Bandit-based Communication-Efficient Client Selection Strategies for Federated Learning

Yae Jee Cho, Samarth Gupta, Gauri Joshi, Osman Yağan

Electrical & Computer Engineering, Carnegie Mellon University, Pittsburgh, PA

Abstract—Due to communication constraints and intermittent client availability in federated learning, only a subset of clients can participate in each training round. While most prior works assume uniform and unbiased client selection, recent work on biased client selection [1] has shown that selecting clients with higher local losses can improve error convergence speed. However, previously proposed biased selection strategies either require additional communication cost for evaluating the exact local loss or utilize stale local loss, which can even make the model diverge. In this paper, we present a bandit-based communication-efficient client selection strategy UCB-CS that achieves faster convergence with lower communication overhead. We also demonstrate how client selection can be used to improve fairness.

Index Terms—distributed optimization, federated learning, fairness, client selection, multi-armed bandits

I. INTRODUCTION

With increasing applications moving from data-center based training to edge-device training, federated learning (FL) [2], [3] has been spotlighted as one of the powerful distributed optimization methods for on-device learning. FL enables a massive distributed network of devices (clients) to participate in the training process without data-sharing. In each communication round, a subset of selected clients perform local model training and send their updated models to a central aggregating server. Most previous works use unbiased client selection and model aggregation. However, due to the inherent data heterogeneity across clients, judicious use of bias in client selection presents an untapped opportunity to improve error convergence.

Recent works have [1], [4] shown that biasing client selection in FL towards clients with higher local loss achieves faster convergence compared to unbiased client selection. However the client selection strategies proposed in previous works either require the server to additionally communicate with clients to retrieve the accurate local loss values or use stale loss values received from selected clients from previous communication rounds. Communication is expensive in FL, and furthermore, we show that using stale loss values can lead to slower error convergence or even divergence.

In this paper, we propose a bandit-based client selection strategy UCB-CS that is communication-efficient and use the observed clients' local loss values more appropriately instead of using stale values. Moreover, we show that biased client

selection can promote fairness, the uniformity of local loss performance across clients. To the best of our knowledge, there has been no work that proposes or evaluates biased client selection strategies in the context of fairness.

II. PROBLEM FORMULATION

A. Federated Learning with Partial Device Participation

The general FL framework FedAvg [2] aims to find the parameter vector $\mathbf{w}$ that minimizes the following objective:

$$F(\mathbf{w}) = \frac{1}{\sum_{k=1}^{K} D_k} \sum_{k=1}^{K} \sum_{\xi \in \mathcal{B}_k} f(\mathbf{w}, \xi) = \sum_{k=1}^{K} p_k F_k(\mathbf{w}) \quad (1)$$

with total K clients, where client k has a local dataset $\mathcal{B}_k$ consisting $|\mathcal{B}_k| = D_k$ data samples. The term $f(\mathbf{w}, \xi)$ is the composite loss function for sample ξ and parameter vector $\mathbf{w}$. The term $p_k = D_k / \sum_{k=1}^{K} D_k$ is the fraction of data at the k-th client, and $F_k(\mathbf{w}) = \frac{1}{|\mathcal{B}_k|} \sum_{\xi \in \mathcal{B}_k} f(\mathbf{w}, \xi)$ is the local objective function of client k.

A central aggregating server optimizes the model parameter $\mathbf{w}$ by selecting a subset of $m = CK$ clients for some fraction $0 < C < 1$ in each communication round (partial-device participation). Each selected client performs τ iterations of local SGD [5], [6], [7] and sends its locally updated model back to the server. Then, the server updates the global model using the local models and broadcasts the global model to a new set of active clients. Formally, we index the local SGD iterations with $t \geq 0$. The set of active clients at iteration t is denoted by $\mathcal{S}^{(t)}$. Since active clients performs τ steps of local update, the active set $\mathcal{S}^{(t)}$ also remains constant for every τ iterations. That is, if $(t+1) \bmod \tau = 0$, then $\mathcal{S}^{(t+1)} = \mathcal{S}^{(t+2)} = \cdots = \mathcal{S}^{(t+\tau)}$. Accordingly, the update rule of FedAvg is as follows:

$$\mathbf{w}_k^{(t+1)} =$$
$$\begin{cases} \frac{1}{m} \sum_{j \in \mathcal{S}^{(t)}} \left(\mathbf{w}_j^{(t)} - \eta_t g_j(\mathbf{w}_j^{(t)}, \xi_j^{(t)}) \right) \triangleq \overline{\mathbf{w}}^{(t+1)} & \text{else} \end{cases}$$
$$(2)$$

where $\mathbf{w}_k^{(t+1)}$ is the local model of client k at iteration t, η_t is the learning rate, and $g_k(\mathbf{w}_k^{(t)}, \xi_k^{(t)}) = \frac{1}{b} \sum_{\xi \in \xi_k^{(t)}} \nabla f(\mathbf{w}_k^{(t)}, \xi)$ is the stochastic gradient over mini-batch $\xi_k^{(t)}$ of size b that is randomly sampled from client k's local dataset $\mathcal{B}_k$. Moreover, $\overline{\mathbf{w}}^{(t+1)}$ is the global model at the server.

B. Biased Client Selection for Faster Convergence

We define a client selection strategy π that maps the parameter vector $\mathbf{w}$ to a specific set of clients $\mathcal{S}(\pi, \mathbf{w})$. The baseline unbiased client selection strategy used in FedAvg [2] chooses clients in proportion to p_k, denoted as π_{rand}. It has been noted in [1], [4] that selecting clients with higher local loss at each communication round leads to faster convergence but incurs an error floor. To attain faster convergence than π_{rand}, [1] proposes the power-of-d client selection scheme $\pi_{\text{pow-d}}$. Under the $\pi_{\text{pow-d}}$ scheme, the central server with $d > m$ clients obtains their local loss $F_k(\mathbf{w})$ for the current global model $\mathbf{w}$. After doing so, the $\pi_{\text{pow-d}}$ selects the m clients with largest local loss values in the next communication round. It has been seen in [1] that this scheme performs much better than the π_{rand} scheme for a variety of ML tasks. A drawback of $\pi_{\text{pow-d}}$ scheme is that it requires additional communication, as the central server is required to poll d clients before selecting clients for the next communication round.

To reduce this additional communication, it is desirable to have a proxy for the local loss $F_k(\mathbf{w})$ of each client k available at the center. Motivated by this, [1] proposes the $\pi_{\text{rpow-d}}$ scheme, where the local loss $F_k(\mathbf{w})$ is approximated by the local loss of the client when it was last selected in the client selection procedure. However, this approximation can be misleading at times as the client loss evaluation can be noisy and stale. Due to these reasons, it was observed that in certain cases the $\pi_{\text{rpow-d}}$ scheme does not have desirable convergence (See Figure 1). In this paper, we design a communication-efficient client selection strategy that enjoys faster convergence performance and robustness to the error floor compared to previously proposed client selection strategies.

C. Fairness in Client Selection

Fairness in FL has been studied only recently in the literature [8], [9], with the main goal of capturing the local accuracy discrepancies across clients for a trained global model. According to (1), clients with larger p_k will intuitively yield lower local loss performance, and vice versa. However, if clients with small p_k perform significantly worse than the other clients, this can be *unfair* [8], [9]. Henceforth, we define *client fairness* as the extent of identical local performance across clients for a single global model. We show that client fairness can be improved by incorporating the estimated local loss values and client's selected frequency to the client selection scheme. Gaining perspective from wireless resource allocation [10], [11], [12], we measure fairness by the Jain's index [10] $J(\mathbf{w})$, $1/K \leq J(\mathbf{w}) \leq 1$ where 1 is when all clients have the same performance. Fairness metric $J(\mathbf{w})$ is defined as:

$$J(\mathbf{w}) = \frac{1}{K} \left[\sum_{k=1}^{K} \left(\frac{F_k(\mathbf{w})}{\sum_{i=1}^{K} F_i(\mathbf{w})} \right)^2 \right]^{-1} \quad (3)$$

With this definition, we show that our proposed UCB-CS gains both fairness and convergence speed compared to the unbiased and previously proposed biased client selection strategies.

III. CLIENT SELECTION WITH DISCOUNTED UCB

In order to achieve faster convergence with low error floor, it is important to select clients with larger local loss (i.e., *exploitation*) as that leads to faster convergence [1]. It is also important to ensure diversity (i.e., *exploration*) in selection to achieve a lower error floor. Motivated by the fact that there is a exploration-exploitation trade-off, we propose the use of Multi-Armed Bandit (MAB) algorithms [13] for the problem of client-selection in FL. Since the local loss values of individual clients are non-stationary during training we make use of discounted MAB algorithms proposed in [14]. We modify the discounted UCB algorithm [14] to balance the exploration-exploitation trade-off in the client selection problem. We view the clients as the arms in the MAB problem and compute discounted cumulative local loss values of each client, $L_t(\gamma, k)$, and a discounted count of the number of times each client has been sampled, $N_t(\gamma, k)$, till communication round t.

Using these, we define the discounted UCB index for each client $k \in [K]$ at communication round t, and select the m clients with largest discounted UCB indices. With $\mathcal{T} = \{\tau, 2\tau, 3\tau, ..., \lfloor t/\tau \rfloor \tau\}$ which are the set of time indices for the communication rounds until t, the discounted UCB indices $A_t(\gamma, k)$ are formally defined as

$$A_t(\gamma, k) = p_k (\underbrace{L_t(\gamma, k)/N_t(\gamma, k)}_{\text{exploitation}} + \underbrace{U_t(\gamma, k)}_{\text{exploration}}) \quad (4)$$

where,

$$L_t(\gamma, k) = \sum_{t' \in \mathcal{T}} \gamma^{t-t'} \mathbb{1}_{\{k \in \mathcal{S}^{(t'-1)}\}} \frac{1}{\tau} \sum_{l=t'-\tau+1}^{t'} \sum_{\xi \in \xi_k^{(l)}} \frac{f(\mathbf{w}_k^{(l)}, \xi)}{b} \quad (5)$$

$$N_t(\gamma, k) = \sum_{t' \in \mathcal{T}} \gamma^{t-t'} \mathbb{1}_{\{k \in \mathcal{S}^{(t'-1)}\}} \quad (6)$$

$$U_t(\gamma, k) = \sqrt{2\sigma_t^2 \log T_t(\gamma)/N_t(\gamma, k)}, \ T_t(\gamma) = \sum_{t' \in \mathcal{T}} \gamma^{t-t'} \quad (7)$$

Here, $0 \leq \gamma \leq 1$ is a hyper-parameter that impacts the importance given to stale values. When $\gamma = 1$, all past local loss samples contribute equally in the calculation of $L_t(\gamma, k)$ and when $\gamma = 0$ only the latest local loss is used to estimate $L_t(\gamma, k)$. For $0 < \gamma < 1$, less weight is put upon stale values of local loss for the calculation of $L_t(\gamma, k)$. By doing so, we compute the estimate for local loss of client in a robust way to escape the noise in the latest evaluation and discount the stale values computed in the past. The exploration term $U_t(\gamma, k)$ gets larger by a factor of σ_t especially for clients that have not been selected recently regardless of their local loss values. This forces UCB-CS to explore other clients that may have relatively smaller exploitation value. This not only can pull the algorithm away from generating the error floor by just selecting clients with estimated larger local loss, but also promotes fairness in terms of the number of times the client is explored. The parameter σ_t is the maximum standard deviation in the local loss computed over the latest update
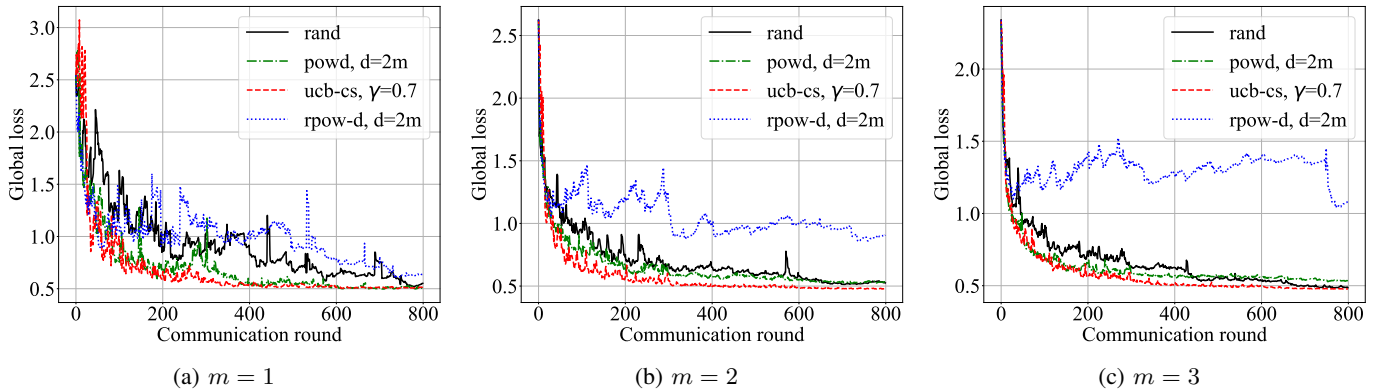
(a) $m = 1$ (b) $m = 2$ (c) $m = 3$

Fig. 1: Global loss for logistic regression on the synthetic dataset, `Synthetic(1,1)`, with π_{rand}, $\pi_{\text{pow-d}}$, $\pi_{\text{rpow-d}}$, and $\pi_{\text{ucb-cs}}$ for $d = 2m$, $\gamma = 0.7$ where $K = 30$, $m \in \{1, 2, 3\}$.

of clients. Note that we multiply the dataset size ratio with the discounted UCB index, as we also want to sample clients proportional to their datasize for fast convergence and lower error floor. Details of UCB-CS are presented in Algorithm 1.

Algorithm 1 Pseudo code for UCB-CS

1: **Input**: m, γ, p_k for $k \in [K]$
2: **Output**: $\mathcal{S}^{(t)}$
3: **Initialize**: empty sets $\mathcal{S}^{(t)}$ and list $\mathcal{A}$ of length K
4: **Global server do**
5:     Receive $\frac{1}{\tau b} \sum_{l=t-\tau+1}^{t} \sum_{\xi \in \xi_k^{(l)}} f(\mathbf{w}_k^{(l)}, \xi)$ from clients $k \in \mathcal{S}^{(t-1)}$ and calculate $A_t(\gamma, k)$ as (4)
6:     Update $\mathcal{A}[k] = A_t(\gamma, k)$
7:     Get $\mathcal{S}^{(t)} = \{m$ clients with largest values in $\mathcal{A}$ (break ties randomly)$\}$
8:     Discount elements in $\mathcal{A}$ by $\mathcal{A} = \gamma \mathcal{A}$
9: **return** $\mathcal{S}^{(t)}$

IV. EXPERIMENT RESULTS

We evaluate the proposed UCB-CS with logistic regression on a heterogeneous synthetic federated dataset, `Synthetic(1,1)` [15], and DNN trained on a non-iid partitioned FMNIST dataset [16]. For logistic regression, we assume in $K = 30$ where the local dataset sizes follow the power law distribution. We set $b = 50$, $\tau = 30$, and $\eta = 0.05$, where η is decayed to $\eta/2$ every 300 and 600 rounds. For DNN, we train a deep multi-layer perceptron network with two hidden layers on the FMNIST dataset. We construct the heterogeneous data partition amongst clients using the Dirichlet distribution $\text{Dir}_K(\alpha)$ [17], where α determines the degree of the data heterogeneity across clients. Smaller α indicates larger data heterogeneity. For all experiments we use $b = 64$, $\tau = 100$, and $\eta = 0.005$, where η is decayed by half for round 150. All experiments are conducted with clusters equipped with one NVIDIA TitanX GPU. The machines communicate amongst each other through Ethernet. The algorithms are implemented by PyTorch. For all results, the hyper-parameters d and γ are tuned for the best performance via grid search.

The training loss performance for the synthetic dataset simulation is presented in Fig. 1. The UCB-CS algorithm, $\pi_{\text{ucb-cs}}$, converges even faster than $\pi_{\text{pow-d}}$ without any error floor, and performs significantly better than π_{rand} and $\pi_{\text{rpow-d}}$. The $\pi_{\text{rpow-d}}$ selection policy performs worse than π_{rand}, showing that using stale local losses for biased client selection can make the performance worse than the unbiased selection strategies. Additionally, in Table I, we show that the biased client selection strategies achieve notable higher fairness than the random selection strategy. While $\pi_{\text{pow-d}}$ is able to achieve higher fairness than $\pi_{\text{ucb-cs}}$, $\pi_{\text{ucb-cs}}$ shows a significant improvement in fairness even with low communication cost and robustness to the error floor in the training curve. Hence we show that $\pi_{\text{ucb-cs}}$ is efficient in the three important factors in FL: loss performance, fairness, and communication-efficiency.

To dive deeper into the difference between $\pi_{\text{pow-d}}$ and $\pi_{\text{ucb-cs}}$, in Fig, 2 we present the local loss distribution across the clients at the end of training for the simulation in Fig. 1(a). We show that both $\pi_{\text{ucb-cs}}$ and $\pi_{\text{pow-d}}$ is able to improve the worst performing client's local loss performance for π_{rand}. While $\pi_{\text{pow-d}}$ is able to keep most of the clients in the approximately average range of performance of the local loss, $\pi_{\text{ucb-cs}}$ allows most of the clients to perform with the lowest local loss, skewing the local loss distribution across clients towards lower loss values. Hence from Fig. 2 we can see that $\pi_{\text{pow-d}}$ is valuing fairness over performance, whereas $\pi_{\text{ucb-cs}}$ is valuing performance slightly over fairness.

In Fig. 3, the test accuracy and training loss for image classification on the FMNIST dataset via DNN are presented. For less data heterogeneity ($\alpha = 2$), both $\pi_{\text{rpow-d}}$ and $\pi_{\text{ucb-cs}}$ perform similarly with higher test accuracy and lower training loss than π_{rand}. However, for larger data heterogeneity ($\alpha = 0.3$), $\pi_{\text{rpow-d}}$ performs worse than $\pi_{\text{ucb-cs}}$, showing that with large τ the estimated local loss values that $\pi_{\text{rpow-d}}$ use becomes very stale, worsening the performance in the presence of large data heterogeneity. On the other hand, $\pi_{\text{ucb-cs}}$ and $\pi_{\text{pow-d}}$ have similar empirical performance, which shows that $\pi_{\text{ucb-cs}}$'s use of discounted and accumulated local losses give an accurate representation of the client's actual local loss value.

TABLE I: Fairness values $J(\overline{\mathbf{w}}^{(T)})$, where T is the last communication round, for the scenarios in Fig. 1.

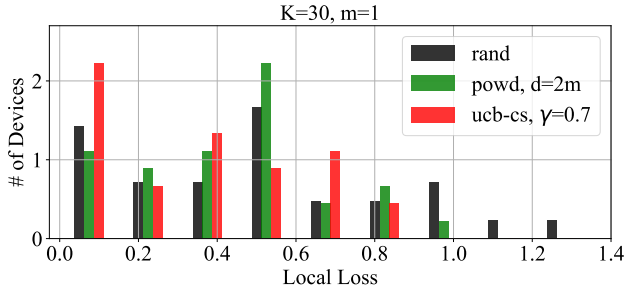|	$m = 1$	$m = 2$	$m = 3$
π_{rand}	0.43	0.29	0.66
$\pi_{\text{pow-d}}$	0.75	0.89	0.91
$\pi_{\text{ucb-cs}}$	0.61	0.61	0.65
$\pi_{\text{rpow-d}}$	0.32	0.52	0.39

Fig. 2: Histogram for the client's individual loss performance after end of training for simulations presented in Fig. 1(a).
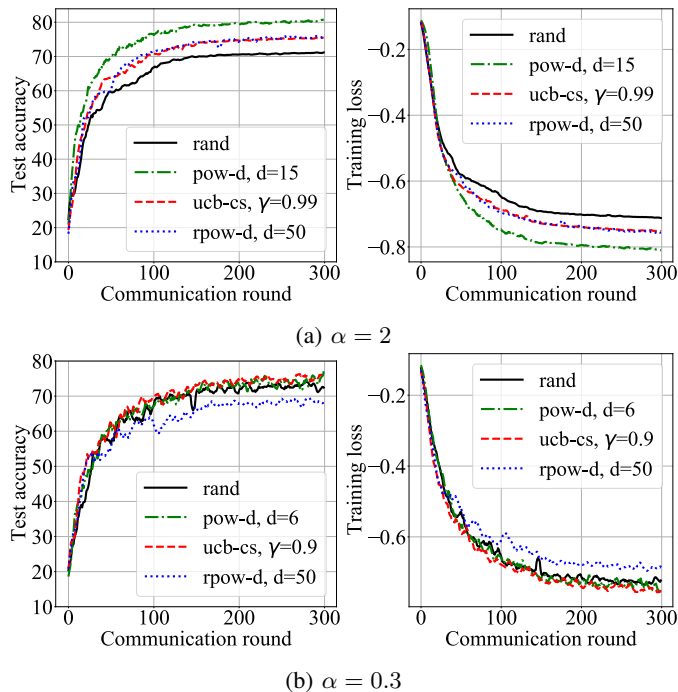

(a) $\alpha = 2$

(b) $\alpha = 0.3$

Fig. 3: Test accuracy and training loss for π_{rand}, $\pi_{\text{pow-d}}$, $\pi_{\text{rpow-d}}$, and $\pi_{\text{ucb-cs}}$ for $K = 100$, $C = 0.03$ on the FMNIST dataset with mini-batch size $b = 64$ and $\tau = 100$.

V. CONCLUDING REMARKS

In this paper, we propose a bandit-based communication-efficient client selection strategy, UCB-CS. It tackles the problem of communication-efficiency, noisy-stale estimates of local loss values, and error floor prevalent in biased client selection strategies from previous literature [1], [4]. We discover that UCB-CS, with no additional communication compared to π_{rand}, is robust to the error floor and gains convergence speed while mitigating the problem of staleness of observed local loss

values. Moreover we show that UCB-CS increases fairness, the uniformity of performance across different clients, compared to other known communication-efficient biased client selection strategies. Throughout this work, we assume that local losses corresponding to a model $\mathbf{w}$ for each clients are independent of each other. However, in reality, similar clients might have similar local losses. We aim to use correlated multi-armed bandit algorithms [18], [19] in such setting to further improve the performance of UCB-CS for future work.

REFERENCES

[1] Y. J. Cho, J. Wang, and G. Joshi, "Client selection in federated learning: Convergence analysis and power-of-choice selection strategies," vol. abs/2010.01243, 2020. [Online]. Available: http://arxiv.org/abs/2010.01243

[2] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-Efficient Learning of Deep Networks from Decentralized Data," *International Conference on Artificial Intelligenece and Statistics (AISTATS)*, Apr. 2017. [Online]. Available: https://arxiv.org/abs/1602.05629

[3] P. Kairouz and H. B. M. et al, "Advances and open problems in federated learning," *arXiv preprint arXiv:1912.04977*, 2019.

[4] J. Goetz, K. Malik, D. Bui, S. Moon, H. Liu, and A. Kumar, "Active federated learning," *ArXiv*, 2019.

[5] S. U. Stich, "Local SGD converges fast and communicates little," *arXiv preprint arXiv:1805.09767*, 2018.

[6] J. Wang and G. Joshi, "Cooperative SGD: A unified framework for the design and analysis of communication-efficient SGD algorithms," *arXiv preprint arXiv:1808.07576*, 2018. [Online]. Available: https://arxiv.org/abs/1808.07576

[7] H. Yu, S. Yang, and S. Zhu, "Parallel restarted SGD for non-convex optimization with faster convergence and less communication," *arXiv preprint arXiv:1807.06629*, 2018.

[8] T. Li, M. Sanjabi, and V. Smith, "Fair resource allocation in federated learning," *arXiv preprint arXiv:1905.10497*, 2019.

[9] M. Mohri, G. Sivek, and A. T. Suresh, "Agnostic federated learning," in *Proceedings of the 36th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, K. Chaudhuri and R. Salakhutdinov, Eds., vol. 97. Long Beach, California, USA: PMLR, 09–15 Jun 2019, pp. 4615–4625.

[10] R. Jain, D. Chiu, and W. Hawe, "A quantitative measure of fairness and discrimination for resource allocation in shared computer systems," vol. cs.NI/9809099, 1998. [Online]. Available: https://arxiv.org/abs/cs/9809099

[11] C. Joe-Wong, S. Sen, T. Lan, and M. Chiang, "Multi-resource allocation: Fairness-efficiency tradeoffs in a unifying framework," in *2012 Proceedings IEEE INFOCOM*, 2012, pp. 1206–1214.

[12] T. Lan, D. Kao, M. Chiang, and A. Sabharwal, "An axiomatic theory of fairness in network resource allocation," in *2010 Proceedings IEEE INFOCOM*, 2010, pp. 1–9.

[13] S. Bubeck and N. Cesa-Bianchi, "Regret analysis of stochastic and nonstochastic multi-armed bandit problems," *Foundations and Trends in Machine Learning*, vol. 5, no. 1, pp. 1–122, 2012.

[14] A. Garivier and E. Moulines, "On upper-confidence bound policies for non-stationary bandit problems," vol. abs/0805.3415, 2008. [Online]. Available: http://arxiv.org/abs/0805.3415

[15] A. K. Sahu, T. Li, M. Sanjabi, M. Zaheer, A. Talwalkar, and V. Smith, "Federated optimization for heterogeneous networks," *https://arxiv.org/abs/1812.06127*, Jan. 2019.

[16] H. Xiao, K. Rasul, and R. Vollgraf, "Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms," *https://arxiv.org/abs/1708.07747*, aug 2017.

[17] T.-M. H. Hsu, H. Qi, and M. Brown, "Measuring the effects of non-identical data distribution for federated visual classification," *https://arxiv.org/abs/1909.06335*, Sep. 2019.

[18] S. Gupta, S. Chaudhari, G. Joshi, and O. Yağan, "Multi-armed bandits with correlated arms," 2020.

[19] S. Gupta, G. Joshi, and O. Yağan, "Correlated multi-armed bandits with a latent random source," *ArXiv*, Aug. 2018. [Online]. Available: https://arxiv.org/abs/1808.05904