

Hierarchical incentive mechanism design for federated machine learning in mobile networks

Lim, Bryan Wei Yang; Xiong, Zehui; Miao, Chunyan; Niyato, Dusit; Yang, Qiang; Leung, Cyril; Poor, H. Vincent

2020

Lim, B. W. Y., Xiong, Z., Miao, C., Niyato, D., Yang, Q., Leung, C., & Poor, H. V. (2020). Hierarchical incentive mechanism design for federated machine learning in mobile networks. *IEEE Internet of Things Journal*, 7(10), 9575-9588. doi:10.1109/JIOT.2020.2985694

<https://hdl.handle.net/10356/144299>

<https://doi.org/10.1109/JIOT.2020.2985694>

© 2020 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works. The published version is available at: <https://doi.org/10.1109/JIOT.2020.2985694>

Downloaded on 30 Sep 2024 22:58:32 SGT

Hierarchical Incentive Mechanism Design for Federated Machine Learning in Mobile Networks

Wei Yang Bryan Lim, Zehui Xiong, Chunyan Miao, *Senior Member, IEEE*, Dusit Niyato, *Fellow, IEEE*, Qiang Yang, *Fellow, IEEE*, Cyril Leung, *Life Member, IEEE*, H. Vincent Poor, *Fellow, IEEE*

Abstract—In recent years, the enhanced sensing and computation capabilities of Internet of Things (IoT) devices have opened the doors to several mobile crowdsensing applications. In mobile crowdsensing, a model owner announces a sensing task following which interested workers collect the required data. However, in some cases, a model owner may have insufficient data samples to build an effective machine learning model. To this end, we propose a Federated Learning based privacy preserving approach to facilitate collaborative machine learning among multiple model owners in mobile crowdsensing. Our system model allows collaborative machine learning without compromising data privacy given that only the model parameters instead of the raw data are exchanged within the federation. However, there are two main challenges of incentive mismatches between workers and model owners, as well as among model owners. For the former, we leverage on the self-revealing mechanism in contract theory under information asymmetry. For the latter, to ensure the stability of a federation through preventing free-riding attacks, we use the coalitional game theory approach that rewards model owners based on their marginal contributions. Considering the inherent hierarchical structure of the involved entities, we propose a hierarchical incentive mechanism framework. Using the backward induction, we first solve the contract formulation and then proceed to solve the coalitional game with the merge and split algorithm. The numerical results validate the performance efficiency of our proposed hierarchical incentive mechanism design, in terms of incentive compatibility of our contract design and fair payoffs of model owners in stable federation formation.

Index Terms—Federated Learning, Incentive Mechanism, Mobile Networks, Mobile Crowdsensing, Artificial Intelligence

I. INTRODUCTION

There are currently close to 7 billion connected Internet of Things (IoT) devices [1] worldwide. These IoT devices are equipped with increasingly enhanced sensing and computation capabilities. As such, mobile crowdsensing tasks [2] are increasingly popular.

WYB. Lim is with Alibaba Group and Alibaba-NTU Joint Research Institute (JRI), Nanyang Technological University (NTU), Singapore. Email: limw0201@e.ntu.edu.sg. Z. Xiong is with Alibaba-NTU JRI, and also with School of Computer Science and Engineering (SCSE), NTU, Singapore. Email: zxiong002@e.ntu.edu.sg. C. Miao is with SCSE, NTU, Singapore, Alibaba-NTU JRI, Joint NTU-UBC Research Centre of Excellence in Active Living for the Elderly (LILY), Singapore. E-mail: ascmiao@ntu.edu.sg. D. Niyato is with SCSE, NTU, Singapore. E-mail: dniyato@ntu.edu.sg. Q. Yang is with Hong Kong University of Science and Technology, Hong Kong, China. Email: qyang@cse.ust.hk. C. Leung is with the Alibaba-NTU JRI, LILY and Department of Electrical and Computer Engineering, The University of British Columbia, Canada. E-mail: cleung@ece.ubc.ca. H. Vincent Poor is with the Electrical Engineering Department, Princeton University, NJ 08544, Email: poor@princeton.edu. *Corresponding author: Zehui Xiong*

Copyright (c) 2020 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

In mobile crowdsensing, a task publisher, i.e., a model owner, announces a sensing task following which interested workers, i.e., participants in the crowdsensing task, will then collect the required data. Together with the advances in Machine Learning, especially in the domain of Deep Learning [3], the wealth of data collected allows effective models to be built, e.g., for medical [4], air quality monitoring, and map navigation applications [5].

However, in some cases, a lone model owner may have insufficient data samples to build an effective model, especially since Deep Learning algorithms usually outperform conventional approaches only when there is an abundance of data available for model training [3]. In addition, the model owner may also not have a comprehensive coverage across data classes to build an effective model [6]. The reason is that mobile crowdsensing involves *location dependency* [7], e.g., the medical data collected by a model owner is restricted to the local community in which its mobile network of workers cover.

Naturally, to build a better inference model, different model owners can collaborate by sharing their data. However, in recent years, the regulations governing data privacy, e.g., General Data Protection Regulation (GDPR) [8], are increasingly stringent. As such, this can potentially prevent the sharing of data across model owners.

To this end, we propose the adoption of a Federated Learning (FL) [9] approach to enable privacy-preserving collaborative Machine Learning across a federation of model owners. In our system model (Fig. 1), the model owner announces a sensing task to interested workers, similar to that in a typical crowdsensing setting. Then, the model owners come together to form different federations based on their profit maximizing objectives. Each federation collaboratively builds a model without the exchange of data. Instead, model training takes place on each model owner's private server, i.e., the data derived from its respective workers remain on the model owner's local database. Then, only the updated model parameters are exchanged and aggregated with that of other model owner's parameters in a trusted third-party server.

Our proposed approach has two key advantages. Firstly, it protects the privacy of workers working under their respective model owners while enabling distributed training on mobile devices [10]. Secondly, it is communication efficient. The reason is that traditional methods of data sharing will require the raw data to be uploaded to an aggregating cloud server. With FL, only the model parameters need to be updated by the model owners [11].

However, there exists two levels of incentive mismatches in

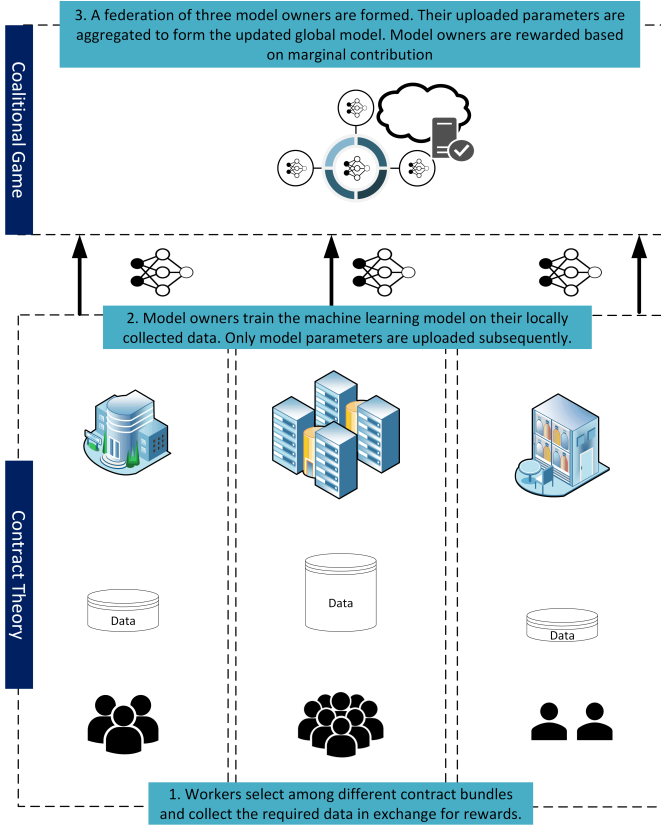


Fig. 1. Our proposed system model involving model parameter aggregation from a federation of multiple model owners. Note that the diagram only shows a single federation. In our system model, a model owner can choose to join between multiple federations based on their profit maximization objectives.

this setting. The lower level is one that is commonly studied in traditional crowdsensing [12], [13], i.e., the workers incur battery consumption and mobile data costs when collecting the data. As such, without an effective incentive mechanism design, the workers may contribute lower quantities or qualities of data. Furthermore, given the information asymmetry between model owners and workers, a poorly designed incentive framework may lead to the situation in which workers conceal their types, e.g., data quality, to maximize their utilities. For example, workers who can collect high quality data may have the incentives to contribute low quality data since the corresponding cost incurred is naturally lower. As such, we propose a contract theoretic approach [14] in incentivizing workers to provide high qualities and quantities of data. The self-revealing mechanism of the contract theory approach enables workers to be rewarded based on their specific types even in the presence of information asymmetry, i.e., when the worker types are not known by a model owner.

The upper level is an incentive mismatch among profit maximizing model owners. Each profit maximizing model owner serves to only maximize its own profits, and not the profits of the federation. Without an effective incentive mechanism design, there may be free-riding attacks [15], i.e., situations in which a model owner does not contribute sufficiently to the federation since high quality data is expensive to collect from the workers. As such, in the upper level, we adopt a coalitional

game approach with fair distribution of payoffs based on each model owner's marginal contribution to the federation. We also study the equilibrium of stable federation formation using the merge and split algorithm proposed in [16], i.e., a model owner can join the federation that allows it to maximize its payoff.

In fact, our proposed incentive mechanism design is hierarchical in nature. In the upper level, the payoff received by a model owner is affected by the decisions of other model owners in the federation. This payoff in turn affects the contract formulation at the level of each individual model owner. For example, when faced with the decision to join either of two federations with differing data qualities, the model owner first designs two distinct hypothetical contracts in response to the two federations. Each contract represents a way to maximize its marginal contribution to a federation while minimizing incentive expenses paid to its workers. Then, the model owner will choose to join the federation that allows it earn more profits, i.e., the federation that offers higher marginal payoffs at a lower incentive expense. This is in contrast to classical works in contract theory, in which the profit will only be affected by the choice variables of the sole model owner alone. Given the hierarchical nature of our mechanism design, we can apply solutions from Stackelberg games [17], [18]. We first implement backward induction by considering each model owner's contract-theoretic incentive mechanism design separately while taking the parameters of other model owners' decisions as given. Then, we explore federation formation as a coalitional game with fair payoffs. The contribution of our study is as follows:

- We propose an FL-assisted privacy-preserving data approach for mobile crowdsensing in the system model in which sole model owners do not have sufficient data to build a viable model.
- We propose a hierarchical incentive mechanism design for FL that considers *multiple* model owners and the formation of multiple federations.
- We provide an analysis of the equilibrium that our system model reaches after iterations of merges and splits, thus giving us an insight of federation structures in the presence of multiple model owners and federations.

The rest of this paper is organized as follows: Section II summarizes related work, Section III discusses the system model and problem formulation, Section IV discusses the optimal contract formulation, Section V discusses the coalition formation game, Section VI presents the performance evaluation and Section VII concludes.

II. RELATED WORK

The issue of incentive mechanism design in mobile crowdsensing is well-studied. Some of the tools that have been explored in previous studies are auctions [19], game theory [20], contract theory [21], and reputation mechanisms [22]. In addition, there are also several surveys in this topic [5], [23], [24].

Game theory has been utilized as a powerful tool to study the incentive mechanisms. The authors in [12], [13] devised the incentive mechanisms with the consideration of social

influence and proposed a novel hierarchical game framework to describe and concisely analyze the decision-making process. In particular, the contract-theoretic approach to mobile crowdsensing is increasingly popular. For example, contract theory has been used to motivate high quality data collection [25], [26] under the incomplete information crowdsensing setting. There have also been studies on incorporating social incentives into contracts [27]. In addition, to protect the privacy of workers, a smart contract on blockchain approach has also been proposed in [28].

Nevertheless, given that FL is a relatively nascent concept, the FL-assisted crowdsensing has not been well explored in the literature. The study in [29] propose privacy-preserving model training for mobile crowdsensing using an FL based framework. However, its focus lies primarily on user privacy protection rather than incentive mechanism design. In fact, most of the aforementioned works assume incentive mechanism design in the setting of only a *sole* model owner. However, in practical scenarios, it is likely that there is a need to consider the cooperation among *multiple* model owners in collaborative machine learning. For example, Nvidia [30] has recently proposed an FL approach for healthcare analytics using patients' data from multiple collaborating hospitals [31], i.e., model owners. Given the potential application of FL in other scenarios, e.g., map data and air quality monitoring, there is a need for a study of incentive mechanism design involving multiple collaborative model owners in crowdsensing.

With the increasing popularity of FL, we can also take reference from the growing literature related to incentive mechanism design for FL. For example, the study in [32], [32] adopts a contract-theoretic approach to motivate workers to contribute more computation resource for efficient FL. On the other hand, the study in [33] formulates the Stackelberg game [34] to analyze the inefficiency in model update transfer. As an extension, the study in [35] uses a Stackelberg game formulation together with Deep Reinforcement Learning to design a learning-based incentive mechanism for FL. For a comprehensive survey in this area, we refer the readers to [36], [37].

Given the self-revealing mechanism of contract theory under the information asymmetry setting of a heterogeneous mobile network, we also adopt the contract theoretic approach in the lower level of our system model, i.e., between each individual model owner and its respective workers. However, similar to that in crowdsensing, most of the studies of incentive mechanism design for FL assumes that there is only a sole model owner with multiple workers under it, i.e., a sole federation. In particular, the contract formulation in [32] assumes a monopolistic, self-sufficient federation with no need for collaboration among model owners. As such, the contract formulation in [32] is affected only by the model owner's individual choice variables. In contrast, in our system model, the payoffs of a model owner is affected by the decisions of *other* model owners. As such, considering the inherent hierarchical structure of the involved entitles, we propose a hierarchical incentive mechanism framework. Our work studies the setting of *multiple* federations and provide both an hierarchical incentive mechanism design and merge-and-split

framework to better understand the structure of federations when more than one model owner exists.

III. SYSTEM MODEL AND PROBLEM FORMULATION

In general, the system model (Fig. 1) comprises workers who collect data in response to crowdsensing tasks published by the model owners, in exchange for contract rewards. The model owners then collaborate with other model owners selectively to form federations that best fulfill their profit maximization objectives. Thereafter, the model owner trains the model on their own locally collected data. Then, only the model parameters are sent to a trusted server for aggregation in a privacy preserving FL approach.

We consider a mobile network with N model owners. Each model owner operates a platform on which p_i participating workers are registered on. Denote $\mathbf{p} = \{p_1, \dots, p_i, \dots, p_N\}$ as the set of all participating workers in the mobile network. The model owners have the basic demographic information of the workers. However, the model owners lack specialized data that have to be collected through participatory crowdsensing. For example, in health crowdsensing, even though hospitals have access to age, weight, and blood type of individuals under their care, they lack non-mandatory data that can only be collected with the individuals' consent, e.g., IoT applications data for health analytics [38], [39].

In an heterogeneous mobile network, there exists differing worker types that can in turn influence the quality and quantity of data collected for model training [23], [40]. A worker of type m belonging to model owner i has its type denoted as θ_{im} where $\theta_{im} \in \{\theta_{i1}, \dots, \theta_{iM}\}$. Similarly, the data quantity contribution is denoted as q_{im} where $q_{im} \in \{q_{i1}, \dots, q_{iM}\}$. The worker type reflects each worker's level of willingness to participate and hence determines the quality of data collected, i.e., workers who are more willing to participate collect higher quality data, e.g., images of higher resolution or datasets with fewer erroneous or missing inputs. Naturally, if a model owner trains the model on more data and higher quality data across all workers, the model performance, e.g., inference accuracy, will be higher [41]. We denote the model performance to be:

$$x_i(\theta_i, Q_i) = 1 - e^{-\phi(\bar{\theta}_i Q_i)^v} \quad (1)$$

where $\bar{\theta}_i$ and Q_i refers to the average data quality and total data quantity used by model owner i to train the model respectively. ϕ and v are weight factors. Following [42], we assume that the global model performance gain has diminishing returns with respect to data quality and quantity. Intuitively, there exists a limit to model performance where $\lim_{\bar{\theta}_i Q_i \rightarrow \infty} (1 - e^{-\phi(\bar{\theta}_i Q_i)^v}) = 1$.

Individually, each model owner usually has insufficient data quantity to build a good predictive and analytics model such as those based on machine learning [43]. Moreover, given the private nature of data involved, each model owner is unable to share or exchange their data with other model owners. To this end, we propose an FL-based [9] approach to facilitate collaborative model training that involves multiple model owners without having the need to share or transfer data to potentially malicious third-parties. A federation S with $|S|$

model owners can be formed where $S \subseteq \mathcal{N}$. In each training iteration, the model owners in the federation first train the local models on their respective collected data. Then, each model owner only uploads the model parameters \mathbf{w}_i to a cloud server for aggregation, e.g., by using the Federated Averaging (*FedAvg*) algorithm [9]. The aggregated model parameters $\mathbf{w} = \cup_{i \in S} \mathbf{w}_i$ are then sent back to the model owners in the federation for the next iteration of training.

The global model owned by the federation can then be sold as a service for profits, of which the profits are dependent on the model performance. For example, a model that performs better, i.e., with higher accuracy, can generate higher profits for the federation. We define the profit function of the federation as:

$$v(S) = \omega(1 - e^{-\phi(\bar{\theta}_S Q_S)^v}) - |S|G. \quad (2)$$

ω is a conversion parameter from model performance to profits which follows:

$$\omega = \begin{cases} \tilde{\omega}, & \text{if } 1 - e^{-\phi(\bar{\theta}_S Q_S)^v} \geq \tilde{x}, \\ 0, & \text{otherwise.} \end{cases} \quad (3)$$

\tilde{x} refers to a threshold model performance that a viable model has to meet for usability. For example, if the model performance is too poor, it is rejected by users. $\bar{\theta}_S$ refers to the quantity-weighted data quality in the coalition S where $\bar{\theta}_S = \sum_{i \in S} \frac{Q_i}{Q_S} \theta_i$ and Q_S refers to total units of data in the federation, i.e., $Q_S = \sum_{i \in S} Q_i$. In addition, when model owners cooperate, there is the coordination cost incurred as well [44]. For example, extra computation resource has to be purchased from the coordinating server to cater to the increased data volume when a new model owner joins and operates within the federation [45]. For simplicity, we denote this cost as G which increases linearly with the federation size $|S|$. Intuitively, as the federation increases in size, coordination cost increases.

After the federation profits from the model, the profits generated are fairly allocated based on the marginal contribution of each model owner. Each individual model owner naturally attempts to achieve the highest share of profits distributed from the federation. However, data collection requires time, energy consumption by IoT sensing devices and data usage costs of workers [24]. To incentivize high quality data contribution from its participating workers, each model owner i designs a contract-theoretic incentive mechanism with self-revealing properties [14]. As such, there exists a tradeoff in the hierarchical incentive design between *marginal contribution maximization* from the federation and *incentive expense minimization*, i.e., while each model owner wants to maximize its share of profits from the federation, it is costly to incentivize workers to collect high quality data. Thus, the model owner will choose to join the federation that allows it to maximize marginal payoffs while minimizing contractual costs.

Similar to solutions in Stackelberg games [17], [46], we first implement backward induction by considering each model owner's contract-theoretic incentive mechanism separately while taking the parameters of other model owners' decisions, i.e., prevailing data quantity and quality in the federation, as

a given. As we have previously elaborated in Section I, the self-revealing mechanisms in contract theory allows incentive compatibility even where there is information asymmetry. Then, we explore federation formation as a coalitional game (\mathcal{N}, v) with transferable utility [47] in which each model owner chooses the federation that enables them to maximize their profits with consideration for incentive expenses.

IV. CONTRACT-THEORETIC INCENTIVE MECHANISM DESIGN

We begin by studying the contract-theoretic incentive mechanism design of a representative model owner. Contrary to classical works in contract theory, the model owner now aims to maximize its payoff from a federation. As such, this payoff is dependent on other model owners in the federation as well. Following the backward induction solution, we first take the prevailing federation parameters as a given, before formulating the optimal contract following [48].

A. Problem Formulation

For ease of notation, we drop the i subscripts and focus on a representative model owner for now. The utility maximization problem of worker type m denoted u_m is as follows:

$$\max_{(R_m, q_m)} u_m = \theta_m R_m - c q_m. \quad (4)$$

θ_m represents willingness to participate [48] for the type m worker, where $m \in \{1, \dots, M\}$ and $\theta_1 < \dots < \theta_m < \dots < \theta_M$. As we have previously established in Section III, a more willing worker provides higher quality data. c represents the cost incurred per unit of data collected and q_m refers to units of data collected. In the contract theoretic formulation, each worker is presented with bundles $\{R_m, q_m\}$ of which it chooses the bundle that best maximizes its utility u_m .

Each model owner aims to maximize its share of payoffs derived from the federation. In fact, maximizing this payoff is equivalent to maximizing its marginal contribution to the federation. The marginal contribution of model owner $\{i\}$ when it joins a federation with $|S| - 1$ members is as follows:

$$\begin{aligned} & v(S) - v(S \setminus \{i\}) \\ &= \left[\omega(1 - e^{-\phi(\bar{\theta}_S Q_S)^v}) - |S|G \right] \\ &\quad - \left[\omega(1 - e^{-\phi(\bar{\theta}_{S \setminus \{i\}} Q_{S \setminus \{i\}})^v}) - (|S| - 1)G \right] \\ &= \omega \left(1 - e^{-\phi(\bar{\theta}_S Q_S)^v} \right) \\ &\quad - \omega \left(1 - e^{-\phi(\bar{\theta}_{S \setminus \{i\}} Q_{S \setminus \{i\}})^v} \right) - G \\ &= \omega \left(1 - e^{-\phi(\bar{\theta}_i Q_i + \sum_{j=1}^{S \setminus \{i\}} \bar{\theta}_j Q_j)^v} \right) \\ &\quad - \omega \left(1 - e^{-\phi(\sum_{j=1}^{S \setminus \{i\}} \bar{\theta}_j Q_j)^v} \right) - G. \end{aligned} \quad (5)$$

For ease of the presentation, we define $\alpha = \sum_{j=1}^{S \setminus \{i\}} \bar{\theta}_j Q_j$. Similarly, we denote $\eta = e^{-\phi(\sum_{j=1}^{S \setminus \{i\}} \bar{\theta}_j Q_j)^v}$. Note that α and η represents the *prevailing* conditions, i.e., of data quality and quantity, in the coalition before model owner i joins. For

TABLE I
TABLE OF COMMONLY USED NOTATIONS.

Notation	Description
θ	Worker's level of willingness/data quality
q	Worker data quantity contribution
Q	Aggregate data quality
$\bar{\theta}$	Aggregate data quantity
ω	Conversion parameter from model performance to profits
ϕ, v	Conversion parameters from data quantity and quality to model performance
$ S $	Number of members in federation
R	Contract rewards
c	Unit cost of data collection

purposes of backward induction, we first take the prevailing parameters, i.e., α and η as given. Following (5), we have the marginal contribution of model owner $\{i\}$ simplified as

$$v(S) - v(S \setminus \{i\}) = \omega(-e^{-\phi(\alpha + \bar{\theta}Q_i)^v} + \eta) - G. \quad (6)$$

As such, the model owner profit maximization function π is as follows:

$$\begin{aligned} \max_{(R_m, q_m)} \pi &= \omega \left(-e^{-\phi(\alpha + \sum_{m=1}^M \theta_m \rho_m p q_m)^v} + \eta \right) \\ &\quad - \sum_{m=1}^M \left(\rho_m p R_m + \gamma p \frac{q_m \psi}{\theta_m} \right) - G - K. \end{aligned} \quad (7)$$

ρ_m refers to the proportion of worker type m where $\sum_{m=1}^M \rho_m = 1$, p refers to number of participating workers, and R_m refers to the expenses incurred for paying each worker of type m for its data collection efforts. Following [32], γ refers to energy cost per iteration of computation, ψ is the coefficient that determines how local data quality affects number of computation iterations required, $\sum_{m=1}^M \gamma \frac{\rho_m p q_m \psi}{\rho_m p \theta_m / p} = \gamma p \frac{q_m \psi}{\theta_m}$ refers to number of computation iterations required for the corresponding aggregate data quality and quantity, and K refers to communication cost incurred in uploading the weight parameters to the aggregating server.

We assume the compute resources, e.g., CPU cycle frequency, contributed by each model owner to be a constant throughout the federation. Instead, we consider computation cost mainly as a function of aggregate data quality and quantity to study how it affects contract formulation. A higher data quality leads to fewer local model iterations [49], [50] whereas higher data quantity is more costly to train. Without loss of generality, we assume the transmission bandwidth, power, and channel gain to be a constant given similar wireless communication environments [45]. As such, communication costs K is a constant regardless of data quantity or quality. This is given that only the fixed-size model parameters \mathbf{w} are uploaded to the aggregating server. For ease of reference, we refer the readers to Table I for commonly used notations.

For feasibility, each contract must satisfy the following constraints.

Definition 1. Individual Rationality (IR): Each worker only participates in data collection when its utility is not less than zero, i.e.,

$$u_m = \theta_m R_m - c q_m \geq 0. \quad (8)$$

Definition 2. Incentive Compatibility (IC): Each worker of type m only chooses the contract designed for its type, i.e., (R_m, q_m) , instead of any other contracts to maximize utility, i.e.,

$$\theta_m R_m - c q_m \geq \theta_m R_z - c q_z, z \neq m. \quad (9)$$

However, this implies that we have to deal with M IR constraints and $M(M-1)$ IC constraints which are all non-convex. As such, we proceed to reduce and relax the conditions that guarantee a feasible contract.

B. Contract Feasibility

Lemma 1. For any feasible contract, we have $R_m \geq R_z$ if and only if $\theta_m \geq \theta_z$, $z \neq m$, $\forall m, z \in \{1, \dots, M\}$

Proof. The lemma is proven by using the IC constraint presented in (2). We first prove the sufficiency, i.e., if $\theta_m \geq \theta_z$, it follows that $R_m \geq R_z$.

From the IC constraints, we have

$$\begin{aligned} \theta_m R_m - c q_m &\geq \theta_m R_z - c q_z \text{ and} \\ \theta_z R_z - c q_z &\geq \theta_z R_m - c q_m. \end{aligned}$$

Then, we add the two inequalities to have

$$\begin{aligned} \theta_m R_m + \theta_z R_z &\geq \theta_m R_z + \theta_z R_m, \\ \theta_m R_m - \theta_z R_m &\geq \theta_m R_z - \theta_z R_z, \\ R_m(\theta_m - \theta_z) &\geq R_z(\theta_m - \theta_z). \end{aligned}$$

Since $\theta_m - \theta_z \geq 0$, it follows that $R_m \geq R_z$. Next, we prove the necessity, i.e., if $R_m \geq R_z$, it follows that $\theta_m \geq \theta_z$. Similarly, from the IC constraints, we have

$$\theta_m(R_m - R_z) \geq \theta_z(R_m - R_z).$$

Since $R_m \geq R_z \geq 0$, it follows that $\theta_m \geq \theta_z$. As such, we have proven that $R_m \geq R_z$ if and only if $\theta_m \geq \theta_z$. \square

Following Lemma 1, workers with higher willingness to participate, i.e., higher θ , will receive more rewards. Since $\theta_1 < \dots < \theta_m < \dots < \theta_M$ and R is a strictly increasing function of q , i.e., the contract bundles are designed such that higher data quantities contributed translate to higher rewards, we have the following monotonicity condition of the contract, thereby forming the necessary conditions for a feasible contract.

Theorem 1. Monotonicity: A feasible contract must meet the following conditions:

$$\begin{cases} 0 \leq R_1 \leq \dots \leq R_i \leq \dots \leq R_M \\ 0 \leq q_1 \leq \dots \leq q_i \leq \dots \leq q_M. \end{cases} \quad (10)$$

C. Devising an Optimal Contract

Next, we proceed to relax the IC and IR constraints with reference to the methods adopted in [51].

Lemma 2. (Reduce IR Constraints): If the IR constraint of worker type-1 is satisfied, the other IR constraints will also hold.

Proof. Following the IC constraints and the condition $\theta_1 < \dots < \theta_m < \dots < \theta_M$, we have

$$\theta_i R_i - cq_i \geq \theta_i R_1 - cq_1 \geq \theta_1 R_1 - cq_1 \geq 0.$$

As such, if the IR constraint of worker type-1 is satisfied, it follows that the other IR constraints automatically hold. Note that worker type-1 refers to the least willing worker, i.e., worker with lowest data quality. \square

Lemma 3. (Reduce IC Constraints): The IC constraints can be reduced into the Local Downward Incentive Constraints (LDIC).

Proof. Consider three worker types where $\theta_{m-1} < \theta_m < \theta_{m+1}$. The two LDICs, i.e., constraints between type m and type $m-1$ workers, are as follows,

$$\begin{aligned} \theta_{m+1} R_{m+1} - cq_{m+1} &\geq \theta_{m+1} R_m - cq_m, \text{ and} \\ \theta_m R_m - cq_m &\geq \theta_m R_{m-1} - cq_{m-1}. \end{aligned}$$

From Lemma 1, we have $R_m \geq R_z$ when $\theta_m \geq \theta_z$. As such, we can rewrite the LDICs as follows:

$$\begin{aligned} \theta_{m+1}(R_m - R_{m-1}) &\geq \\ \theta_m(R_m - R_{m-1}) &\geq c(q_m - q_{m-1}) \text{ and} \\ \theta_{m+1} R_{m+1} - cq_{m+1} &\geq \\ \theta_{m+1} R_m - cq_m &\geq \theta_{m+1} R_{m-1} - cq_{m-1}. \end{aligned}$$

As such, we have

$$\theta_{m+1} R_{m+1} - cq_{m+1} \geq \theta_{m+1} R_{m-1} - cq_{m-1}.$$

Hence, if the IC constraint holds for type- m worker, it will also hold for type $m-1$ worker. This process can be extended downward from type $m-1$ to type 1 worker, i.e., all DICs hold, as follows:

$$\begin{aligned} \theta_{m+1} R_{m+1} - cq_{m+1} &\geq \theta_{m+1} R_{m-1} - cq_{m-1} \\ &\geq \dots \\ &\geq \theta_{m+1} R_1 - cq_1, \\ M &> m \geq 1. \end{aligned}$$

A similar procedure can be taken to show that if the Local Upward Incentive Constraint (LUIC) holds, all UICs are also satisfied. Given the monotonicity condition in Definition 1, i.e., $R_m \geq R_{m-1}$, the LDIC also implies the Local Upward Incentive Constraint (LUIC) as follows.

$$\theta_{m-1} R_m - cq_m \leq \theta_{m-1} R_{m-1} - cq_{m-1}.$$

As such, we have proven that the IC constraints can be reduced to the LDIC constraint, since it also ensures that all UIC and DIC constraints hold. \square

D. Contract Optimality

With the constraints relaxed, we can rewrite the optimization problem as follows:

$$\begin{aligned} \max_{(R_m, q_m)} \quad & \pi = \omega \left(-e^{-\phi(\alpha + \sum_{m=1}^M \theta_m \rho_m p q_m)^v} + \eta \right) - \\ & \sum_{m=1}^M \left(\rho_m p R_m + \gamma p \frac{q_m \psi}{\theta_m} \right) - G - K, \\ \text{s.t.} \quad & \theta_1 R_1 - cq_1 \geq 0, \\ & \theta_m R_m - cq_m \geq \theta_m R_{m-1} - cq_{m-1}, \\ & 0 \leq q_1 \leq \dots \leq q_i \leq \dots \leq q_M, \\ & \sum_{m=1}^M q_m \leq q_{max}, \\ & \forall m \in \{1, \dots, M\}. \end{aligned} \tag{11}$$

Note that the last constraint refers to data quantity constraint specified for each worker. For example, in a classification task, each worker may belong to a specific class. The model does not benefit from training on duplicated inputs from the same worker of the same class [52]. In addition, this also serves as an indirect budget constraint. Following this, we first establish the dependence of optimal rewards \mathbf{R} on quantity of data provided \mathbf{q} . Thereafter, we solve the problem in (11) with \mathbf{q} only. Specifically, we obtain the optimal data rewards $R_m^*(\mathbf{q})$ ($1 \leq m \leq M$) given a set of feasible data contribution from each worker $\mathbf{q} = \{q_1, q_2, \dots, q_M\}$ which satisfies the monotonicity constraint $0 \leq q_1 \leq \dots \leq q_i \leq \dots \leq q_M$. The optimal rewarding scheme can be summarized in the following theorem.

Theorem 2. For a known set of data quantity \mathbf{q} satisfying $0 \leq q_1 \leq \dots \leq q_i \leq \dots \leq q_M$ in a feasible contract, the optimal reward is given by

$$R_m^* = \begin{cases} \frac{1}{\theta_m} cq_m, & \text{if } m = 1, \\ R_{m-1} - \frac{1}{\theta_m} cq_{m-1} + \frac{1}{\theta_m} cq_m, & \text{if } m = 2, 3, \dots, M. \end{cases} \tag{12}$$

Proof. We adopt the proof by contradiction to validate this theorem. We first assume that there exist some \mathbf{R}^\dagger in the feasible contract that yields greater profit for the model owner, meaning that the theorem is not correct, i.e., $\pi(\mathbf{R}^\dagger) > \pi(\mathbf{R}^*)$.

For simplicity, we need to consider only the rewards portion of the model owner's profit function in this proof, i.e., $\sum_{m=1}^M R_m^\dagger < \sum_{m=1}^M R_m^*$. This implies there exists at least a $t \in \{1, 2, \dots, M\}$ that satisfies the inequality $R_t^\dagger < R_t^*$. According to the LDIC conditions presented in Lemma (3), we have

$$R_t^\dagger \geq R_{t-1}^\dagger - \frac{1}{\theta_t} cq_{t-1} + \frac{1}{\theta_t} cq_t, \tag{13}$$

and

$$R_t^* = R_{t-1}^* - \frac{1}{\theta_t} cq_{t-1} + \frac{1}{\theta_t} cq_t. \tag{14}$$

From (13) and (14), we can deduce that $R_{t-1}^\dagger \leq R_{t-1}^*$. Continuing this process, we eventually obtain $R_1^\dagger \leq R_1^* = \frac{1}{\theta_1} cq_1$.

However, this violates the IR constraint we present in Lemma (2). Therefore, there does not exist the rewards R^\dagger in the feasible contract that yields greater profit for the model owner. Intuitively, the model owner chooses the lowest reward that satisfies the IR and IC constraints for profit maximization. The proof is now completed. \square

The optimal rewards given in (12) can be expressed as

$$R_T^* = \frac{1}{\theta_1} c q_1 + \sum_{t=1}^T \Delta_t, \quad (15)$$

where $\Delta_1 = 0$, $\Delta_t = -\frac{1}{\theta_t} c q_{t-1} + \frac{1}{\theta_t} c q_t$, and $T = 2, \dots, M$. To group all the same quantity types together, we follow the approach in [53] by rewriting the rewards as:

$$\begin{aligned} & \sum_{m=1}^M \rho_m p R_m \\ &= \sum_{m=1}^M \rho_m p \left(\frac{1}{\theta_1} c q_1 + \sum_{t=1}^m \Delta_t \right) \\ &= \sum_{m=1}^M \left(\rho_m p \frac{1}{\theta_m} c q_m + \Lambda_m \sum_{t=m+1}^M \rho_t p \right), \end{aligned}$$

where $\Lambda_m = \frac{1}{\theta_m} c q_m - \frac{1}{\theta_{m+1}} c q_m$ and $\Lambda_M = 0$.

Thereafter, this yields the following optimization problem,

$$\begin{aligned} & \text{maximize} \quad \sum_{m=1}^M G_m(q_m) \\ & \text{subject to} \quad 0 \leq q_1 \leq \dots \leq q_i \leq \dots \leq q_M, \end{aligned}$$

where

$$\begin{aligned} G_m &= \omega \left(-e^{-\phi(\alpha + \sum_{m=1}^M \theta_m \rho_m p q_m)^v} + \eta \right) - \\ & \sum_{m=1}^M \left(\rho_m p \frac{1}{\theta_m} c q_m + \Lambda_m \sum_{t=m+1}^M \rho_t p + \gamma p \frac{q_m \psi}{\theta_m} \right) - G - K. \end{aligned} \quad (16)$$

As the objective function $G_m(q_m)$ is structurally separate from different data quantities q_m , i.e., $G_m(q_m)$ is independent of $G_z(q_z)$, $\forall m, z \in \{1, \dots, M\}$, $z \neq m$, and thereby $G_m(q_m)$ is only associated with q_m . As such, the variable of each data quantity q_m can be derived by separately optimizing each $G_m(q_m)$ as follows:

$$\begin{aligned} q_m^* &= \arg \max_{q_m} \omega \left(-e^{-\phi(\alpha + \theta_m \rho_m p q_m)^v} + \eta \right) - \\ & \left(\rho_m p \frac{1}{\theta_m} c q_m + \Lambda_m \sum_{t=m+1}^M \rho_t p + \gamma p \frac{q_m \psi}{\theta_m} \right) - G - K, \end{aligned} \quad (17)$$

subject to the feasibility constraint where $0 \leq q_1 \leq \dots \leq q_i \leq \dots \leq q_M$. Using convex optimization tools such as *cvxpy* [54], we can first solve the relaxed problem by dropping the monotonicity condition, and then check whether the solution satisfies the monotonicity condition or not while also taking into account the quantity constraint of the model owner. If the solutions satisfy the monotonicity conditions, they are the optimal solutions. Otherwise, we can solve the infeasible subsequences using an iterative adjusted algorithm presented

in Algorithm 1, i.e., the Bunching and Ironing algorithm [55]. The algorithm iteratively adjusts the results such that it satisfies the monotonicity conditions. As the concavity of (17) is guaranteed, the solutions obtained are also globally optimal.

Algorithm 1 “Bunching and Ironing” Adjusted Algorithm

- 1: **Initialization:** Let $q_m^* = \arg \max_{q_m} G_m(q_m)$, $\forall m \in \{1, \dots, M\}$
 - 2: **while** The set of $\mathbf{q}^* = \{q_m^*\}$ violates the monotonicity constraint, **do**
 - 3: Find an infeasible sub-sequence $\{q_i^*, q_{i+1}^*, \dots, q_j^*\}$, where $q_i^* \leq \dots \leq q_j^*$ and $i < j$;
 - 4: Set $q_l^* = \arg \max_q \sum_{t=i}^j G_t(q)$, $\forall l \in \{i, i+1, \dots, j\}$;
 - 5: **end while**
 - 6: **Return** The feasible set $\mathbf{q}^* = \{q_m^*\}$, $m \in \{1, \dots, M\}$
-

For ease of presentation and derivation, we now re-express the optimization output with subscripts in our formulation. The optimal contract by model owner i results in the following aggregate data quantity and quality, where quality is computed as the weighted average across M worker types. Note that in our formulation, we have taken the prevailing federation conditions α and η as a given. When choosing between different federations to join subsequently, the contract formulation changes, i.e., we may vary the two variables.

$$\begin{cases} Q_i = \sum_{m=1}^M q_m^* \\ \bar{\theta}_i = \sum_{m=1}^M \frac{q_m^* \theta_m}{Q} \end{cases} \quad (18)$$

V. FEDERATION FORMATION

We now proceed to model the problem of federation formation as a coalitional game. In particular, we study the formation of stable federations of model owner given the different characteristics of data quantity and quality they have. Note that we will use the terms “*federation*” and “*coalition*” interchangeably in this section.

A. Coalitional Game Formulation and Properties

Property 1. The federation formation process is a coalitional game that has a *transferable* utility because the value of the federation, i.e., profit generated from model performance $v(S)$, can be arbitrarily apportioned [16] among players in the federation.

Property 2. The coalitional formation game does not always result in the formation of a grand coalition. Instead, independent disjoint coalitions will form in the network.

Proof. Following the procedure in [16], we first show the non-superadditivity of the coalitional game. Then, we proceed to prove that the core of the proposed game is empty. For purpose of this proof, we assume there exists two disjoint federations S_1 and S_2 where $S_1, S_2 \subset \mathcal{N}$. The two disjoint federations can choose to exist non-cooperatively, or form the grand coalition $S_1 \cup S_2$.

A coalitional game is superadditive if $v(S_1 \cup S_2) \geq v(S_1) + v(S_2)$. However, we observe from (2) that as $|S|$ increases,

i.e., the number of model owners in the federation increases, the marginal cost of coordination increases linearly whereas marginal benefits from gains in model performance falls. This is due to the diminishing returns to data quantity and quality. Thus, the exception to rule will be some $|S_1 \cup S_2|$ in which $\omega(1 - e^{-\phi(\bar{\theta}_S Q_S)^v}) - \omega(1 - e^{-\phi(\bar{\theta}_{S \setminus \{i\}} Q_{S \setminus \{i\}})^v}) \approx 0$, which implies $v(S) - v(S \setminus \{i\}) = -G < 0$. This further implies $v(S_1 \cup S_2) < v(S_1) + v(S_2)$, i.e., the coalitional game is non-superadditive since marginal gains from adding a member is negative.

Next, we proceed to prove the emptiness of the core. Following the definition in [16], a payoff vector $\mathbf{z} = (z_1, \dots, z_N)$ is an imputation if it is (i) group rational, i.e., $\sum_{i=1}^N z_i = v(N)$ and (ii) individually rational where $z_i \geq v(\{i\}) \forall i$, i.e., players can obtain benefit no less than acting alone. Following this definition, the *core* of a coalition refers to a set ζ of stable imputations such that

$$\zeta = \left\{ \mathbf{z} : \sum_{i \in N} z_i = v(N) \text{ and } \sum_{i \in S} z_i \geq v(S) \forall S \subset N \right\}$$

As we have previously established in this proof, the marginal gain of adding members into the coalition can be negative. This implies the violation of the individual rationality condition of an imputation, i.e., $z_i < v(S)$. Notably, the grand coalition does not form if there exists some combination of model owners with large data quantities and qualities $\bar{\theta} Q_S$. In the case, other model owners are better off existing in disjoint federations. The proof is now complete. \square

Following property 2, the grand coalition does not form. Instead, disjoint federations are expected to form in the network. We now proceed to establish two useful definitions from [47] that can be used to study the subsequent formation of a stable partition of coalitions.

Definition 3. Denote S as the set $S = \{S_1, \dots, S_l, \dots, S_L\}$ of mutually disjoint coalitions $S_l \in \mathcal{N}$. A collection of disjoint coalitions that spans all players in N , i.e., $\bigcup_{l=1}^L S_l = N$, is called a *partition* of N . In our system model, a partition is simply a federation of model owners. For discussion, we use the two terms interchangeably.

Definition 4. We define a *comparison relation* \triangleright for the purpose of comparing two collections, e.g., $R = \{R_1, \dots, R_l\}$ and $S = \{S_1, \dots, S_k\}$, both of which are partitions of the same subset $A \subseteq \mathcal{N}$. $R \triangleright S$ implies that the partition R is preferred to that of S based on the Pareto order we further define below.

Given the nature of our transferable utility game involving profit maximizing model owners, we adopt the *Pareto order* [56] that bases preferences on the payoff of individual model owners rather than the combined coalition value. The Pareto order is as follows

$$R \triangleright S \iff \{v_i(R) \geq v_i(S) \forall i \in R, S\},$$

with at least one strict equality ($>$) for a player k . The Pareto order $R \triangleright S$ can be interpreted as a preference relation in which partition R is preferred over S if at least one model owner in S

can improve its payoff through forming R , without decreasing the payoffs of other model owners throughout R and S . The latter condition ensures the stability of a partition, i.e., if any one model owner has lower utility, it will reject the formation of R .

B. Coalition Formation Algorithm

Following the merge and split algorithm proposed in [57] and applied in [56], we define two operations that can be used to modify a partition based on the Pareto order

- **Merge Rule** - Merge any set of coalitions $\{S_1, \dots, S_l\}$ where $\left\{ \bigcup_{j=1}^l S_j \right\} \triangleright \{S_1, \dots, S_l\}$, therefore $\{S_1, \dots, S_l\} \rightarrow \left\{ \bigcup_{j=1}^l S_j \right\}$
- **Split Rule** - Split any coalition $\bigcup_{j=1}^l S_j$ where $\{S_1, \dots, S_l\} \triangleright \left\{ \bigcup_{j=1}^l S_j \right\}$, therefore, $\left\{ \bigcup_{j=1}^l S_j \right\} \rightarrow \{S_1, \dots, S_l\}$

The merge and split rule also fulfills the internal and external stability notions presented in [58]. For an in-depth discussion of stability notions, we refer the interested readers to the studies conducted in [57].

We present the federation formation pseudocode in Algorithm 2. Denote an initial network partition as $T = \{T_1, \dots, T_l, \dots, T_L\}$, with each partition's member, i.e., model owner, denoted as T_{l_i} . The prevailing conditions of each partition, e.g., aggregate data quality and quantity, are known to model owners in the network (line 1) for example through reputation mechanisms [59].

We illustrate the merge mechanism from the perspective of a representative partition $T_l \in T$. The partition T_l considers a merge with other partitions $T_{l'}, l' \neq l, T_{l'} \in T$. During this process, each model owner $T_{l_i} \in T_l$ take as given the prevailing conditions in $T_{l'}$, e.g., current data quality and data quantity, i.e., $\theta_{T_{l'}}$ and $Q_{T_{l'}}$, respectively. Using this information, each model owner T_{l_i} formulates a hypothetical contract and computes the resulting payoff for the $L - 1$ possible partitions it can join (lines 6-10). By the Pareto order, the merge occurs where $\{T_l \cup T_{l'}\} \triangleright \{T_l, T_{l'}\}$. Specifically, the Pareto order is fulfilled under two conditions as follows (i) if the potential payoff from the merge exceeds the current payoff for at least one model owner and (ii) no other model owners in the current partitions are made worse off. If the two conditions are fulfilled, the merge will occur and a new partition $T^\dagger = \{T_l \cup T_{l'}\}$ is formed (lines 11-12). Then, the new partition's conditions and all model owners' payoffs are updated following (18) (line 13). In the next iteration, the new partition T^\dagger reconsiders a merge with another existing partition again. On the other hand, if there is no benefit to gain from merging, the partition T_l remains unchanged and we proceed to analyze the next partition's, i.e., T_{l+1} , merge decision (line 14-15). The merge mechanism stops when all partitions have been considered.

Following the termination of the merge process, we denote the set of modified partitions to be $\tilde{T} = \{\tilde{T}_1, \dots, \tilde{T}_j, \dots, \tilde{T}_J\}$ where each \tilde{T}_j denotes the resulting partition derived from the merge mechanism. The split mechanism is considered next.

Algorithm 2 Merge and Split algorithm for federation formation among model owners

```

1: Initialization: Random network partition
    $T = \{T_1, \dots, T_l, \dots, T_L\}$ , prevailing partition conditions
    $\{\theta_{T_l}, Q_{T_l}\}$ , model owner  $T_{l_i}$  characteristics
2: loop
3: Merge mechanism
4:  $\bar{l} = 1$ 
5: if  $\bar{l} \leq L$  then
6:   for  $T_{\bar{l}}$  in  $T$  do
7:     for each model owner  $T_{l_i}$  in  $T_{\bar{l}}$  do
8:       for  $T_{l'}$  in  $T$  do
9:         Formulate optimal contract
10:        Compute model owner payoff  $u_{T_{l_i}, T_{l'}}$ 
11:        if  $\{T_{\bar{l}} \cup T_{l'}\} \triangleright \{T_{\bar{l}}, T_{l'}\}$  then
12:           $T_{\bar{l}} \leftarrow T_{\bar{l}} \cup T_{l'}$ 
13:          Update  $\theta_{T_{\bar{l}}}, Q_{T_{\bar{l}}}, u_{T_{l_i}, T_{l'}}$ 
14:        else
15:           $\bar{l} = \bar{l} + 1$ 
16:        return  $\tilde{T} = \{T_1, \dots, \tilde{T}_j, \dots, \tilde{T}_J\}$ 
17: Split mechanism
18: for  $\tilde{T}_j$  in  $\tilde{T}$  do
19:    $r = |\tilde{T}_j| - 1$ 
20:   if  $r \geq 1$  then
21:     Compute split permutations  $\tilde{T}_{jz}^-$ 
22:     for each model owner  $\tilde{T}_{jz_i}^-$  in  $\tilde{T}_{jz}^-$  do
23:       Compute payoff
24:       if  $\tilde{T}_{jz_i}^- \triangleright \tilde{T}_j$  then,
25:          $\tilde{T}_j \leftarrow \tilde{T}_{jz_i}^-$ 
26: end loop
27: return  $T^* = \{T_1^*, \dots, T_a^*, \dots, T_A^*\}$ 

```

We first consider a split of the partition \tilde{T}_j into a partition with $|\tilde{T}_j| - 1$ members. Denote $\tilde{T}_j^- = \{\tilde{T}_{j1}^-, \dots, \tilde{T}_{jz}^-, \dots, \tilde{T}_{jZ}^-\}$ as the set of possible splits of a partition and $\tilde{T}_{jz_i}^-$ refers to a model owner i belonging in the split partition \tilde{T}_{jz}^- . The payoff for each split is considered (lines 19-22).

If the split fulfills the conditions of the Pareto order (lines 23-24), new partitions will be formed and updated. For splits into partitions involving more than one model owner, we let the prevailing condition of partitions be defined by the members with more data. As an illustration, we consider the scenario in which two members in partition \tilde{T}_{l_i} namely \tilde{T}_{l_i} and $\tilde{T}_{l_{i+1}}$ considers a split to form a two member partition. If $Q_{\tilde{T}_{l_i}} > Q_{\tilde{T}_{l_{i+1}}}$, \tilde{T}_{l_i} shall first formulate the contract and set the prevailing conditions for $\tilde{T}_{l_{i+1}}$ to follow. Similarly, if there are more than two members, the initial conditions shall be set by the members with more data first, i.e., in a descending order. In practice, this is realistic since members with too little data usually are unable to exist alone. As such, it is likely that they follow other members who can initiate the split with more data. Then, we iterate down the updated partitions to consider more possible and smaller splits until the split process terminates, i.e., the stable partition $T^* = \{T_1^*, \dots, T_a^*, \dots, T_A^*\}$ is returned.

VI. PERFORMANCE EVALUATION

In this section, we perform numerical experiments to evaluate our designed incentive mechanism. The network parameters are set as follows: $\tilde{\omega} = 10,000$, $\tilde{x} = 0.5$, $\phi = 0.01$, $v = 0.5$, $c = 0.1$, $\gamma = 1$, $\psi = 0.001$, $G = 5000$, $q_{max} = 10$ and $K = 0.1$. We also consider that the user parameter θ follows a normal distribution $\theta \sim \mathcal{N}(0.7, 0.1)$ unless otherwise stated. Then, we adopt the k -means clustering method to derive $M = 5$ clusters of users. Based on the above settings, we first evaluate the performance of our contract design for a representative model owner. Then, we proceed to study the static and dynamic federation formation.

A. Contract Optimality

In considering contract optimality, we first disregard the quantity constraints q_{max} from the optimization problem. This serves to enhance the visualization and evaluation of our contract design. Note that under the condition in which quantity constraints are in place, the contract is still optimal and the evaluations from this section still holds. The key difference lies in that if the constraint is binding, we will expect more user types to contribute the same data quantities.

We observe from Fig. 3 that both the quantity of data contributed and rewards for user types increase as the user's data quality increases. This implies that our designed contract satisfies the monotonicity constraint, as is validated in Theorem 1. Then, we study the utility of each user type in Fig. 4. Firstly, we observe that all the user utilities are positive, thus satisfying the Individual Rationality constraint that we provide in Definition 1. Secondly, we can observe that all users can achieve utility maximization only if they choose the contract items designed for their types. This is consistent with the Incentive Compatibility constraint that we define in Definition 2. The results in the Fig. 3 and Fig. 4 show the self-revealing properties of our optimal contract. Note that the bunching and ironing algorithm was not required in our numerical experiment because the current network parameters returned feasible sequences at its first attempt.

B. Static Federation Formation

We study the federation formation preferences of the representative model owner in a static setting, i.e., where there is only one iteration of merging from the perspective of a representative model owner. In our subsequent experiments, we present a representative model owner with existing federations that have varying data quantities and qualities. Then, we compute the potential profit the model owner can derive from the respective federations and study the preference of the model owner.

In Fig. 5, the model owner is presented with existing federations with prevailing data quantities ranging from 1000 to 5000 and an homogeneous average data quality of 0.9. Note that we also consider the scenario in which the model owner exists alone, i.e., without joining any federation. Then, we compute the profit that the model owner can potentially derive from joining each respective federations. If the model owner




Fig. 2. Data quantity vs. User types.




Fig. 3. Total rewards vs. User types.




Fig. 4. User utility vs. contract items.




Fig. 5. Profit vs. prevailing data quantity.




Fig. 6. Profit vs. prevailing data quality.

does not cooperate, we observe that it has zero profit, i.e., it chooses not to collect data and train a model since it does not have the capabilities to build a model that can generate positive profit. Similarly, if the model owner joins a federation with 1000 data units currently, the profit is negative, i.e., the resulting federation formed in this cooperation is not capable of covering the costs of cooperation since the combined model is not viable. We note that the maximum profit that can be derived comes from joining a federation with 2000 prevailing data units. If the model owner joins a larger federation, we observe a drop in its profit. The reason is that the marginal contribution it brings to a large federation is diminishing as the prevailing size of the federation increases.

In Fig. 6, we consider federations with a range of varying data qualities from 0.1 to 0.9. The federations can take on two prevailing data quantities, i.e., 5000 and 10,000. Then, we compute the profit that a model owner can potentially derive from joining the different federations with the respective combinations of data quantities and qualities. For federations with 5000 units of data, the model owner can derive maximum profit from the cooperation if the federation has an average data quality of 0.7. In contrast, for federations with 10,000 units of data, the model owner can derive maximum profit

from cooperation if the federation has an average data quality of 0.1.

The above experiments show that a model owner prefers to join federations with a *minimum* threshold of data quantity and quality such that a viable combined model can be built. On one hand, joining a federation with prevailing quantities and qualities below this threshold leads to losses since the model performance is insufficient to cover the cooperation costs. On the other hand, joining a federation that has data quantity and quality above the threshold implies that the model owner will have a lower payoff due to its declining marginal contribution. This insight serves to aid us in interpreting the merge and split decisions of the model owners subsequently.

C. Dynamic Federation Formation

We consider a dynamic federation formation, in which all the model owners considered can merge or split till an equilibrium stable partitions is achieved. We first initialize six model owners with 100, 200, 500, 750, 1500, and 2000 workers labeled model owners 1 to 6 respectively. For ease of illustration, we consider that each model owner has an average data quality of 0.7. Note that our subsequent results can also generalize to situations in which quality is heterogeneous.




Fig. 7. Evolution of federations with iterations. The solid lines imply a merge decision whereas dashed lines imply no changes.

We implement the proposed merge and split algorithm presented in Algorithm 1. Fig 7 shows the result of the merge and split. In fact, the federations are stable after 3 iterations of merges and did not require any splits. In each iteration, the model owner makes a merge or split decision based on the Pareto order presented in Section V-B.

In iteration 1, we observe that model owners 1 and 5 merge to form a federation whereas model owner 6 benefits more from being alone. The merge decision of 1 and 5 comes naturally. 1 requires the merge with a larger model owner, i.e., 5 or 6, to build a viable model. It prefers 5 since it is able to derive a higher marginal payoff. Similarly, 5 chooses 1 to also maximize its payoff, i.e., it does not require a merge with a larger model owner to be viable.

We observe that model owner 2 does not merge with any other model owners in the first iteration. This observation can be supported in our prior findings. Model owner 2 is unable to profit from a cooperation with 3 or 4 *alone* since the collaboration does not result in the production of a viable model, i.e., the resulting federation is still too small. On the other hand, it is unable to merge with 5 since 5 prefers forming a federation with 1. In iteration 2, model owner 2 finally merges with 3 and 4 since a viable model can now be built when it joins the relatively larger federation involving two other model owners.

In Fig. 8, we observe that the merge and split decisions follow the Pareto order, i.e., no model owner should be made worse off by a merge or split decision. It also supports our proof in V that the grand coalition does not form, due to the existence of large model owners, i.e., 5 and 6 in particular.

Fig. 8 also shows that our designed incentive mechanism does not incentivize model owners to hide their types. For example, the profits of model owner 1 is the lowest due to its cooperation with a large model owner, i.e., 5. However, it does not have the ability to join the federation with smaller model owners, i.e., 2, 3, 4. On the other hand, model owner 2 does not have the incentives to hide its type, e.g., pretend to be model owner 1, since its benefits from joining a federation comprised of smaller model owners imply that its marginal payoffs are




Fig. 8. Profit of each model owner for every iteration. Note that profit for iteration zero, i.e., initialization, is not computed.

eventually higher than that of 1. The same conclusion applies for the other model owners.

The numerical experiment also suggests to us the federation formation equilibrium in our proposed mechanism design. The largest model owners tend to stay alone, or cooperate with small model owners that are unable to build a viable model alone. On the other hand, medium sized model owners are likely to band together to maximize their marginal payoffs.

VII. CONCLUSIONS

In this paper, we have proposed a hierarchical incentive mechanism design for an FL-based crowdsensing network involving multiple model owners and federations. Using contract theory, we first propose an incentive design for model owners to incentivize high quality and quantity data from different worker types in the presence of information asymmetry. The contract is also designed to maximize the model owner's payoff in the federation. Then, we use the merge and split algorithm to study federation formation in the system model.

For our future works, we can consider more conditions that affect the collaboration of model owners. For example, some model owners may not be in close proximity to others, thus resulting in heterogeneous cooperation costs. In addition, we can also consider the possibility of competing or malicious model owners that can affect the performance of a federation adversely. For example, malicious model owners may intentionally provide erroneous inputs to corrupt the global model. In this case, reputation mechanisms [60] may be considered to exclude these model owners from joining the federation.

ACKNOWLEDGMENT

This research is supported, in part, by the National Research Foundation (NRF), Singapore, under Singapore Energy Market Authority (EMA), Energy Resilience, NRF2017EWTEP003-041, NRF2015-NRF-ISF001-2277, Singapore NRF National Satellite of Excellence, Design Science and Technology for Secure Critical Infrastructure NSOE DeST-SCI2019-0007, A*STAR-NTU-SUTD Joint Research Grant

Call on Artificial Intelligence for the Future of Manufacturing RGANS1906, WASP/NTU M4082187 (4080), AI Singapore Programme AISG-GC-2019-003, NRF-NRFI05-2019-0002, Singapore MOE Tier 2 MOE2014-T2-2-015 ARC4/15, MOE Tier 1 2017-T1-002-007 RG122/17 and U.S. National Science Foundation under Grant CCF-1908308. This research is also supported, in part, by the Alibaba-NTU JRI (Alibaba-NTU-AIR2019B1), NTU, Singapore. Qiang Yang also thanks the support of Hong Kong CERF grants 16209715 and 16244616.

REFERENCES

- [1] K. L. Lueth, "State of the IoT 2018: Number of IoT devices now at 7b – market accelerating," *IOT Analytics*, 2018.
- [2] B. Guo, Z. Yu, X. Zhou and D. Zhang, "From participatory sensing to mobile crowd sensing," in *2014 IEEE International Conference on Pervasive Computing and Communication Workshops (PERCOM WORKSHOPS)*. IEEE, 2014, pp. 593–598.
- [3] Y. LeCun, Y. Bengio and G. Hinton, "Deep learning," *nature*, vol. 521, no. 7553, pp. 436, 2015.
- [4] R. Pryss, M. Reichert, J. Herrmann, B. Langguth and W. Schlee, "Mobile crowd sensing in clinical and psychological trials—a case study," in *2015 IEEE 28th International Symposium on Computer-Based Medical Systems*. IEEE, 2015, pp. 23–24.
- [5] R. K. Ganti, F. Ye and H. Lei, "Mobile crowdsensing: current state and future challenges," *IEEE Communications Magazine*, vol. 49, no. 11, pp. 32–39, 2011.
- [6] M. Buda, A. Maki and M. A. Mazurowski, "A systematic study of the class imbalance problem in convolutional neural networks," *Neural Networks*, vol. 106, pp. 249–259, 2018.
- [7] T. Luo, S. S. Kanhere, J. Huang, S. K. Das and F. Wu, "Sustainable incentives for mobile crowdsensing: Auctions, lotteries, and trust and reputation systems," *IEEE Communications Magazine*, vol. 55, no. 3, pp. 68–74, 2017.
- [8] B. Custers, A. Sears, F. Dechesne, I. Georgieva, T. Tani and S. van der Hof, *EU Personal Data Protection in Policy and Practice*, Springer, 2019.
- [9] H. B. McMahan, E. Moore, D. Ramage, S. Hampson et al., "Communication-efficient learning of deep networks from decentralized data," *arXiv preprint arXiv:1602.05629*, 2016.
- [10] K. B. Letaief, W. Chen, Y. Shi, J. Zhang and Y.-J. A. Zhang, "The roadmap to 6G: AI empowered wireless networks," *IEEE Communications Magazine*, vol. 57, no. 8, pp. 84–90, 2019.
- [11] K. Yang, T. Jiang, Y. Shi and Z. Ding, "Federated learning via over-the-air computation," *IEEE Transactions on Wireless Communications*, 2020.
- [12] J. Nie, J. Luo, Z. Xiong, D. Niyato and P. Wang, "A stackelberg game approach toward socially-aware incentive mechanisms for mobile crowdsensing," *IEEE Transactions on Wireless Communications*, vol. 18, no. 1, pp. 724–738, 2018.
- [13] J. Nie, J. Luo, Z. Xiong, D. Niyato, P. Wang and M. Guizani, "An incentive mechanism design for socially aware crowdsensing services with incomplete information," *IEEE Communications Magazine*, vol. 57, no. 4, pp. 74–80, 2019.
- [14] P. Bolton, M. Dewatripont et al., *Contract theory*, MIT press, 2005.
- [15] H. Kim, J. Park, M. Bennis and S.-L. Kim, "On-device federated learning via blockchain and its latency analysis," *arXiv preprint arXiv:1808.03949*, 2018.
- [16] W. Saad, Z. Han, M. Debbah and A. Hjørungnes, "A distributed merge and split algorithm for fair cooperation in wireless networks," in *IEEE International Conference on Communications Workshops*. IEEE, 2008, pp. 311–315.
- [17] X. Kang and Y. Wu, "Incentive mechanism design for heterogeneous peer-to-peer networks: A stackelberg game approach," *IEEE Transactions on Mobile Computing*, vol. 14, no. 5, pp. 1018–1030, 2014.
- [18] Z. Xiong, J. Kang, D. Niyato, P. Wang and V. Poor, "Cloud/edge computing service management in blockchain networks: Multi-leader multi-follower game-based admm for pricing," *IEEE Transactions on Services Computing*, 2019.
- [19] Z. Duan, W. Li and Z. Cai, "Distributed auctions for task assignment and scheduling in mobile crowdsensing systems," in *2017 IEEE 37th International Conference on Distributed Computing Systems (ICDCS)*. IEEE, 2017, pp. 635–644.
- [20] X. Liu, K. Ota, A. Liu and Z. Chen, "An incentive game based evolutionary model for crowd sensing networks," *Peer-to-Peer Networking and Applications*, vol. 9, no. 4, pp. 692–711, 2016.
- [21] M. Dai, Z. Su, Y. Wang and Q. Xu, "Contract theory based incentive scheme for mobile crowd sensing networks," in *2018 International Conference on Selected Topics in Mobile and Wireless Networking (MoWNeT)*. IEEE, 2018, pp. 1–5.
- [22] D. G. Dobakhshari, P. Naghizadeh, M. Liu and V. Gupta, "A reputation-based contract for repeated crowdsensing with costly verification," *IEEE Transactions on Signal Processing*, vol. 67, no. 23, pp. 6092–6104, 2019.
- [23] X. Zhang, Z. Yang, W. Sun, Y. Liu, S. Tang, K. Xing and X. Mao, "Incentives for mobile crowd sensing: A survey," *IEEE Communications Surveys & Tutorials*, vol. 18, no. 1, pp. 54–67, 2015.
- [24] L. G. Jaimes, I. J. Vergara-Laurens and A. Raij, "A survey of incentive techniques for mobile crowd sensing," *IEEE Internet of Things Journal*, vol. 2, no. 5, pp. 370–380, 2015.
- [25] M. Li, J. Lin, D. Yang, G. Xue and J. Tang, "Quac: quality-aware contract-based incentive mechanisms for crowdsensing," in *2017 IEEE 14th International Conference on Mobile Ad Hoc and Sensor Systems (MASS)*. IEEE, 2017, pp. 72–80.
- [26] J. Wang, J. Tang, D. Yang, E. Wang and G. Xue, "Quality-aware and fine-grained incentive mechanisms for mobile crowdsensing," in *2016 IEEE 36th International Conference on Distributed Computing Systems (ICDCS)*. IEEE, 2016, pp. 354–363.
- [27] G. Yang, S. He, Z. Shi and J. Chen, "Promoting cooperation by the social incentive mechanism in mobile crowdsensing," *IEEE Communications Magazine*, vol. 55, no. 3, pp. 86–92, 2017.
- [28] D. Chatzopoulos, S. Gujar, B. Faltings and P. Hui, "Privacy preserving and cost optimal mobile crowdsensing using smart contracts on blockchain," in *2018 IEEE 15th International Conference on Mobile Ad Hoc and Sensor Systems (MASS)*. IEEE, 2018, pp. 442–450.
- [29] Y. Liu, Z. Ma, X. Liu, S. Ma, S. Nepal and R. Deng, "Boosting privately: Privacy-preserving federated extreme boosting for mobile crowdsensing," *arXiv preprint arXiv:1907.10218*, 2019.
- [30] K. Powell, "Nvidia clara federated learning to deliver ai to hospitals while protecting patient data," *Nvidia Blog*.
- [31] T. S. Brisimi, R. Chen, T. Mela, A. Olshevsky, I. C. Paschalidis and W. Shi, "Federated learning of predictive models from federated electronic health records," *International journal of medical informatics*, vol. 112, pp. 59–67, 2018.
- [32] J. Kang, Z. Xiong, D. Niyato, H. Yu, Y.-C. Liang and D. I. Kim, "Incentive design for efficient federated learning in mobile networks: A contract theory approach," in *2019 IEEE VTS Asia Pacific Wireless Communications Symposium (APWCS)*. IEEE, 2019, pp. 1–5.
- [33] S. Feng, D. Niyato, P. Wang, D. I. Kim and Y.-C. Liang, "Joint service pricing and cooperative relay communication for federated learning," in *IEEE Green Computing and Communications (GreenCom)*. IEEE, 2019, pp. 815–820.
- [34] Z. Xiong, S. Feng, D. Niyato, P. Wang, A. Lessem and Z. Han, "Joint sponsored and edge caching content service market: A game-theoretic approach," *IEEE Transactions on Wireless Communications*, vol. 18, no. 2, pp. 1166–1181, 2019.
- [35] Y. Zhan, P. Li, Z. Qu, D. Zeng and S. Guo, "A learning-based incentive mechanism for federated learning," *IEEE Internet of Things Journal*, 2020.
- [36] W. Y. B. Lim, N. C. Luong, D. T. Hoang, Y. Jiao, Y.-C. Liang, Q. Yang, D. Niyato and C. Miao, "Federated learning in mobile edge networks: A comprehensive survey," *arXiv preprint arXiv:1909.11875*, 2019.
- [37] J. Kang, Z. Xiong, D. Niyato, Y. Zou, Y. Zhang and M. Guizani, "Reliable federated learning for mobile networks," *IEEE Wireless Communications*, 2020.
- [38] L. G. Jaimes and R. Steele, "Incentivization for health crowdsensing," in *2017 IEEE 15th Intl Conf on Dependable, Autonomic and Secure Computing, 15th Intl Conf on Pervasive Intelligence and Computing, 3rd Intl Conf on Big Data Intelligence and Computing and Cyber Science and Technology Congress (DASC/PiCom/DataCom/CyberSciTech)*. IEEE, 2017, pp. 139–146.
- [39] P. Vepakomma, O. Gupta, A. Dubey and R. Raskar, "Reducing leakage in distributed deep learning for sensitive health data," *arXiv preprint arXiv:1812.00564*, 2019.
- [40] J. Nie, Z. Xiong, D. Niyato, P. Wang and J. Luo, "A socially-aware incentive mechanism for mobile crowdsensing service market," in *IEEE Global Communications Conference (GLOBECOM)*. IEEE, 2018, pp. 1–7.
- [41] G. Cong, W. Fan, F. Geerts, X. Jia and S. Ma, "Improving data quality: Consistency and accuracy," in *Proceedings of the 33rd international*

conference on Very large data bases. VLDB Endowment, 2007, pp. 315–326.

- [42] Q. Liu, S. Huang, J. Opadere and T. Han, “An edge network orchestrator for mobile augmented reality,” in *IEEE INFOCOM 2018-IEEE Conference on Computer Communications*. IEEE, 2018, pp. 756–764.
- [43] X.-W. Chen and X. Lin, “Big data deep learning: challenges and perspectives,” *IEEE Access*, vol. 2, pp. 514–525, 2014.
- [44] T. T. Anh, N. C. Luong, D. Niyato, D. I. Kim and L.-C. Wang, “Efficient training management for mobile crowd-machine learning: A deep reinforcement learning approach,” *IEEE Wireless Communications Letters*, 2019.
- [45] N. H. Tran, W. Bao, A. Zomaya and C. S. Hong, “Federated learning over wireless networks: Optimization model design and analysis,” in *IEEE INFOCOM 2019-IEEE Conference on Computer Communications*. IEEE, 2019, pp. 1387–1395.
- [46] Z. Xiong, S. Feng, D. Niyato, P. Wang, Y. Zhang and B. Lin, “A stackelberg game approach for sponsored content management in mobile data market with network effects,” *IEEE Internet of Things Journal*, 2020.
- [47] W. Saad, Z. Han, M. Debbah, A. Hjørungnes and T. Basar, “Coalitional games for distributed collaborative spectrum sensing in cognitive radio networks,” in *IEEE INFOCOM 2009*. IEEE, 2009, pp. 2114–2122.
- [48] Y. Zhang, L. Song, W. Saad, Z. Dawy and Z. Han, “Contract-based incentive mechanisms for device-to-device communications in cellular networks,” *IEEE Journal on Selected Areas in Communications*, vol. 33, no. 10, pp. 2144–2155, 2015.
- [49] T. Nishio and R. Yonetani, “Client selection for federated learning with heterogeneous resources in mobile edge,” in *IEEE International Conference on Communications (ICC)*. IEEE, 2019, pp. 1–7.
- [50] J. Konečný, H. B. McMahan, D. Ramage and P. Richtárik, “Federated optimization: Distributed machine learning for on-device intelligence,” *arXiv preprint arXiv:1610.02527*, 2016.
- [51] Y. Zhang and Z. Han, *Contract Theory for Wireless Networks*, Springer, 2017.
- [52] C. Huang, Y. Li, C. Change Loy and X. Tang, “Learning deep representation for imbalanced classification,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 5375–5384.
- [53] D. H. Nguyen, Y. Zhang and Z. Han, “A contract-theoretic approach to spectrum resource allocation in wireless virtualization,” in *IEEE Global Communications Conference (GLOBECOM)*. IEEE, 2016, pp. 1–6.
- [54] S. Diamond and S. Boyd, “Cvxpy: A python-embedded modeling language for convex optimization,” *The Journal of Machine Learning Research*, vol. 17, no. 1, pp. 2909–2913, 2016.
- [55] L. Gao, X. Wang, Y. Xu and Q. Zhang, “Spectrum trading in cognitive radio networks: A contract-theoretic modeling approach,” *IEEE Journal on Selected Areas in Communications*, vol. 29, no. 4, pp. 843–855, 2011.
- [56] W. Saad, Z. Han, M. Debbah, A. Hjørungnes and T. Basar, “Coalitional game theory for communication networks: A tutorial,” *arXiv preprint arXiv:0905.4057*, 2009.
- [57] K. R. Apt and T. Radzik, “Stable partitions in coalitional games,” *arXiv preprint cs/0605132*, 2006.
- [58] D. Niyato, P. Wang, W. Saad and A. Hjørungnes, “Coalition formation games for bandwidth sharing in vehicle-to-roadside communications,” in *2010 IEEE Wireless Communication and Networking Conference*. IEEE, 2010, pp. 1–5.
- [59] P. Michiardi and R. Molva, “Core: a collaborative reputation mechanism to enforce node cooperation in mobile ad hoc networks,” in *Advanced communications and multimedia security*, pp. 107–121. Springer, 2002.
- [60] J. J. Jaramillo and R. Srikant, “A game theory based reputation mechanism to incentivize cooperation in wireless ad hoc networks,” *Ad Hoc Networks*, vol. 8, no. 4, pp. 416–429, 2010.



Wei Yang Bryan Lim graduated with double First Class Honours in Economics and Business Administration (Finance) from the National University of Singapore (NUS) in 2018. He is currently an Alibaba PhD candidate with the Alibaba Group and Alibaba-NTU Joint Research Institute, Nanyang Technological University, Singapore. His research interests include Federated Learning and Edge Intelligence.



Zehui Xiong (S'17) received his B.Eng degree with the highest honors in Telecommunication Engineering from Huazhong University of Science and Technology, Wuhan, China, in Jul 2016. From Aug 2016 to Oct 2019, he pursued the Ph.D. degree in the School of Computer Science and Engineering, Nanyang Technological University, Singapore. Since Nov 2019, he has been with Alibaba-NTU Singapore Joint Research Institute. He was a visiting scholar with Department of Electrical Engineering at Princeton University from Jul to Aug 2019. He was also a visiting scholar with BCCR lab in Department of Electrical and Computer Engineering at University of Waterloo from Dec 2019 to Jan 2020. His research interests include resource allocation in wireless communications, network games and economics, blockchain, and edge intelligence.



Chunyan Miao received the BS degree from Shandong University, Jinan, China, in 1988, and the MS and PhD degrees from Nanyang Technological University, Singapore, in 1998 and 2003, respectively. She is currently a professor in the School of Computer Science and Engineering, Nanyang Technological University (NTU), and the director of the Joint NTU-UBC Research Centre of Excellence in Active Living for the Elderly (LILY). Her research focus on infusing intelligent agents into interactive new media (virtual, mixed, mobile, and pervasive media) to create novel experiences and dimensions in game design, interactive narrative, and other real world agent systems.



Dusit Niyato (M'09-SM'15-F'17) is currently a professor in the School of Computer Science and Engineering and, by courtesy, School of Physical & Mathematical Sciences, at the Nanyang Technological University, Singapore. He received B.E. from King Mongkut's Institute of Technology Ladkrabang (KMUTL), Thailand in 1999 and Ph.D. in Electrical and Computer Engineering from the University of Manitoba, Canada in 2008. He has published more than 380 technical papers in the area of wireless and mobile networking, and is an inventor of four US and German patents. He has authored four books including “Game Theory in Wireless and Communication Networks: Theory, Models, and Applications” with Cambridge University Press. He won the Best Young Researcher Award of IEEE Communications Society (ComSoc) Asia Pacific (AP) and The 2011 IEEE Communications Society Fred W. Ellersick Prize Paper Award. Currently, he is serving as a senior editor of IEEE Wireless Communications Letter, an area editor of IEEE Transactions on Wireless Communications (Radio Management and Multiple Access), an area editor of IEEE Communications Surveys and Tutorials (Network and Service Management and Green Communication), an editor of IEEE Transactions on Communications, an associate editor of IEEE Transactions on Mobile Computing, IEEE Transactions on Vehicular Technology, and IEEE Transactions on Cognitive Communications and Networking. He was a guest editor of IEEE Journal on Selected Areas on Communications. He was a Distinguished Lecturer of the IEEE Communications Society for 2016-2017. He was named the 2017, 2018, 2019 highly cited researcher in computer science. He is a Fellow of IEEE.



Qiang Yang (F'09) is the head of AI at WeBank (Chief AI Officer) and Chair Professor at the Computer Science and Engineering (CSE) Department of Hong Kong University of Science and Technology (HKUST), where he was a former head of CSE Department and founding director of the Big Data Institute (2015-2018). His research interests include artificial intelligence, machine learning and data mining, especially in transfer learning, automated planning, federated learning and case-based reasoning. He is a fellow of several international societies,

including ACM, AAAI, IEEE, IAPR and AAAS. He received his PhD from the Computer Science Department in 1989 and MSc in Astrophysics in 1985, both from the University of Maryland, College Park. He obtained his BSc in Astrophysics from Peking University in 1982. He had been a faculty member at the University of Waterloo (1989-1995) and Simon Fraser University (1995-2001). He was the founding Editor in Chief of the ACM Transactions on Intelligent Systems and Technology (ACM TIST) and IEEE Transactions on Big Data (IEEE TBD). He served as the President of International Joint Conference on AI (IJCAI, 2017-2019) and an executive council member of Association for the Advancement of AI (AAAI, 2016 - 2020). Qiang Yang is a recipient of several awards, including the 2004/2005 ACM KDDCUP Championship, the ACM SIGKDD Distinguished Service Award (2017) and AAAI Innovative AI Applications Award (2016, 2019). He was the founding director of Huawei's Noah's Ark Lab (2012-2014) and a co-founder of 4Paradigm Corp, an AI platform company. He is an author of several books including **Transfer Learning** (Cambridge Press), **Federated Learning** (Morgan Claypool), **Intelligent Planning** (Springer), **Crafting Your Research Future** (Morgan Claypool) and **Constraint-based Design Recovery for Software Engineering** (Springer).



Cyril Leung received the B.Sc. (First Class Hons.) degree from Imperial College, University of London, U.K., and the M.S. and Ph.D. degrees in electrical engineering from Stanford University. He has been an Assistant Professor with the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, and the Department of Systems Engineering and Computing Science, Carleton University. Since 1980, he has been with the Department of Electrical and Computer Engineering, University of British Columbia (UBC),

Vancouver, Canada, where he is a Professor and currently holds the PMC-Sierra Professorship in Networking and Communications. He served as an Associate Dean of Research and Graduate Studies with the Faculty of Applied Science, UBC, from 2008 to 2011. His research interests include wireless communication systems, data security and technologies to support ageless aging for the elderly. He is a member of the Association of Professional Engineers and Geoscientists of British Columbia, Canada.



H. Vincent Poor (S'72, M'77, SM'82, F'87) received the Ph.D. degree in EECS from Princeton University in 1977. From 1977 until 1990, he was on the faculty of the University of Illinois at Urbana-Champaign. Since 1990 he has been on the faculty at Princeton, where he is the Michael Henry Strater University Professor of Electrical Engineering. From 2006 until 2016 he also served as Dean of Princeton's School of Engineering and Applied Science. Dr. Poor's research interests are in the areas of information theory, signal processing and machine

learning, and their applications in wireless networks, energy systems, and related fields. Among his publications in these areas is the forthcoming book *Advanced Data Analytics for Power Systems* (Cambridge University Press, 2020).

Dr. Poor is a member of the National Academy of Engineering and the National Academy of Sciences, and is a foreign member of the Chinese Academy of Sciences, the Royal Society, and other national and international academies. Recent recognition of his work includes the 2017 IEEE Alexander Graham Bell Medal, the 2019 ASEE Benjamin Garver Lamme Award, a D.Sc. honoris causa from Syracuse University, awarded in 2017, and a D.Eng. honoris causa from the University of Waterloo, awarded in 2019.