Economics of Investment and Use of Shared Network Infrastructures

Iordanis Koutsopoulos Angeliki Anastopoulou Merkourios Karaliopoulos
Department of Informatics
Athens University of Economics and Business
Athens, Greece

***Abstract*—We study the interaction of a set of self-interested, budget-limited Service Providers (SPs), each of which invests a portion of its budget towards building a common network infrastructure, which is then shared among SPs for offering services. SPs pay an infrastructure Operator (IO) for leasing resource slices and then obtain revenue from services. The problem faced by each SP is how to split their budget into a portion invested in the infrastructure, and a portion for using it. The IO may control the allocation of resource slices and the shares of total revenue to each SP. The sum of SP investments determines the total available resource slices, and the portion of slices allocated to an SP for providing services affects SP revenue. In shared infrastructures, an important metric is the total attained revenue of SPs, as it is linked to sustainability of the IO and the infrastructure, and quality of services. We show that if the IO allocates to each SP a share of the total revenue equal to its Shapley value, then the selfish investment policy that maximizes revenue share of each SP, maximizes total revenue as well, and the global optimal policy is a Nash equilibrium of the investment game. If revenue sharing is not possible, the IO may allocate resource slices so as to maximize total revenue in the presence of SP strategic investments, and thus minimize price of anarchy. Numerical results show the benefits of our approach in attained total revenue and price of anarchy.**

I. Introduction

Shared crowdsourced networks are formed based on crowdsourcing, which dictates that different stakeholders or users contribute infrastructure, monetary or network resources towards building a shared infrastructure which is used collectively by them. There are two classes of crowdsourced networks. In the first one, access point (AP) owners make their APs available for connection to all users in exchange for payment or extended coverage through APs of others. A network operator regulates pricing schemes, user incentive payments and subscription modes. FON [1] is the world's largest crowdsourced network, currently counting more than 17 million members.

The second and perhaps larger class of shared networks includes the so-called *wireless community networks*, in which the network equipment is contributed by end-users and professional actors such as service providers (SPs) [2]. The latter wish to build a shared infrastructure and deploy services, they are interested in revenue out of services, and they invest money in building an infrastructure so that services with good quality are deployed. Services are vital for SP revenue and ultimately for network sustainability. An infrastructure operator (IO) regulates the interaction of stakeholders and builds the infrastructure. Guifi [3] is a community network in Spain with over $33,000$ nodes, both wireless and wireline (fiber) infrastructure, several professional actors and a public-administration IO.

We target the second class above. A network infrastructure is built out of monetary contributions of some budget-limited SPs. Each SP invests some money to a common pool, and an IO builds a network infrastructure out of these contributions. The infrastructure is used in a shared fashion among SPs. Namely, each SP pays the IO for leasing network resource slices in order to provide services to its clients. Slices are leased for a time duration according to the budget spent. The service is provided for the same duration and brings revenue based on its quality, which depends on the amount of allocated slices to the SP.

We study the interaction of a set of budget-limited, self-interested SPs, each of which aims to best utilize its own budget and maximize its revenue. Each SP invests a portion of its budget towards building a common network infrastructure, which is then shared among SPs for services. SPs pay an infrastructure Operator (IO) the remaining portion of their budget for leasing resource slices and obtain revenue from services. The SP problem is how to split their budget into a portion invested in the infrastructure and a portion for using it. SPs compete with each other, and the investment policy of each SP affects all SPs: the sum of SP investments determines the available resource slices, and the portion of these slices allocated to the SP for providing services affects SP revenue.

In shared infrastructures, a metric to maximize is total SP revenue, since it implies network sustainability, high profit for the IO, and good quality of service to users. The IO may control the allocation of resource slices to each SP and/or the shares of the total accrued revenue to each SP. We ask the following questions: (*i*) Given that SPs invest so as to maximize their own revenue, can the IO induce an SP investment policy such that the total revenue of SPs under the Nash equilibrium strategy is the same as the maximum total revenue obtained under the global optimal investment policy (which would require SP cooperation), or at least as close as possible to that? (*ii*) Can the IO achieve this goal through a revenue sharing or slice

I. Koutsopoulos acknowledges the support from the AUEB internal project "Original scientific publications". The research work of M. Karaliopoulos received funding from the Hellenic Foundation for Research and Innovation (HFRI) and the Hellenic General Secretariat for Research and Technology (GSRT), under grant agreement No 892.

allocation policy?

We show that the answers are affirmative. Thus, SPs are incentivized and encouraged to invest in the infrastructure, and although each of them aims at maximizing its own revenue, the total incurred revenue is maximum, and the shared network is best operated. Our contributions are as follows:

- We devise a model for investment and use of a shared network by SPs and derive an expression for revenue.
- We show that if the IO allocates to each SP a share of the total revenue equal to its Shapley value, then the SP selfish investment strategy that maximizes its revenue share, maximizes total revenue as well, and the global optimal strategy is a Nash equilibrium of the investment game, i.e. the Nash equilibrium is efficient.
- If revenue sharing is not possible, we show how the IO may allocate resource slices to SPs so as to maximize the total revenue in the presence of SP strategic investments, and thus minimize the price of anarchy.

II. Model

A set $\mathcal{S}$ of n budget-limited Service Providers (SPs) are interested in investing money towards building a common network infrastructure that will be used by them in a shared fashion to provide services to their customers. Each SP i has a unit budget and may invest a portion x_i of budget towards contributing to building the infrastructure, and a portion $(1-x_i)$ for using the infrastructure and providing services.

An infrastructure operator (IO) entity is responsible for: (*i*) collecting SP investments and building the infrastructure out of them, (*ii*) allocating network resource slices to SPs for their services, (*iii*) charging SPs for the use of the infrastructure, and (*iv*) sharing revenue among SPs (Fig. 1).

Given an SP investment vector $\mathbf{x} = (x_1, \ldots, x_n)$, the IO builds an infrastructure. This can be abstracted as a set of resources such base stations/small cells or routers for network access, network links of certain bandwidth for network traffic transport, server clusters of certain computational capacity for computational tasks, or caches for caching content. In a most general form, the infrastructure is a network graph with nodes (e.g. routers, base stations) of some computational and/or cache capacity, and edges (e.g. optical fiber ones) of some bandwidth.

In this work, we abstract the infrastructure as just one type of resource. Let $b(\cdot)$ be a continuous function that maps the investment vector $\mathbf{x}$ into an amount of that resource. Thus, for computational resources, $b(\mathbf{x}) = \frac{1}{c}\sum_{i=1}^n x_i$ would denote the amount of computational capacity available as a result of total investment $\sum_{i=1}^n x_i$, where c is the cost per unit of computational capacity. Function $b(\cdot)$ is assumed to be known to SPs. Thus, SPs know the collective impact of their investments on the network infrastructure and network resources. We assume unit cost, $c = 1$, thus $b(\mathbf{x}) = \sum_{i=1}^n x_i$.

A. Slice allocation

After the infrastructure is built, the IO allocates resource slices to SPs to deploy and run their services. Without loss of generality, we assume that each SP i deploys one service, and all slices have the same resource requirements. Thus, the total number of slices created after SP investments is $\frac{b(\mathbf{x})}{C}$, where C is the amount of resource requirements per slice.

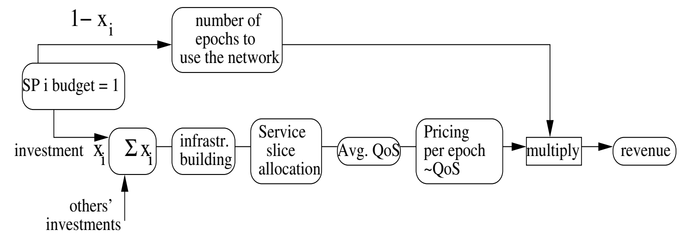

Fig. 1. An SP i invests a portion x_i of its budget towards building the infrastructure, and the rest, $1-x_i$, for using the infrastructure in a shared fashion for service provisioning. Infrastructure building, service slice allocation and pricing are performed by the IO.

The IO has a rule $\mathbf{s} = (s_1, \ldots, s_n)$ for allocating service slices to SPs, where s_i is the portion of the total number of slices assigned to SP i, with $\sum_{i=1}^n s_i = 1$. Thus, SP i gets

$$k_i(\mathbf{x}) = \frac{s_i}{C}\sum_{i=1}^n x_i \tag{1}$$

slices. For instance, slices could be split equally among SPs, i.e. $s_i = \frac{1}{n}$ for all i; or a proportional allocation rule could be applied, where the number of allocated slices to each SP i is proportional to its demand i.e. $s_i = d_i / (\sum_{j=1}^n d_j)$.

B. Use of the network by SPs

Each SP i uses the non-invested portion $(1-x_i)$ of its budget to pay the IO for using the infrastructure so as to offer services to its customers. Our payment model is inspired from the realistic situation where slices are leased for some amount of time e.g. months, for some fee f per unit of time, called an *epoch*. The number of epochs for which slices are leased is $(1-x_i)/f$. We assume a unit fee, $f = 1$ per epoch for each SP, thus SP i will lease slices for $(1-x_i)$ epochs.

Let d_i be the per-epoch anticipated user demand for the service of SP i. SP i can provide service with a given average quality level. We assume that the average quality of service (QoS) level provisioned by SP i is uniform across users according to a continuous function $q_i(\cdot)$ which depends on the amount of allocated slices $k_i(\mathbf{x})$ (and thus on investment vector $\mathbf{x}$), and on d_i. This QoS level translates to a money flow for the SP through a service charging model per epoch and per unit of demand, which includes a pricing function $p_i(q_i(\cdot))$. The total revenue $u_i(\cdot)$ out of total demand over all epochs is

$$u_i(\mathbf{x}) = (1-x_i) d_i p_i(q_i(\mathbf{x}))\,. \tag{2}$$

Proposition 1. *Function $u_i(\cdot)$ in (2) is a concave function of x_i if: (i) $q_i(\cdot)$ is decreasing and convex function of x_i, and $p_i(\cdot)$ is decreasing, and convex or linear function of $q_i(\cdot)$; or if (ii) $q_i(\cdot)$ is increasing and concave function of x_i, and $p_i(\cdot)$ is increasing, and concave or linear function of $q_i(\cdot)$.*

Proof. The second partial derivative of $u_i(\mathbf{x})$ with respect to x_i is

$$\frac{\partial^2 u_i(\mathbf{x})}{\partial x_i^2} = -2p_i'(q_i(\mathbf{x}))\frac{\partial q_i(\mathbf{x})}{\partial x_i} +$$

$$+(1-x_i)\left[p_i''(q_i(\mathbf{x}))\left(\frac{\partial q_i(\mathbf{x})}{\partial x_i}\right)^2 + p_i'(q_i(\mathbf{x}))\frac{\partial^2 q_i(\mathbf{x})}{\partial x_i^2}\right].$$

Using basic calculus about the sign of the first derivative (which is positive for increasing, and negative for decreasing functions) and the sign of the second derivative (which is positive for convex, negative for concave, and zero for linear functions), we get that $\frac{\partial^2 u_i(\mathbf{x})}{\partial x_i^2} < 0$ for the cases (*i*) or (*ii*) above. Thus, $u_i(\cdot)$ is a concave function of x_i. □

C. Examples

1) Example 1: Let s_i be the portion of computational capacity slices allocated to SP i, and $k_i(\mathbf{x})$ be the total computational capacity of SP i (in cycles/sec), given by (1). If the stream of service requests for SP i obeys the Poisson distribution with rate d_i requests per unit of time, and the request size is exponentially distributed (e.g., with unit mean), a QoS metric is the *average delay per request* for the M/M/1 service queue above,

$$q_i(\mathbf{x}) = \frac{1}{k_i(\mathbf{x}) - d_i}\,. \tag{3}$$

A pricing function $p_i(\cdot)$ charges customers according to average delay. This can be a convex decreasing function of q_i, e.g. a piecewise-linear convex function, consisting of line segments of different slopes, each of which denotes the rates of decrease of price per unit of delay increase. In its simplest form, this piece-wise linear convex function has two segments,

$$p_i(q_i(\mathbf{x})) = \begin{cases} a_i - D_i q_i(\mathbf{x}), & \text{if } q_i(\mathbf{x}) \le a_i/D_i \\ 0, & \text{otherwise,} \end{cases} \tag{4}$$

where a_i is a reference price for zero delay, and a_i/D_i is the maximum delay for which there is nonzero charge. If the epoch duration is τ, the average revenue per epoch for all demand is $d_i \tau p_i(q_i(\mathbf{x}))$ and the total average revenue over all epochs is:

$$u_i(\mathbf{x}) = (1-x_i) d_i \tau p_i(q_i(\mathbf{x}))\,. \tag{5}$$

2) Example 2: Let s_i be the portion of a link's bandwidth slices allocated to SP i, and let $k_i(\mathbf{x})$ be the total bandwidth of SP i (in bits/sec). Let QoS be a concave function of transmission rate, e.g. $q_i(\mathbf{x}) = \gamma_i \log(1 + k_i(\mathbf{x}))$, where $\gamma_i > 0$ is a multiplicative factor. Pricing can be proportional to $q_i(\cdot)$, i.e. $p_i(q_i(\mathbf{x})) = \beta_i q_i(\mathbf{x})$, and total revenue is $u_i(\mathbf{x}) = (1 - x_i) d_i \beta_i q_i(\mathbf{x})$, where demand d_i denotes the number of traffic sessions per epoch that are transported through the link.

D. Revenue sharing and the Shapley value

The IO may allocate the total incurred revenue among SPs, after possibly withholding a certain percentage as a commission. Revenue allocation needs to be done so that investments are encouraged. *Shapley value* [4] is an appropriate mechanism for allocating the worth of a coalition among its participants, where the worth of a coalition is the total maximum revenue obtained by its members through cooperation. In our case, the worth of a coalition $\mathcal{S}$ of SPs is

$$v(\mathcal{S}) = \max_{\mathbf{x}\ge \mathbf{0}} \sum_{i\in\mathcal{S}} u_i(\mathbf{x}), \text{ such that } u_i(\mathbf{x}) \ge 1 \text{ for all SPs } i\,. \tag{6}$$

The constraint above is a participation incentive for each SP, since the resulting revenue for each SP after investing and using the network should be at least as much as its initial budget.

Among the class of revenue distribution mechanisms, the Shapley-value one is the only one that satisfies desirable properties such as fairness, efficiency, symmetry and strong monotonicity [5], [6]. The Shapley value for SP i is

$$\phi_i^v(\mathcal{S}) = \frac{1}{n!}\sum_{\pi\in\Pi} \Delta_i(v, \mathcal{S}(\pi, i))\,, \tag{7}$$

where Π is the set of all $n!$ orderings of $\mathcal{S}$, $\mathcal{S}(\pi, i)$ is the set of SPs preceding i in ordering π, and

$$\Delta_i(v, \mathcal{A}) = v(\mathcal{A}\cup\{i\}) - v(\mathcal{A}) \tag{8}$$

is the marginal contribution of SP i to subset $\mathcal{A}$. The Shapley value is interpreted as the expected marginal contribution of SP i to various subsets of SPs that precede i in a uniformly distributed random ordering of $\mathcal{S}$. Due to efficiency of the mechanism, it is

$$\sum_{i=1}^{n} \phi_i^v(\mathcal{S}) = v(\mathcal{S})\,. \tag{9}$$

Alternatively, the Shapley value can be written as,

$$\phi_i^v(\mathcal{S}) = \frac{|\mathcal{A}|!(n-|\mathcal{A}|-1)!}{n!}\sum_{\mathcal{A}\subseteq\mathcal{S}\setminus\{i\}} (v(\mathcal{A}\cup\{i\}) - v(\mathcal{A}))\,, \tag{10}$$

which is the average marginal contribution of SP i over all possible permutations in which coalition $\mathcal{S}$ can be formed. For $n = 2$ SPs, the Shapley value-shares are

$$\phi_1^v(\mathcal{S}) = \frac{1}{2}[v(\{1\}) + v(\{1,2\}) - v(\{2\}] \tag{11}$$

and

$$\phi_2^v(\mathcal{S}) = \frac{1}{2}[v(\{2\}) + v(\{1,2\}) - v(\{1\}] \tag{12}$$

In order to simplify notation in the next section, we drop "v" from the notation of Shapley value.

III. Problem statement

SPs are strategic in the portion of budget they invest towards building the common infrastructure. On the one hand, an SP would like to conserve its budget and invest as little as possible to the common infrastructure. In that case, it would be able to reap the benefits of the shared infrastructure for service provisioning for a larger number of epochs. On the other hand, with a larger invested amount, the total amount of resource slices created in the infrastructure is larger, and the SP's share of resource slices will be larger as well. A larger number of resource slices implies better QoS, and thus larger revenue per epoch from provisioned services. Overall, the tradeoff in deciding how much to invest amounts to deciding between providing the service for more epochs but possibly at lower quality (and thus with fewer earnings per epoch), versus providing better quality services (and thus with more earnings per epoch) but for smaller number of epochs.

A key metric is the total revenue of SPs. It shows the total earnings of SPs and reflects the utility of end-users that pay SPs to enjoy services. Furthermore, the IO is also interested in high total SP revenue if it withholds some percentage of it.

If SPs jointly coordinate their strategies, they seek the investment policy $\mathbf{x}$ that maximizes total revenue, i.e. they solve

$$\max_{\mathbf{x}\geq \mathbf{0}} \sum_{i=1}^{n} u_i(\mathbf{x}), \text{ such that } u_i(\mathbf{x}) \geq 1 \text{ for all SPs } i\,. \tag{13}$$

We denote this global optimal solution by $\mathbf{x}^* = (x_1^*, \dots, x_n^*)$.

A. Investment game and revenue sharing by the IO

We assume that the resource slice allocation rule $\mathbf{s}$ is fixed, and that the IO applies a total revenue sharing rule among SPs, which is announced apriori to them. Each SP aims to invest an amount that maximizes its own revenue share. We are interested in finding the revenue sharing rule $\boldsymbol{\phi} = (\phi_1, \dots \phi_n)$, such that the selfish SP investment policy results in the global optimal investment solution. In other words, if each SP i aims at maximizing its own share ϕ_i, the collective selfish investment strategy of SPs should maximize total revenue as well.

In order to stress the dependence of the revenue share of each SP i on the investment strategy $\mathbf{x}$, we also write $\boldsymbol{\phi}_i(\mathcal{S}, \mathbf{x})$. For the same reason, we write the worth of coalition $\mathcal{S}$ as $v(\mathcal{S}, \mathbf{x})$. We write $\mathbf{x}_{-i}$ to denote the investment policy of all SPs except i. Given a revenue sharing rule $\boldsymbol{\psi}$, each SP solves the following optimization problem:

$$\max_{x_i\geq 0} \psi_i(\mathcal{S}, x_i, \mathbf{x}_{-i}), \text{ such that } u_i(\mathbf{x}) \geq 1\,. \tag{14}$$

A Nash equilibrium $\mathbf{x}^0(\boldsymbol{\psi}) = (x_1^0(\boldsymbol{\psi}), \dots, x_n^0(\boldsymbol{\psi}))$ is an investment vector such that for each SP i, it is $\psi_i(\mathcal{S}, x_i^0(\boldsymbol{\psi}), \mathbf{x}_{-i}^0(\boldsymbol{\psi})) \geq \psi_i(\mathcal{S}, x_i, \mathbf{x}_{-i}^0(\boldsymbol{\psi}))$ for any other investment strategy $x_i \neq x_i^0(\boldsymbol{\psi})$, given that all other SPs do not change their strategies. In a Nash equilibrium, no SP has an incentive to unilaterally deviate from its chosen strategy.

Proposition 2. *If revenue sharing to SPs is performed based on the Shapley value of each SP i, i.e. if $\boldsymbol{\psi} = \boldsymbol{\phi}$, then if each SP i applies the optimal investment policy x_i^*, it maximizes its Shapley value share, i.e. it is*

$$x_i^* = \arg\max_{x_i\geq 0} \phi_i(\mathcal{S}, x_i, \mathbf{x}_{-i})\,. \tag{15}$$

Proof. Following a proof from [6], the marginal contribution of SP i to subset $\mathcal{A}$ of SPs such that $i \notin \mathcal{A}$ is:

$$\begin{aligned}\Delta_i(v(\mathcal{A}, x_i^*, \mathbf{x}_{-i}), \mathcal{A}) &= v(\mathcal{A}\cup\{i\}, x_i^*, \mathbf{x}_{-i}) - v(\mathcal{A}, x_i^*, \mathbf{x}_{-i}) \\ &= v(\mathcal{A}\cup\{i\}, x_i^*, \mathbf{x}_{-i}) - v(\mathcal{A}, x_i, \mathbf{x}_{-i}) \\ &\geq v(\mathcal{A}\cup\{i\}, \mathbf{x}) - v(\mathcal{A}, \mathbf{x}) = \Delta_i(v(\mathcal{A}, \mathbf{x}), \mathcal{A})\,,\end{aligned}$$

where the second equality holds because $v(\mathcal{A}, x_i^*, \mathbf{x}_{-i}) = v(\mathcal{A}, x_i, \mathbf{x}_{-i})$ since $i \notin \mathcal{A}$, and the inequality holds because by optimizing $v(\mathcal{A}\cup\{i\}, \mathbf{x})$ over only x_i, the outcome cannot be smaller. Therefore, we get that $\Delta_i(v(\mathcal{A}, x_i^*, \mathbf{x}_{-i}), \mathcal{A}) \geq \Delta_i(v(\mathcal{A}, \mathbf{x}), \mathcal{A})$ for all $\mathcal{A} \subseteq \mathcal{S}\setminus\{i\}$. By the strong monotonicity property of the Shapley value, we get that $\phi_i(\mathcal{S}, x_i^*, \mathbf{x}_{-i}) \geq \phi_i(\mathcal{S}, \mathbf{x})$, and the proof is completed. ☐

Thus, by adopting the optimal investment strategy x_i^*, an SP i optimizes its own revenue share, if sharing is performed according to its Shapley value. Further, we show the following:

Proposition 3. *Under the Shapley value revenue-sharing mechanism, an optimal investment strategy $\mathbf{x}^*$ is a Nash equilibrium, i.e. it is $\mathbf{x}^* = \mathbf{x}^0(\boldsymbol{\phi})$.*

Proof. Assume that an optimal investment strategy $\mathbf{x}^*$ is not a Nash equilibrium. Then, there would exist an SP i that could change its strategy from x_i^* to x_i and achieve a share $\phi_i(\mathcal{S}, x_i, \mathbf{x}_{-i}^*) > \phi_i(\mathcal{S}, x_i^*, \mathbf{x}_{-i}^*)$. This is a contradiction, since x_i^* is an optimal investment strategy for SP i. ☐

Therefore, under the Shapley-value revenue-sharing rule, the selfish revenue-share maximization strategy of each SP coincides with the global optimal investment strategy that maximizes total revenue.

B. Investment game and resource slicing by the IO

Now we assume that revenue sharing cannot be applied, either because SPs may be unwilling to abide to the IO rules, or because the IO cannot enforce such a sharing rule. In that case, each SP will select its investment strategy so as to maximize its own revenue, through solving

$$\max_{x_i\geq 0} u_i(\mathbf{x}), \text{ such that } u_i(\mathbf{x}) \geq 1\,. \tag{16}$$

A Nash equilibrium $\tilde{\mathbf{x}}$ is an investment vector such that for each SP i, it is $u_i(\tilde{x}_i, \tilde{\mathbf{x}}_{-i}) \geq u_i(x_i, \tilde{\mathbf{x}}_{-i})$ for any other investment strategy $x_i \neq \tilde{x}_i$. The Nash equilibrium can be computed numerically as follows. We write the Lagrangian for problem (16) for each SP i. We apply the necessary and sufficient KKT conditions for each Lagrangian, and we solve the resulting system of equations. Since $u_i(\mathbf{x})$ is concave in x_i, the Nash equilibrium is unique.

The *Price of Anarchy (PoA)* for the game above is

$$\text{PoA} = \frac{\sum_{i=1}^n u_i(\mathbf{x}^*)}{\sum_{i=1}^n u_i(\tilde{\mathbf{x}})} \geq 1\,. \tag{17}$$

1) Example: Nash equilibrium for $n = 2$ SPs: Consider example 1 above, and a fixed resource slice allocation rule given by $s_1 \in (0, 1)$ and $s_2 = 1 - s_1$. The necessary and sufficient conditions for $\tilde{\mathbf{x}} = (\tilde{x}_1, \tilde{x}_2)$ to be a Nash equilibrium give

$$\frac{\partial u_1(\tilde{\mathbf{x}})}{\partial x_1} = 0, \text{ and } \frac{\partial u_2(\tilde{\mathbf{x}})}{\partial x_2} = 0\,, \tag{18}$$

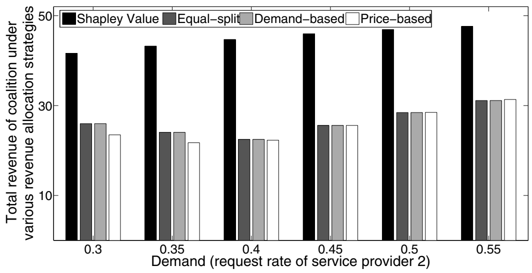

Fig. 2. Example 1: Total revenue of coalition under different revenue allocation (sharing) strategies.

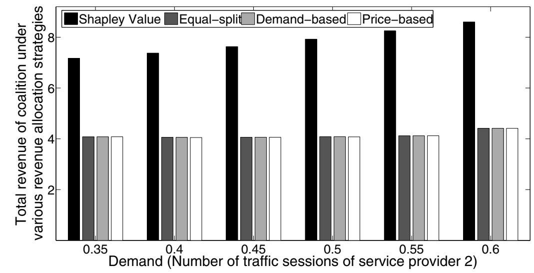

Fig. 3. Example 2: Total revenue of coalition under different revenue allocation (sharing) strategies.

which lead to the system of equations:

$$a_1 g_1^2(x_1,x_2) - D_1 C g_1(x_1,x_2) - (1-x_1)s_1 D_1 C = 0,$$
$$a_2 g_2^2(x_1,x_2) - D_2 C g_2(x_1,x_2) - (1-x_2)s_2 D_2 C = 0,$$

where $g_1(x_1,x_2) = s_1(x_1+x_2) - Cd_1$, $g_2(x_1,x_2) = s_2(x_1+x_2) - Cd_2$. The system solution gives the Nash equilibrium.

2) Finding the optimal slice allocation rule: The IO wishes to alleviate the negative effect of selfish SP investment on total revenue through a resource slice allocation rule $\mathbf{s}$. We write the Nash equilibrium point as $\tilde{\mathbf{x}}(\mathbf{s})$ to stress dependence on $\mathbf{s}$. The IO wishes to minimize the PoA, namely it needs to solve:

$$\min_{\mathbf{s}} \frac{\sum_{i=1}^n u_i(\mathbf{x}^*(\mathbf{s}))}{\sum_{i=1}^n u_i(\tilde{\mathbf{x}}(\mathbf{s}))} \text{ such that } \sum_{i=1}^n s_i = 1\,. \tag{19}$$

IV. Numerical results

In order to evaluate the performance of the proposed mechanisms, we conduct a numerical investigation with two service providers, SP1 and SP2, each with unit budget. We study the 2 examples discussed in subsection II-C, where the resource slices concern (*i*) for example 1, computational power and (*ii*) for example 2, link bandwidth.

The parameter values are; for Example 1, $C = 0.6$, $a_1 = 10$, $D_1 = 2$, $a_2 = 12$, $D_2 = 3$; for Example 2, $C = 0.05$, $\gamma_1 = 2.3$, $\beta_1 = 4.1$, $\gamma_2 = 1.8$, $\beta_2 = 3.3$. The demand of SP1 is kept fixed and equal to $d_1 = 0.4$, and we vary in x-axes in the following figures the demand d_2 of SP2, so that $d_2 \in [0.35, 0.6]$.The simulation was performed in a discrete-event simulator (MATLAB).

A. Revenue Sharing

The IO announces to SPs the applied revenue allocation (sharing) strategy. Then, SPs negotiate over the values of investment x_1 and x_2 in a best-response fashion. Namely, for fixed x_2, SP1 announces its investment strategy x_1 that maximizes her own revenue, and vice versa. The negotiation process is a sequence of best-responses applied by SP1 and SP2, until convergence to the equilibrium investment strategies.

In Figures 2 and 3, we compare the total revenue of the coalition in the equilibrium under different revenue sharing strategies: (i) Shapley value mechanism, (ii) equal-split among SPs, (iii) revenue allocation proportionally to demand of each SP, and (iv) revenue allocation proportionally to the price that each SP charges for its service. For both examples, we assume that the number of slices allocated to SP1 and SP2, s_1 and s_2, is fixed and proportional to their demand. Under Shapley-value based revenue sharing, the total revenue is consistently higher than that under the other strategies, by $50-60\%$ and $75-100\%$ in examples 1 and 2 respectively. As demand d_2 increases, the total revenue of the coalition slightly increases.

TABLE I
Revenue sharing among SP1 and SP2 under various revenue allocation strategies.

	Revenue Sharing			
	Example 1		Example 2	
	SP1	SP 2	SP 1	SP 2
Shapley value	22.31	23.65	4.44	3.49
Equal-split	12.79	12.79	2.04	2.04
Demand-based	12.04	13.54	1.82	2.27
Price-based	12.12	13.46	2.40	1.67

Table I depicts how the revenue is shared among SP1 and SP2 in the equilibrium under the different revenue allocation strategies, for $(d_1, d_2) = (0.4, 0.5)$. For both examples, SP1 and SP2 are better off under the Shapley value mechanism.

B. Slice allocation

The IO chooses the slice allocation policy that minimizes PoA as in (19). We find the optimal slice allocation $(s_1^*, s_2^*) \equiv (s_1^*, 1-s_1^*)$ numerically: we let s_1 run from 0 to 1 with increment of 0.01. For each s_1, the optimal policy $(x_1^*(s_1), x_2^*(s_1))$ and equilibrium policy $(\tilde{x}_1(s_1), \tilde{x}_2(s_1))$ are computed, and the value of PoA(s_1) is found. Next, the value of s_1 that achieves the smallest value of PoA out of these 100 values is found.

In Figures 4 and 5, the value of PoA is depicted under optimal slice allocation, and under the different slice allocation rules, (i) equal-split of slices among SPs, (ii) slice allocation proportionally to SP demand, and (iii) slice allocation proportionally to SP investment. The optimal slice allocation rule leads to the lowest PoA, while investment-based slice allocation appears to be the second best in terms of PoA. For example 1, we observe that as the demand d_2 increases, i.e. the average delay per request increases, the total revenue of the coalition decreases, and this leads to higher PoA. On the other hand, for

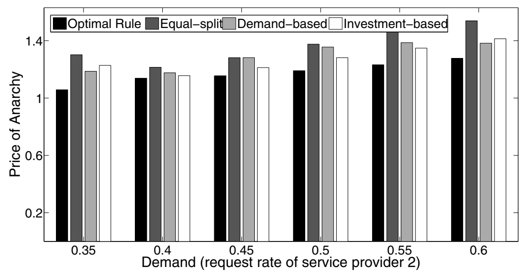

Fig. 4. Example 1: Price of anarchy under different slice allocation rules.

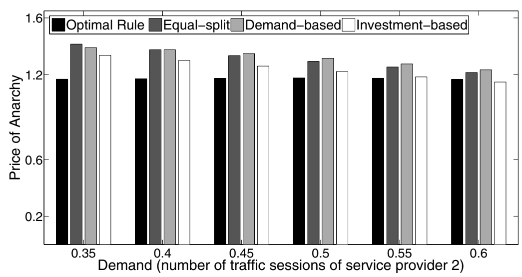

Fig. 5. Example 2: Price of anarchy under different slice allocation rules.

Example 2, the quality of service is not directly related to the demand. Thus as d_2 increases, the total revenue of coalition also increases, which leads to lower PoA.

V. Related work

Infrastructure sharing has been studied through the lens of network economics. In [7], the authors design resource-sharing and payment mechanisms through optimal auction and mechanism design so as to maximize social efficiency. Participants are strategic in revealing private information about their resource needs, and resource-sharing mechanisms incentivize participants so that they contribute to the infrastructure and cover their costs. A different setting in resource sharing is considered in [8], where players compete for location-specific resources. The authors consider the limit of large number of players and employ mean-field theory to show that the equilibrium has a threshold structure: a player switches to a different location based on the numbers of resources and players at the current location. In [6], the authors consider regulation of routing and connecting/ peering strategies of selfish internet SPs through a profit-sharing mechanism based on Shapley value. The work [9] studies the interaction between infrastructure and service providers with respect to investment for different ownership scenarios of resources.

A comprehensive survey of community networks is presented in [10], with emphasis on sustainability and incentives of different stakeholders for engaging with the network. The authors in [11] study contract design in FON-like networks in the presence of user hidden information such as quality of AP connectivity, over which users strategize. The contract includes AP pricing and subscription fee. Optimal infrastructure sharing of co-existing mobile network operators (MNOs) is studied in [12]. Each MNO decides whether to deploy more base stations and share them with other MNOs so as to maximize provisioned rate to its users. The global optimal strategy of maximizing total rate is derived through a mixed-integer linear program, while the individual perspective of each MNO is studied through non-transferable-utility coalition games.

VI. Conclusion

We studied the investment decision problem of SPs that share a network infrastructure which emerges out of their investments. Our main findings are that: (*i*) the selfish revenue-share maximization investment strategy coincides with the global optimal one that maximizes total SP revenue, if total revenue is shared among SPs according to their Shapley values, (*ii*) SP selfishness can be alleviated as much as possible by a resource slice allocation rule that minimizes PoA.

Our model could be extended in various ways. First, the demand may be shaped by user behavior in terms of which SP to choose. The users' SP selection rules may be affected by resource availability (and thus by the investment vector), or by price. More elaborate models could involve diverse resources (e.g. base stations, computing clusters) and different service quality for a user based on the number of other users using the same resource. In that case, user competition would need to be modeled as well. Another aspect is that SPs request resource slices by the IO by declaring their demand and strategize over these declarations. The sharing rule devised by the IO should then discourage dishonest reporting of demand.

References

[1] FON website: https://corp.fon.com/en.

[2] Wikipedia: http://en.wikipedia.org/wiki/List of wireless community networks by region.

[3] R. Baig, R. Roca, F. Freitag, and L. Navarro, "Guifi.net, a crowdsourced network infrastructure held in common", *Computer Networks,* vol.90, pp.150-165, July 2015.

[4] A. Roth, "The Shapley Value: Essays in Honor of Lloyd S. Shapley", Cambridge Univ. Press, 1988.

[5] E. Winter, "The Shapley Value", in Handbook of Game Theory, R. J. Aumann and S. Hart, Eds. Amsterdam, North- Holland, 2002.

[6] R. T. B. Ma, D. M. Chiu, J. C. S. Lui, V. Misra, and D. Rubenstein, "Internet Economics: The Use of Shapley Value for ISP Settlement", *IEEE/ACM Trans. on Networking,* vol.18, no.3, pp.775-787, June 2010.

[7] C. Courcoubetis and R. Weber, "Economic Issues in Shared Infrastructures", *IEEE/ACM Trans. on Networking,* vol.20, no.2, pp.594-608, April 2012.

[8] P. Yang, K. Iyer, and P. Frazier, "Mean Field Equilibria for Competitive Exploration in Resource Sharing Settings", *in Proc. WWW,* 2016.

[9] Y. Zhang, C. Jiang, L. Song, W. Saad, Z. Dawy, and Z. Han, "Complementary Investment of Infrastructure and Service Providers in Wireless Network Virtualization", *in Proc. GLOBECOM,* 2016.

[10] P. Micholia, M. Karaliopoulos, I. Koutsopoulos, L. Navarro, R. Baig, D. Boucas, M. Michalis, P. Antoniadis, "Community Networks and Sustainability: a Survey of Perceptions, Practices, and Proposed Solutions", *IEEE Comm. Surveys and Tutorials,* vol.20, no.4, pp.3581-3606, Q4, 2018.

[11] Q. Ma, L. Gao, Y.-F. Liu, and J. Huang, "A Contract-Based Incentive Mechanism for Crowdsourced Wireless Community Networks", *in Proc. WiOpt,* 2016.

[12] L. Cano, A. Capone, G. Carello, M. Cesana, and M. Passacantando, "On Optimal Infrastructure Sharing Strategies in Mobile Radio Networks", *IEEE Trans. Wireless Commun.,* vol.16, no.5, pp. 3003 - 3016, May 2017.