Provider-Customer Coalitional Games

Chandramani Singh*, Saswati Sarkar, Alireza Aram

***Abstract*—Efficacy of commercial wireless networks can be substantially enhanced through large scale cooperation among involved entities such as providers and customers. The success of such cooperation is contingent upon the design of judicious resource allocation strategies that ensure that the individuals payoffs are commensurate to the resources they offer to the coalition. The resource allocation strategies depend on which entities are decision-makers, and whether and how they share their aggregate payoffs. Initially, we consider the scenario where the providers are the only decision-makers, and they do not share their payoffs. We formulate the resource allocation problem as a nontransferable payoff coalitional game and show that there exists a cooperation strategy that leaves no incentive for any subset of providers to split from the grand coalition, i.e., the *core* is nonempty. To compute this cooperation strategy and the corresponding payoffs, we subsequently relate this game and its core to an exchange market setting, and its equilibrium which can be computed by several efficient algorithms. Next, we investigate cooperation when customers are also decision makers and decide which provider to subscribe to, based on whether there is cooperation. We formulate a coalitional game in this setting and show that it has a nonempty core. We extend the formulations and results to the cases, where players assume more general payoff sharing relations, that is, sub-groups of providers can share the revenue among each other, and the benefits are modeled as "vector payoff functions", where only some components can be shared. Finally, we also consider multi-hop networks.**

I. INTRODUCTION

A. Motivation

We have witnessed a significant growth in commercial wireless services in the past few years, and the trend is likely to continue in the foreseeable future. This growth has been in part fueled by demand for new services such as network games and multimedia transmissions. These services are taxing the available transmission resources which are either limited (e.g., spectrum, transmission energy), or costly (e.g., infrastructure). Cooperation among service providers has the potential to substantially improve the resource utilization, and should therefore facilitate the proliferation of wireless services.

To serve its customers, each provider uses (i) wireless spectrum that it acquires either directly from central regulators such as the FCC or in secondary markets from other providers that have already licensed this spectrum from the regulators, and (ii) infrastructure such as base station, access points, mesh points (which we refer to as service units) that it deploys in its coverage area. Cooperation between providers entails pooling and sharing some of these resources to ultimately better serve the common pool of customers. This pooling and sharing of resources can improve coverage and throughput, which can in turn lead to higher customer satisfaction and higher revenues for the providers.

We now describe the benefits of cooperation among service providers. When different providers cooperate, their resources such as spectrum and infrastructure are likely to be optimally utilized. For example, if a provider's resources exceed traffic demands of its customers, it can use the underutilized portion to serve customers of other providers in its coalition, and enhance its profit. Similarly, even when its resources are congested, owing to poor propagation quality in the spectrum it owns, or temporary demand overload, it can deliver the desired quality of service to its customers using the resources of its collaborators. Such sharing turns out to be mutually beneficial as different providers are unlikely to experience poor quality of transmission and overload at the same time. Similarly, providers can augment their coverage areas by utilizing each others service units. Thus, overall, the providers can substantially enhance their net payoffs by cooperating.

B. Research challenges and Contributions

The success of this setup, however, is contingent on whether providers, as selfish entities, find the cooperation worthwhile. More specifically, a provider expects to receive a payoff commensurate to the resources such as service units and channels it offers the coalition, and the wealth it generates. The cooperation strategy of each provider involves the determination of which providers to cooperate with and how to share resources (i.e., the allocations of the service units and the spectrum to the customers). Design of rational cooperation strategies is imperative to motivate providers to participate in such cooperation. In particular, different choices of these decision variables determine the individual payoffs and the efficacy of cooperation. Also, collaborating providers may be able to share their profits. However, some providers may only be willing to collaborate so as to enhance individual profits, but may not be willing to share the profits, due to lack of trust, the nontransferable nature of the profit, or other reasons. In general, providers' total utilities could be a function of different types of payoffs, and they may be willing to share some types but not the rest. Coalition formation also depends on whether and how providers share their resulting payoffs. Finally, cooperation among providers could have negative effects on the customer base of some providers. A successful cooperation strategy may as well be required to guarantee that such potential downside of cooperation does not outweigh its

*C. Singh is in the Dep. of Electrical Communications Eng., Indian Institute of Science, Bangalore, and his research has been supported by an INRIA Associates program DAWN, and also by the Indo-French Centre for the Promotion of Advanced Research (IFCPAR), Project No. 4000-IT-A.

S. Sarkar & A. Aram are in the Dep. of Electrical and Systems Eng. and Wharton School at the University of Pennsylvania. Their contributions have been supported by NSF grants NCR- 0238340, CNS-0721308, ECS-0622176.

upside.

We introduce our communication model, notion of coalitions and feasible actions and payoffs of providers in a coalition and relate them to realistic wireless systems (Section III). We present a coalitional game framework [1] for cooperation among providers in a single-hop network using tools from cooperative game theory (Section IV-A). In particular, we first investigate providers' cooperative resource allocation in the scenario, where providers do not engage in payoff sharing. Using this framework, we show that there exists an operating point and corresponding payoffs that renders it optimal for all providers to cooperate. Specifically, if a subset of providers leave the grand coalition, regardless of how they cooperate among each other, at least one provider will be worse off (Section IV-A). In the cooperative game theory terminology, this is equivalent to saying that the core of the game is nonempty. To compute such an operating point, we next construct an "exchange market" setting, where service providers are considered to be agents in the market, and service units and channels are the goods (Section IV-B). Agents will then trade goods so as to maximize their own benefits. We show that in this setting, market equilibrium exists. Furthermore, we show that the allocation of goods in the economy given by the market equilibrium can be translated to a cooperation strategy among providers with the corresponding payoffs in the core. As a result, we can compute an element of the core, by computing the market equilibrium, which is possible by using several available algorithms. This result is also of independent interest, as it links two different concepts in this context.

Next we study cooperation in a scenario, where customers are also decision makers. Particularly, customers can subscribe to the provider of their choice, and that choice can depend on providers' cooperation decisions (section V). We propose a cooperation model and show that the core of this game is nonempty. Subsequently, we examine an algorithm to obtain a core element in this game.

Finally, we generalize our framework to accommodate a) more general payoff sharing rules, such as when there are groups of providers, and providers within each group would share their payoffs, while those in different groups would not, b) vector payoff functions that are comprised of mixed transferable and nontransferable components of different types, and c) multi-hop wireless networks where providers use ciutomers as potential relays. We formulate three coalitional games with the above three generalizations (section VI). We show that the previous results extend to these scenarios as well. Considering different payoff functions, we numerically evaluate the providers' and customers' payoff increases resulting from cooperation for a range of available demands (customers) and assets (base stations, spectrum) (Section VII).

II. Related Work

Interactions between different entities in wireless networks have primarily been investigated from the following extreme perspectives. In the first, each entity is assumed to select its actions so as to maximize its individual incentive without coordinating with others, e.g., [2]–[8]. This scenario, which has been investigated using noncooperative game theory, in general suffers from inefficient utilization of resources [9]. The other perspective has been to assume that entities selflessly choose their actions so as to optimize a global utility function even when such actions may deteriorate individual incentives of some entities (e.g., [2], [6]). We investigate interactions among providers assuming that each provider would be willing to cooperate and coordinate its actions with others when such cooperation enhances its individual incentives.

We obtained optimal cooperation schemes using the framework of cooperative game theory. This choice of tools allowed us to combine the desirable features of the extreme approaches studied in the existing literature, that of allowing entities to choose their actions guided by selfish objectives, and of maximizing global utility functions. Surprisingly, cooperative game theory has seen only limited use in wireless context so far. For bandwidth allocation among mobiles in heterogeneous wireless access environments, a Shapley value based algorithm is proposed in [10]. Nash bargaining solutions have been proposed for power control and spectrum sharing among multiple users [11]. Coalitional games have been used recently for modeling cooperation among nodes in the physical layer [12], [13], rate allocation in multiple access channels (MAC) [14], and rate allocation among mobiles and admission control in heterogeneous wireless access environments [10]. All these works use the framework of transferable utility coalitional games as they assume that players can share their aggregate payoffs in any arbitrary way. Non-transferable utility coalitional games have been used to model cooperation between single antenna receivers and transmitters in an interference channel [15], and to study collaborative sensing by secondary users in cognitive radio networks [16]. For an overview of applications of coalitional game theory in communication networks, see [17].

We also adopt the framework of non-transferable utility coalitional games. However, our problem formulation, solution techniques, and results significantly differ from the above owing to the difference in contexts - our focus is on cooperation among providers and customers at the network and MAC layers. Cooperation among providers has been previously studied in [18], using transferable payoff coalitional games. In the current work, we generalize formulations and results in [18], so as to consider (i) more general payoff sharing rules, (ii) the scenario, where customers are also players in the game, and (iii) vector payoff functions with components of various types. In order to obtain these results, we use different analytical tools and concepts, e.g., nontransferable payoff coalitional games and exchange market.

III. System Model

We present the communication model in Section III-A and describe in Section III-B how the formulations capture the essence of existing wireless technologies. We describe providers' coalitions and their payoffs in Section III-C.

A. Communication Model

Consider a network with a set of providers $\mathcal{N}$. Each provider i deploys a set of base stations (or access points) in order to serve its set of customers $\mathcal{M}_i$. Each base station has access to a certain set of channels (frequency bands)[1], and each base station-channel pair is referred to as a service unit. Thus, a provider's resources are its service-units. Let $\mathcal{B}_i$ be the set of service units of provider i, $\mathcal{B}_i \cap \mathcal{B}_j = \emptyset$ and $\mathcal{M}_i \cap \mathcal{M}_j = \emptyset$ for $i \neq j$. For a $\mathcal{S} \subseteq \mathcal{N}$, let $\mathcal{B}_{\mathcal{S}}$ and $\mathcal{M}_{\mathcal{S}}$ denote the set of service units and customers associated with providers in $\mathcal{S}$. Thus $\mathcal{B}_{\mathcal{N}}$ and $\mathcal{M}_{\mathcal{N}}$ are the sets of all service units and all customers in the network, respectively; we also define $\mathcal{M} = \mathcal{M}_{\mathcal{N}}$. We assume, unless mentioned otherwise, that the service units and the customers communicate via single-hop links. We assume that the achievable rates of a customer-service unit pair do not depend on communications of other customers and service units. Each customer or a service unit may be involved in at most one communication at a given time (*time sharing*). The system can be represented by a complete bipartite graph ($\mathcal{G} = (\mathcal{V}, \mathcal{E})$) where the customers and the service units represent the nodes and there exists a link between every customer-service unit node pair. Any customer-service unit assignment corresponds to a matching (a set of links such that no two links have a common node) in the above graph.

For ease of exposition, we consider only downlink communications in our model (the results easily extend to the case where communications involve both uplinks and downlinks). We consider frequency selective fading and assume that when customer j is served by service unit k, j receives at a rate r_{jk}, a random variable which is a function of the location of customer j and the state of channel k both of which can vary randomly. Let ω represent a network state (customer location, channel qualities resulting from fading and channel access of primary users), Ω be the collection of all ωs and $\mathbb{P}(\omega)$ be the probability that the network state is ω. The rates r_{jk} are functions of ω and are denoted as $r_{jk}(\omega)$. We assume that $|\Omega|$ is finite, since (i) feasible service rates in any practical communication system belong to a finite set, and (ii) we can partition the service region in such a way that the service rates received by the customers inside a member of the partition do not depend on the locations of the customers.

B. How the formulations relate to existing wireless networks

We now illustrate via examples how our framework can be used to model specific communication systems. Consider elastic data transfers in the downlink of a CDMA cellular system (e.g., used for internet access of cellular subscribers) [19, Chapter 5] with provider set $\mathcal{N}$. Owing to simplicity of physical layer implementations, a base station k always transmits at a pre-determined fixed power P_k (which may be different for different base stations). This happens even when no mobiles associated with it require downlink transmission [20]. Each base station has access to only one band and thus the base stations are the service units. Customers in a cell are served on *time-sharing* basis, i.e., a base station transmits to at most one customer at a given time. Also, at any given time, a customer receives transmissions from at most one base station. Then, $\{\alpha_{jk}(\omega)\}$ represent the fraction of times customers are served by different base stations. When base station k transmits to customer j and the channel gain realization is ω, the achievable rate $r_{jk}(\omega)$ from k to j is a function of the downlink SINR $\text{SINR}_{jk}(\omega)$ [19, Chapter 5], where

$$\text{SINR}_{jk}(\omega) = \frac{h_{jk}(\omega) P_k}{\sum_{i' \in \mathcal{B}_{\mathcal{N}} \setminus \{k\}} h_{ji'}(\omega) P_{i'} + N_0 W},$$

$h_{jk}(\omega)$ are the channel gains between customer-base station pairs, N_0 is the power spectral density of the additive noise and W is the spectrum bandwidth[2]. Thus, $\text{SINR}_{jk}(\omega)$ and hence $r_{jk}(\omega)$, is independent of which customers are being served by other base stations.

In a variant of the above service discipline (*power sharing*), a base station distributes its total power among the downlink transmissions in its cell. Orthogonal codes and chip synchronous transmissions can ensure that the intra-cell interference for a customer is negligible. As in the earlier case, the inter-cell interference remains fixed. Then $r_{jk}(\omega)$ is the fixed peak rate between customer j and base station k, and is achieved when k uses its entire power to transmit to j. The variables $\{\alpha_{jk}(\omega)\}$ account for the fractional power allocation[3]. We formulate the characteristic functions considering time-sharing, and point out the modifications required for incorporating power sharing.

Next, consider downlink communications in a multi-cell OFDMA system [19, Chapter 6]. Different providers acquire non-overlapping bands and the bandwidth acquired by a provider is divided into several channels (sub-carriers in OFDM terminology) (For small-scale providers, some of these channels can be secondary access channels or spectrum white-spaces acquired from primary users). In order to manage interference, each provider partitions the set of sub-carriers into reuse groups, assigning one such group of sub-carriers to each base-station so as to ensure that inter-cell interference to simultaneous transmissions in other base-station-sub-carrier pairs is negligible. At any given time, a base station assigns a sub-carrier to only one customer, but more than one sub-carrier can be assigned to a customer (*multiple allocation*). With such reuse partitioning and spatial allocation of sub-carriers we can assume that the interference (both inter-cell and intra-cell) is zero. Also assume that each base-station, in

[1]We assume that each base station has a separate radio available for every channel. Most of our formulations and all our results go through even when some base stations have fewer radios than channels - wherever applicable we mention the necessary changes in the formulations in this case.

[2]This SINR expression assumes that all base stations use the same band. This facilitates smooth hand-overs but provides poor SINR to the mobiles at cell boundaries owing to high interference from neighboring base stations. Note that CDMA technology can provide acceptable rates even in presence of low SINRs. Nevertheless, in some implementations, neighboring base stations are allocated different bands. In that case, we sum over all co-channel base stations to obtain the aggregate interference in the denominator.

[3]In the low SNR regime, the rates are proportional to the SNR, and thus the peak rates are shared among the mobiles in the same proportion as the overall power.

each state ω, assigns a fixed transmit power to each of its carriers. thus, the downlink rate that a customer gets from a service unit (which denotes a base-station and sub-carrier pair) to which it is assigned depends only on the channel gain from the corresponding base-station-sub-carrier pair to itself, channel usage of primary users as applicable, and not on the assignments of other customers and service units. The communication model presented in Section III-A captures all these attributes except the multiple allocation condition. We will point out the modifications required for allowing multiple allocation while formulating the feasible allocations in the next subsection.

C. Coalitions and Payoffs

Definition III.1. *A coalition $\mathcal{S} \subseteq \mathcal{N}$ is a subset of providers who cooperate. We refer to $\mathcal{N}$ as the grand coalition.*

A service unit can serve a customer only when either both are associated with the same provider, or the providers associated with them are in a coalition. Consider a network state ω. Let $\alpha_{jk}(\omega) \in [0,1]$ be the fraction of time service unit k serves customer j. Now consider a coalition $\mathcal{S}$. A joint allocation for $\mathcal{S}$ is $\{\alpha_{jk}(\omega), j \in \mathcal{M}_{\mathcal{S}}, k \in \mathcal{B}_{\mathcal{S}}, \omega \in \Omega\}$. Let $y^{\mathcal{S}}_{jk}(\omega)$ denote the rate a customer $j \in \mathcal{M}_{\mathcal{S}}$ receives from a service unit $k \in \mathcal{B}_{\mathcal{S}}$; $y^{\mathcal{S}}_{jk}(\omega) = \alpha_{jk}(\omega) r_{jk}(\omega)$. In coalition $\mathcal{S}$, provider i receives a benefit (payoff in economics terminology) $f_i(\mathbf{y}^{\mathcal{S}}_i(\omega))$ where $\mathbf{y}^{\mathcal{S}}_i(\omega) = (y^{\mathcal{S}}_{jk}(\omega), j \in \mathcal{M}_i, k \in \mathcal{B}_{\mathcal{S}})$ is the rate vector of its customers; $\mathbf{y}_i(\omega) = \mathbf{y}^{\mathcal{N}}_i(\omega)$. The payoff of a provider may be the revenue it earns from its customers or may reflect any other benefits it incurs by serving the customers, e.g., reputation, social welfare, etc. The payoff functions $f_i(\cdot)$s are assumed to be concave since customers would pay in accordance with their satisfactions, which are usually concave functions of rates [21], [22]. Usually (and especially for the revenue connotation), $f_i(\mathbf{y}_i)$ are additive functions of different components, i.e., $f_i(\mathbf{y}_i) = \sum_{j \in \mathcal{M}_i} h_{ij}\left(\sum_k y_{jk}\right)$, where $h_{ij}(\cdot)$ is a concave (either strict or linear) revenue function chosen by provider i for customer j. We therefore allow the revenue functions to be different for different customers of the same provider, though mostly, $h_{ij}(x) = h_{ik}(x)$ for all $j,k \in \mathcal{M}_i$, i.e., a provider uses the same revenue functions for all of its customers. The expected payoff provider $i \in \mathcal{S}$ earns will be $\sum_{\omega \in \Omega} \mathbb{P}(\omega) f_i(\mathbf{y}^{\mathcal{S}}_i(\omega))$[4]. We assume that the payoff functions $f_i(.)$s are decided apriori (based on governmental regulations, customer charging policies etc.), and do not investigate the optimal selections of these functions. Thus, in our setup, any joint allocation $\{\alpha_{jk}(\omega)\}$ uniquely determines the payoffs of all the providers.

Providers in a coalition $\mathcal{S}$ have to decide how to schedule service units to customers, i.e., select the variables $\alpha_{jk}(\omega)$s, for each $\omega \in \Omega$, based on the payoff functions $f_i(.)$s, and the service unit to customer rates $r_{jk}(\omega)$s so as to maximize their payoffs. For a coalition $\mathcal{S}$, a feasible allocation $\{\alpha_{jk}(\omega), j \in \mathcal{M}_{\mathcal{S}}, k \in \mathcal{B}_{\mathcal{S}}, \omega \in \Omega\}$ satisfies the following conditions.

1) $\sum_{j \in \mathcal{M}_{\mathcal{S}}} \alpha_{jk}(\omega) \leq 1, k \in \mathcal{B}_{\mathcal{S}}, \omega \in \Omega$
2) $\sum_{k \in \mathcal{B}_{\mathcal{S}}} \alpha_{jk}(\omega) \leq 1, j \in \mathcal{M}_{\mathcal{S}}, \omega \in \Omega$
3) $\alpha_{jk}(\omega) \geq 0, j \in \mathcal{M}_{\mathcal{S}}, k \in \mathcal{B}_{\mathcal{S}}, \omega \in \Omega$.

Constraints (1) ensure that for all service units $k \in \mathcal{B}_{\mathcal{S}}$, their service times are upper bounded by 1[5]. Constraints (2) ensure that for all $j \in \mathcal{S}$, the fraction of time customer j is served is at most 1. Incidentally, constraints (1), (2) arise from the time-sharing model [6], but for the multiple allocation model, only constraints (1) suffice - all results presented below extend even in absence of constraint (2).

Let $\mathcal{A}(\mathcal{S})$ denote the set of feasible joint actions of coalition $\mathcal{S}$. Now, consider a joint action $\alpha \in \mathcal{A}(\mathcal{S})$. Define $\mathcal{F}^{\mathcal{S}}(\alpha) \in \mathbb{R}^{\mathcal{S}}$ to be the payoff vector generated by the joint action α, i.e., the ith component of $\mathcal{F}^{\mathcal{S}}(a)$ is $\sum_{\omega \in \Omega} \mathbb{P}(\omega) f_i(\mathbf{y}^{\mathcal{S}}_i(\omega)),\ \forall i \in \mathcal{S}$, where $(\mathbf{y}^{\mathcal{S}}_i(\omega), i \in \mathcal{S}, \omega \in \Omega)$ are the rate vectors resulting from the joint action α.

Associated with each coalition $\mathcal{S}$, there is a set of feasible payoff profiles, $v(\mathcal{S})$, defined as:

$$v(\mathcal{S}) = \{\mathbf{x} \in \mathbb{R}^{\mathcal{S}} : \mathbf{x} \leq \mathcal{F}^{\mathcal{S}}(\alpha) \text{ for some } \alpha \in \mathcal{A}(\mathcal{S})\}. \quad (1)$$

In other words, $v(\mathcal{S})$ contains all payoff profiles which are less than or equal to the payoff vector generated by some feasible joint action. Now the stage is set for the following definition.

IV. Provider Coalitional Game

A. An NTU Game Formulation

Definition IV.1. *A nontransferable payoff cooperative (NTU) game consists of a pair $< \mathcal{N}, v >$, where $\mathcal{N}$ is the set of players, and $v(\mathcal{S}),\ \forall \mathcal{S} \subseteq \mathcal{N}$ is the set of feasible payoff profiles satisfying*

1) *For each $\mathcal{S}$, $v(\mathcal{S})$ is a closed set.*
2) *If $\mathbf{z} \in v(\mathcal{S})$ and $\mathbf{x} \in \mathbb{R}^{\mathcal{S}}$ with $\mathbf{x} \leq \mathbf{z}$, then $\mathbf{x} \in v(\mathcal{S})$.*
3) *The set of vectors in $v(\mathcal{S})$ in which each player in $\mathcal{S}$ receives no less than the maximum that he can obtain by himself is a nonempty, bounded set.*

When the providers cooperate can take any feasible joint action, and consequently, achieve any payoff profile in $v(\mathcal{N})$. Thus, cooperation has the potential to enhance the quality of service to customers, which in turn can increase providers' payoffs. However, there is a need for a criterion that determines which payoff profile in $v(\mathcal{N})$ will be acceptable to the providers. We use a well known solution concept in coalitional

[4]One can also define the payoff of provider i to be a function of the expected service rates, i.e., $f_i(\sum_{\omega \in \Omega} \mathbb{P}(\omega) \mathbf{y}^{\mathcal{S}}_i(\omega))$ - all results extend to this type of payoff functions as well.

[5]This condition can be modified to capture the scenario when a service unit has access to multiple channels with only 1 radio, as follows. The modified Constraint (1) for a service unit, bounds the sum of $\alpha_{jk}(\omega)$ over customers $j \in \mathcal{M}_{\mathcal{S}}$, and channels k accessed by that service unit, by 1. It can be shown that all the subsequent results extends to this scenario.

[6]Note that for each ω, $\{\alpha_{jk}(\omega)\}$ comprise a feasible allocation of service units to customers if and only if there exists a corresponding collection of matchings $L_1, L_2, \ldots$ and a collection of non-negative real numbers $\gamma_1, \gamma_2, ..$ such that (i) $\sum_i \gamma_i = 1$, $\gamma_i \geq 0$ and (ii) if the service unit - customer allocation follows matching L_i for γ_i fraction of time for each i, then service unit k transmits to customer j for $\alpha_{jk}(\omega)$ fraction of time for all j,k. Constraints (1), (2) provide the necessary and sufficient condition for feasibility of $\{\alpha_{jk}(\omega)\}$ for each ω [23].

game theory,*core*, to provide a rational basis for choosing a payoff profile. The idea behind the core in a cooperative game is analogous to that behind a Nash equilibrium of a noncooperative game: an outcome is stable if no deviation is profitable.

Definition IV.2. *A payoff profile $\mathbf{x} \in v(\mathcal{N})$ is said to be* blocked *by coalition $\mathcal{S} \subseteq \mathcal{N}$, if there is a payoff profile $\mathbf{z} \in v(\mathcal{S})$ such that $\mathbf{z}_i > \mathbf{x}_i$ for all $i \in \mathcal{S}$, i.e., $\mathbf{z}$ makes every provider in $\mathcal{S}$ better off.*

Definition IV.3. *The core $\mathcal{C}$ of the game $< \mathcal{N}, v >$ is the set of all feasible payoff profiles which can not be blocked by any coalition. That is,*

$$\mathcal{C} = \{\mathbf{x} \in v(\mathcal{N}) : \forall \mathcal{S},\ \nexists \mathbf{z} \in v(\mathcal{S}) \text{ such that } \mathbf{z}_i > \mathbf{x}_i,\ \forall i \in \mathcal{S}\} \tag{2}$$

The significance of the core comes from the fact that every payoff profile in the core renders the grand coalition stable. To see this, let providers form the grand coalition and select a joint action that results in a payoff profile $\mathbf{x} \in \mathcal{C}$. Now, suppose a set of providers $\mathcal{S} \subset \mathcal{N}$ leave the grand coalition and choose a joint action and the corresponding payoff profile $\mathbf{z} \in v(\mathcal{S})$. They, however, would do so only if all of them receive a higher payoff than what they could in the grand coalition, i.e., $\mathbf{z}_i > \mathbf{x}_i,\ \forall i \in \mathcal{S}$. But this contradicts the fact that $\mathbf{x} \in \mathcal{C}$. Therefore, the grand coalition is stable. *This is a globally desirable outcome, since the grand coalition has the potential to achieve higher network-wide efficiency.*

We now elucidate $v(\cdot)$ and $\mathcal{C}$ using a simple example.

Example IV.1. *Consider a network with $\mathcal{N} = \{1, 2, 3\}$, and $|\Omega| = 1$. Suppose each provider has one service unit and one customer; $\mathcal{B}_i = \mathcal{M}_i = \{i\}, i = 1, 2, 3$. Let $r_{ii} = 1, i = 1, 2, 3$, $r_{21} = r_{32} = r_{13} = 3$, and $r_{31} = r_{12} = r_{23} = 2$, and the payoff of each provider equal the aggregate service rate of its customers. When providers do not cooperate, each attains a payoff of 1. In other words, $v(\{i\}) = [0, 1],\ \forall i$. For the coalition $\{1, 2\}$, we can similarly specify the feasible payoff profiles as $v(\{1, 2\}) = \{(x_1, x_2) : x_1 \le 2, x_2 \le 3\}$. Finally, for the grand coalition we have $v(\{1, 2, 3\}) = \{(x_1, x_2, x_3) : x_i \le 3,\ \forall i\}$. Note that the feasible payoff profile $(3, 3, 3)$ is not blocked by any coalition, and hence is in the core. This profile enhances the payoff of each provider by 200% compared to when they operate individually - rate diversity (customers have different rates from different service units) leads to a better match between customers and service units in presence of cooperation (e.g., 1's service units provides the highest rate to 2's customer).*

Though cooperation enhances providers' payoffs, it may not be advantageous for all the customers. The following example illustrates this.

Example IV.2. *Consider a network with $\mathcal{N} = \{1, 2\}$, and $|\Omega| = 1$. $\mathcal{B}_i = \{i\}, i = 1, 2, \mathcal{M}_1 = \{1, 2\}$ and $\mathcal{M}_2 = \{3\}$. Let $r_{11} = r_{32} = P$, $r_{22} = r_{31} = Q$ and $r_{jk} = 0$ otherwise. Suppose $P < Q$. Suppose that the payoff of each provider equals the aggregate service rate of its customers. Then, $v(\{1\}) = v(\{2\}) = [0, P]$ and $v(\{1, 2\}) = \{(x_1, x_2) : x_1 \le Q, x_2 \le Q\}$. Clearly, customers 2 and 3 get better rates, while customer 1 does not get any service in the coalition.*

B. Nonemptyness of The Core

In several coalitional games the core is empty, i.e., the grand coalition can not be stabilized [1, Example 260.3] [15, Example 20], and in general it is NP-hard to determine whether the core of a coalitional game is nonempty [24]. Nevertheless, in this section we show that the game $< \mathcal{N}, v >$ always has a nonempty core.

Definition IV.4. *A collection of coalitions $\mathcal{I} \subset 2^{\mathcal{N}} \backslash \emptyset$ is called balanced if there exist nonnegative weights $(\lambda_{\mathcal{S}}, \mathcal{S} \in \mathcal{I})$ such that*

$$\sum_{\mathcal{S} \in \mathcal{I}:\ i \in \mathcal{S}} \lambda_{\mathcal{S}} = 1,\ \forall i \in \mathcal{N}.$$

Note that the balancedness condition for a collection of coalitions does not depend on the players' payoff functions, but only on memberships of the coalitions in the collection. Intuitively, a collection is balanced if the players are distributed "uniformly" among the coalitions in the collection. For example, a collection, where each player is in the same number (say k) of coalitions, is balanced. Then, let $\lambda_{\mathcal{S}} = 1/k$ for each $\mathcal{S} \in \mathcal{I}$, and since exactly k coalitions include the player i, for each i, $\sum_{\mathcal{S} \in \mathcal{I}:\ i \in \mathcal{S}} \lambda_{\mathcal{S}} = k \times (1/k) = 1$. On the other hand, if one player, say player 1, belongs in all the coalitions in the collection, and every other player belongs in only one coalition each, the coalition is not balanced. To see this, note that for a player i ($i \ne 1$), $\sum_{\mathcal{S} \in \mathcal{I}:\ i \in \mathcal{S}} \lambda_{\mathcal{S}}$ equals only one $\lambda_{\mathcal{S}}$, the one that contains i, and hence this $\lambda_{\mathcal{S}}$ must equal 1 if this collection is to be balanced. Now, for player 1, $\sum_{\mathcal{S} \in \mathcal{I}:\ 1 \in \mathcal{S}} \lambda_{\mathcal{S}}$ equals $\sum_{\mathcal{S} \in \mathcal{I}} \lambda_{\mathcal{S}}$, which exceeds 1 since each $\lambda_{\mathcal{S}}$ equals 1.

Example IV.3. *Let $\mathcal{N} = \{1, 2, 3\}$. Then $\mathcal{I}_1 = \{\{1, 2\}, \{2, 3\}, \{1, 3\}\}$ is balanced; every player is in exactly two coalitions, so $\lambda_{\mathcal{S}} = \frac{1}{2}$ is the balancing weight for each coalition $\mathcal{S} \in \mathcal{I}_1$. On the other hand $\mathcal{I}_2 = \{\{1, 2\}, \{2, 3\}\}$ is not balanced, since there do not exist nonnegative λ_1 and λ_2 such that $\lambda_1 = 1$, $\lambda_1 + \lambda_2 = 1$, and $\lambda_2 = 1$. Note that here player 2 belongs in both coalitions, whereas players $1, 3$ belong in one coalition each.*

Definition IV.5. *A game is balanced if for every balanced collection of coalitions $\mathcal{I}$, if $u \in \mathbb{R}^{\mathcal{N}}$ and $u^{\mathcal{S}} \in v(\mathcal{S})$ for all $\mathcal{S} \in \mathcal{I}$, then $u \in v(\mathcal{N})$*[7].

Thus, balancedness of a game depends on the payoff functions and feasible payoff profiles. For any coalitional game, balancedness provides a sufficient condition for non-emptiness of the core [25].

Theorem IV.1. *A balanced game always has a nonempty core.*

Here is the main result.

[7]For any $u \in \mathbb{R}^{\mathcal{N}}$, we denote by $u^{\mathcal{S}} \in \mathbb{R}^{\mathcal{S}}$, the sub-vector of u corresponding to the coalition $\mathcal{S}$, i.e., $u^{\mathcal{S}}_i = u_i, \forall i \in \mathcal{S}$

Theorem IV.2. *The coalitional game among providers, $<\mathcal{N}, v>$, is balanced and hence has a nonempty core.*

Proof: Consider a balanced collection of coalitions $\mathcal{I}$. Let $(\lambda_{\mathcal{S}}, \mathcal{S} \in \mathcal{I})$ be the corresponding balancing weights. Also, let $u \in \mathbb{R}^{\mathcal{N}}$ be such that $u^{\mathcal{S}} \in v(\mathcal{S})$ for all $\mathcal{S} \in \mathcal{I}$, i.e, there exists a joint action $\{\alpha_{jk}^{\mathcal{S}}(\omega), j \in \mathcal{M}_{\mathcal{S}}, k \in \mathcal{B}_{\mathcal{S}}, \omega \in \Omega\}$ for each $\mathcal{S} \in \mathcal{I}$ such that

1) $\{\alpha_{jk}^{\mathcal{S}}(\omega), j \in \mathcal{M}_{\mathcal{S}}, k \in \mathcal{B}_{\mathcal{S}}, \omega \in \Omega\}$ satisfy feasibility constraints (1) - (3) in Section III, for each $\mathcal{S} \in \mathcal{I}$.
2) $u_i \leq \sum_{\omega\in\Omega} \mathbb{P}(\omega) f_i(\mathbf{y}_i^{\mathcal{S}}(\omega)), \ \forall i \in \mathcal{S}$, where $\mathbf{y}_i^{\mathcal{S}}(\omega)$ denotes the rate vector corresponding to joint action $\{\alpha_{jk}^{\mathcal{S}}(\omega), j \in \mathcal{M}_{\mathcal{S}}, k \in \mathcal{B}_{\mathcal{S}}, \omega \in \Omega\}$.

We next show that $u \in v(\mathcal{N})$. Define a joint action $\{\alpha_{jk}(\omega), j \in \mathcal{M}_{\mathcal{N}}, k \in \mathcal{B}_{\mathcal{N}}, \omega \in \Omega\}$ as follows

$$\alpha_{jk}(\omega) = \sum_{\mathcal{S}\in\mathcal{I}: \substack{j\in\mathcal{M}_{\mathcal{S}} \\ k\in\mathcal{B}_{\mathcal{S}}}} \lambda_{\mathcal{S}} \alpha_{jk}^{\mathcal{S}}(\omega). \tag{3}$$

The rest of the proof consists of two steps.

Step 1: We show that $\{\alpha_{jk}(\omega), j \in \mathcal{M}_{\mathcal{N}}, k \in \mathcal{B}_{\mathcal{N}}, \omega \in \Omega\}$ satisfy feasibility constraints (1) in Section III.

$$\begin{aligned}
\sum_{j\in\mathcal{M}_{\mathcal{N}}} \alpha_{jk}(\omega) &= \sum_{j\in\mathcal{M}_{\mathcal{N}}} \sum_{\mathcal{S}\in\mathcal{I}: \substack{j\in\mathcal{M}_{\mathcal{S}} \\ k\in\mathcal{B}_{\mathcal{S}}}} \lambda_{\mathcal{S}} \alpha_{jk}^{\mathcal{S}}(\omega) \\
&= \sum_{\mathcal{S}\in\mathcal{I}:\ k\in\mathcal{B}_{\mathcal{S}}} \lambda_{\mathcal{S}} \sum_{j\in\mathcal{M}_{\mathcal{S}}} \alpha_{jk}^{\mathcal{S}}(\omega) \\
&\leq \sum_{\mathcal{S}\in\mathcal{I}:\ k\in\mathcal{B}_{\mathcal{S}}} \lambda_{\mathcal{S}} \\
&= \sum_{\mathcal{S}\in\mathcal{I}:\ i\in\mathcal{S}} \lambda_{\mathcal{S}} \ \text{(where } k \in \mathcal{B}_i\text{)} \\
&= 1.
\end{aligned}$$

The first equality follows from (3). The inequality follows from feasibility of $\{\alpha_{jk}^{\mathcal{S}}(\omega), j \in \mathcal{M}_{\mathcal{S}}, k \in \mathcal{B}_{\mathcal{S}}, \omega \in \Omega\}$ for $\mathcal{S} \in \mathcal{I}$ and Constraints (1) in Section III.

Similarly, one can show that feasibility constraints (2) are also satisfied. Constraints (3) are trivial. Thus, the joint action $\{\alpha_{jk}(\omega), j \in \mathcal{M}_{\mathcal{N}}, k \in \mathcal{B}_{\mathcal{N}}, \omega \in \Omega\}$ is feasible.

Step 2: We show that $u_i \leq \sum_{\omega\in\Omega} \mathbb{P}(\omega) f_i(\mathbf{y}_i(\omega)), \forall i \in \mathcal{N}$, where $(\mathbf{y}_i(\omega), i \in \mathcal{S}, \omega \in \Omega)$ are the rate vectors resulting from the joint action $\{\alpha_{jk}(\omega), j \in \mathcal{M}_{\mathcal{N}}, k \in \mathcal{B}_{\mathcal{N}}, \omega \in \Omega\}$.

Using (3), it can be easily verified that

$$\mathbf{y}_i(\omega) = \sum_{\mathcal{S}\in\mathcal{I}:i\in\mathcal{S}} \lambda_{\mathcal{S}} \mathbf{y}_i^{\mathcal{S}}(\omega), \ \forall i \in \mathcal{N}, \omega \in \Omega, \tag{4}$$

i.e., $\mathbf{y}_i(\omega)$ are convex combinations of $\{\mathbf{y}_i^{\mathcal{S}}(\omega), \mathcal{S} \in \mathcal{I} : i \in \mathcal{S}\}$. Since $f_i(\cdot)$s are concave, for each provider i we have:

$$\begin{aligned}
\sum_{\omega\in\Omega} \mathbb{P}(\omega) f_i(\mathbf{y}_i(\omega)) &\geq \sum_{\omega\in\Omega} \mathbb{P}(\omega) \sum_{\mathcal{S}\in\mathcal{I}:i\in\mathcal{S}} \lambda_{\mathcal{S}} f_i(\mathbf{y}_i^{\mathcal{S}}(\omega)) \\
&= \sum_{\mathcal{S}\in\mathcal{I}:i\in\mathcal{S}} \lambda_{\mathcal{S}} \sum_{\omega\in\Omega} \mathbb{P}(\omega) f_i(\mathbf{y}_i^{\mathcal{S}}(\omega)) \\
&\geq \sum_{\mathcal{S}\in\mathcal{I}:i\in\mathcal{S}} \lambda_{\mathcal{S}} u_i \\
&= u_i
\end{aligned}$$

Thus, the game is balanced. It then follows from Theorem IV.1 that the core of the game is nonempty. ■

C. Computation of a Payoff Profile in the Core

We first construct an ”exchange market” setting, a concept borrowed from micro-economics (Section IV-C1), and show that a *market equilibrium* in this setting, if exists, corresponds to a payoff profile in the core of our NTU game (Section IV-C2). Note that the fact that two different concepts are equivalent in this context, can itself be of independent interest. Finally, we show that a market equilibrium exists (Section IV-C3), and can be computed without requiring exchange of confidential information among the providers (Section IV-C4).

1) Exchange Market Preliminaries: We now introduce the concept of ”exchange market” from micro-economic theory. Consider a market with a set of agents $\mathcal{N}$. Let $\mathcal{L}$ denote the set of goods in the market. Each agent i has a positive initial endowment of the goods given by the vector $\mathbf{e}_i = (e_i^l, l \in \mathcal{L})$. Associated with each agent i is a utility function $u_i : \mathbb{R}_+^{\mathcal{L}} \rightarrow \mathbb{R}$; $u_i(\mathbf{x}_i)$ represents the satisfaction level of agent i from the allocation of goods $\mathbf{x}_i = (x_i^l, l \in \mathcal{L})$. Now let vector $\mathbf{p} = (p_l, l \in \mathcal{L})$ denote the prices of goods in the market. The agents will then try to maximize their utilities through trading of goods according to prices $\mathbf{p}$. We now present the definition of the market equilibrium [26, pp. 579].

Definition IV.6. *An allocation $\mathbf{x}^*$ and a price vector $\mathbf{p}$ constitute a market equilibrium if*

i) $\mathbf{x}_i^* \in \arg\max_{\mathbf{x}_i\in\mathbb{R}_+^{\mathcal{L}}} u_i(\mathbf{x}_i)$ *subject to* $\mathbf{p}.\mathbf{x}_i \leq \mathbf{p}.\mathbf{e}_i, \ \forall i \in \mathcal{N}$. *Note that $\mathbf{p}.\mathbf{x}_i^*$ is the value of agent i's allocation, which clearly can not be larger than the value of his initial endowment (budget constraint).*

ii) $\sum_{i\in\mathcal{N}}(\mathbf{x}_i^* - \mathbf{e}_i) = 0$, *that is, it is possible to attain the agents' desired allocations, just by using the total initial endowments in the market (market clearing).*

Such an allocation $\mathbf{x}^*$ is called a market equilibrium allocation.

We now elucidate equilibrium prices and equilibrium allocations using a simple example.

Example IV.4. *Consider a market with two agents and two goods. Initial endowments of goods are $\mathbf{e}_1 = (1, 0)$ and $\mathbf{e}_2 = (0, 1)$. Utility functions are $u_i(x_i^1, x_i^2) = \beta_{i1} \log(1 + x_i^1) + \beta_{i2} \log(1 + x_i^2), i = 1, 2$. For identical utility functions, i.e., with $\beta_{i1} = \beta, \beta_{i2} = 1$, $((\frac{\beta}{1+\beta}, \frac{\beta}{1+\beta}), (\frac{1}{1+\beta}, \frac{1}{1+\beta}))$ is an* *equilibrium allocation with $(\beta, 1)$ as an equilibrium price vector. For non-identical utility functions, i.e., with $(\beta_{11}, \beta_{12}) = (1, \beta)$ and $(\beta_{21}, \beta_{22}) = (\beta, 1)$, $((q_\beta, 1 - q_\beta), (1 - q_\beta, q_\beta))$ is a market equilibrium allocation, where $q_\beta = 0$ if $\beta > 2$, $q_\beta = \frac{2-\beta}{1+\beta}$ if $\beta \in (1/2, 2)$ and $q_\beta = 1$ otherwise. Here, $(1, 1)$ is an equilibrium price vector.*

We explain the above equilibrium allocations assuming that $\beta > 1$. For identical utility functions, both agents value agent 1's initial possession more. The second agent is able to obtain only a small part $(1/(1 + \beta))$ of this more valuable good by ceding a large part $(\beta/(1 + \beta))$ of its initial possession. Thus, the first (second, resp.) ends up with a higher (lower, resp.) share $\beta/(1 + \beta)$ $(1/(1 + \beta))$ of both goods. Note that the equilibrium price is higher for the first good. For non-identical utility functions, each agent values the other agent's initial possession more, and hence they engage in an even exchange after which each possesses greater amount of the good it values more. In fact, for $\beta > 2$, agents exchange their initial possessions fully, and each possesses the good it values more, in its entirety. The goods have equal equilibrium prices.

The following theorem provides a sufficient condition for a market equilibrium to exist [26, pp. 585].

Theorem IV.3. *Suppose that for every agent $i \in \mathcal{N}$, $u_i(\cdot)$ is continuous, strictly concave, and strictly increasing. Also suppose that $\sum_{i\in\mathcal{N}} \mathbf{e}_i \in \mathbb{R}^{\mathcal{L}}_{++}$. Then a market equilibrium exists, with the property that the price vector is strictly positive, i.e., $\mathbf{p} \in \mathbb{R}^{\mathcal{L}}_{++}$.*

Now suppose instead of trading, agents pool their goods and reallocate them among themselves. The amount of goods allocated to each agent has to be commensurate with agents initial endowments, or some agents would not agree to it. Consequently, one can use the concept of *core of the market* as a policy to determine the allocation of goods among agents. An allocation $\mathbf{x}^* = (\mathbf{x}^*_i, i \in \mathcal{N}) \in \mathbb{R}^{\mathcal{L}\times\mathcal{N}}_+$ is in the core of the market, $\mathfrak{C}$, if is can not be blocked by any coalition of agents $\mathcal{S} \subseteq \mathcal{N}$, i.e., for any $\mathcal{S} \subseteq \mathcal{N}$, there does not exist an allocation $(\mathbf{x}_i, i \in \mathcal{S}) \in \mathbb{R}^{\mathcal{L}\times\mathcal{S}}_+$ with the properties:

i) $\sum_{i\in\mathcal{S}} \mathbf{x}_i \le \sum_{i\in\mathcal{S}} \mathbf{e}_i$.
ii) $u_i(\mathbf{x}_i) > u_i(\mathbf{x}^*_i),\ \forall i \in \mathcal{S}$.

Example IV.5. *Consider the setting of Example IV.4. The first condition for an allocation $(y^1_1, y^2_1, y^1_2, y^2_2)$ to be in the core of the market, $\mathfrak{C}$, is that $y^1_1 + y^1_2 = e^1_1 + e^1_2 = 1$ and $y^2_1 + y^2_2 = e^2_1 + e^2_2 = 1$, and hence $y^1_2 = 1 - y^1_1$ and $y^2_2 = 1 - y^2_1$. The second condition is that either $\beta \log(1 + y^1_1) + \log(1 + y^2_1) \ge \beta \log(1 + x^1_1) + \log(1 + x^2_1)$, or $\beta \log(2 - y^1_1) + \log(2 - y^2_1) \ge \beta \log(2 - x^1_1) + \log(2 - x^2_1)$, for any $x^1_1, x^2_1 \in [0, 1]$.*

The following theorem states the relation between market equilibria and the core of the market [26, pp. 654].

Theorem IV.4. *Any market equilibrium allocation is in the core of the market $\mathfrak{C}$.*

2) Provider Coalition as an Exchange Market: Consider the NTU game defined in Section III. Think of the set of providers $\mathcal{N}$ as the agents in the market. The combinations of service units and network realizations, $\mathcal{B}_\mathcal{N} \times \Omega$, then constitute the goods in the economy. Specifically, an agent i with allocation $\mathbf{x}_i$ can access service unit k at most $x^k_i(\omega)$ fraction of time, when the realization is ω. Consequently, the initial endowments of the providers are the full access to the set of service units they own; for a provider i, for all $\omega \in \Omega$, $e^k_i(\omega)$ is 1 if $k \in \mathcal{B}_i$ and 0, otherwise.

Now consider an allocation of goods $\mathbf{x}$ in this setup. We define the providers' corresponding utilities to be the maximum payoff they can obtain by serving their own customers using their access levels of service units. In other words,
$u_i(\mathbf{x}_i) = \max \sum_{\omega\in\Omega} \mathbb{P}(\omega) f_i(\mathbf{y}_i(\omega))$
subject to:

1) $\mathbf{y}_i(\omega) = (y_{jk}(\omega), j \in \mathcal{M}_i, k \in \mathcal{B}_\mathcal{N}), \quad \omega \in \Omega$
2) $y_{jk}(\omega) = \alpha_{jk}(\omega) r_{jk}(\omega), \quad j \in \mathcal{M}_i, k \in \mathcal{B}_\mathcal{N}, \omega \in \Omega$
3) $\sum_{k\in B_\mathcal{N}} \alpha_{jk}(\omega) \le 1, \quad j \in \mathcal{M}_i, \omega \in \Omega.$
4) $\sum_{j\in\mathcal{M}_i} \alpha_{jk}(\omega) \le x^k_i(\omega), \quad k \in \mathcal{B}_\mathcal{N}, \omega \in \Omega.$
5) $\alpha_{jk}(\omega) \ge 0, \quad j \in \mathcal{M}_i, k \in \mathcal{B}_\mathcal{N}, \omega \in \Omega.$

Example IV.6. *We now obtain the exchange market for the setting of Example IV.1. Each provider constitutes an agent, and each service unit constitutes a good. Provider i possesses the ith good entirely, initially. Consider the payoff profile of $(3, 3, 3)$. This is attained when the first (second, third, resp.) service unit serves agent 2's (3's, 1's) customers respectively for the entire time. This corresponds to allocations of goods $\mathbf{x}_i$ to agents as follows: $\mathbf{x}_1 = (0, 0, 1), \mathbf{x}_2 = (1, 0, 0)$ and $\mathbf{x}_3 = (0, 1, 0)$ - thus, the first agent gets the whole of the third good and so on. The utilities of the agents can be obtained by solving the above optimizations.*

We next show how an allocation in the core of the market can be used to obtain a payoff profile in the core of the NTU game.

Theorem IV.5. *Consider any allocation $\mathbf{x}^*$ belonging to $\mathfrak{C}$. Then $(u_i(\mathbf{x}^*_i), i \in \mathcal{N}) \in \mathcal{C}$.*

Proof: First, let $\{\alpha^*_{jk}(\omega), j \in \mathcal{M}_i, k \in \mathcal{B}_\mathcal{N}, \omega \in \Omega\}$ be an optimal solution of the optimization defining $u_i(\mathbf{x}^*_i)$. Note that $\{\alpha^*_{jk}(\omega), j \in \mathcal{M}_\mathcal{N}, k \in \mathcal{B}_\mathcal{N}, \omega \in \Omega\}$ constitute a feasible joint action of providers in $\mathcal{N}$. Thus $(u_i(\mathbf{x}^*_i), i \in \mathcal{N}) \in v(\mathcal{N})$.

We now prove the claim by contradiction. Suppose $(u_i(\mathbf{x}^*_i), i \in \mathcal{N}) \notin \mathcal{C}$. Then there exist a coalition $\mathcal{S}$ and a payoff profile $\mathbf{z} \in v(\mathcal{S})$ such that $\mathbf{z}_i > u_i(\mathbf{x}^*_i)$ for all $i \in \mathcal{S}$. We argue that there exists an allocation $(\mathbf{x}_i, i \in \mathcal{S})$ such that (i) $\sum_{i\in\mathcal{S}} \mathbf{x}_i \le \sum_{i\in\mathcal{S}} \mathbf{e}_i$ and (ii) $u_i(\mathbf{x}_i) \ge \mathbf{z}_i\ \forall i \in \mathcal{S}$, as follows. Consider the joint action $\{\alpha_{jk}(\omega), j \in \mathcal{M}_\mathcal{S}, k \in \mathcal{B}_\mathcal{S}, \omega \in \Omega\}$, corresponding to payoff profile $\mathbf{z}$. Now define $\mathbf{x}_i = (\sum_{j\in\mathcal{M}_i} \alpha_{jk}(\omega), k \in \mathcal{B}_\mathcal{S}, \omega \in \Omega)$. Since $\{\alpha_{jk}(\omega), j \in \mathcal{M}_\mathcal{S}, k \in \mathcal{B}_\mathcal{S}, \omega \in \Omega\}$ is a feasible joint action of providers in $\mathcal{S}$, it satisfies Constraint (1) in Section III. This fact, together with the definitions of $\mathbf{x}_i$ and $\mathbf{e}_i$, implies that $\sum_{i\in\mathcal{S}} x^k_i(\omega) = \sum_{j\in\mathcal{M}_\mathcal{S}} \alpha_{jk}(\omega) \le 1 = \sum_{i\in\mathcal{S}} e^k_i(\omega)$ for all $k \in \mathcal{B}_\mathcal{S}, \omega \in \Omega$. Thus (i) holds. Also, since $\{\alpha_{jk}(\omega), j \in \mathcal{M}_i, k \in \mathcal{B}_\mathcal{S}, \omega \in \Omega\}$ is a feasible solution and $\mathbf{z}_i$ is the corresponding value of the optimization defining $u_i(\mathbf{x}_i)$, (ii) immediately follows. As a consequence of (ii), $u_i(\mathbf{x}_i) > u_i(\mathbf{x}_i^*), \forall i \in \mathcal{S}$. This is in contradiction with $\mathbf{x}^* \in \mathfrak{C}$. ∎

Theorem IV.6. *If $\mathbf{x}^*$ is a market equilibrium allocation in the exchange market, the corresponding payoff profile $(u_i(\mathbf{x}_i^*), i \in \mathcal{N})$ is in the core of the NTU game.*

Proof: Using Theorems IV.4 and IV.5, the claim immediately follows. ∎

3) Existence of The Market Equilibrium: In this section, we establish the existence of the market equilibrium in our model. We make the following technical assumptions.

1) The functions f_is are strictly concave, strictly increasing, and smooth functions (i.e., the first two derivatives exist and are continuous).
2) For any feasible allocation of interest, $(\mathbf{x}_i, \in \mathcal{N})$, Constraints (3) in the optimizations defining $u_i(\cdot)$s are never binding.

We originally considered the functions f_is to be concave; assumption (1) imposes stronger conditions. Assumption (2), on the other hand, can be motivated by considering the number of customers high enough so that it is always sub-optimal to serve any one customer all the the time.

Using Assumption (2) we can rewrite the provider i's utility function $u_i(\cdot)$ as

P: $u_i(\mathbf{x}_i) = \max \sum_{\omega \in \Omega} \mathbb{P}(\omega) f_i(\mathbf{y}_i(\omega))$

subject to:

1) $\mathbf{y}_i(\omega) = (y_{jk}(\omega), j \in \mathcal{M}_i, k \in \mathcal{B}_{\mathcal{N}}), \quad \omega \in \Omega$
2) $y_{jk}(\omega) = \alpha_{jk}(\omega) r_{jk}(\omega), \quad j \in \mathcal{M}_i, k \in \mathcal{B}_{\mathcal{N}}, \omega \in \Omega$
3) $\sum_{j \in \mathcal{M}_i} \alpha_{jk}(\omega) \leq x_i^k(\omega), \quad k \in \mathcal{B}_{\mathcal{N}}, \omega \in \Omega.$
4) $\alpha_{jk}(\omega) \geq 0, \quad j \in \mathcal{M}_i, k \in \mathcal{B}_{\mathcal{N}}, \omega \in \Omega.$

It follows from Assumptions (1) and (2) that functions $u_i(\cdot)$s are continuous, strictly increasing, strictly concave, and smooth. Then Theorem IV.3 implies that a market equilibrium exists.

We give a simple example to illustrate the above notions.

Example IV.7. *Consider a network with $\mathcal{N} = \{1, 2\}$, and $|\Omega| = 1$. $\mathcal{B}_i = \{i\}, i = 1, 2$, and $\mathcal{M}_i = \{2i-1, 2i\}, i = 1, 2$. Let $r_{ji} = P, j \in \mathcal{M}_i$, $r_{j1} = Q, j \in \mathcal{M}_2$, and $r_{j2} = 2Q, j \in \mathcal{M}_1$. A provider's payoff is a function of rate received by each of its customers; $f_i(\mathbf{y}_i^{\mathcal{S}}) = \sum_{j \in \mathcal{M}_i} \log(\sum_{k \in \mathcal{B}_{\mathcal{S}}} y_{jk}^{\mathcal{S}}), i = 1, 2$. The set of all market equilibrium prices and allocations, $(\mathbf{p}, \mathbf{x})$, is*

$$\{(\beta, 1), ((0, \beta), (1, 1-\beta)) : \beta \in [\tfrac{P}{2Q}, 1]\}$$
$$\cup \{(1, \beta), ((1-\beta, 1), (\beta, 0)) : \beta \in [\tfrac{P}{Q}, 1]\}.$$

Each provider allocates its share equally among its customers. Finally, we obtain the following set of payoff profiles that are in the core.

$$\{(2\log(\beta Q), 2\log(\tfrac{Q+(1-\beta)P}{2}) : \beta \in [\tfrac{P}{2Q}, 1]\}$$
$$\cup \{(2\log(\tfrac{(1-\beta)P+2Q}{2}), 2\log(\tfrac{\beta Q}{2})) : \beta \in [\tfrac{P}{Q}, 1]\}.$$

4) Computation of the Market equilibrium:

By Theorem IV.6, a payoff profile in the core of the NTU game can be obtained by computing a market equilibrium, which can be computed as follows. For a price vector $\mathbf{p}$, define the demand vector of agent (provider) i as $\mathbf{d}_i(\mathbf{p}) = \arg\max_{\mathbf{x}_i \in \mathbb{R}_+^{\mathcal{L}}} u_i(\mathbf{x}_i)$ subject to $\mathbf{p}.\mathbf{x}_i \leq \mathbf{p}.\mathbf{e}_i$, i.e., an allocation of goods to agent i, that maximizes his utility, subject to his budget constraint. Then the aggregate excess demand in the market is the function $\xi : \mathbb{R}_+^{\mathcal{L}} \to \mathbb{R}^{\mathcal{L}}$ given by $\xi(\mathbf{p}) = \sum_{i \in \mathcal{N}} (\mathbf{d}_i(\mathbf{p}) - \mathbf{e}_i)$, i.e., the aggregate demand minus the total endowment. From Definition IV.6, $\mathbf{p}^*$ is an equilibrium price vector if $\xi(\mathbf{p}^*) = 0$. This equation can be solved using the *global Newton method* [27]. The market equilibrium allocation $(\mathbf{x}_i^*, i \in \mathcal{N})$ then is $\mathbf{d}_i(\mathbf{p}^*)$. Then, by Theorem IV.6, the corresponding payoff vector, $(u_i(\mathbf{x}_i^*), i \in \mathcal{N})$, is in the core of the NTU game.

The above computation can be executed using a central controller but without requiring any provider i to reveal its own benefit functions $f_i(\cdot)$ and the service-unit-customer rates $r_{jk}(\omega)$ for its customers $j \in \mathcal{M}_i$ and any service unit $k \in \mathcal{B}_{\mathcal{N}}$ to other providers (or the central controller). The central controller only needs to know the total number of service units of the providers and the number of network states. The need for limited access to global information ensures confidentiality of operations. At each iteration, the central controller selects an initial price vector $\mathbf{p}^0$ arbitrarily, and broadcasts it to all providers. Each provider calculates its demand vector $\mathbf{d}_i(\mathbf{p}^0)$ (as in the last paragraph), using $P(\omega)$ for each $\omega \in \Omega$, $f_i(\cdot)$ and the service-unit-customer rates $r_{jk}(\omega)$ for its customers $j \in \mathcal{M}_i$ and any service unit $k \in \mathcal{B}_{\mathcal{N}}$ and at each ω. Once the central controller receives the demand vectors from the providers, it determines the excess for this price vector as $\xi(\mathbf{p}^0) = \sum_{i \in \mathcal{N}} (\mathbf{d}_i(\mathbf{p}^0) - \mathbf{e}_i)$, and updates the price vector based on the value of this excess[8], and communicates the new price vector to the providers. The process is repeated, using the new price vector at each step, until the excess is 0.

The computation time of this algorithm is dominated by that of the *Global Newton Method*. This method, despite not guaranteed to run in polynomial time, in practice, is known to terminate fast for large problems (i.e., grows polynomially with problem size) [29, pp. 670]. But, when the customer locations and channel states are random, the number of network states $|\Omega|$, and therefore the problem size, and hence the computation times can grow very fast (exponentially) with increase in the number of service units and customers. This may not however pose a major challenge as the computations are done off-line using large work-stations and at a slower time-scale (only when the network state statistics change or the coalitions are assessed). In addition, the above algorithm is efficient (i.e., attains low computation time) for small $|\Omega|$, i.e., for deterministic and pseudo-deterministic systems in which customers are (almost) static and the qualities of channels are also (almost) fixed.

Remark IV.1. *We can solve the optimization problem P by solving separate problems for each $\omega \in \Omega$. Each problem*

[8]The *global Newton method* updates the prices as: $\mathbf{p}^{n+1} = \mathbf{p}^n - \mathrm{adj}\,(\mathrm{J}\xi(\mathbf{p}^n))\,\xi(\mathbf{p}^n)$, where, $\mathrm{J}(\cdot)$, $\mathrm{adj}(\cdot)$ are the *Jacobian* of a multivariate function [28, Section 4.5] and the *adjont* of a matrix [28, Section 0.8.2] respectively. In practice, the *Jacobian* is replaced by its finite difference approximation; $(\mathrm{J}\xi(\mathbf{p}^n))_{jk} = \dfrac{\xi_j(\mathbf{p}^n + \epsilon \mathbf{e}_k) - \xi_j(\mathbf{p}^n)}{\epsilon}$.

yields $\{x_i^k(\omega), \alpha_{jk}(\omega), i \in \mathcal{N}, j \in \mathcal{M}_{\mathcal{N}}, k \in \mathcal{B}_{\mathcal{N}}\}$*. The number of variables in these problems is* $\frac{1}{|\Omega|}$ *times that in the original problem. This substantially reduces the computation complexity because the computation time for a convex program is polynomial in the number of variables, and* $|\Omega|$*, typically, is large.*

V. Provider-Customer Coalitional Game

We have so far assumed that the customer subscriptions are determined apriori. However, cooperation among providers may make some of their customers worse off (see Example IV.2). Thus, customers can strategically subscribe to the providers of their choices depending upon providers' cooperation decisions. Then, cooperation may enhance some providers' customer bases, and reduce others'. The following example further elucidates this point.

Example V.1. *Consider two providers with one service unit each, i.e.,* $\mathcal{B}_i = \{i\}$*. Customer* 1 *(say* $C1$*) would subscribe to one of them. Let the network have two states, i.e.,* $\Omega = \{\omega_1, \omega_2\}$*,* $P(\omega_1) = P(\omega_2) = 1/2$ *and* $r_{11}(\omega_1) = 0, r_{11}(\omega_2) = H, r_{12}(\omega_1) = r_{12}(\omega_2) = L$*, where* $H > L > 1$*.* $C1$*'s expected satisfaction is* $\sum_{\omega\in\Omega} P(\omega)\log(x(\omega))$ *if its rate is* $x(\omega)$ *in state* ω*. In the non-cooperative regime,* $C1$ *must choose* 2 *as with* 1 *its expected satisfaction is* $-\infty$ *since* $x(\omega_1) = 0$*, but with* 2 *it may be* $\log(L) > 0$ *instead (since potentially* $x(\omega_1) = x(\omega_2) = L > 1$*.) But,* $C1$*'s decision* may *be different if providers cooperate. Then, if* $C1$ *subscribes to* 2*, its expected utility may still be (at most)* $\log(L)$*. This is because at* ω_2*, provider* 1 *may not serve* 2*'s customers since* 2 *may not be able to reciprocate (owing to low transmission quality). But, if* $C1$ *subscribes to* 1*,* 1 *may have* 2 *serve* $C1$ *at* ω_1 *(possibly by serving at* ω_1 *one of* 2*'s customers to whom it may have a high rate) and offer rate* H *to* $C1$ *at* ω_2*. Thus,* $C1$*'s expected utility is* $(\log(L) + \log(H))/2$ *which exceeds* $\log(L)$*.*

It is therefore important to understand how (and whether) coalitions will be formed when both providers and customers are decision-makers. We formulate the interactions among providers and customers as a nontransferable payoff coalitional game, and show that this game has a nonempty core (Section V-A). Thus, the grand coalition is optimal in the sense that it generates at least one payoff profile for providers and customers that can not be blocked by any coalition. We then investigate how to compute such a profile and prove that it is polynomial time computable when the payoffs are linear functions and the network states are deterministic (Section V-B).

A. An NTU Game Formulation

We first redefine a coalition as follows.

Definition V.1. *A coalition* $(\mathcal{S}, \mathcal{T}), \mathcal{S} \subseteq \mathcal{N}, \mathcal{T} \subseteq \mathcal{M}$*, is a subset of providers and customers who cooperate, that is each customer in* $\mathcal{T}$ *agrees to subscribe to one of the providers in* $\mathcal{S}$*, and providers in* $\mathcal{S}$ *jointly serve customers in* $\mathcal{T}$*. The* grand coalition *now refers to* $(\mathcal{N}, \mathcal{M})$*.*

Consider a network realization ω, and also a coalition $(\mathcal{S}, \mathcal{T})$. Let $y_{jk}^{\mathcal{ST}}(\omega)$ denote the rate a customer $j \in \mathcal{T}$ receives from a service unit $k \in \mathcal{B}_{\mathcal{S}}$; $y_{jk}^{\mathcal{ST}}(\omega) = \alpha_{jk}(\omega) r_{jk}(\omega)$. Define customer j's rate vector $\mathbf{y}_j^{\mathcal{ST}}(\omega) = (y_{jk}^{\mathcal{ST}}(\omega), k \in \mathcal{B}_{\mathcal{S}})$, and its rate vector from a provider $i \in \mathcal{S}$, $\mathbf{y}_{ij}^{\mathcal{ST}}(\omega) = (y_{jk}^{\mathcal{ST}}(\omega), k \in \mathcal{B}_i)$. For serving customer j, provider i receives a payoff (e.g., revenue from j) $f_{ij}(\mathbf{y}_{ij}^{\mathcal{ST}}(\omega))$, while customer j attains a payoff (satisfaction) $g_j(\mathbf{y}_j^{\mathcal{ST}}(\omega))$, which is a function of j's received rate. Such revenue and satisfaction functions ($f_{ij}(\cdot)$s and $g_j(\cdot)$s, resp.) are widely assumed to be concave [21], [22]. Thus, the expected payoffs of provider i and customer j will be $\sum_{\substack{j\in\mathcal{T}\\ \omega\in\Omega}} \mathbb{P}(\omega) f_{ij}(\mathbf{y}_{ij}^{\mathcal{ST}}(\omega))$ and $\sum_{\omega\in\Omega} \mathbb{P}(\omega) g_j(\mathbf{y}_j^{\mathcal{ST}}(\omega))$, respectively, and are determined once the service-unit-customer allocations $\{\alpha_{jk}(\omega)\}$ are decided.

Similar to that in Section III, we define a feasible joint action of providers and customers in coalition $(\mathcal{S}, \mathcal{T})$ as an allocation $\{\alpha_{jk}(\omega), j \in \mathcal{T}, k \in \mathcal{B}_{\mathcal{S}}, \omega \in \Omega\}$ that satisfies the following conditions.

1) $\sum_{j\in\mathcal{T}} \alpha_{jk}(\omega) \le 1, k \in \mathcal{B}_{\mathcal{S}}, \omega \in \Omega$
2) $\sum_{k\in\mathcal{B}_{\mathcal{S}}} \alpha_{jk}(\omega) \le 1, j \in \mathcal{T}, \omega \in \Omega$
3) $\alpha_{jk}(\omega) \ge 0, j \in \mathcal{T}, k \in \mathcal{B}_{\mathcal{S}}, \omega \in \Omega$.

Note that for any feasible joint action $\{\alpha_{jk}(\omega), j \in \mathcal{T}, k \in \mathcal{B}_{\mathcal{S}}, \omega \in \Omega\}$, there is a schedule that allocates service units to customers, ensuring that for all $j \in \mathcal{T}, k \in \mathcal{B}_{\mathcal{S}}, \omega \in \Omega$, service unit k serves customer j for $\alpha_{jk}(\omega)$ fraction of time [23]. Let $\mathcal{A}(\mathcal{S}, \mathcal{T})$ denote the joint action space of coalition $(\mathcal{S}, \mathcal{T})$. For a joint action $\alpha \in \mathcal{A}(\mathcal{S}, \mathcal{T})$, let $\mathcal{F}^{\mathcal{ST}}(\alpha) \in \mathbb{R}^{\mathcal{S}\cup\mathcal{T}}$ be the resulting payoff vector. We now define the set of feasible payoff profiles, $v(\mathcal{S}, \mathcal{T})$, as follows:

$$v(\mathcal{S}, \mathcal{T}) = \{\mathbf{x} \in \mathbb{R}^{\mathcal{S}\cup\mathcal{T}} : \mathbf{x} \le \mathcal{F}^{\mathcal{ST}}(\alpha) \text{ for some } \alpha \in \mathcal{A}(\mathcal{S}, \mathcal{T})\}. \tag{5}$$

That is, $v(\mathcal{S}, \mathcal{T})$ is the set of all payoff profiles which are achievable through some feasible joint action of coalition $(\mathcal{S}, \mathcal{T})$, and all payoff profiles lower than those. Now, according to Definition IV.1, $< (\mathcal{N}, \mathcal{M}), v >$ is a well defined NTU game. The core of the game is defined as follows.

$$\mathcal{C} = \{\mathbf{x} \in v(\mathcal{N}, \mathcal{M}) : \forall(\mathcal{S}, \mathcal{T}), \nexists \mathbf{z} \in v(\mathcal{S}, \mathcal{T}) \text{ such that } \mathbf{z}_i > \mathbf{x}_i, \mathbf{z}_j > \mathbf{x}_j,\ \forall i \in \mathcal{S}, j \in \mathcal{T}\} \tag{6}$$

Note that every payoff profile in the core renders the grand coalition stable. To see this, let providers and customers form the grand coalition and select a joint action that results in a payoff profile $\mathbf{x} \in \mathcal{C}$. Now, suppose a set of providers and customers $(\mathcal{S}, \mathcal{T}) \subset (\mathcal{N}, \mathcal{M})$ leave the grand coalition and choose a joint action and the corresponding payoff profile $\mathbf{z} \in v(\mathcal{S} \cup \mathcal{T})$. They, however, would do so only if all of them receive a higher payoff than what they could in the grand coalition, i.e., $\mathbf{z}_i > \mathbf{x}_i, \mathbf{z}_j > \mathbf{x}_j,\ \forall i \in \mathcal{S}, j \in \mathcal{T}$. But this contradicts the fact that $\mathbf{x} \in \mathcal{C}$. Therefore, the grand coalition is stable. Also note that the condition for a payoff profile to be in the core, as given in (6), does not depend on which provider a customer subscribes to. Therefore, once the grand coalition has been formed, a customer can not improve his payoff by changing subscription.

Allowing customers also to be decision makers is likely to improve the global utility, as the following example shows.

Example V.2. *Let $\mathcal{N} = \{1,2\}$, $\mathcal{B}_i = \{i\}, i = 1,2$, and $\mathcal{M} = \{1,2,3\}$. Let $r_{1k} = P$ and $r_{jk} = Q$ for $j > 1$, for all $k \in \mathcal{B}_{\mathcal{N}}$. Suppose $P < Q$. Let each provider's payoff equal the sum of the service rates it provides to the customers, and each customer's payoff equal its service rate. Then the only payoff allocation in the core provides payoff Q to each provider and Q to customers $2,3$ each and 0 to customer 1. This allocation corresponds to customers $2,3$ being served full-time from the service units and customer 1 receiving no service whatsoever. We now prove that this payoff profile is the only allocation in the core. Note that neither the providers nor customers $2,3$ can enhance their payoffs in any other coalition. If customer 1 receives any service (say from provider 1), then either customer 2 or 3 receives lower payoffs (as they can no longer receive full-time service), and then provider 1 (and 2) generates revenue at rate P for the fraction of time it serves customer 1. Thus, provider 1's (or 2's or for both) payoff decreases below Q. Then the provider whose payoff decreases can split along with one of the customers whose payoff decreases and each can obtain a payoff of Q - thus this payoff blocks any allocation that provides customer 2 a positive payoff. Now, note that the payoff profile in the core provides an aggregate rate of $2Q$ to the customers. This is the maximum possible aggregate rate to the customers. In fact, if customers had subscribed apriori, say $\mathcal{M}_1 = \{1\}, \mathcal{M}_2 = \{2,3\}$ (coalition game in Section IV), and providers' payoffs had been the sum of their customers rates, $v(\{1\}) = [0,P], v(\{2\}) = [0,Q]$, and the only payoff in the core would have been (P,Q) - this is the only payoff that provides a payoff of at least P (Q, resp.) to provider 1 (2, resp.). This corresponds to serving provider 1's (2's, resp.) customer using provider 1's (2's, resp.) service unit respectively. Thus, effectively, there is no cooperation! Also, the customers obtain a lower aggregate rate, $P+Q$.*

In general, cooperation is expected to enhance the participants' payoffs. But, in provider customer games, as the following example illustrates, it can reduce the payoffs of some of the providers and customers. This is because cooperation provides more options to the providers and customers.

Example V.3. *Consider a provider customer game with $\mathcal{N} = \mathcal{M} = \{1,2,3,4\}$, $\mathcal{B}_i = \{i\}$, $i = 1,2,3,4$ and $|\Omega| = 1$. Let $r_{11} = r_{44} = P$, $r_{21} = r_{43} = Q$ and $r_{jk} = 0$ otherwise. Suppose $P < Q$. Suppose that the payoff of each provider (customer, resp.) equals the aggregate rate it provides (receives, resp.). The unique core payoff for providers and customers in the grand coalition (of all providers and customers) are $(Q,0,Q,0)$ and $(0,Q,0,Q)$, respectively. In coalitions with 1 provider and 1 customer each, i.e., $\mathcal{S}_i = \mathcal{T}_i = \{i\}$ (providers do not cooperate and customer i has subscribed to provider i without examining other options), providers' and customers' payoffs are $(P,0,0,P)$. Clearly, cooperation reduces the payoffs of provider 4 and customer 1 - but the aggregate providers' (customers', resp.) payoffs improve!*

We now prove that the core of the above game is nonempty.

Theorem V.1. *The nontransferable payoff game $< (\mathcal{N},\mathcal{M}), v >$ is balanced and hence has a nonempty core.*

Proof: Consider a balanced collection of coalitions $\mathcal{I} \subset 2^{\mathcal{N}\cup\mathcal{M}}\backslash\emptyset$ and the corresponding balancing weights $(\lambda_{\mathcal{ST}}, (\mathcal{S},\mathcal{T}) \in \mathcal{I})$. Also, let $u \in \mathbb{R}^{\mathcal{N}\cup\mathcal{M}}$ be such that $u^{\mathcal{ST}} \in v(\mathcal{S},\mathcal{T})$ for all $(\mathcal{S},\mathcal{T}) \in \mathcal{I}$[9], i.e., there exists a joint action $\{\alpha^{\mathcal{ST}}_{jk}(\omega), j \in \mathcal{T}, k \in \mathcal{B}_{\mathcal{S}}, \omega \in \Omega\}$ for each $(\mathcal{S},\mathcal{T}) \in \mathcal{I}$ such that

a) $\{\alpha^{\mathcal{ST}}_{jk}(\omega), j \in \mathcal{T}, k \in \mathcal{B}_{\mathcal{S}}, \omega \in \Omega\}$ satisfy feasibility constraints (1) - (3) introduced in this section, for each $(\mathcal{S},\mathcal{T}) \in \mathcal{I}$.

b) $u_i \leq \sum_{\substack{j\in\mathcal{T}\\ \omega\in\Omega}} \mathbb{P}(\omega) f_{ij}(\mathbf{y}^{\mathcal{ST}}_{ij}(\omega)),\ \forall i \in \mathcal{S}$, where $\mathbf{y}^{\mathcal{ST}}_j(\omega)$ denotes customer j's rate vector corresponding to joint action $\{\alpha^{\mathcal{ST}}_{jk}(\omega), j \in \mathcal{T}, k \in \mathcal{B}_{\mathcal{S}}, \omega \in \Omega\}$.

c) $u_j \leq \sum_{\omega\in\Omega} \mathbb{P}(\omega) g_j(\mathbf{y}^{\mathcal{ST}}_j(\omega)),\ \forall j \in \mathcal{T}$.

We next show that $u \in v(\mathcal{N},\mathcal{M})$. The procedure is similar to that in the proof of Theorem IV.2. First, define a joint action $\{\alpha_{jk}(\omega), j \in \mathcal{M}, k \in \mathcal{B}_{\mathcal{N}}, \omega \in \Omega\}$ as follows

$$\alpha_{jk}(\omega) = \sum_{(\mathcal{S},\mathcal{T})\in\mathcal{I}:\substack{j\in\mathcal{T}\\ k\in\mathcal{B}_{\mathcal{S}}}} \lambda_{\mathcal{ST}} \alpha^{\mathcal{ST}}_{jk}(\omega). \tag{7}$$

The following two steps, concludes the proof.

Step 1: We need to show that $\{\alpha_{jk}(\omega), j \in \mathcal{M}, k \in \mathcal{B}_{\mathcal{N}}, \omega \in \Omega\}$ satisfy feasibility constraints (1) - (3). The argument is similar to that in Step 1 of the proof of Theorem IV.2, and is omitted for brevity.

Step 2: Using concavity of $f_{ij}(\cdot)$s and $g_j(\cdot)$s, it is straightforward to show that

i) $u_i \leq \sum_{\substack{j\in\mathcal{M}\\ \omega\in\Omega}} \mathbb{P}(\omega) f_{ij}(\mathbf{y}^{\mathcal{NM}}_{ij}(\omega)),\ \forall i \in \mathcal{N}$, and

ii) $u_j \leq \sum_{\omega\in\Omega} \mathbb{P}(\omega) g_j(\mathbf{y}^{\mathcal{NM}}_j(\omega)),\ \forall j \in \mathcal{M}$,

where $(\mathbf{y}^{\mathcal{NM}}_j(\omega), j \in \mathcal{M}, \omega \in \Omega)$ are the rate vectors resulting from the joint action $\{\alpha_{jk}(\omega), j \in \mathcal{M}, k \in \mathcal{B}_{\mathcal{N}}, \omega \in \Omega\}$ (Refer to Step 2 of the proof of Theorem IV.2 for analogous arguments).

Thus the game is balanced, and nonemptyness of the core follows from Theorem IV.1. ■

B. Computation of a Payoff Profile in the Core

The coalitional game $< (\mathcal{N},\mathcal{M}), v >$ can not be related to the exchange market setting that we constructed in Section IV-B. But, we obtain a payoff profile in the core, with arbitrary precision using a different technique. We first introduce the concept of *approximate core*.

[9]For any $u \in \mathbb{R}^{\mathcal{N}\cup\mathcal{M}}$, we denote by $u^{\mathcal{ST}} \in \mathbb{R}^{\mathcal{S}\cup\mathcal{T}}$, the sub-vector of u corresponding to the coalition $(\mathcal{S},\mathcal{T})$, i.e., $u^{\mathcal{ST}}_i = u_i, \forall i \in \mathcal{S}$ and $u^{\mathcal{ST}}_j = u_j, \forall j \in \mathcal{T}$

A feasible payoff profile is said to be in the approximate core of the game $<(\mathcal{N},\mathcal{M}),v>$, $\mathcal{C}_\epsilon$, if it can not be blocked by at least a margin of ϵ, by any coalition. Formally,

$$\begin{aligned}\mathcal{C}_\epsilon = \{\mathbf{x}\in v(\mathcal{N},\mathcal{M}) : \forall(\mathcal{S},\mathcal{T}), \nexists \mathbf{z}\in v(\mathcal{S},\mathcal{T}) \text{ such that} \\ \mathbf{z}_i > \mathbf{x}_i+\epsilon, \mathbf{z}_j > \mathbf{x}_j+\epsilon,\ \forall i\in\mathcal{S}, j\in\mathcal{T}\}\end{aligned} \tag{8}$$

It is straightforward to check that for $\epsilon<0$, $\epsilon=0$, and $\epsilon>0$, $\mathcal{C}_\epsilon$ is a subset of, equal to, and superset of $\mathcal{C}$, respectively. Here we are naturally interested in the approximate core for an strictly positive value of ϵ. It is also evident from the definition of the approximate core (8), that by letting ϵ go to 0, payoff profiles in the approximate core get closer to those in the core of the game, hence the name approximate core.

Let us fix an $\epsilon>0$. In the following two steps, we present an algorithm to compute a payoff profile in the approximate core, $\mathcal{C}_\epsilon$, of our NTU game.

Step 1: For each coalition $(\mathcal{S},\mathcal{T})\subseteq(\mathcal{N},\mathcal{M})$, construct a finite sequence of payoff profiles in $v(\mathcal{S},\mathcal{T})$, $V(\mathcal{S},\mathcal{T}) = \{u^{1,\mathcal{ST}}, u^{2,\mathcal{ST}},\ldots,u^{k_{\mathcal{ST}},\mathcal{ST}}\}$, such that any payoff profile in $v(\mathcal{S},\mathcal{T})$ is dominated by at least one profile in $V(\mathcal{S},\mathcal{T})$ after $\epsilon/2$ is added to each component of the latter. In other words,

$$\begin{aligned}\forall\mathbf{x}\in v(\mathcal{S},\mathcal{T}),\ \exists u^{m,(\mathcal{ST})}\in V(\mathcal{S},\mathcal{T}) \text{ such that} \\ \mathbf{x}-u^{m,(\mathcal{ST})} \le \frac{\epsilon}{2}\mathbb{1}_{1\times(\mathcal{S}\cup\mathcal{T})}.\end{aligned} \tag{9}$$

We next show how to construct $V(\mathcal{S},\mathcal{T})$, for any $(\mathcal{S},\mathcal{T})\subseteq(\mathcal{N},\mathcal{M})$. Note that we can restrict our search to the Pareto-optimal payoff profiles in $v(\mathcal{S},\mathcal{T})$, i.e., those such that no other payoff profile in $v(\mathcal{S},\mathcal{T})$ can give every one in $(\mathcal{S},\mathcal{T})$ a strictly better payoff. Also, every Pareto-optimal payoff profile $\mathbf{x}$ in $v(\mathcal{S},\mathcal{T})$ can be obtained as a solution of the following optimization, OPT($\lambda^{\mathcal{ST}}$), for different choices of $\lambda^{\mathcal{ST}}$; $\lambda^{\mathcal{ST}}\neq 0$ are sets of nonnegative weights [30, Section 2.6.3].
OPT($\lambda^{\mathcal{ST}}$) : $\max\sum_{i\in\mathcal{S}}\lambda_i^{\mathcal{ST}}\mathbf{x}_i+\sum_{j\in\mathcal{T}}\lambda_j^{\mathcal{ST}}\mathbf{x}_j$
subject to:

1) $\mathbf{x}_i=\sum_{j\in\mathcal{T}, \omega\in\Omega}\mathbb{P}(\omega)f_{ij}(\mathbf{y}_{ij}^{\mathcal{ST}}(\omega))$
2) $\mathbf{x}_j=\sum_{\omega\in\Omega}\mathbb{P}(\omega)g_j(\mathbf{y}_j^{\mathcal{ST}}(\omega))$
3) $y_{jk}^{\mathcal{ST}}(\omega)=\alpha_{jk}(\omega)r_{jk}(\omega),\quad j\in\mathcal{T},k\in\mathcal{B}_\mathcal{S},\omega\in\Omega$
4) $\sum_{j\in\mathcal{T}}\alpha_{jk}(\omega)\le 1,\quad k\in\mathcal{B}_\mathcal{S},\omega\in\Omega$
5) $\sum_{k\in\mathcal{B}_\mathcal{S}}\alpha_{jk}(\omega)\le 1,\quad j\in\mathcal{T},\omega\in\Omega$
6) $\alpha_{jk}(\omega)\ge 0,\quad j\in\mathcal{T},k\in\mathcal{B}_\mathcal{S},\omega\in\Omega$

Let $\mathbf{x}(\lambda^{\mathcal{ST}})$ be the solution of OPT($\lambda^{\mathcal{ST}}$). Note that the function $\mathbf{x}(\lambda^{\mathcal{ST}})$ is continuous in $\lambda^{\mathcal{ST}}$. Also, since scaling $\lambda^{\mathcal{ST}}$ does not change the solution of the above optimization, we can set $\sum_{i\in\mathcal{S}}\lambda_i^{\mathcal{ST}}+\sum_{j\in\mathcal{T}}\lambda_j^{\mathcal{ST}}=1$. As a result, we have a continuous function over a bounded set $\{\lambda^{\mathcal{ST}}:\sum_{i\in\mathcal{S}}\lambda_i^{\mathcal{ST}}+\sum_{j\in\mathcal{T}}\lambda_j^{\mathcal{ST}}=1\}$, whose range covers the set of Pareto-optimal feasible payoff profiles. It then follows that if we select an appropriate collection of weights $\{\lambda^{1,\mathcal{ST}},\lambda^{2,\mathcal{ST}},\ldots,\lambda^{k_{\mathcal{ST}},\mathcal{ST}}\}$, the set of feasible payoff profiles obtained by solving the above optimization will be the desired set $V(\mathcal{S},\mathcal{T})$.

Step 2: Consider the discrete coalitional game $<(\mathcal{N},\mathcal{M}),V>$, where $V(\mathcal{S},\mathcal{T})$ is as defined in Step 1. Let $\hat{\mathcal{C}}_{\epsilon/2}$ denote the $\epsilon/2$ approximate core of $<(\mathcal{N},\mathcal{M}),V>$. In the following, we show that $\hat{\mathcal{C}}_{\epsilon/2}\neq\emptyset$, and an element of $\hat{\mathcal{C}}_{\epsilon/2}$ can be computed in finite time using a brute force search. Next, we show that $\hat{\mathcal{C}}_{\epsilon/2}\subset\mathcal{C}_\epsilon$, and consequently, the result of the brute force search belongs to $\mathcal{C}_\epsilon$.

Theorem V.2. *$\hat{\mathcal{C}}_{\epsilon/2}\neq\emptyset$. Furthermore, $\hat{\mathcal{C}}_{\epsilon/2}\subset\mathcal{C}_\epsilon$.*

Proof: Consider a payoff profile $\mathbf{x}\in\mathcal{C}$, the core of the original game. Such a payoff profile exists, by Theorem V.1. It follows from the construction of $V(\mathcal{M},\mathcal{N})$ that there exists a $\hat{\mathbf{x}}\in V(\mathcal{N},\mathcal{M})$ satisfying $\mathbf{x}_i-\hat{\mathbf{x}}_i\le\epsilon/2, \mathbf{x}_j-\hat{\mathbf{x}}_j\le\epsilon/2,\ \forall i\in\mathcal{N}, j\in\mathcal{M}$. We prove via contradiction that $\hat{\mathbf{x}}\in\hat{\mathcal{C}}_{\epsilon/2}$. Suppose it is not true. Then there exists a coalition $(\mathcal{S},\mathcal{T})$ and a payoff profile $\hat{\mathbf{z}}\in V(\mathcal{S},\mathcal{T})$ such that $\hat{\mathbf{z}}_i>\hat{\mathbf{x}}_i+\epsilon/2, \hat{\mathbf{z}}_j>\hat{\mathbf{x}}_j+\epsilon/2,\ \forall i\in\mathcal{S}, j\in\mathcal{T}$. This implies that $\hat{\mathbf{z}}_i>\mathbf{x}_i, \hat{\mathbf{z}}_j>\mathbf{x}_j,\ \forall i\in\mathcal{S},j\in\mathcal{T}$, i.e., $\mathbf{x}$ is blocked by the coalition $(\mathcal{S},\mathcal{T})$ through the payoff profile $\hat{\mathbf{z}}$. This is in contradiction with $\mathbf{x}\in\mathcal{C}$. Thus $\hat{\mathbf{x}}\in\hat{\mathcal{C}}_{\epsilon/2}$, and so $\hat{\mathcal{C}}_{\epsilon/2}\neq\emptyset$.

We now show that $\hat{\mathcal{C}}_{\epsilon/2}\subset\mathcal{C}_\epsilon$. Consider some $\hat{\mathbf{x}}\hat{\mathcal{C}}_{\epsilon/2}$. Suppose $\hat{\mathbf{x}}\notin\mathcal{C}_\epsilon$. Then by (8), there exists a coalition $(\mathcal{S},\mathcal{T})$ and a payoff profile $\mathbf{z}\in v(\mathcal{S},\mathcal{T})$ such that $\mathbf{z}_i>\hat{\mathbf{x}}_i+\epsilon, \mathbf{z}_j>\hat{\mathbf{x}}_j+\epsilon,\ \forall i\in\mathcal{S},j\in\mathcal{T}$. It follows from the construction of $V(\mathcal{S},\mathcal{T})$ that there exists a $\hat{\mathbf{z}}\in V(\mathcal{S},\mathcal{T})$ satisfying $\hat{\mathbf{z}}_i\ge\mathbf{z}_i-\epsilon/2, \hat{\mathbf{z}}_j\ge\mathbf{z}_j-\epsilon/2,\ \forall i\in\mathcal{S},j\in\mathcal{T}$. Putting these inequalities together, we have $\hat{\mathbf{z}}_i>\hat{\mathbf{x}}_i+\epsilon/2, \hat{\mathbf{z}}_j>\hat{\mathbf{x}}_j+\epsilon/2,\ \forall i\in\mathcal{S},j\in\mathcal{T}$. This contradicts the fact that $\hat{\mathbf{x}}\in\hat{\mathcal{C}}_{\epsilon/2}$. Thus $\hat{\mathbf{x}}\in\mathcal{C}_\epsilon$ and the claim follows. ■

This method requires us to quantize the set of payoff profiles in a finite set (Step 1), which and therefore the computation time) grows exponentially in the dimension of the problem and parameter ϵ. In addition, the quantization and the subsequent brute-force search is centralized in that a central unit needs to know all the payoff functions $f(\cdot),g(\cdot)$. The algorithm therefore essentially provides a proof of concept, that an allocation in the core can be computed to arbitrary precision, in finite time.

A payoff profile in the core may however be computed in polynomial time using distributed computations in an important special case: (i) the network states are deterministic (i.e., $|\Omega|=1$) and (ii) providers and customers have linear payoff functions, i.e., $f_{ij}(\mathbf{y}_{ij})=\beta_i\sum_{k\in\mathcal{B}_i}y_{jk}$ and $g_j(\mathbf{y}_j)=\gamma_j\sum_{k\in\mathcal{B}_\mathcal{S}}y_{jk}$, $\beta_i,\gamma_j>0$ for all $i\in\mathcal{N},j\in\mathcal{M}$. Also, no provider i (customer j, resp.) is required to reveal its revenue (satisfaction) per unit throughput β_i (γ_j, resp.) to any other provider (provider and customer, resp.).

Let the non-zero link rates r_{jk} be in $\{r_1,r_2,\ldots,r_L\}$ where $r_1>r_2>\ldots>r_L$, and $\mathcal{E}_i$ be the set of service-unit-customer links with rate r_i, $\mathcal{G}^{(i)}$ be the bipartite graph $(\mathcal{V},\mathcal{E}_i)$ (i.e., the service-unit-customer graph $\mathcal{G}=(\mathcal{V},\mathcal{E})$ but with only the links in $\mathcal{E}_i$). We describe a service-unit-customer allocation algorithm for the grand coalition, **ALG** (elucidated in Fig. 1):

1) Let $i=1$, $\mathcal{E}_0'=\phi$.
2) Assign service units to customers as per any maximal

matching[10] in $\mathcal{G}^{(i)} = (\mathcal{V}, \mathcal{E}_i \setminus \mathcal{E}'_{i-1})$. Let $\mathcal{E}'_i$ be the set of links incident to service units or customers matched (allocated service) so far.

3) If $\mathcal{E}_{i+1} \setminus \mathcal{E}'_i$ is empty, terminate; else increment i, and go to the previous step.

Theorem V.3. *For networks with linear payoff functions and deterministic states, the payoff profile corresponding to* **ALG**, **ALG-PAYOFF**, *is in the core.*

Proof: We prove that **ALG-PAYOFF**, denoted as $\mathbf{x}$, is in $\mathcal{C}$. Suppose not. Then there exists sets $\mathcal{S}$ and $\mathcal{T}$ of providers and customers respectively, and a Pareto optimal payoff profile in $v(\mathcal{S}, \mathcal{T})$ that blocks $\mathbf{x}$. Every such payoff profile in $v(\mathcal{S}, \mathcal{T})$ can be obtained as a solution of OPT($\lambda^{\mathcal{ST}}$). However, in the case of linear payoff functions, OPT($\lambda^{\mathcal{ST}}$) becomes a linear optimization problem with totally unimodular [31] constraint matrix, and hence assumes only integral solutions [32] (i.e., $\alpha_{jk} \in \{0,1\}, j \in \mathcal{T}, k \in \mathcal{B}_{\mathcal{S}}$), Thus all Pareto optimal payoff profiles in $v(\mathcal{S}, \mathcal{T})$ correspond to matchings.[11] In the following we prove that no payoff profile in $V(\mathcal{S}, \mathcal{T})$, that corresponds to a matching, can block $\mathbf{x}$. Otherwise, there exists a service unit $k \in \mathcal{B}_{\mathcal{S}}$ that serves in the coalition $(\mathcal{S}, \mathcal{T})$ at strictly higher rate than in the grand coalition. Assume that its new rate is r_{l^*}, and k serves customer $j \in \mathcal{T}$ at this rate. Let customer j and service unit k obtain rates $r_{l'}$ and $r_{l''}$ respectively, in the grand coalition; $l'' > l^*$. Now, there are two possibilities.

1) $l' > l^*$: this implies that till the l^*th stage in **ALG** neither customer j nor service unit k were allocated (matched), and link (j,k) was not selected at the l^*th stage. This contradicts the algorithm.
2) $l' \leq l^*$: this is equivalent to $r_{l'} \geq r_{l^*}$, and so contradicts the fact that $(\mathcal{S}, \mathcal{T})$ can block $\mathbf{x}$.

■

ALG can be computed in $|\mathcal{B}_{\mathcal{N}}|^2 |\mathcal{M}_N|^2$ time. This is because **ALG** needs at most $|\mathcal{E}|$ iterations, in each of which it executes (i) $|\mathcal{E}|$ operations and (ii) computes a maximal matching. A maximal matching can be computed in a graph with $|\mathcal{E}|$ edges in $O(\log(|\mathcal{E}|))$ time [33]. The result follows by observing that $|\mathcal{E}| = |\mathcal{B}_{\mathcal{N}}||\mathcal{M}_N|$.

C. Provider Customer Coalition as a Competitive Market

The coalitional game $< (\mathcal{N}, \mathcal{M}), v >$ can not be related to the exchange market setting (Section IV-B). However, interaction among providers and customers when both are decision makers can be modelled as a two-sided market which has been studied in [34] ("pairing model" in the terminology of [34]). In a pairing model, providers as well as customers are agenets, while service units and customers themselves can be considered as goods; the providers' goods are paired against the customers' goods. Our setup is more general than the one in [34] in the following two ways:

[10] A matching is maximal if it is not a proper subset of any other matching.

[11] Note that all the payoff profiles in the core are Pareto optimal in $v(\mathcal{N}, \mathcal{M})$, and hence, following the arguments in the proof, correspond to matchings. Thus, if payoff functions are linear, fractional associations are not needed to achieve a payoff profile in the core.

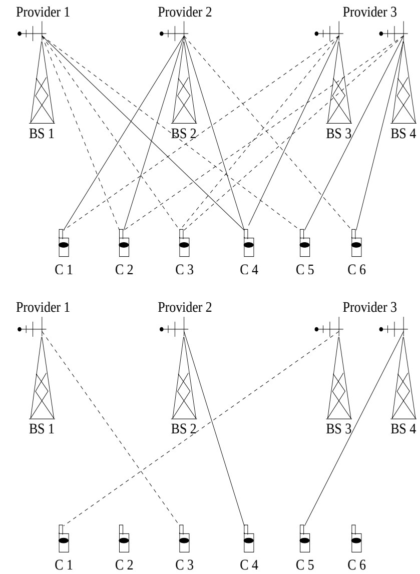

Fig. 1: The top figure shows an example network. Each base station has access to only one channel and therefore constitutes one service unit. The solid and dashed links have rates 200 Kbps and 100 Kbps respectively; no link between a service unit-customer pair denotes 0 rate. The bottom figure shows the service-unit-customer allocations generated by **ALG**: the solid and dashed links are selected in iterations 1 and 2 respectively.

i) we allow multiple goods per agent (more than one service units for a provider).
ii) we also consider fractional assignments ([34] only allowed integral matchings).

In a two-sided market, an allocation, to be a market equilibrium, must satisfy a certian *reciprocity* condition (see [34, Section 3]) in addition to *budget constraint* and *market clearing*. For the provider-customer competitive market, this condition takes the following form.

reciprocity: if customer j is assigned α_{jk} fraction of a service unit $k \in \mathcal{B}_i$, the provider i gets $\sum_{k \in \mathcal{B}_i} \alpha_{jk}$ fraction of customer j (considered as a good now).

It is shown in [34] (by means of a counter-example) that a competitive equilibrium need not exist in a two-sided competitive market, even when the market has a non-empty core.

Nevertheless, in order to ensure consistency with the overall setup we use, we choose to position the provider-customer interactions in a cooperative context. This positioning provides following important advantages. It allows us to analyze the impact of composition and size of coalitions on the payoffs of providers and customers using coalitional game theory - a tool that may not port to the competition interpretation. Specifically, we do not conclude apriori that any customer can be assigned to any provider, but allow the possibility that the assignments will occur in smaller sets (coalitions). But, we show, that the grand coalition is optimal (or stable) in the sense that there exists one service unit-customer allocation such that the corresponding payoff profile can not be blocked by that of any smaller coalition, and thus no set of providers and customers has incentive to split from the grand coalition(Theorem V.1). We also present a fully distributed and polynomial time algorithm to compute such an allocation and the corresponding payoff profile for linear payoff functions and deterministic network states, that can be executed without participants having to reveal confidential information such as payoff functions to each other (Section V-B).

Finally, we use the above theoretical framework to discover several interesting artifacts of the system through examples. For instance, as Example V.3 shows, some providers and customers *may* be worse off in the grand coalition, than when they operate in smaller groups, since the grand coalition offers more choices to everybody, including the potential collaborators of the participants. More importantly, the participants may not be able to circumvent this loss by separating from the grand coalition, as they need to persuade others to leave the coalition with them so as to enhance their payoffs (i.e., a customer needs to persuade a provider to leave with it in order not to have 0 rate and therefore 0 satisfaction, and similarly for a provider). However, our numerical evaluations demonstrate how pervasive this phenomenon is (SectionVII-B). We see that providers' payoffs usually increase due to coalitions and most of the customers' payoffs and their aggregate payoffs typically increase, but payoffs of some customers may decrease as well.

VI. Generalizations

A. Generalized Payoff Sharing

We have so far assumed that players (providers and customers) do not share their payoffs. When a group of providers, for instance, agree to share their payoffs, instead of each one trying to maximize its own payoff, they attempt to maximize their aggregate payoff, which is generally higher than the sum of the maximized individual payoffs - the increase in the aggregate may lead to increase in individual shares. The following example illustrates this phenomenon.

Example VI.1. *Consider the provider coalitional game in Section IV (i.e., customers have subscribed apriori) and the setting of Example V.2: $\mathcal{M}_1 = \{1\}, \mathcal{M}_2 = \{2, 3\}$, and provider's payoffs (revenues) equal the sum of their customers rates. Recall that when the providers do not share the aggregate payoffs, the only payoff in the core is (P, Q). But, the providers can generate an aggregate revenue of $2Q$ by serving only provider 2's customers, which can for example be shared as $(\frac{P+Q}{2}, Q + \frac{Q-P}{2})$. Then, provider 1's (2's, resp.) payoff is higher (lower, resp.) than its aggregate customer revenue; but, each provider's payoff (and hence the aggregate payoff) strictly improves due to payoff sharing (compare with (P, Q)).*

But, again, players may (i) refuse to share payoffs owing to mutual distrust (and the need to disclose individual payoffs) (ii) not be able to share payoffs as not all types of payoffs can be shared, since they may not have monetary equivalence and may have individual satisfaction connotations.

Let $\mathcal{P} = \{\mathcal{N}_l, l \in \mathcal{L}\}$ be a partition of the set of providers $\mathcal{N}$. Now consider a coalition $(\mathcal{S}, \mathcal{T})$ of providers and customers. Define $\mathcal{S}_l = \mathcal{S} \cap \mathcal{N}_l,\ \forall l \in \mathcal{L}$. We consider a generalized payoff sharing among providers in which for each $l \in \mathcal{L}$ providers in the same $\mathcal{S}_l$ share their payoffs, but those in different $\mathcal{S}_l$ do not share their payoffs. For example, if $\mathcal{P} = \mathcal{N}$, all providers in a coalition share their payoffs, while if $\mathcal{P} = \{\{i\}, i \in \mathcal{N}\}$ no provider shares its payoff with others. The feasibility constraints, and consequently the set of feasible joint actions remain the same as in Section V. But for a payoff profile to be in the set of feasible payoff profiles $v(\mathcal{S}, \mathcal{T})$, we only need the *aggregate payoff* of each sharing group to be less than or equal to that in the payoff vector resulting from some feasible joint action. Formally,

$$\begin{aligned} v(\mathcal{S}, \mathcal{T}) = \{\mathbf{x} \in \mathbb{R}^{\mathcal{S}\cup\mathcal{T}} : &\sum_{i\in\mathcal{S}_l} \mathbf{x}_i \leq \sum_{i\in\mathcal{S}_l} \mathbf{z}_i, \forall l \in \mathcal{L}, \\ &\mathbf{x}_j \leq \mathbf{z}_j, \forall j \in \mathcal{T},\ \text{where} \\ &\mathbf{z} = \mathcal{F}^{\mathcal{S}\mathcal{T}}(\alpha) \text{ for some } \alpha \in \mathcal{A}(\mathcal{S}, \mathcal{T})\}. \end{aligned} \tag{10}$$

With this definition, the coalitional game $< (\mathcal{N}, \mathcal{M}), v >$ is now well defined. The definition of the core of this game will be the same as (6). Using a similar technique as that used in the proof of Theorems IV.2 and V.1, one can show that this game is balanced. It then follows from Theorem IV.1 that the core of this game is nonempty, and thus the grand coalition is stable. Also, a payoff profile in the approximate core of this game can be computed by an algorithm similar to the one discussed in Section V-B. All the formulations and results extend to the scenario where there are groups of customers, and only those within the same group share their payoffs.

B. Vector Payoff Functions

We have so far focused on coalitional games with scalar payoff functions. In this section, we examine the scenario where players have vector payoff functions. Such functions can have several payoff components of different types. For instance, a provider's payoff can be a vector of its total revenue, its competitive power in the market, fairness in the network, reputational issues and social welfare, among others. A customer's payoff, on the other hand, may consist of its service rate and cost, power consumption, the size of the network, and so on. Note that it is possible to have payoff vectors with mixed transferable and nontransferable payoff components. Then, which components can be shared does not only depend on the players in a coalition, but also on the types of components. Specifically, there could be groups of players, and players within each group would share the transferable components of payoffs, whenever they are in a coalition.

In this section, we investigate cooperation among providers in presence of vector payoff functions with two components; one transferable and another nontransferable. We consider the scenario where all providers in a coalition would share the transferable component. The formulations and results can extend to more general cases, where payoff functions have several components, and for given groups of providers, sharing of transferable components happens only among providers within the same group.

Consider a coalition $\mathcal{S}$ and a provider $i \in \mathcal{S}$. Assume that provider i's transferable and nontransferable payoff components be $f_i^t(\mathbf{y}_i^{\mathcal{S}}(\omega))$ and $f_i^n(\mathbf{y}_i^{\mathcal{S}}(\omega))$, respectively. Functions $f_i^t(\cdot)$s and $f_i^n(\cdot)$s are concave. The definitions of a rate vector $\mathbf{y}_i^{\mathcal{S}}(\omega)$, a feasible payoff profile $\{\alpha_{jk}(\omega), j \in \mathcal{M}_S, k \in \mathcal{B}_{\mathcal{S}}, \omega in \Omega\}$, and a joint action space $\mathcal{A}(\mathcal{S})$ are exactly the same as those in Section III. Now consider a joint action $\alpha \in \mathcal{A}(\mathcal{S})$. We define $\mathcal{F}^{\mathcal{S}}(\alpha) \in \mathbb{R}^{\mathcal{S}} \times \mathbb{R}^{\mathcal{S}}$ to be the payoff vector corresponding to the joint action α, i.e., $\mathcal{F}^{\mathcal{S}}(a) = (\mathbf{x}^t, \mathbf{x}^n)$ where $\mathbf{x}_i^t = \sum_{\omega \in \Omega} \mathbb{P}(\omega) f_i^t(\mathbf{y}_i^{\mathcal{S}}(\omega))$ and $\mathbf{x}_i^n = \sum_{\omega \in \Omega} \mathbb{P}(\omega) f_i^n(\mathbf{y}_i^{\mathcal{S}}(\omega))$ for all $i \in \mathcal{S}$, and $(\mathbf{y}_i^{\mathcal{S}}(\omega), i \in \mathcal{S}, \omega \in \Omega)$ are the rate vectors resulting from the joint action α.

We now define the set of feasible payoff profiles, $v(\mathcal{S})$, as follows.

$$\begin{aligned} v(\mathcal{S}) = \{(\mathbf{x}^t, \mathbf{x}^n) \in \mathbb{R}^{\mathcal{S}} \times \mathbb{R}^{\mathcal{S}} &: \mathbf{x}^t \cdot \mathbb{1}_{1\times\mathcal{S}} \leq \mathbf{z}^t \cdot \mathbb{1}_{1\times\mathcal{S}}, \\ &\mathbf{x}^n \leq \mathbf{z}^n, \\ &\text{where } (\mathbf{z}^t, \mathbf{z}^n) = \mathcal{F}^{\mathcal{S}}(\alpha) \text{ for some } \alpha \in \mathcal{A}(\mathcal{S})\}. \end{aligned} \tag{11}$$

In words, $v(\mathcal{S})$ is the set of all payoff profiles in which the nontransferable utilities of providers in $\mathcal{S}$, as well as the sum of their transferable utilities, are either equal to or less than those in the payoff vector generated by some feasible joint action.

We now define the core of this game. To do this, an order relation between two different payoff profiles of a provider is necessary, i.e., we need to know which of the two payoffs, $(\mathbf{x}_i^t, \mathbf{x}_i^n)$ and $(\mathbf{z}_i^t, \mathbf{z}_i^n)$, provider i prefers. We consider a lexicographic ordering, in which a provider prefers the payoff that offers him higher nontransferable utility. In case there are several payoffs with this property, the one that offers the highest transferable utility is preferred. With this lexicographic order relation in place, the core of the game is defined to be the set of payoff profiles that can not be blocked lexicographically, by any coalition, i.e.,

$$\begin{aligned} \mathcal{C} = \{(\mathbf{x}^t, \mathbf{x}^n) \in v(\mathcal{N}) : \forall \mathcal{S}, \nexists (\mathbf{z}^t, \mathbf{z}^n) \in v(\mathcal{S}) \text{ such that} \\ \mathbf{z}_i^n > \mathbf{x}_i^n, \text{ or } \mathbf{z}_i^n = \mathbf{x}_i^n \text{ and } \mathbf{z}_i^t > \mathbf{x}_i^t, \ \forall i \in \mathcal{S}\}. \end{aligned} \tag{12}$$

We seek to show that $\mathcal{C}$ is nonempty. Since functions $f_i^t(\cdot)$s and $f_i^n(\cdot)$s are concave, using a similar technique as in the proof of Theorem IV.2, it is straightforward to verify that this game is balanced. But since the coalitional games considered in [25] have scalar payoff functions, it is not evident whether balancedness leads to nonemptiness of the core, in presence of vector payoff functions. However, with a slight twist in the definition of the core, we can use Theorem IV.1 and derive interesting results. Towards this end, we first define the *approximate core* for this game as follows

$$\begin{aligned} \mathcal{C}_\epsilon = \{(\mathbf{x}^t, \mathbf{x}^n) \in v(\mathcal{N}) : \forall \mathcal{S}, \nexists (\mathbf{z}^t, \mathbf{z}^n) \in v(\mathcal{S}) \text{ such that} \\ \mathbf{z}_i^n > \mathbf{x}_i^n + \epsilon, \text{ or } \mathbf{z}_i^n = \mathbf{x}_i^n + \epsilon \text{ and } \mathbf{z}_i^t > \mathbf{x}_i^t, \ \forall i \in \mathcal{S}\}. \end{aligned} \tag{13}$$

In words, $\mathcal{C}_\epsilon$ is the set of all payoff profiles that can not be blocked lexicographically by a margin of ϵ, by any coalition. Here is the main result.

Theorem VI.1. *For any $\epsilon > 0$, $\mathcal{C}_\epsilon$ is nonempty.*

Proof: Suppose that instead of lexicographic ordering, providers use a linear ordering to compare different payoffs, i.e., for some given $\lambda > 0$, $(\mathbf{x}_i^t, \mathbf{x}_i^n)$ is preferred over $(\mathbf{z}_i^t, \mathbf{z}_i^n)$ if $\lambda \mathbf{x}_i^t + \mathbf{x}_i^n > \lambda \mathbf{z}_i^t + \mathbf{z}_i^n$. In other words, we can assume that providers have scalar payoff functions given by $\lambda f_i^t(\cdot) + f_i^n(\cdot)$ for all $i \in \mathcal{N}$. We redefine the sets of feasible payoff profiles, $\hat{v}(\mathcal{S}), \mathcal{S} \subseteq \mathcal{N}$, according to the new payoff functions. We can then define the core, $\hat{\mathcal{C}}(\lambda)$, as follows.

$$\begin{aligned} \hat{\mathcal{C}}(\lambda) = \{(\mathbf{x}^t, \mathbf{x}^n) \in v(\mathcal{N}) : \forall \mathcal{S}, \nexists (\mathbf{z}^t, \mathbf{z}^n) \in v(\mathcal{S}) \text{ such that} \\ \lambda \mathbf{z}_i^t + \mathbf{z}_i^n > \lambda \mathbf{x}_i^t + \mathbf{x}_i^n, \ \forall i \in \mathcal{S}\}. \end{aligned} \tag{14}$$

It is straightforward to verify that the coalitional game $< \mathcal{N}, \hat{v} >$ is balanced. Here, Theorem IV.1 applies, and we conclude that $\hat{\mathcal{C}}(\lambda)$ is nonempty for all $\lambda > 0$..

We next claim that $\hat{\mathcal{C}}(\lambda) \subset \mathcal{C}_\epsilon$ if $\lambda = \epsilon / (\max_{i \in \mathcal{N}, (\mathbf{x}^t, \mathbf{x}^n) \in v(\mathcal{N})} \mathbf{x}_i^t)$. We prove this claim via contradiction. Consider a payoff profile $(\mathbf{x}^t, \mathbf{x}^n) \in \hat{\mathcal{C}}(\lambda)$. Suppose that $(\mathbf{x}^t, \mathbf{x}^n) \notin \mathcal{C}_\epsilon$. Then by (13), there exists a coalition $\mathcal{S}$ and a payoff profile $(\mathbf{z}^t, \mathbf{z}^n) \in v(\mathcal{S})$ such that either of the following holds for all $i \in \mathcal{S}$.

i) $\mathbf{z}_i^n > \mathbf{x}_i^n + \epsilon$
ii) $\mathbf{z}_i^n = \mathbf{x}_i^n + \epsilon$ and $\mathbf{z}_i^t > \mathbf{x}_i^t$

Since $(\mathbf{x}^t, \mathbf{x}^n) \in \hat{\mathcal{C}}(\lambda)$, it can not be blocked in linear ordering sense, by coalition $\mathcal{S}$. Thus, there exits an $i \in \mathcal{S}$ such that

$$\lambda \mathbf{z}_i^t + \mathbf{z}_i^n \leq \lambda \mathbf{x}_i^t + \mathbf{x}_i^n, \tag{15}$$
$$\text{or,} \quad \mathbf{z}_i^n - \mathbf{x}_i^n \leq \lambda(\mathbf{x}_i^t - \mathbf{z}_i^t),$$
$$\text{or,} \quad \mathbf{z}_i^n - \mathbf{x}_i^n \leq \epsilon. \tag{16}$$

The last inequality follows since $\lambda = \epsilon / (\max_{i \in \mathcal{N}, (\mathbf{x}^t, \mathbf{x}^n) \in v(\mathcal{N})} \mathbf{x}_i^t)$. Clearly, (15) implies that (ii) can not hold for i. On the other hand, (16) implies that (i) also can not hold for i. These are in contradiction with $(\mathbf{x}^t, \mathbf{x}^n) \notin \mathcal{C}_\epsilon$.

Thus, the claim and subsequently, the theorem follows. ■

Remark VI.1. *Note that Theorem VI.1 does not imply that $\mathcal{C}$ is nonempty. However, the approximate core can be made as close to the core as required, by selecting ϵ appropriately.*

We now discuss computing a payoff profile in the approximate core. As $\hat{\mathcal{C}}(\lambda) \subset \mathcal{C}_\epsilon$, we can obtain a payoff profile in $\mathcal{C}_\epsilon$ by finding one in $\hat{\mathcal{C}}(\lambda)$. Since $\hat{\mathcal{C}}(\lambda)$ is the core of a coalitional game with scalar payoff functions, a payoff profile in $\hat{\mathcal{C}}(\lambda)$ can be computed by the algorithm discussed in Section IV-C4.

C. Cooperation In Multi-hop Networks

In this section, we investigate cooperation among providers in multi-hop networks. Similar to the single-hop scenario, providers in multi-hop networks can cooperate by pooling their resources, such as service units and spectrum. In addition, they have the possibility to share their communication routes, which is a component specific to these types of networks. Having access to a larger set of routes, providers then have the potential to redirect their traffic through possibly better routes, which in turn can enhance their transmission power efficiency. In the following, we generalize the framework of Section III to accommodate cooperation in such networks. We then formulate cooperation as another coalitional game. Finally, we show that this game has a nonempty core and therefore, the grand coalition is stable .

Consider a network in which customers can communicate with service units via potentially multi-hop routes. In particular, customers can act as relays and carry packets of other customers to their destination or other relays, in return for service discounts, for example. In such a network, a provider has to determine the allocation of its service units to customers as well as the communication routes. If now a set of providers cooperate and pool their resources (which in this case include service units and customers/relays), besides gaining access to others' service units, they also enjoy a larger set of relays. This, in turn, can enhance customers' service rates and quality of service, and can also improve power efficiency in the network.

As in Section III, let $\mathcal{N}$ be the set of providers. Let $\mathcal{B}_i$ and $\mathcal{M}_i$ be the sets of provider i's service units and customers, respectively. As before, we consider downlink communications. We assume that each service unit (likewise, each customer) has access to a single channel (for transmission). In addition, we assume that no two service units in a vicinity have access to the same channel. We also assume that a pair of customers can communicate with each other (to relay packets) without interfering with the communications of other customer-customer or service unit-customer pairs (owing to appropriate channel allocation for example). Therefore, the necessary and sufficient condition for the simultaneous transmissions to be successful is that the set of transmitter-receiver pairs form a matching. Similar communication models have extensively been assumed in related contexts [22], [35].

A sufficient condition for a schedule to be feasible is that the fraction of time each service unit or customer communicates be below θ, where θ is a constant in $(0, 1]$ and depends on the network topology. For bipartite networks, for instance, $\theta = 1$, which is also a necessary condition [23]. It has been shown that in general, $\theta = \frac{2}{3}$ is a sufficient but not a necessary condition [23]. We assume that the network operates in a way that this condition always holds. This assumption can be motivated by the fact that operating the network at full capacity raises the delay which is not desirable.

Suppose now that a service unit or customer j can transmit to another customer k at a rate equal to r_{jk}, a random variable which is a function of the location of customer j and the state of channel k. Let Ω be the state space of the channels' states and customers' locations. We assume $|\Omega|$ is finite. Let ω be an outcome in this state space and $\mathbb{P}(\omega)$ be its probability.

A service unit and a customer, or two customers, can communicate only when both are associated with the same provider or the providers associated with them are in a coalition. Let random variable $\alpha_{jk} \in [0, 1]$ be the fraction of time, service unit or customer k transmits packets to customer j. α_{jk}s are determined by the allocation scheme.

Now consider a coalition $\mathcal{S}$. When the providers associated with customer j and service unit or customer k are in $\mathcal{S}$ and the network realization is ω, j receives a service rate $y^{\mathcal{S}}_{jk}(\omega) = \alpha_{jk}(\omega) r_{jk}(\omega)$ from k. For a provider $i \in \mathcal{S}$, define the rate vector $\mathbf{y}^{\mathcal{S}}_i(\omega) = (y^{\mathcal{S}}_{jk}(\omega), y^{\mathcal{S}}_{lj}(\omega), j \in \mathcal{M}_i, k \in \mathcal{B}_{\mathcal{S}} \cup \mathcal{M}_{\mathcal{S}}, l \in \mathcal{M}_{\mathcal{S}})$. In words, $\mathbf{y}^{\mathcal{S}}_i(\omega)$ is comprised of all the rates received by customers of provider i (either form service units or other customers) as well as the rates these customers deliver to other customers. Provider i receives a payoff equal to $f_i(\mathbf{y}^{\mathcal{S}}_i(\omega))$, which is the difference between the revenue i receives from its customers, and the costs (e.g., power consumption) it incurs by serving the customers in the coalition. Functions $f_i(\cdot)$s are assumed to be concave. The expected payoff provider $i \in \mathcal{S}$ earns will be $\sum_{\omega \in \Omega} \mathbb{P}(\omega) f_i(\mathbf{y}^{\mathcal{S}}_i(\omega))$.

For a coalition $\mathcal{S}$, a feasible joint action (allocation) $\{\alpha_{jk}(\omega), j \in \mathcal{M}_{\mathcal{S}}, k \in \mathcal{B}_{\mathcal{S}} \cup \mathcal{M}_{\mathcal{S}}, \omega \in \Omega\}$ should satisfy the following conditions.

1) $\sum_{j \in \mathcal{M}_{\mathcal{S}}} \alpha_{jk}(\omega) \leq \theta, \quad k \in \mathcal{B}_{\mathcal{S}}, \omega \in \Omega$
2) $\sum_{k \in \mathcal{B}_{\mathcal{S}} \cup \mathcal{M}_{\mathcal{S}}} \alpha_{jk}(\omega) + \sum_{l \in \mathcal{M}_{\mathcal{S}}} \alpha_{lj}(\omega) \leq \theta, \quad j \in \mathcal{M}_{\mathcal{S}}, \omega \in \Omega$
3) $\sum_{l \in \mathcal{M}_{\mathcal{S}}} \alpha_{lj}(\omega) r_{lj}(\omega) \leq \sum_{k \in \mathcal{B}_{\mathcal{S}} \cup \mathcal{M}_{\mathcal{S}}} \alpha_{jk}(\omega) r_{jk}(\omega), \quad j \in \mathcal{M}_{\mathcal{S}}, \omega \in \Omega$
4) $\alpha_{jk} \geq 0, \quad j \in \mathcal{M}_{\mathcal{S}}, k \in \mathcal{B}_{\mathcal{S}} \cup \mathcal{M}_{\mathcal{S}}, \omega \in \Omega$

Constraints (1) and (2) ensure that each node (service unit or customer) communicates for at most θ fraction of time. By constraints (3), each customer transmits packets at most at a rate less than or equal to that of receiving packets. Let $\mathcal{A}(\mathcal{S})$ denote the joint action space of coalition $\mathcal{S}$.

Consider a joint action $\alpha \in \mathcal{A}(\mathcal{S})$ and let $\mathcal{F}^{\mathcal{S}}(\alpha)$ be the resulting payoff vector. We can now define the sets of feasible payoff profiles $(v(\mathcal{S}), \mathcal{S} \subseteq \mathcal{N})$ and the core of the game $\mathcal{C}$, similar to (1) and (2). Then our NTU game $< \mathcal{N}, v >$ is well defined. Using Theorem IV.1 for balanced games, we can prove the following.

Theorem VI.2. *The nontransferable payoff game $< \mathcal{N}, v >$ is balanced, and hence has a nonempty core.*

Proof: The proof is similar to those of Theorems IV.2 and V.1 and is omitted. ∎

VII. Quantitative Evaluations

We evaluate the benefits of cooperation in context of the provider coalitional game (Section IV) as well as provider-customer game (Section V). We consider logarithmic and linear payoff functions $f(\cdot)$ ($g(\cdot)$) for providers (customers, resp.). Such functions are concave, increasing and assume non-negative values, and have been widely used as satisfaction functions of customers and therefore constitute good candidates for the revenues they pay (and hence for the payoffs the providers obtain) [21], [22]. We allow the rates r_{jk} to be uniformly distributed over the set $\{0, 100, 200\}$ Kbps, and to be independent across service-unit-customer pairs j, k. *In each* *figure, the legends appear in the same order as the plots (e.g., x_5 is the topmost plot in the left sub-plot of Figure 2).*

A. Provider Coalitional Game

We have mostly considered systems where providers do not share their payoffs (except Section VI-A). Example VI.1 reveals that the lack of payoff sharing may be inefficient as payoffs of each provider (and hence the aggregate) may increase in presence of sharing. But, again, such lack of payoff sharing may be enforced due to mutual distrust among the providers. We now examine whether such inefficiencies are pathological or pervasive. Note that the maximum aggregate payoff and the corresponding service unit-customer allocations (referred to as the *optimal allocation* henceforth) can be computed by solving a concave maximization with the linear feasibility constraints (1), (2), (3) (see Section IV-A). The question then is does the core of the non-transferable utility coalition game (i.e., where payoffs can not be shared) usually have at least one payoff profile that maximizes the aggregate payoffs of the providers - more specifically does the payoff profile corresponding to the optimal allocation belong in the core? Also, what are the payoff gains due to cooperation for each provider for such a payoff profile (if it is in the core)? We seek to answer these questions considering logarithmic payoff functions for the providers: $f_i^{\mathcal{S}}(\mathbf{y}_i) = \sum_{j\in\mathcal{M}_i} \log\left(1 + (\sum_{k\in\mathcal{B}_\mathcal{S}} y_{jk})\right)$. Note that $\log\left(1 + (\sum_{k\in\mathcal{B}_\mathcal{S}} y_{jk})\right)$ is the payoff (revenue) provider i earns from its customer j when in coalition $\mathcal{S}$.

We first assume a symmetric scenario with n providers, each provider having one base station and $5k$ customers. We consider $n = 1, 2, 3, 5$ and vary k from 1 to 20. Owing to the symmetry, under the optimal allocation, all providers get equal payoffs (x_n for n providers) and percentage payoff gains (as compared to when each operates individually, i.e., x_1), which we plot as a function of k for different n in Figure 2. In each case, the payoff profile corresponding to the optimal allocation belongs in the core[12]. Thus, the lack of sharing of payoffs does not introduce any inefficiency in these cases. The payoff gains are significant (in the range of $20\% - 40\%$ for each provider), and both payoffs and payoff gains increase with increase in n, k. Thus cooperation becomes more beneficial with increase in the size of the grand coalition, and the number of customers (demand) of each provider.

Now we consider an asymmetric scenario with $N = 3$, $B_1 = B_2 = B_3 = 1$ and $M_1 = 3k, M_2 = 4k, M_3 = 5k$ where k ranges from 1 to 20. (Note that Example VI.1 demonstrated the inefficiency due to lack of sharing using an asymmetric setting). Now, due to asymmetry, the payoff profile resulting from the optimal allocation provides different payoffs to different providers (Figure 3). The third sub-figure shows that $\sum_{i\in\mathcal{S}} x_i^* \geq \bar{v}(\mathcal{S})$ for all $\mathcal{S} \subset \mathcal{N}$, for all the chosen parameter values. Thus, in each case, this payoff profile is in the core and thus the lack of sharing did not introduce any inefficiency in any of these cases. Note that providers' payoffs and payoff gains are in increasing order of their number of customers.

B. Provider Customer Coalitional Game

Example V.3 demonstrated that cooperation in provider customer games, can reduce the payoffs of some of the providers and customers. The question then is how pervasive this phenomenon is, and also whether cooperation in general improves the aggregate payoffs of providers (customers, resp.)?

we seek to answer the above considering logarithmic profit functions as before: $f_{ij}(\mathbf{y}_{ij}) = \log\left(1 + (\sum_{k\in\mathcal{B}_i} y_{jk})\right)$ and $g_j(\mathbf{y}_j) = \log\left(1 + (\sum_{k\in\mathcal{B}_\mathcal{S}} y_{jk})\right)$. We consider a provider customer game in the symmetric scenario described in the second paragraph of Section VII-A. The service unit-customer allocation that maximizes the aggregate payoffs of all participants (providers' and customers') is referred to as socially-optimal. The payoff profile for providers and customers corresponding to the socially optimal allocation have been plotted in Figure 4. Owing to symmetry, all providers (customers, resp.) receive equal payoffs - in the figures, x_n (y_n, resp.) is the payoff of a provider (customer, resp.) when it is in a coalition with n providers and $5 * n * k$ customers. In each case, the payoff profile turns out to be in the core[13]. Note that for any given k, x_n, y_n increase with n - as the size of the coalition increase, each participant's payoff increases. Thus, cooperation becomes more beneficial as both the resources (service units) and the demands (customers) increase. Customers' payoffs decrease with increase in k for any given n as each customer needs to contend with more customers for sharing the same amount of resource. The trend is the opposite for providers' payoffs as demand increases with increase in k. The aggregate providers' payoffs increase with increase in n, k. The aggregate customers' payoffs (the third sub-figure in Figure 4), and therefore the overall (i.e., over all participants) aggregate payoffs also increase with increase in n, k.

Next, we consider an asymmetric setting and linear payoff functions $f_{ij}(\mathbf{y}_{ij}) = \sum_{k\in\mathcal{B}_i} y_{jk}$ and $g_j(\mathbf{y}_j) = \sum_{k\in\mathcal{B}_\mathcal{S}} y_{jk}$ (as in Example V.3). Also, following Example V.3, we consider deterministic network realizations, where each realization is obtained as per the service unit-customer rate distributions considered so far. A payoff profile in the core is chosen as

[12]To verify whether a payoff profile $(x_1^*, x_2^*, \ldots, x_n^*)$ is in the core, it suffices to check that $\sum_{i\in\mathcal{S}} x_i^* \geq \bar{v}(\mathcal{S})$ for all $\mathcal{S} \subset \mathcal{N}$, where $\bar{v}(\mathcal{S})$ is the maximum aggregate payoff of providers in $\mathcal{S}$. In the symmetric case we check the above for $x_i^* = x_n$ for all i. Since $\bar{v}(\mathcal{S})$ equals $|\mathcal{S}|x_{|\mathcal{S}|}$, we only need to ensure that $x_n \geq x_{n-1} \geq \ldots \geq x_1$, which holds as per Figure 2.

[13]Let us assume that the grand coalition has n providers and $5 * n * k$ customers (n, k fixed). Let $(\mathbf{x}^*, \mathbf{y}^*)$ be the payoff profile corresponding to the socially optimal allocation. Now, Consider a generic coalition $(\mathcal{S}, \mathcal{T})$ with $n' \leq n$ providers and $m' \leq m$ customers ($(n', m') \neq (n, 5 * n * k)$). There are two possibilities.

1) $m' \leq 5 * m' * k$: The first sub-figure in Figure 4 reveals that the aggregate subplot of all the n' providers is less in $(\mathcal{S}, \mathcal{T})$ than in the grand coalition. Thus no feasible payoff profile can make all of them happier, and so $(\mathcal{S}, \mathcal{T})$ can not block $(\mathbf{x}^*, \mathbf{y}^*)$.
2) $m' > 5 * m' * k$: The second sub-figure in Figure 4 reveals that not all m' customers can obtain better payoffs in $(\mathcal{S}, \mathcal{T})$. Thus, again, $(\mathcal{S}, \mathcal{T})$ can not block $(\mathbf{x}^*, \mathbf{y}^*)$.

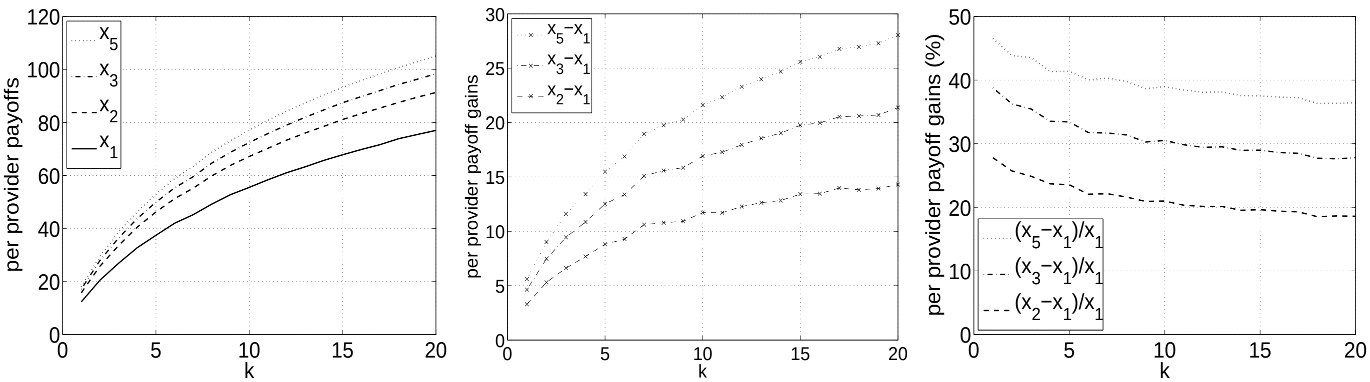

Fig. 2: The left, middle and right sub-plots respectively show a provider's payoff, payoff gains and percentage payoff gains as functions of the number of customers. The payoff of a provider, when n of them cooperate, is denoted by x_n.

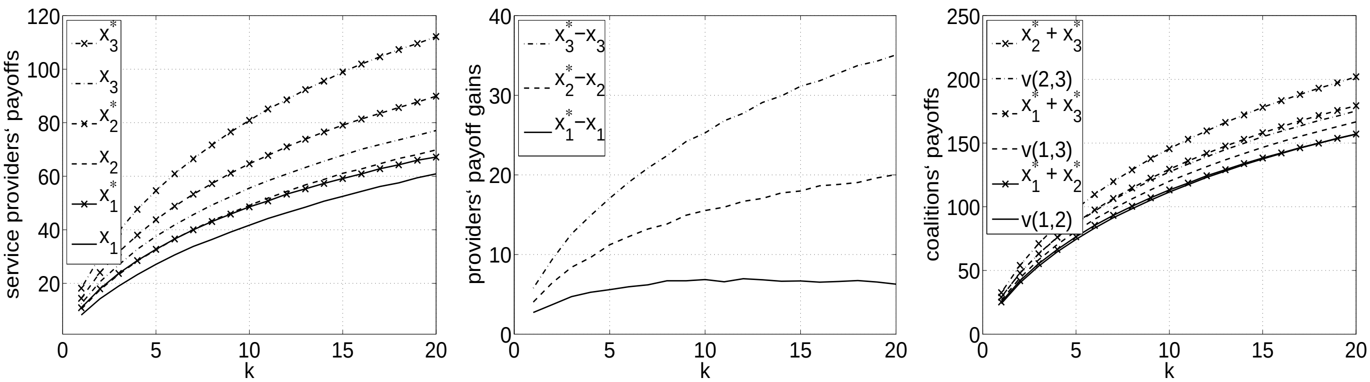

Fig. 3: The left and middle sub-plots respectively show providers' payoffs and payoff gains as functions of the number of customers: the three providers have $3k, 4k$ and $5k$ customers, respectively. The payoff of provider i is x_i when it is operating alone, and x_i^* in the grand coalition. The last plot shows aggregate payoffs under the globally optimal allocation and the maximum aggregate payoffs for coalitions $\mathcal{S} \in \{\{1,2\},\{1,3\},\{2,3\}\}$.

per **ALG** (Section V-B) in each realization and the payoffs are averaged over all the realizations. Here, $N = 3, B_1 = B_2 = B_3 = 3, M = 6k$, and we vary k from 1 to 10. We consider the payoffs of providers and customers, (i) when all participants are in the grand coalition (grand-coalition payoffs), and (ii) when each provider i operates individually with M_i subscribed customers (no-cooperation payoff) (Figure 5). Here, $M_1 = k, M_2 = 2k, M_3 = 3k$. For ease of reference, we denote the customers in $\mathcal{M}_i$ as those of provider i (though in the grand coalition they can seek service from and pay to any provider).

Owing to symmetry, providers (and customers) attain equal grand-coalition payoff; no-cooperation payoffs are however higher for providers with larger number of customers and lower for customers with larger number of contenders (i.e., larger k) (Figure 5). The grand-coalition payoff of a provider however exceeds its no-cooperation payoff, suggesting that by and large providers enhance their payoffs through cooperation. As k increases, each provider's (both grand-coalition and no-cooperation) payoff initially increases rapidly but subsequently saturates (once the number of customers becomes large enough to allow the utilization of the service units at highest rates most of the time).

Note that customers of provider 1 (2, 3, resp.), attain mostly lower (higher, resp.) grand-coalition payoff than no-cooperation payoff. This is because the aggregate resource (number of service units) to demand (number of customers) is higher (equal, lower, resp.) for customers of provider 1 (2, 3, resp.) in the no cooperation case. Customers of provider 2 attain slightly higher payoff in grand coalition owing to rate diversity (in the grand coalition they find service units with high rates even when their rates from provider 2's service units are low). Thus, cooperation may reduce the rates of certain classes of customers (which they can not circumvent since they can not persuade providers to leave the grand coalition as they are gaining from cooperation).

Finally, the choice of the payoff functions ensures that the aggregate of providers payoffs equals that of the customers payoffs for any service-unit-customer allocation. Figure 5 reveals that under **ALG** the aggregate grand coalition payoff of customers (and hence of providers) is (i) identical to the maximum aggregate grand-coalition payoff, and (ii) exceeds their aggregate no-cooperation payoff. Thus, **ALG** is efficient in the above sense, and ensures that the customers are overall

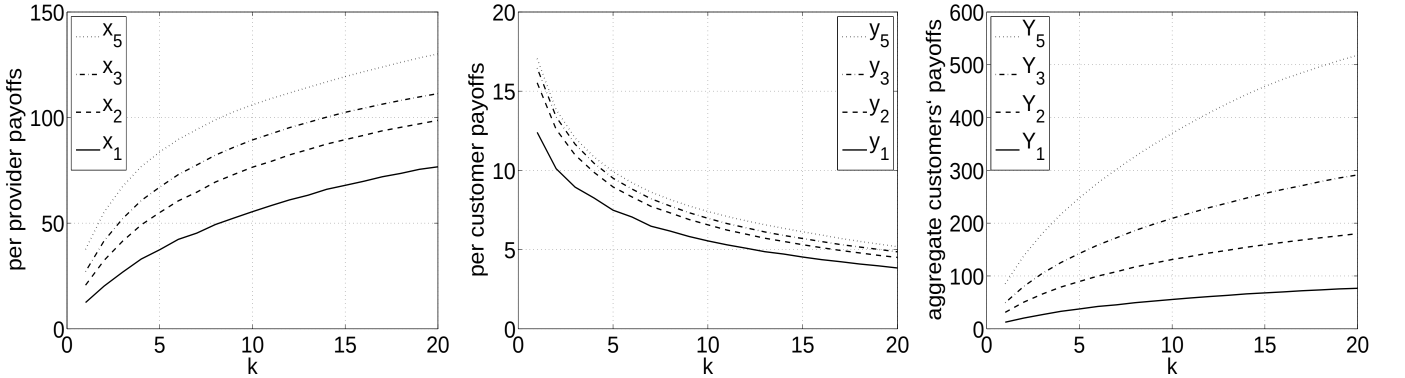

Fig. 4: The left and middle sub-plots respectively show a provider's and a customer's payoffs. The payoff of a provider (customer), when the coalition consists of n providers and $5 * n * k$ customers, is denoted by x_n (y_n). The right sub-plots show aggregate customers' payoffs $Y_n = 5 * n * k * y_n$.

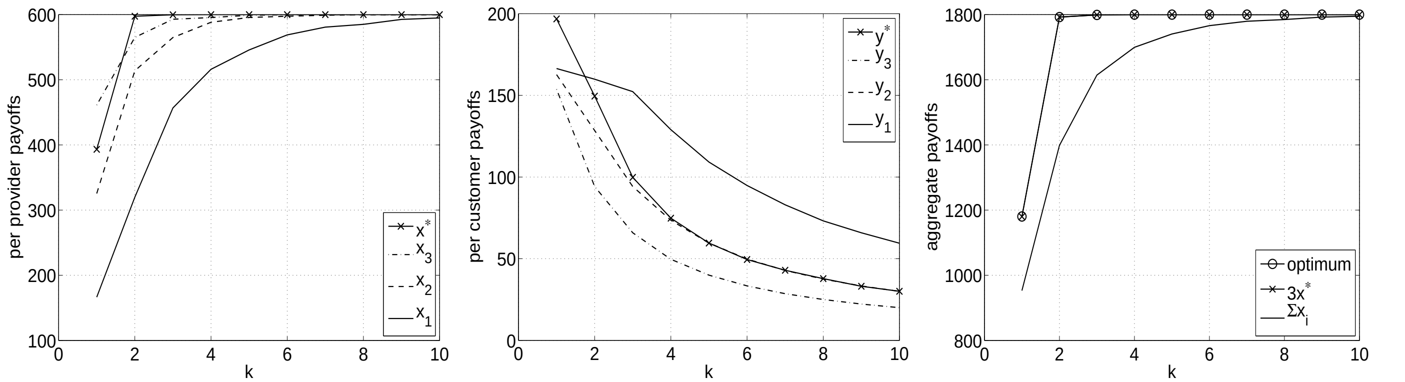

Fig. 5: The left and middle sub-plots respectively show providers' and customers' payoffs as functions of the number of customers: the three providers have $k, 2k$ and $3k$ customers respectively. The payoff of provider i (customers i) is x_i (y_i) in the absence of cooperation. Each provider (customer) gets x^* (y^*) in the grand coalition. The right sub-plot shows aggregate payoffs of customers (and hence of providers) with and without cooperation, and also their maximum aggregate payoffs.

better off after cooperation.

VIII. Conclusion and Future work

We formulated interaction among cooperating service providers and customers in wireless networks as nontransferable utility coalitional games. We showed nonemptyness of cores in various scenarios (see Theorems IV.2, V.1 etc.) implying that cooperation is globally desirable. We used the concept of market equilibrium to obtain a payoff profile in the core when customer subscriptions are fixed (Section IV-C2). When customers strategically select their subscriptions, we showed how to compute a payoff profile in the core using distributed polynomial time computations for linear payoff functions and deterministic network states; we presented a technique to obtain a payoff profile in the approximate core when the above assumptions must be relaxed (Section V-B). Many of our algorithms do not require participants to reveal their confidential information (such as payoff functions) to each other. Our numerical investigations indicate that usually cooperation enhances providers' payoffs, aggregate providers' (and customers') payoffs substantially, whereas individual customer's payoffs may either increase or decrease (mostly increases) depending on the resources and demands of the coalitions and the nature of the payoff functions (Section VII).

Computing a payoff profile in the core using distributed polynomial time solutions in systems with arbitrary payoff functions and random state evolution remains open. Next, in practice, coalition formation can incur overheads, e.g., it can lead to increased loads on the call processors and billing systems. Finally, we have assumed that participants are rational in that they do not separate from coalitions unless they can improve their payoffs by separating. But, a participant may adopt a more detrimental attitude towards others, in that, if its payoff is significantly less than that of the others it would not cooperate even at the cost of reducing its payoff by not doing so. Investigating the stability of the grand coalition considering the coalition formation overheads and the payoff parity objectives constitute interesting open problems.

References

[1] M. Osborne and A. Rubinstein, *A Course in Game Theory*. The MIT press, 1999.

[2] E. Altman, A. Kumar, C. Singh, and R. Sundaresan, "Spatial sinr games combining base station. placement and mobile association," in *Proc. of IEEE INFOCOM*, (Rio de Janeiro, Brazil), April 2009.
[3] T. Alpçan, T. Başar, R. Srikant, and E. Altman, "CDMA uplink power control as a noncooperative game," *Wireless Networks*, vol. 8, pp. 659–670, 2002.
[4] E. Hossein, D. Niyato, and Z. Han, *Dynamic Spectrum Access and Management in Cognitive Radio Networks*. Cambridge, UK: Cambridge University Press, 2009.
[5] Y. Song, C. Zhang, and Y. Fang, "Joint channel and power allocation in wireless mesh networks: A game theoretical perspective," *IEEE J. Select. Areas Commun*, vol. 26, pp. 1149–1159, September 2008.
[6] T. Himsoon, W. P. Siriwongpairat, Z. Han, and K. J. R. Liu, "Lifetime maximization via cooperative nodes and relay deployment in wireless networks," *IEEE J. Select. Areas Commun*, vol. 25, pp. 306–317, February 2007.
[7] L. Grokop and D. N. Tse, "Spectrum sharing between wireless networks," in *Proc. of IEEE INFOCOM*, (Phoenix, USA), pp. 735–743, April 2008.
[8] W. Saad, Z. Han, M. Debbah, A. Hjorungnes, and T. Basar, "A game-based self-organizing uplink tree for voip services in ieee 802.16j networks," in *Proc. of IEEE ICC*, (Dresden, Germany), June 2009.
[9] H. Kameda and E. Altman, "Inefficient noncooperation in networking games of common-pool resources," *IEEE J. Select. Areas Commun*, vol. 26, pp. 1260–1268, September 2008.
[10] D. Niyato and E. Hossain, "A cooperative game framework for bandwidth allocation in 4g heterogeneous wireless networks," in *Proc. of IEEE ICC*, (Istanbul, Turkey), June 2006.
[11] J. E. Suris, L. A. DaSilva, Z. Han, and A. B. MacKenzie, "Cooperative game theory for distributed spectrum sharing," in *Proc. of IEEE ICC*, (Glasgow, Scotland), June 2007.
[12] Z. Han and H. V. Poor, "Coalition games with cooperative transmission: A cure for the curse of boundary nodes in selfish packet-forwarding wireless networks," *IEEE Trans. Comm.*, January 2009.
[13] S. Mathur, L. Sankaranarayanan, and N. Mandayam, "Coalitional games in cooperative radio networks," *Fortieth Asilomar Conference on Signals, Systems and Computers*, pp. 1927–1931, Oct.-Nov. 2006.
[14] R. La and V. Anantharam, "A game theoretic look at the gaussian multiaccess channel," in *Proc. of the DIMACS Workshop on Network Information Theory*, (New Jersey, NY, USA), Mar. 2003.
[15] S. Mathur, L. Sankaranarayanan, and N. Mandayam, "Coalitions in cooperative wireless network," *IEEE J. Select. Areas Commun.*, vol. vol. 26, pp. 1104–1115, Sep. 2008.
[16] W. Saad, Z. Han, M. Debbah, A. Hjorungnes, and T. Basar, "Coalitional games for distributed collaborative spectrum sensing in cognitive radio networks," in *Proc. of IEEE INFOCOM*, (Rio de Janeiro, Brazil), April 2009.
[17] W. Saad, Z. Han, M. Debbah, A. Hjorungnes, and T. Basar, "Coalitional game theory for communication networks: A tutorial," *IEEE Signal Processing Magazine*, vol. 26, pp. 77–97, September 2009.
[18] A. Aram, C. Singh, S. Sarkar, and A. Kumar, "Cooperative profit sharing in coalition based resource allocation in wireless networks," in *Proc. of IEEE INFOCOM*, (Rio de Janeiro, Brazil), Apr. 2009.
[19] A. Kumar, D. Manjunath, and J. Kuri, *Wireless Networks*. Morgan Kaufmann Publishers, 2008.
[20] L. Jiang, S. Parekh, and J. Walrand, "Base station association game in multi-cell wireless networks," in *Proceedings of IEEE WCNC*, (Las Vegas, NV), pp. 1616–1621, 2008.
[21] F. Kelly, A. Maulloo, and D. Tan, "Rate control in communication networks: shadow prices, proportional fairness and stab," *Journal of the Operational Research Society*, vol. 49, pp. 237–252, 1998.
[22] X. Lin and N. Shroff, "The impact of imperfect scheduling on cross-layer rate control in multihop wireless networks," in *Proceedings of INFOCOM*, (Miami, Fl), March 2005.
[23] B. Hajek and G. Sasaki, "Link scheduling in polynomial time," *IEEE Transactions on Information Theory*, vol. 34, no. 5, 1988.
[24] V. Conitzer and T. Sandholm, "Complexity of determining nonemptiness of the core," Technical Reprt CS-02-137, CMU, 2002.
[25] H. E. Scarf, "The core of an n person game," *Econometrica*, vol. 35, pp. 50–69, January 1967.
[26] A. Mas-colell, M. D. Whinston, and J. R. Green, *Microeconomic Theory*. Oxford University Press, June 1995.
[27] S. Smale, "A convergent process of price adjustment," *Journal of Mathematical Economics*, vol. 3, pp. 107–120, 1976.
[28] R. A. Horn and C. R. Johnson, *Matrix Analysis*. Cambridge, UK: Cambridge University Press, 1990.
[29] M. Ferris and J. Pang, "Engineering and economic applications of complementarity problems," *SIAM Review*, vol. 39, pp. 669–713, 1997.
[30] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, UK: Cambridge University Press, 2004.
[31] I. Heller and C. B. Tompkins, *Linear Inequalities and Related Systems*, vol. 38 of *Annals of Mathematics Studies*, ch. An extension of a theorem of dantzigs, pp. 247–254. Princeton University Press, 1956.
[32] A. J. Hoffman and J. B. Kruskal, *Linear Inequalities and Related Systems*, vol. 38 of *Annals of Mathematics Studies*, ch. Integral boundary points of convex polyhedra, pp. 223–246. Princeton University Press, 1956.
[33] A. Israeli and A. Itai, "A fast and simple randomized parallel algorithm for maximal matching," *Information Processing Letters*, vol. 22, pp. 77–80, February 1986.
[34] M. Quinzii, "Core and competitive equilibria with indivisibilities," *International Journal of Game Theory*, vol. 13, no. 1, pp. 41–60, 1984.
[35] X. Wu and R. Srikant, "Regulated maximal matching: a distributed scheduling algorithm for multihop wireless networks with nodeexclusive spectrum sharing," in *Proceedings of IEEE CDC-ECC05*, (Seville, Spain), Dec 2005.