheduled transmissions in the next communication round (Fig. 1).
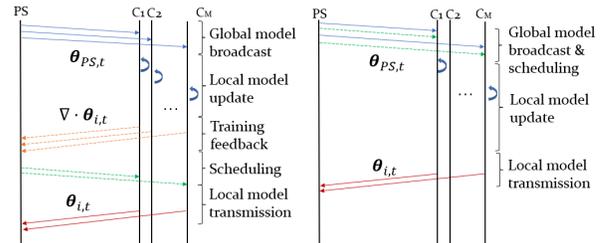
Fig. 1: Comparison between importance-aware scheduling protocols in FL. On the left side (a), classic energy-inefficient approaches require training feedback from all network clients every communication round. As the number of clients in a network grows, this becomes highly inefficient from the computational and energetic point of view. On the right side (b), importance-aware scheduling via UGRL requires only scheduled users to perform local training.

In this work, we propose a novel energy-efficient, importance-aware FL metric based on graph representation learning, which leads to effective scheduling of client devices without any need for collecting training feedback information.

A. State of the Art & Contributions

Over the last years, different metrics for FL client scheduling have been proposed and discussed. In [14], by leveraging the concept of Age of Information (AoI), a metric termed age of update (AoU) is introduced, which takes into account the staleness of the received parameters. In [10], [11], [14], [15] channel conditions experienced by different clients are considered during the scheduling decision. In [15], a channel prediction algorithm based on Gaussian process regression is incorporated into the scheduling process when dealing with imperfect channel state information. Authors from [11], instead, exploit both diversity in multiuser channels and diversity in the importance of the edge devices' local learning updates. Importance, in this case, is measured by the local parameter update's gradient divergence, which must be reported to the PS downstream of the training of all client devices. A similar training-feedback metric of importance is introduced in [10], where the significance of the model updates at the devices is captured by the L2-norm of the model update.

Under the frequentist setting, training data constitute a fundamental part of the inductive bias of a model. In this context, in line with the idea of importance-aware communications, knowledge about data distribution among devices would suffice for driving proper FL scheduling decisions. However, as previously stated, this approach is unfeasible in FL systems. Nevertheless, clients exhibit relations and correlations in a network setting, especially in the context of massive Internet-of-Things (IoT) and ultra-dense 6G networks. Here, we argue that the scheduling of FL client devices is a problem with a rich relational structure and, as a consequence, there is a need to tackle this problem effectively by taking node correlations into account. Relations among clients, which relate to local data distribution too, can be learned and inferred by encompassing network geometry and relational representation learning, while at the same time preserving users' privacy. Graphs, generally, are a representation that supports arbitrary relational structures, and computations over graphs afford a strong relational inductive bias [16]. By considering networks of clients as graphs, we introduce this bias in the clients' scheduling process by leveraging Unsupervised Graph Representation Learning (UGRL). As results show, this effectively makes up for the impossibility of selecting users based on their data, while at the same time aiming for an energy and computationally-efficient scheduling protocol.

The main contributions of this work are listed hereafter:

- We consider network geometry in the form of node embeddings obtained via UGRL as a fundamental new metric for driving efficient FL scheduling decisions in the context of non-i.i.d. and spatially correlated data. We aim to show it is possible to make up for the absence of explicit knowledge information about clients' data by introducing a relational inductive bias into the scheduling process.
- We compare the performance of the proposed scheduling metric with respect to baseline metrics recently proposed in the literature.
- We discuss the range of applicability of our proposed solution with respect to different kinds of data distributions with application to IoT and 6G networks.

II. SYSTEM MODEL

A. Network Scenario and Propagation Channel

Let us consider a system comprised of one Access Point (AP) co-located with a PS and multiple client devices with local data and computation capabilities, as depicted in Fig. 2. Physically running on the PS, there are two software entities responsible for model aggregation and radio resource management: the *Federated Aggregator* and the *Graph Scheduler*, respectively described in the following subsections.
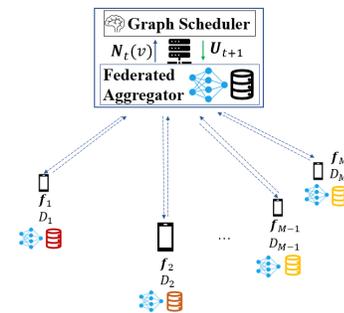
Fig. 2: System Model

Each client $m \in \mathcal{M}$ holds a local data set $D_m = \{\mathbf{x}_m \in \mathcal{R}^d, \mathbf{y}_m \in \mathcal{R}\}$ with cardinality $|D_m|$, such that $\sum_{m \in \mathcal{M}} |D_m| = |D|$, and is equipped with a single isotropic antenna. On the PS side, we consider an antenna with a directive gain of 15 dBi. All clients are randomly uniformly distributed within a radius R of the PS. The PS broadcasts the global model to the selected clients with a transmit power of 15 dBm, while the latter send their local updates with a transmit power of 10 dBm.

As further detailed in the next subsections, we consider the two cases of non-i.i.d. spatially correlated data and spatially correlated/uncorrelated non-i.i.d clusters of i.i.d. data. Both cases are artificially reproduced in our experiments by spatially distributing MNIST digit labels among neighbor clients.

Within the network area, we consider an Orthogonal Frequency Division Multiple Access (OFDMA) scheme with perfect equalization, where $M = |\mathcal{M}|$ clients share the same spectrum and can be assigned one of the $K < M$ set of orthogonal sub-channels for model parameters transmission. Moreover, we assume a slow fading propagation model, where each model transmission from device $m \in \mathcal{M}$ to the PS is shorter than the channel coherence time T_c:

$$T_{D,m} < T_c \quad \text{for } m \in \mathcal{M}, \tag{1}$$

where $T_{D,m}$ is the model transmission time for client m. With the aforementioned assumptions, the channel impulse response $h_{m,PS}(f, t)$ from device m to the PS loses its time and

frequency dependency within a block transmission duration $T_{D,m}$ (2):

$$h_{m,PS}(f,t) \rightarrow h_{m,PS} \in \mathbf{H}_{M+1}, \qquad (2)$$

where $\mathbf{H}_{M+1}$ is the channel matrix of dimension $M+1$. Finally, we consider the Okumura-Hata model for the median path loss, and the Nakagami-m distribution for the fading propagation model, as it provides a flexible formulation to characterize Rician and Rayleigh fading.

B. Federated Learning Framework

The goal of the *Federated Aggregator* is to learn a ML model by offloading and aggregating the training to the set $\mathcal{M}$ of distributed clients with local data. The federated training process involves a number of iterations, namely communication rounds, until convergence. Each client $m \in \mathcal{M}$, upon receiving a global model $\boldsymbol{\theta}_{PS}$ from the PS at the beginning of a new round, executes multiple Stochastic Gradient Descent (SGD) updates to minimize the model's loss function with respect to its local dataset (3):

$$F_m(\boldsymbol{\theta}, D_m) = \frac{1}{|D_m|} \sum_{\{x_i, y_i\} \in D_m} f(\boldsymbol{\theta}, \{x_i, y_i\}), \qquad (3)$$

where F_m is the m-th client loss function and $f(\boldsymbol{\theta}, \{x_i, y_i\})$ indicates the task-dependent loss (e.g., Mean Square Error (MSE), categorical cross-entropy, etc.) for every training example $\{x_i, y_i\}$. At every j-th local SGD iteration, each client m updates its local model according to (4):

$$\boldsymbol{\theta}_m^{j+1}(t) = \boldsymbol{\theta}_{PS}^j(t) - \alpha(t) \cdot \mathbf{g}_m, \qquad (4)$$

where $\alpha(t)$ denotes the learning rate scheduled by the PS for communication round t and $\mathbf{g}_m := \nabla(F_m(\boldsymbol{\theta}_m, D_m)$ is the m-th client's gradient of the local model's weights. Once the training is terminated, each client selected for scheduling must forward its local model $\boldsymbol{\theta}_m(t)$ to the PS, which will update the global model upon aggregation of all received clients' models. Here, we refer to *FedAvg* algorithm [1], for which the aggregation is a weighted sum described by (5):

$$\boldsymbol{\theta}_{PS}(t+1) = \frac{1}{K} \sum_{m \in \mathcal{K}} \frac{|D_m|}{\sum_{m \in \mathcal{K}} |D_m|} \cdot \boldsymbol{\theta}_m(t), \qquad (5)$$

where we denote by $\mathcal{K}$ ($|\mathcal{K}| = K$) the set of scheduled clients for communication. For model evaluation, we consider a separate centralized test set $D_{test,PS}$ locally residing on the PS.

C. Graph Scheduler

The Graph Scheduler is the entity responsible for the procedures depicted in Fig. 3: dynamic graph creation, UGRL, and distance maximization.

- *Sensing phase and reporting*: a sensing phase is performed by all client devices, and it should be repeated with a periodicity that depends on environment dynamicity. Following, all nodes v report a list of sensed neighbors $N_t(v)$ (Fig. 2) to the graph scheduler.
- *Dynamic graph formation*: the scheduler builds a dynamic graph from each node's list of sensed neighbors. In
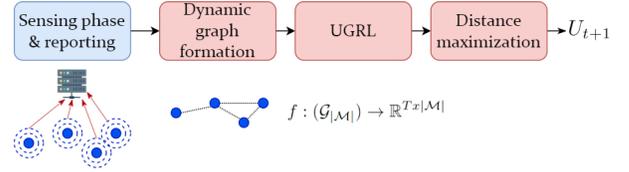
Fig. 3: Graph scheduler block scheme. The scheduling sequence $\mathbf{U}_{t+1}$ is computed downstream Distance Maximization scheduler and the UGRL block. The UGRL block is responsible for the transformation f from a graph space $\mathcal{G}_{|\mathcal{M}|}$ to the vectorial space $\mathbb{R}^{Tx|\mathcal{M}|}$ of node embeddings.

our approach, we retain the strongest-K neighbors, as it allows to adapt to the density of nodes' deployment. Note that alternative approaches, such as retaining adjacencies based on a received power threshold, might also be used.

- *UGRL*: graph representation learning involves the transformation via an encoding function $f : (\mathcal{G}_{|\mathcal{M}|}) \rightarrow \mathbb{R}^{Tx|\mathcal{M}|}$ from the graph-structured node representation $\mathcal{G}_{|\mathcal{M}|}$, to a vectorial space of T-dimensional node embeddings. In this procedure, node embeddings from $\mathcal{G}_{|\mathcal{M}|}$ are efficiently computed in an unsupervised way with the use of random walks procedures. In our scenario, we make use of Node2Vec algorithm [17]. Further details are discussed in section IV.
- *Distance maximization*: the final step involves the computation of the scheduling sequence $\mathbf{U}_{t+1}$. This is based on the distance maximization between node embeddings retained in a context window of tunable dimension. Additional details are provided in section IV.

D. Data Distribution

The intuition behind the proposed approach is that the clients' data distribution reflects the structural relation of nodes in a graph. Consequently, this method does not apply to the trivial case of i.i.d. data. Vice versa, it is possible to think of a plethora of applications and AI-driven network procedures foreseen for 6G in which this condition holds: localization, tracking, integrated sensing and communication, channel estimation and measures of a physical quantity from a sensors network is just a non-exhaustive list of examples where clients' data are non-i.i.d. and spatially correlated. Another vertical of great importance for future 6G networks is Industrial Internet-of-Things (IIoT), where typically nodes are arranged into clusters, and nodes within a cluster might hold similar kinds of measurements (e.g., monitoring sensors inside automatic machines in a warehouse). To this end, we consider in this work typical kinds of client data distributions that find use in many real-world 6G applications:

- Non-i.i.d. and spatially correlated data distribution (Fig. 4a).
- Spatially correlated/uncorrelated clusters of i.i.d. data (Fig. 4b, 4c).

For benchmarking purposes, we reproduced the three scenarios described above in our simulations by distributing MNIST data arranged by labels to a set of randomly distributed clients. According to the considered scenario, clients are distributed, at the beginning of every new simulation, a random
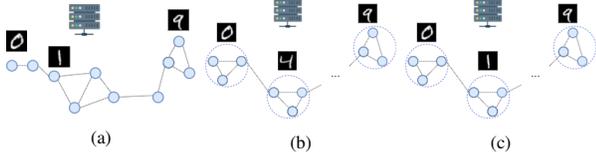
Fig. 4: Typical kinds of clients data distribution in AI-native network procedures and IoT: (a) Non-i.i.d, spatially correlated data distribution, (b) Non-i.i.d, spatially uncorrelated clusters of i.i.d. data and (c) Non-i.i.d, spatially correlated clusters of i.i.d. data

number of examples according to: a) their respective position to the PS (for the case depicted in Fig. 4a), or b) their belonging cluster (for the case depicted in Fig. 4b, 4c).

III. PROBLEM FORMULATION

The observation space can be represented as a graph composed of nodes (client devices), edges (adjacency matrix), and node features $\mathcal{F}$ (FL metrics). In the most general formulation, each edge can also be associated with a weight $|h_{i,j}|$, corresponding to the module of the complex channel impulse response $h_{i,j}$ between client i and j, drawn from a $(M+1)$-dimensional channel matrix $\mathbf{H}_{M+1}$. Nevertheless, considering the case of orthogonal resources assignment (i.e., no inter-users interference), and assuming link reciprocity, allows for a simplification of the problem formulation, since matrix $\mathbf{H}_{M+1}$ reduces to an M-dimensional vector $\mathbf{H} = [|h_{1,PS}|, \ldots, |h_{M,PS}|]$. Hence, its elements can be represented as node features instead of edges.

$$S_t = \begin{cases} \mathbf{F} = \big[(\mathcal{F}_1, |h_{1,PS}|), ..., \\ \qquad\quad (\mathcal{F}_M, |h_{M,PS}|)\big], \quad \text{with } |h_{i,PS}| \in \mathbb{R} \\[2ex] \mathbf{A} = \begin{bmatrix} a_{11} & \ldots & a_{1M} \\ \vdots & \ddots & \vdots \\ a_{M1} & \ldots & a_{MM} \end{bmatrix}_M, \quad \text{with } a_{ij} \in \{0,1\} \end{cases} \quad (6)$$

In (6), $\mathbf{F}$ is the feature space of every node, which includes the FL-metrics $\mathcal{F}$ used for decision-making during the scheduling procedure, and $\mathbf{A}$ is the adjacency matrix, which is obtained downstream of the dynamic graph formation block of Fig. 3. The feature space is composed of the following metrics:

- Age of Information (AoI): a scalar indicator, introduced in [14], describing the number of rounds that elapsed since the client was last scheduled for model transmission.
- Path loss $|h_{i,j}|$: the path loss value in dB related to the client-PS link, assuming link reciprocity.
- L2-norm of model update $||\boldsymbol{\theta}_{i,t+1} - \boldsymbol{\theta}_{PS,t}||_2$: an importance metric indicating the L2 norm of the i-th client model update.

Notably, no explicit informative content can be collected about the data of the clients, as this would violate the privacy-preserving nature of FL.

IV. PROPOSED ALGORITHM

In the formulation above, nodes connected by edges have similar data, but not necessarily similar FL metric features. In fact, AoI, for instance, does not show any spatial correlation property among nodes, as it purely depends on the scheduling mechanism. In turn, this, together with the inability to have explicit features about nodes' data, depicts a situation where client features don't necessarily reflect the structure of the graph. Moreover, the nature of the problem is naturally unsupervised, as nodes don't have any label. To this end, we make use of UGRL via random walks in place of common supervised methods with Graph Neural Network (GNN)s to incorporate information about the structure of the graph into the decision-making process. By employing Node2Vec [17], we are able to compute the graph encoding function $f(\mathcal{G})$ without the assistance of any node labels and features, but rather by maximizing the log-likelihood of the 2nd order biased random walks $N_s(u)$ conditioned by $f(u)$, for $u \in \mathcal{G}$ as per (7) [17].

$$\arg\max_f \sum_{u \in \mathcal{G}(V)} \log\left(\text{Prob}(N_s(u)|f(u))\right) \quad (7)$$

A. Distance Maximization via UGRL

Once $f(\mathcal{G})$ is obtained, clients' scheduling relies on the distance maximization of the nodes in the embedding space. This has the effect of increasing data heterogeneity of the client devices during consecutive communication rounds. The distance among a pair of nodes (v, u) in the graph $\mathcal{G}$ is evaluated as the normalized dot product (cosine similarity) $s(v, u)$ of their node embeddings $\mathbf{z}_v, \mathbf{z}_u$ (8):

$$s(v, u) = \frac{\mathbf{z}_v^T \cdot \mathbf{z}_u}{||\mathbf{z}_v|| \cdot ||\mathbf{z}_u||}. \quad (8)$$

To introduce memory of the past actions in the scheduling process, nodes scheduled in previous communication rounds are stored in a context window W_L of tunable length $L = \max(|W_L|)$. Accordingly, the similarity scores of the nodes $v \in W_L$ are computed and collected in a matrix $\mathbf{S}_{v,u}$ of dimension $|W_L| \times (M - |W_L|)$:

$$\mathbf{S}_{v,u} = \begin{bmatrix} s(v_1, u_{L+1}) & \ldots & s(v_1, u_M) \\ \vdots & \ddots & \vdots \\ s(v_L, u_{L+1}) & \ldots & s(v_L, u_M) \end{bmatrix}. \quad (9)$$

A scheduling decision is finally determined by equation (10).

$$\arg\min_{u \in \mathcal{G} \setminus \{v \in W_L\}} \sum_{v \in W_L} s(v, u). \quad (10)$$

The latter is equivalent to summing all elements of the matrix $\mathbf{S}_{v,u}$ by column, and selecting the next scheduled client as the argument corresponding to the column holding the minimum sum value, i.e., the node with maximum distance with respect to all previously scheduled nodes in W_L.

V. SIMULATION METHODOLOGY

Experiments and evaluation were conducted on a simulator based on Tensorflow Federated Core API. The logical steps of the designed FL framework, namely *"FederatedEnv"*, are reported in Algorithm 1.

After the initialization of clients' positions $(r_m, \phi_m)_M$ and datasets D_m, the algorithm loops over a fixed number of communication rounds. Each round can be subdivided into 7 logical steps:

1) The module of the client-PS Downlink (DL) channel impulse response $|h_{PS,m}|$ is computed, assuming perfect CSI, by summing fading and shadowing contributions ($f \sim \mathfrak{F}, s \sim N$) to the median path loss PL. The signal-to-noise ratio γ is thereby computed with a noise floor N_{dB} of -115 [dBm] and an AP gain G_{tx} of 15 [dBi].

2) The PS model $\boldsymbol{\theta}_{PS,t}$ is broadcasted to the clients. At the receiver side, Gaussian noise with standard deviation $\sigma_{DL} = \sqrt{E[\mathbf{N^2}]} = \sqrt{E[\boldsymbol{\theta}^2]/\gamma_{DL,m}}$ is added to the PS's model weights, where $\boldsymbol{\theta}$ is the discrete signal of weights of every model's layer.

3) Each client performs a local update of its model weights, controlled by the learning rate $\alpha(t)$, using stochastic gradient descent for N_{ep} epochs, as per (4).

4) After local model updates, a new round of Uplink (UL) channel estimation is performed in the same way as for step 1, and the corresponding UL signal-to-noise ratio γ_{UL} is computed.

5) A binary mask of the scheduled users $\mathbf{U} = [u_1, \ldots, u_M]$ is applied to the vector of updated model weights $\boldsymbol{\Theta}'_{\mathrm{M}} = [\boldsymbol{\theta}'_1, \ldots, \boldsymbol{\theta}'_M]^T$. During the aggregation phase, only models belonging to scheduled clients will be retained and included in the aggregation process.

6) The noisy models from the scheduled clients are aggregated as per (5).

7) The model is evaluated on a test set. We denote by $f(D_{test}|\boldsymbol{\theta}_{PS}(t+1))$ any metric computed over D_{test}.

Algorithm 1 FederatedEnv: logical steps

---

for each episode **do**
  for $m \in \mathcal{M}$ **do**
    $(r_m, \phi_m) = (r_m \sim U(0,R), \phi_m \sim U(0, 2\pi))$
    $|D_m| = |D| \cdot (u_m \sim U(0,1)/\sum u_m)$
    $D_m = \text{sample\_data}(D, |D_m|, \phi_m)$
  end for
  for each round t in T and for every $m \in \mathcal{M}$ **do**
    Step 1: DL Channel State Estimation
      $|h_{PS,m}|_{dB} = -\text{PL}_m + s \sim N(0, \sigma_s) + f \sim \mathfrak{F}(k, \omega)$
      $\gamma_{DL,m} = P_{tx,DL} + G_{tx} + |h_{PS,m}|_{dB} - N_{dB}$
    Step 2: Model Broadcast
      $\boldsymbol{\theta}_m(t) \leftarrow \boldsymbol{\theta}_{PS}(t) + \mathbf{n} \sim N(0, \sigma_{\propto \gamma_{DL,m}})$
    Step 3: Client Update
      $\boldsymbol{\theta}_m^{j+1}(t) = \boldsymbol{\theta}_m^j(t) - \alpha(t) \cdot \mathbf{g}_m$
    Step 4: UL Channel State Estimation
      $|h_{m,PS}|_{dB} = -\text{PL}_m + s \sim N(0, \sigma_s) + f \sim \mathfrak{F}(k, \omega)$
      $\gamma_{UL,m} = P_{tx,UL} + G_{tx} + |h_{m,PS}|_{dB} - N_{dB}$
    Step 5: Client Scheduling
      $\boldsymbol{\Theta}'_{\mathrm{M}} = \boldsymbol{\Theta}'_{\mathrm{M}} \cdot \mathbf{U}^T = [\boldsymbol{\theta}'_1, \ldots, \boldsymbol{\theta}'_M]^T \cdot [u_1, \ldots, u_M]$
    Step 6: Model update
      $\boldsymbol{\theta}_{PS}(t+1) = \frac{1}{K} \sum_{m \in \mathcal{K}} \frac{|D_m|}{\sum_{m \in \mathcal{K}} |D_m|} \cdot \boldsymbol{\theta}_m(t) + \mathbf{n}_m$
    Step 7: Performance Evaluation
      $f(D_{test}|\boldsymbol{\theta}_{PS}(t+1))$
  end for
end for

VI. RESULTS

In this section, we present and comment on the obtained results. We consider a scenario where $M = 50$ clients are randomly distributed among equally populated clusters (Fig. 4b, 4c). The distance maximization scheduling was tested against baseline policies (Table I) based on the feature nodes metrics introduced in section III. All policies have been tested on MNIST classification when transmitting a shallow neural network with 1 hidden layer, achieving 0.91 accuracy in a centralized setting.

POLICY	DESCRIPTION.
Max Age of Information (AoI)	nodes with max AoI
Random (RND)	nodes chosen randomly
Round robin (RR)	nodes chosen in a round-trip fashion
Best Channel (BC)	nodes with the best channel condition
Oracle (OCL)	explicit data knowledge information - maximize label heterogeneity
Max L2-norm (L2N)	nodes with maximum L2-norm of their local model update
Distance maximization (DM)	our proposed scheduling policy

TABLE I: Baseline policies

Fig. 5 and Fig. 6 show performance comparison in terms of training accuracy and energy (Fig. 6a) vs. communication efficiency (Fig. 6b), respectively. The three metrics are evaluated as follows:

- *Accuracy*: the sparse categorical cross accuracy on a centralized test set.
- *Communication efficiency*: the number of communication rounds r to achieve an accuracy of 0.8.
- *Energy efficiency*: the number of rounds r, multiplied by the number of training devices per round (K, or M, depending on the policy).

In all three metrics, the proposed solution outperforms the baselines and approaches the performance of the oracle, which is an empirical upper bound. In particular, the performance gap between the proposed solutions increases as the number of schedulable users (K) decreases. It is of particular interest to compare our proposed DM policy with respect to the L2N importance-aware policy. Since the former does not require any form of training feedback (Fig. 1), it is much more energy efficient, as the number of clients training per round is reduced from M to K (i.e., only the scheduled users perform local training). Moreover, results show that, in the proposed scenario, DM outperforms L2N even in terms of model accuracy and communication efficiency (Fig. 5 and 6b). In fact, even though the L2N policy is successful in scheduling the users holding the most significant models every round (i.e., those contributing more significantly to the global model, according to the L2-norm of the local updates [10]), when nodes show spatial relation, this may result in the selection of nearby clients in space holding similar data. On the opposite, our proposed policy aims to achieve data heterogeneity by maximizing nodes' distance in the graph space, making it more suitable for the considered scenario. For the same reason, policies such as AoI and random scheduling achieve better performance than round-robin, since they inherently increase
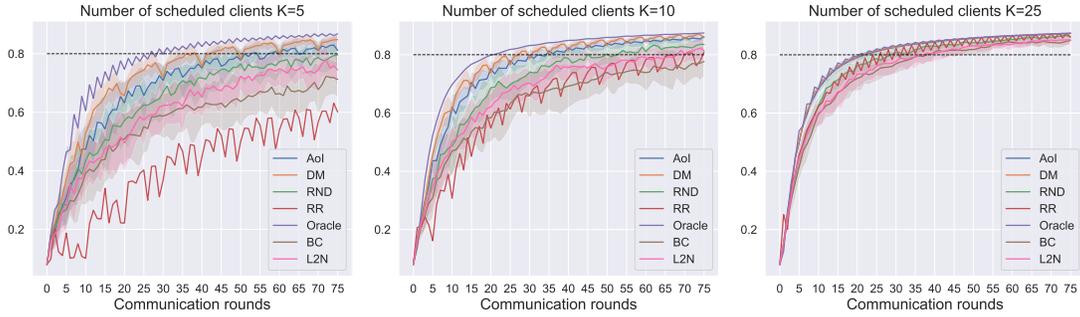
Fig. 5: Training results: sparse categorical cross accuracy vs number of communication rounds for different number of schedulable users per round: (a) $K = 5, M = 50$, (b) $K = 10, M = 50$ and $K = 25, M = 50$

data heterogeneity among the clients scheduled every round. Results show that DM can achieve an average 10% gain in model accuracy with respect to L2N at the end of the training while increasing energy efficiency by 17 times for $K = 5$. Moreover, we register an average gain of 4% accuracy at the end of the training and 31% in communication and energy efficiency with respect to the second-best performing policy (AoI). Finally, it is interesting to notice how the choice of K generates a tradeoff between communication and energy efficiency. Indeed, for our considered model and simulation parameters, if the aim of the designer is to maximize the system's communication efficiency, then $K = 10$ yields a gain of $27, 9\%$ in communication efficiency, but a loss of $30, 7\%$ in energy efficiency with respect to $K = 5$. Therefore, this is an indicator that for energy-sensitive applications, like IoT, the maximization of the number of schedulable clients is not always the best design choice.

(a) Energy efficiency

(b) Communication efficiency
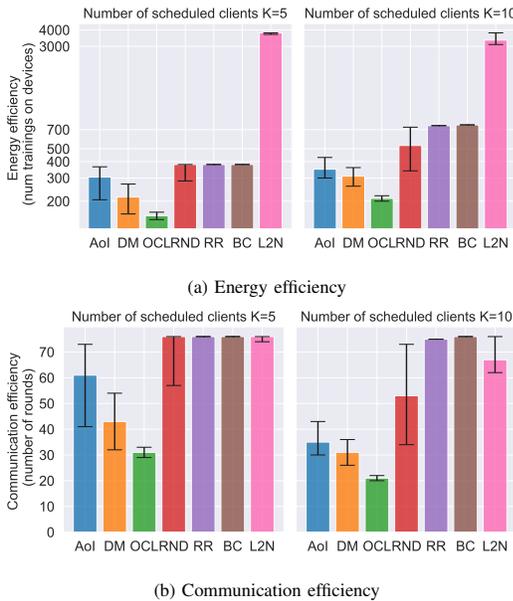
Fig. 6: Training results: energy (a) vs communication efficiency (b).

VII. CONCLUSION

In this study, we present a novel metric for the scheduling of client devices in FL applications leveraging the use of UGRL.

With respect to state-of-the-art importance-aware scheduling methods, our solution does not require any training feedback from client devices. Hence, it provides a much more computationally and energy-efficient solution. Our results indicate that, when tested against baseline importance-aware policies, our solution achieves a gain of up to 10% in model accuracy, while requiring up to 17 times fewer local training phases on client devices.

REFERENCES

[1] H.B. McMahan *et al.*, "Communication-efficient learning of deep networks from decentralized data," in *AISTATS*, 2017.

[2] K.B. Letaief *et al.*, "Edge artificial intelligence for 6G: Vision, enabling technologies, and applications," *IEEE J. Select. Areas Commun.*, 2022.

[3] M. Chen *et al.*, "Communication-efficient federated learning," *Proc. Natl. Acad. Sci. U.S.A.*, vol. 118, no. 17, p. e2024789118, 2021.

[4] E. Calvanese Strinati *et al.*, "6G networks: Beyond shannon towards semantic and goal-oriented communications," *Comput. Netw.*, vol. 190, p. 107930, 2021.

[5] H. Hellström *et al.*, "Wireless for machine learning: A survey," *Found. Trends Signal Process.*, vol. 15, no. 4, pp. 290–399, 2022.

[6] H.H. Yang *et al.*, "Scheduling policies for federated learning in wireless networks," *IEEE Trans. Commun.*, vol. 68, no. 1, pp. 317–333, 2019.

[7] E. Rizk *et al.*, "Federated learning under importance sampling," *IEEE Trans. Signal Process.*, pp. 1–15, 2022.

[8] E. Rizk *et al.*, "Optimal importance sampling for federated learning," in *Proc. IEEE Int. Conf. Acoustics, Speech, and Signal Processing (ICASSP)*, 2021, pp. 3095–3099.

[9] W. Chen *et al.*, "Optimal client sampling for federated learning," *arXiv preprint arXiv:2010.13723*, 2020.

[10] M.M. Amiri *et al.*, "Convergence of update aware device scheduling for federated learning at the wireless edge," *IEEE Trans. Commun.*, vol. 20, no. 6, pp. 3643–3658, 2021.

[11] J. Ren *et al.*, "Scheduling for cellular federated edge learning with importance and channel awareness," *IEEE Trans. Commun.*, 2020.

[12] A. Aral *et al.*, "Staleness control for edge data analytics," *Proc. ACM Meas. Anal. Comput. Syst.*, vol. 4, no. 2, 2020.

[13] D. Wen *et al.*, "An overview of data-importance aware radio resource management for edge machine learning," *J. Commun. Netw.*, vol. 4, no. 4, pp. 1–14, 2019.

[14] H.H. Yang *et al.*, "Age-based scheduling policy for federated learning in mobile edge networks," in *Proc. IEEE Int. Conf. Acoustics, Speech, and Signal Processing (ICASSP)*, 2020, pp. 8743–8747.

[15] M.M. Wadu *et al.*, "Federated learning under channel uncertainty: Joint client scheduling and resource allocation," in *Proc. IEEE Wireless Commun. and Networking Conf. (WCNC)*, 2020, pp. 1–6.

[16] P.W. Battaglia *et al.*, "Relational inductive biases, deep learning, and graph networks," *arXiv preprint arXiv:1806.01261*, 2018.

[17] A. Grover *et al.*, "node2vec: Scalable feature learning for networks," *Proc. of the 22nd ACM SIGKDD Int. Conf. on Knowledge Discovery and Data Mining*, 2016.