

Federation of virtualized infrastructures: sharing the value of diversity

Panayotis Antoniadis, Serge Fdida, Timur Friedman
UMPC Sorbonne Universités and CNRS

Vishal Misra
Columbia University

ABSTRACT

By federating virtualized computing and network resources one can significantly increase their value thanks to gains from statistical multiplexing and increases in resource diversity (more distinct locations, technologies, etc.). Successful federation depends upon resource providers being able to agree on policies: how to share the profit generated by external customers and/or how to allocate the resources contributed by the federation participants to their affiliated users. This paper's main contribution is a method that enables organizers of a federation to evaluate the relative importance of the resources contributed by each participant. We build on coalitional game theory concepts and formulate a generic economic model of federation that captures the notion of diversity, which is relevant for a variety of overlay services, and notably the networking research experiments that are running today on PlanetLab. Based on this model, we propose the Shapley value as a means for participants to share the value of federation. We show how this approach can help in the design of policies that encourage infrastructure owners to federate.

1. INTRODUCTION

1.1 Problem statement

Virtualization technology has augmented the potential for sharing computing and network resources. Bandwidth, storage, CPU, and services belonging to different organizations could in principle become part of a global commodity infrastructure that can be used for a wide variety of network services and applications. Cloud computing, distributed experimental facilities, and the Grid are examples of virtualization-enabled sys-

tems that can augment their value to users by pooling resource contributions from multiple organizations.

We call the process of aggregating administratively isolated resources into a common pool, *federation*. Federation enables more efficient utilization of resources through statistical multiplexing. In many cases it is also the only way for certain demanding services to be deployed at a global scale. The Internet is an archetypal example of such a global service: it functions only thanks to federation agreements between pairs or groups of the more than 30,000 autonomous systems.

Federation of computing and network resources provides two benefits, generally speaking, to users: *increased capacity* and *increased diversity*. Increased capacity reduces the completion time of computation-intensive tasks and enables faster communications. Diversity provides value through the complementary characteristics of resources, such as their geographic location, technology, attached users or services, reliability against natural disasters through redundancy, etc.

The core question is how to share the *extra value* that federation brings to its participants. The answer will come in the form of a federation policy whose parameters depend on the special characteristics of the system: the types of resources it offers, the locations of the customers, the levels and type of demand, etc. If properly conceived, this policy will provide incentives to potential members that encourage them to share their resources.

This paper examines the case of the PlanetLab experimental facility [1], which has recently been structured to operate as a federation of regional PlanetLab testbeds (PlanetLab in North America, PlanetLab Europe, etc.) that each provide their respective users with full access to the global system [2]. This federated system provides a geographically distributed set of computing and network resources. On top of PlanetLab, users can deploy a network overlay to run a peer-to-peer service or a content distribution network, experiment with novel Internet architectures in realistic conditions, perform location-sensitive measurements, etc.

PlanetLab is unlike Grid systems, which have been

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ACM CoNEXT 2010, November 30 – December 3 2010, Philadelphia, USA.

Copyright 2010 ACM 1-4503-0448-1/10/11 ...\$10.00.

the focus of existing literature in this area [3, 4, 5, 6]. Whereas the value from federating grids comes principally from shared computational resources, PlanetLab federation provides value in the form of the *multiple locations* at which shared resources reside. Our federation model extends recent work on coalitional game theory for communication networks [7, 8, 9, 10, 11, 12] by capturing this diversity dimension.

Using this model we demonstrate that simple proportional sharing mechanisms can lead to solutions that are considerably different from the optimal ones dictated by coalitional game theory. Our objective is to apply in practice this and other insights gained from our numerical analysis in the design of realistic policies for the emerging PlanetLab federation, as discussed in Sec. 4.

While PlanetLab is our main target, our proposal can be extended to other federation scenarios for network connectivity, content distribution networks, and cloud computing. In fact, our model serves as a generalization of these scenarios that can help in comparing the properties of the coalitional game theory approach in different cases and understanding in more depth the factors that should influence the design of federation policies at large.

1.2 The case of PlanetLab

PlanetLab is the largest geographically-distributed experimental facility yet constructed for computer networking. It is at its base a federated system since its infrastructure is built from the individual contributions of participating universities and research institutions, the *sites*, which are each required to contribute at least two servers, or *nodes*, connected to the Internet. When an institution fulfills this requirement it obtains the right to deploy up to ten *slices* across the facility. A slice consists of one virtual machine on each of a set of nodes. Each slice benefits from a short-term fair allocation policy (i.e., competing slices acquire a fair share of the available resources at each node and in each time slot).

Much of the value of PlanetLab comes from its geographic diversity. Its approximately 1,000 nodes are widely situated, allowing experimenters to deploy overlays on the Internet at a global scale. This value is attested to by, for example, one particularly high profile user of PlanetLab: Google pays an annual fee and uses PlanetLab to check the accessibility and performance of its services from different geographic locations.

Until recently, all nodes were controlled by the PlanetLab operations center at Princeton University, which acted as the sole manager of the system’s complex trust and security issues. Today, the European part of PlanetLab, PlanetLab Europe (PLE), is part of the OneLab experimental facility [2] and PlanetLab Central (PLC) and PLE are federated through the direct exchange of user credentials and node descriptions.

PLC and PLE each run almost identical instances of the MyPLC operations server software and they each follow the same local provision and allocation policies. This makes the implementation of the current federation policy (a simple peering agreement) rather straightforward. But the PlanetLab federation is rapidly becoming complex: PLC and PLE are being joined by emerging regional authorities such as PlanetLab Japan (PLJ), and other testbeds (e.g., G-Lab, EmanicsLab, and VINI) are joining the federation through the regional authorities.

More sophisticated federation policies will soon be needed for PlanetLab to succeed in becoming a hierarchical federated system that offers both increased resource capacity and diversity to its users. Currently, the main focus is on federation at the top level of this hierarchy (PLC-PLE-PLJ, etc.). We will treat this as a single layer by making the simplifying assumption that each authority has full control over the amount of resources that it brings to the federation. In future work, we will study the interdependencies between local and global federation policies.

1.3 Our contribution

This paper formulates a simple but expressive economic model that captures an interesting characteristic of the evolving federation ecosystem: the value of diversity. Based on this model, we compute the Shapley value [13] of the individual members of a federation facing a certain type and volume of demand. Our study leads us to the following conclusions:

- Economic models considering a single resource dimension and/or simple proportional sharing mechanisms are not enough to capture the diversity dimension and lead in general to inefficient outcomes.
- The required value sharing policies highly depend on the type and volume of demand.
- The Shapley value provides a suitable basis for sharing the additional value that federation brings.

This paper is organized as follows. Sec. 2 formulates our economic model by defining the main actors in federation and its value. Sec. 3 presents the high-level federation game and the different theoretical optimization problems that need to be solved in order to compute the optimal policy based on the model’s assumptions. Sec. 4 presents some interesting insights using different numerical examples since the general problem is not analytically tractable. Sec. 5 discusses related work and places our contributions in context. We summarize these insights and discuss the possible future directions of our work in Sec. 6.

2. ECONOMIC MODELING

There are two possible scenarios that benefit from the federation of virtualized infrastructures.

- *Commercial scenario*: the value for the federation members corresponds to a monetary profit generated by the payments of external customers wishing to use the federated infrastructure.
- *P2P scenario*: the value is strictly related to the utility acquired by the users, who gain access to the infrastructure thanks to their affiliation with a contributing facility.

Both scenarios apply to PlanetLab, which currently has external industrial customers participating by paying a fixed annual fee and internal users, mostly researchers, affiliated with the contributing research institutions and universities.

Federation is meaningful only if this extra value minus the cost of federation and local resource provision is greater than the value acquired when participants act alone. In general, the assumed utility and cost functions, and the characteristics of demand are the main factors that will affect the estimation of the value of federation.

An individual facility’s contribution to a federation has many characteristics, among them: capacity (CPU, physical memory, hard disk, bandwidth), geographical or network position of nodes (location or AS number), scale (the number of nodes), technology, availability (the portion of time that it can be used), reliability (how long it remains available without interruption), and average load (congestion level). Their effect on the total value of a federation depends on the types of experiments and services that are to be run on the federated infrastructure and the corresponding utility and cost functions. If we wish to have a tractable model, we must simplify by sacrificing some details as described below.

2.1 Resource providers

Consider a set of facilities \mathcal{N} participating in a federation, with $N = |\mathcal{N}|$. Each facility i contributes resources (e.g., nodes having certain capabilities), which reside in one or more of a set \mathcal{L} of locations, indexed with l . We denote $L = |\mathcal{L}|$ the total number of possible locations and $\mathcal{L}_i \subseteq \mathcal{L}$ the subset of these locations for which facility i provides nodes.

Then the facility’s contribution to the geographic diversity of a federation could be expressed through the number of distinct locations at which it provides resources, $L_i = |\mathcal{L}_i| \leq L$. This implicitly assumes that all distinct locations are equally valuable to the users, which is not unrealistic for a facility designed to support networking experiments.

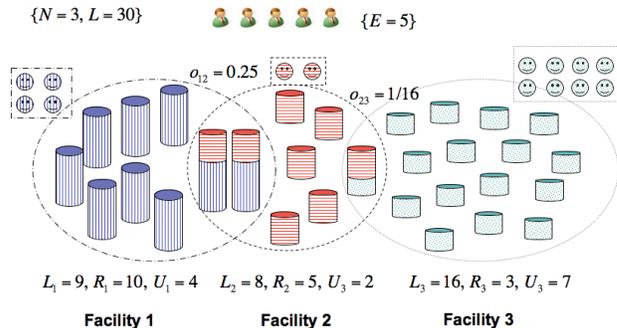


Figure 1: Federation model

In addition, each facility i will need to assign an amount of resources at each of its locations $l \in \mathcal{L}_i$. In the general case, this amount should be expressed as a combination of values for CPU, storage, bandwidth, and other resource characteristics, as described above. A possible simplification that is meaningful in a scenario that focuses on the value of diversity is to aggregate the different resource dimensions into a single resource capacity variable R_i^l , which expresses the number of different experiments that can run at location l thanks to the contribution of facility i . In other words, R_i corresponds to whatever happens to be the bottleneck resource. Finally, in the general case, the resources of each facility could be made available only for a subset of time $T_i \in [0, 1]$.¹

Under this model, some locations can host resources from multiple facilities. We can capture this by introducing the probability of overlap o_{ij} between the sets \mathcal{L}_i and \mathcal{L}_j . For simplicity, we could assume that these probabilities are independent but more realistic overlapping models can be easily developed if needed.

Resource providers are also characterized by their affiliated users and/or customers $U_i = |\mathcal{U}_i|$. In a commercial setting, this number will be part of their contribution to the total profit generated, as for example in the model of Aram et al. [8]. In contrast, in a P2P setting, this number will correspond to their contribution to the federation cost (because of congestion). In addition, affiliated users will represent the means for a facility to acquire a subset of the total federation value generated, which will need to be shared fairly among providers as in the case of profit.

There is also a possible set of external customers, denoted by \mathcal{E} , who do not belong to any facility but pay for the acquired service to a central authority responsible for sharing the total profit among the contributing facilities. This is the case that we will consider in our examples below since it is easier to analyze without any

¹In our analysis in this paper we will assume that, $\forall i, T_i = 1$, and express resource contributions only through R .

loss of generality in terms of the federation value and/or profit sharing function that needs to be put in place.

Fig. 1 depicts a federation of three facilities ($N = 3$) providing different resource units each on 30 distinct locations. Notice that at locations where there is overlapping the total available resources are the sum of the resources per location provided by each facility.

2.2 Users

As already stressed, the utility acquired by a user of an experimental facility depends not only on the total amount of resources acquired (as often assumed in the literature on grids [3, 6]) but also on the number of different locations and/or technologies providing these resources. This means that the value that a researcher/customer acquires through an experiment k will depend on the number of geographically distributed nodes that will have enough capacity to host this experiment upon its arrival. So, an experiment can be characterized by three different demand attributes:

- The set of individual **locations**, or for simplicity just a range for a required number of distinct locations $l_k \in [l_k, \bar{l}_k]$.
- The **resources per location** $r_{kl} \in [r_{kl}, \bar{r}_{kl}]$.
- The **holding time per location** $t_{kl} \in [t_{kl}, \bar{t}_{kl}]$, which could depend or not on the corresponding amount of resources.

For example, the following are some realistic values for three different types of experiments/services running today on PlanetLab:

- A P2P experiment: $l = 40, \bar{l} = \infty, r = 1, t = 0.1$.
- A CDN service: $l = 100, \bar{l} = 500, r = 4, t = 1$.
- A measurement experiment: $l = 500, \bar{l} = \infty, r = 2, t = 0.4$.

Notice that capacity-hungry services (e.g., computationally intensive jobs) have lower holding times as the amount of resources made available to them increases. On the other hand, diversity-hungry network experiments require as many network nodes as possible but their holding time remains a wholly independent variable. This property of holding time can significantly affect the level of statistical multiplexing achieved under different federation scenarios. In our model, we take this into account by assuming that holding time is part of the request and independent of the available capacity in the system. This assumption captures an important differentiating factor of networking experimental facilities compared to grids and Internet services.

2.3 The value of federation

In order to quantify the gains of a federation we need to define the utility and cost functions of the players involved.

2.3.1 Utility

The utility function expresses the satisfaction of a user acquiring a certain amount of resources to run her experiment. It is typical to assume that it is a concave function of the acquired resources, which is often a critical assumption for the analysis to be tractable.

In our context, the number of different requested locations is the most important dimension of the resources assigned to an experiment. The utility of a networking experiment is typically zero if it cannot be run in a sufficiently large number of locations. So it is reasonable to assume a minimum threshold level of locations required by a user.

A simple and intuitive way to express this type of utility would be with a simple step function, which could be linear, concave, or convex after a certain threshold. This is the utility function that we will consider in our examples below:

$$u_k(x_k) = \begin{cases} x_k^d, & \text{if } x_k \geq L_k \\ 0, & \text{if } x_k < L_k \end{cases} \quad (1)$$

where x_k is the number of distinct locations assigned to the experiment, L_k is the experiment-specific threshold value, and d defines the shape of the utility function as depicted in Fig. 2.

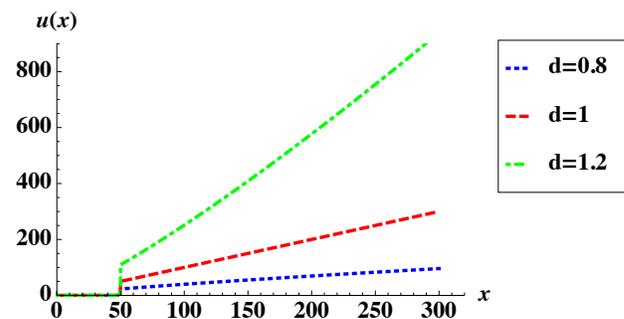


Figure 2: Utility functions for $l = 50$.

For simplicity we assume that d is the same for all users and that $\bar{l}_k = \infty$, since in practice the maximum requirements of experiments are much higher than the total number of available locations.

Based on this utility function the contribution L_i affects the value per experiment generated by a federated infrastructure while the contribution R_i affects the number of different experiments that can be hosted at each location. So, using a rather simple model we manage to capture both dimensions of federation: capacity and diversity.

2.3.2 Cost

Certain resource attributes are easier to provide than others. For example, many organizations have just a few locations and cannot easily increase the geographical diversity that they provide. Other attributes depend on long-term investments such as the number of nodes and their resource capabilities. Finally, there are attributes such as reliability that depend on investing ongoing effort. In general, the cost function of a facility i depends with different intensity on all three L_i, R_i, T_i . $c_i(L_i, R_i, T_i) = \alpha L_i + \beta R_i + \gamma T_i$, with usually $\alpha < \beta < \gamma$.

In this paper we will ignore resource provision costs, assuming that facilities have invested in their infrastructure prior to federation. In practice many of these costs are subsidized by research funds (e.g., [14, 15]) and in general are similar for different institutions. So, our solutions for dividing the value will not be significantly affected by the actual costs involved.

Note that there is also a fixed federation cost c_F related to the administrative, technical, and legal aspects. But this cost will only affect the final decision about the benefit for forming a federation, which again in practice will depend on often external factors and will not influence the properties of the different profit sharing policies.

3. THE FEDERATION GAME

3.1 Resource allocation

The first decision that needs to be made by a federation is the way resources will be allocated to different users when demand exceeds the available capacity.

In the commercial scenario the common objective is to choose the set of experiments that maximize the total profit. In the context of our model, assuming for simplicity that $r_{kl} = 1$ and $t_{kl} = 1, \forall k \in \mathcal{E}$, and $l \in \mathcal{L}$, this decision would correspond to the solution of the following optimization problem:

$$\begin{aligned} & \underset{a_{kl}}{\text{maximize}} \left[\sum_{k \in \mathcal{E}} u_k \left(\sum_{l \in \mathcal{L}} a_{kl} \right) \right] \\ & \text{subject to} \quad \sum_{k \in \mathcal{E}, l \in \mathcal{L}} a_{kl} < \sum_{i \in \mathcal{N}, l \in \mathcal{L}} R_i^l \end{aligned} \quad (2)$$

where $a_{kl} \in \{0, 1\}$ equals 1 if experiment k is assigned to location l and 0 otherwise. Then the total number of nodes assigned to experiment k , $x_k = \sum_{l \in \mathcal{L}} a_{kl}$. The total profit, denoted P , will be a function of the total utility acquired depending on the properties of the underlying market. That is $P = \mu \sum_k u_k(x_k)$, where $\mu \leq 1$ when that utility is measured in monetary units. Then given the profit P of a coalition \mathcal{N} , facilities need to agree on a vector $\mathbf{s} = \{s_1, \dots, s_N\}$ that will determine the part of the total profit p_i assigned to each one

of them as follows:

$$p_i = s_i P, \text{ s.t. } \sum_i s_i = 1.$$

The focus of this paper is precisely on how to decide on the values s_i and not on how to maximize the total profit. Thus the above simplifications for profit maximization do not change the qualitative results that we obtain. We could equally well use any predefined price function that depends on external factors (e.g., competing services provided by cloud computing providers).

In the P2P scenario it is through resource allocation schemes that the total value will be shared. There are different ways for the system designer to differentiate the quality perceived by users of different facilities based on contribution [16]. In the following, we will constrain ourselves to those that affect the core parameters of our model: capacity (R) and diversity (L).

In this scenario each experiment k belongs to a certain facility i ($k \in \mathcal{U}_i$). So, an amount of locations x_i are assigned first to facilities which then are responsible for redistributing them internally to the candidate experiments of their affiliated users ($x_i = \sum_{k \in \mathcal{U}_i} x_k$). Then the resource allocation decision should incorporate the value sharing policy since it is through their users that facilities acquire value:

$$u_i^f(x_i) = \sum_{k \in \mathcal{U}_i} u_k(x_k).$$

So, in this case resource allocation decisions are more complex since the objective is not only to maximize the total value generated in the system but at the same time to share it in an incentive-compatible way, depending on the contributions of the different facilities.

This means that while in the commercial scenario it is in the interest of the facilities to maximize the total utility of the customers (and thus the profit), the requirement for incentive compatibility might force a coalition to a suboptimal solution in terms of total utility generated. Our optimization problem in this case takes the following form:

$$\begin{aligned} & \underset{a_{il}}{\text{maximize}} \left[\sum_{i \in \mathcal{N}} u_i^f \left(\sum_{l \in \mathcal{L}} a_{il} \right) \right] \\ & \text{subject to} \quad \sum_{i \in \mathcal{N}} a_{il} r_{il} < \sum_{j \in \mathcal{N}} R_j^l, \forall l \\ & \quad \quad \quad u_i^f \left(\sum_{l \in \mathcal{L}} a_{il} \right) \geq u_i^f(L_i) \end{aligned} \quad (3)$$

The value sharing decision is incorporated into the result of the above optimization problem since

$$s_i = u_i^f(x_i^*) / \sum_{j \in \mathcal{N}} u_j^f(x_j^*),$$

where $x_i^* = \sum_{l \in \mathcal{L}} a_{il}^*$.

Notice that now the decision variable a_{il} refers to a facility i and not to individual experiments. The term r_{il} is then added at the capacity constraint due to the fact that more than one experiment belonging to a single facility i (each with $r_{kl} = 1$ as assumed before) might happen to be assigned to the same location l . The second constraint of (3) ensures that each facility receives a higher value than when acting alone.

This is sufficient to ensure that all facilities will have the appropriate incentives to participate in the federation, but only if the *grand coalition*, the one in which all candidate members participate, is the only possible federation. If facilities have also the option to form smaller coalitions, one needs to ensure that their value share exceeds that which could be achieved in all different possible coalitions. This leads us to the notion of the core and other important concepts, developed in the context of coalitional game theory, discussed next.

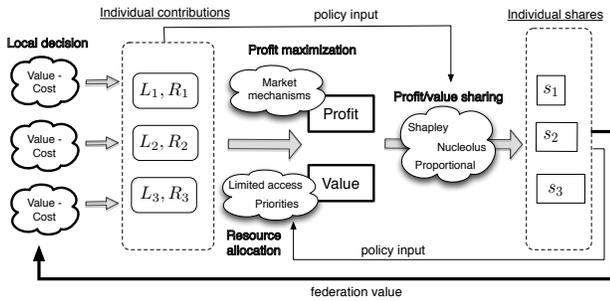


Figure 3: The federation game

3.2 Sharing profit and value

3.2.1 The core

The fundamental question that a federation of providers needs to answer is: “what you get for what you give?” In other words, having collected an amount of money and/or potential value, how should this be shared amongst the different contributors?

Our model defines a transferable utility coalitional game [17]. The *core* C of the game is a well defined concept in coalitional game theory: it is the set of solutions that have the property that no player has an incentive to depart from the grand coalition either alone or by forming smaller coalitions $\mathcal{S} \subset \mathcal{N}$. Formally, given a cooperative game with total possible value $V(\mathcal{N})$ and the vector $\mathbf{v} = \{v_1, \dots, v_N\}$ such that $v_i = s_i V(\mathcal{N})$,

$$C = \left\{ \mathbf{v} : \sum_{i \in \mathcal{N}} v_i = V(\mathcal{N}) \text{ and } \sum_{i \in \mathcal{S}} v_i \geq V(\mathcal{S}), \forall \mathcal{S} \subseteq \mathcal{N} \right\}.$$

Clearly, the existence of the core for a given game is of significant importance and it has been studied extensively in the literature (see Saad et al. [12] for a recent survey of applications of this concept to communication networks). The most important properties of the

game that determine the existence of the core are super-additivity and convexity. Both depend significantly on the utility function assumed.

For example, it is easy to see that if our utility function (1) is strictly concave and continuous with no minimum diversity threshold and no possibility for statistical multiplexing (i.e., $d < 1$, $L_k = 0$, and $t_k = 1$, $\forall k$), the corresponding game is not super-additive and thus not convex either. Notably, for concave utility functions convexity depends highly on the levels of statistical multiplexing and the value of diversity.

More specifically, the smaller the t_k s, the more chances for the game to be super-additive and convex. Note that this is the typical case in which Grid federation would be meaningful [6]. In our scenario, it is also the importance of diversity that significantly affects the convexity of the game: As \underline{l} grows, the more small coalitions are of zero value and thus the comparative value of the grand coalition increases, turning the core nonempty. On the other hand, when $d \geq 1$ the core always exists.²

The existence of the core ensures the feasibility for forming the grand coalition but not all different solutions in the core are equally desirable since they could be based on different notions of fairness (e.g., equity, max-min, proportional) and different ways to evaluate the individual contribution of a member of the coalition. In our case, facilities that offer resources in locations with low overlapping will contribute more to the total value generated. How can one quantify the relative contribution of different facilities in a testbed federation?

3.2.2 Shapley value

The Shapley value, originally proposed by Lloyd Shapley [13], provides an elegant means to evaluate the individual contribution of the members of a coalition. It expresses the ‘importance’ or ‘uniqueness’ of each member of a coalition with regard to the value of its output. Applied to the case of ISP interconnection, for example, the share of each provider depends on the number of similar providers, which could be those providing the same content or those covering the same region [7]. The less overlapping, the more valuable one’s contribution.

Formally, given a coalition \mathcal{S} , the Shapley value determines how the value of the coalition, captured by $V(\mathcal{S})$, should be shared among the players in \mathcal{S} . Specifically, for each player i and coalition \mathcal{S} , the Shapley value of player i is denoted by $\varphi_i(\mathcal{S}, V)$ and is uniquely defined by the following three axioms

Axiom 1 (EFFICIENCY). $\sum_{i \in \mathcal{S}} \varphi_i(\mathcal{S}, V) = V(\mathcal{S})$.

²Although values of $d > 1$ are not realistic for an unlimited amount of resources, one could assume that $d > 1$ for a certain range of $[\underline{l}, \bar{l}]$.

Axiom 2 (SYMMETRY). *If for all $S' \subseteq \mathcal{S} \setminus \{i, j\}$,*

$$V(\mathcal{S}' \cup \{i\}) = V(\mathcal{S}' \cup \{j\})$$

then $\varphi_i(\mathcal{S}, V) = \varphi_j(\mathcal{S}, V)$.

Axiom 3 (FAIRNESS/BALANCED CONTRIBUTION).

For any $i, j \in \mathcal{S}$, j 's contribution to i equals i 's contribution to j , or, in other words

$$\varphi_i(\mathcal{S}, V) - \varphi_i(\mathcal{S} \setminus \{j\}, V) = \varphi_j(\mathcal{S}, V) - \varphi_j(\mathcal{S} \setminus \{i\}, V).$$

The efficiency axiom states that the total of the revenues assigned to each player equals the actual profit created by their coalition. In other words, the mechanism does not contribute or receive extra profit. The symmetry axiom requires that if two players contribute the same to every subset of other players, they should receive the same amount of revenue. Finally, the balanced contribution axiom addresses the fairness between any pair of players. It may be illustrated on a two-player system where $\mathcal{N} = \{1, 2\}$. By efficiency we have that, for a coalition of a single player, $\varphi_i(\{i\}, V) = V(\{i\})$. The fairness axiom states that the gain (or loss) of revenue from cooperation, as seen by players 1 and 2, should be the same: $\varphi_1(\mathcal{N}, V) - V(\{1\}) = \varphi_2(\mathcal{N}, V) - V(\{2\})$. In that case it means that the global gain of cooperation, defined as $V(\mathcal{N}) - V(\{1\}) - V(\{2\})$, is split evenly among players. The balanced contribution axiom preserves and generalizes this egalitarian property [18].

Based on the axioms above, one can show that the Shapley value φ can be computed as follows [13]:

$$\forall i \in \mathcal{S}, \varphi_i(\mathcal{S}, V) = \frac{1}{|\mathcal{S}|!} \sum_{\pi \in \Pi} \Delta_i(V, S(\pi, i)), \quad (4)$$

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

The Shapley value of a player i can thus be interpreted as the *expected* marginal contribution $\Delta_i(V, S')$ where S' is the set of players in \mathcal{S} preceding i in a uniformly distributed random ordering of \mathcal{S} . We denote $\hat{\varphi}_i$ the *normalized* Shapley value,

$$\hat{\varphi}_i = \varphi_i / V(\mathcal{N}), \quad (5)$$

which expresses the percentage of the total value of the federation attributed to facility i . That is, following the Shapley value scheme, $s_i = \hat{\varphi}_i$ and $v_i = s_i V(\mathcal{N})$.

There are two main criticisms of the use of the Shapley value in practice. First, its centralized computation relies on the willingness of the interested parties to share often sensitive private information, as, for example, the level of local demand and profits [7, 8]. However, in the case of PlanetLab, federation is based on mutual trust between all top-level authorities and federated resources are to be accessed through an open distributed architecture (namely, the Slice-based Federation Architecture,

or SFA, currently under development [19]). So, all required information for the computation of the Shapley value, as in our numerical analysis, is available to all federation members.

Second, the problems defined above are not analytically tractable in the general case. However, it is easy to compute the Shapley value for specific realistic scenarios. This is especially so in our case since PlanetLab federation has a hierarchical nature and the number of providers at each level of the hierarchy is not expected to be large. So, even though our numerical examples are simple, they are meaningful and the insights they provide could be used for the design of practical policies, like the ones required for the top level federation between PLC, PLE, and PLJ. For example, $\hat{\varphi}_i$ s can be computed off-line and used as heuristic evaluators of the individual contributions of facilities, given the mixture of expected users.

3.2.3 Other solutions

A simpler notion of contribution is the one considering only the amount of resources shared by different facilities either in terms of availability or consumption. Proportional sharing schemes based on this measure of contribution have been recently studied in the context of scientific grids (see, for example, [6]).

We denote $\hat{\pi}_i$ the *proportionally fair share* of facility i , defined as

$$\hat{\pi}_i = (L_i R_i) / \sum_{k=1}^N L_k R_k. \quad (6)$$

Then $\pi_i = \hat{\pi}_i V(\mathcal{N})$. Alternatively, one could consider the proportionally fair share based on consumed resources $\hat{\rho}_i$ ($\rho_i = \hat{\rho}_i V(\mathcal{N})$),

$$\hat{\rho}_i = \frac{\sum_{j \in \mathcal{N}, l \in \mathcal{L}_i} a_{jl} r_{jl}}{\sum_{j \in \mathcal{N}, l \in \mathcal{L}} a_{jl} r_{jl}}, \quad (7)$$

which demand is low differs significantly from $\hat{\pi}_i$.

An alternative to the Shapley value and the simpler proportional fair schemes presented above, is the *nucleolus function*, which is a more complex mechanism [17]. Simply put, the nucleolus function achieves a max-min fair allocation of profit among all possible coalitions. It achieves this by recursively maximizing the excess value³ of the worst treated coalitions.

It can be proven that if the core of a game is nonempty, the nucleolus solution is in the core [17]. However, in this case the profit share of the members would be completely independent of their contribution, which means that the incentives for resource provision will not in general be correctly aligned. The same concerns exist

³Excess value is defined as the sum of the final allocations of the members of a coalition \mathcal{S} , computed for the grand coalition \mathcal{N} , minus the total value of coalition \mathcal{S} .

for simple fairness schemes such as the equity approach which always divides the profit equally among all the members of the grand coalition. Thus, in the following we will focus on comparing Shapley value with the two above types of proportionally fair profit sharing. In Sec. 5 we discuss how our work is related to market-based approaches (e.g., [4, 3, 20, 21, 22]).

3.3 Decision making and equilibrium

Faced with an agreed profit sharing or resource allocation policy, facilities decide the quantity of resources they will bring to the federation based on the trade-off between the extra value/profit generated and their internal costs. Fig. 3 illustrates the different stages of the federation game and their interdependencies. The fact that more sophisticated schemes like the Shapley value do not have a closed form makes it very challenging to analytically study the evolution and possible equilibria of this game in the general case.

In this paper we don't fully address this challenging question but define one of its most important elements: the estimation of the relational importance of the individual contributions of the federation members. We present numerical results that help us understand the dynamics of the game and the properties of the different solution concepts.

A way to generalize and formalize our initial analysis toward characterizing the equilibrium of the global federation game is a very interesting research avenue with significant potential impact on the theory and practice of coalitional game theory for communication networks. We believe that our model and first observations constitute a good starting point.

4. NUMERICAL ANALYSIS

In this section, we compare profit sharing schemes and demonstrate the important roles of the utility function and the level of demand using simple numerical examples as our guide. In particular, we focus on the differences between the Shapley value approach, where $s_i = \hat{\varphi}_i$, defined in (5), and the resource-based proportional sharing strategy, where $s_i = \hat{\pi}_i$, defined in (6).

We do this in the context of the commercial scenario, assuming that the profit generated by a coalition equals the sum of the utility acquired by the users as defined by (1): $P = V = \sum_{k \in \mathcal{E}} u_k(x_k)$. Although we study only subsets of the model defined in Sec. 2, our analysis can provide a practical guide for sharing profit in a way that will be considered fair and that will provide independent facilities with suitable incentives for each to contribute as many resources as possible to the federation.

Notice that sharing P efficiently is an issue that already arises in the PlanetLab context, as subscription fees are paid by industrial users of the system, such as Google and HP. The default policy at present is for

each top-level authority (PLC, PLE, eventually PLJ) to retain the totality of the fees that it brings in.

4.1 The threshold value

Let there be three facilities, $\mathcal{N} = \{1, 2, 3\}$. We compute the Shapley value of a facility i as follows:

$$\varphi_i = (2V(\{i\}) + (V(\{i, j\}) - V(\{i\})) + (2(V(\mathcal{N}) - V(\{j, k\})) + (V(\{i, k\}) - V(\{k\}))/6,$$

where j and k are the other two facilities.

Assume that: the facilities make available resources $L_1 = 100$, $L_2 = 400$, and $L_3 = 800$; there is a single experiment in the system; both the resources and holding time per location are 1 ($\forall l, r_l = t_l = 1$); the utility function (1) is linear ($d = 1$); and the only condition for the experiment to be accepted is that there should be x locations available, such that $x \geq \underline{l}$, whereas otherwise it is blocked.

For this case, it is easy to see that $V(\{i\}) = u(L_i)$, $V(\{i, j\}) = u(L_i + L_j)$, and $V(\mathcal{N}) = u(L_1 + L_2 + L_3)$. For $\underline{l} = 500$, for example, we find the following values for the different possible coalitions: $V(\{1\}) = 0$, $V(\{2\}) = 0$, $V(\{3\}) = 800$, $V(\{1, 2\}) = 500$, $V(\{2, 3\}) = 1300$, $V(\mathcal{N}) = 1300$. Then for facility 2, for example, the normalized Shapley value $\hat{\varphi}_2 = 2/13$, while the corresponding proportional share $\hat{\pi}_2 = 4/13$.

Notice that for $\underline{l} = 0$, each $\hat{\varphi}_i$ and $\hat{\pi}_i$ are equal. The differences arise when $\underline{l} > 0$ because of the possible inability of certain coalitions to serve the experiment.

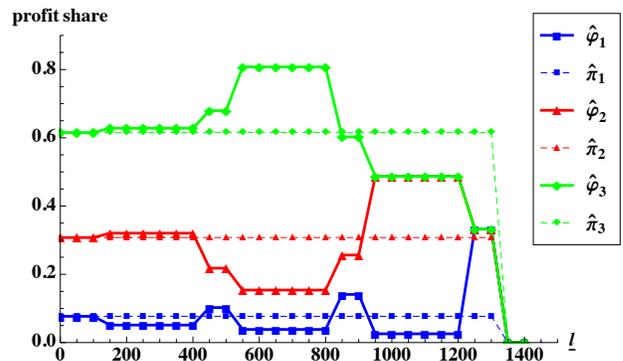


Figure 4: Profit shares with respect to \underline{l}

Fig. 4 depicts $\hat{\varphi}_i$ and $\hat{\pi}_i$, $\forall i$, as \underline{l} increases. It is instructive to follow the corresponding changes of the Shapley value.⁴ The first change occurs when the threshold hits $L_1 = 100$. Above this point, facility 1 has zero payoff when acting alone and this is the reason why its share decreases. The next change occurs at $\underline{l} = L_2 = 400$. Facility 2 is the one that henceforth

⁴Fig. 4 shows $\hat{\varphi}_i$ and $\hat{\pi}_i$ as \underline{l} increases in increments of 50. The lines in this and subsequent figures are added only to better illustrate the trends.

becomes incapable of serving the customer alone and thus its share decreases as well. Facilities 1 and 2 are still able to serve the customer in coalition, but above $\underline{l} = L_1 + L_2 = 500$ they lose this ability and so both lose share to facility 3. Facility 3 in turn loses a significant part of its importance above $\underline{l} = L_3 = 800$, but it can still serve the customer in coalition with facility 1. Above $L_1 + L_3 = 900$, this coalition is no longer possible, and above $L_2 + L_3 = 1200$ only the grand coalition will work. In the grand coalition, all facilities receive an equal share even if their resource contributions are very different! Finally, above $\underline{l} = L_1 + L_2 + L_3 = 1300$, no coalition can meet customer demand.

4.2 The shape of the utility function

We now show how the shape d of the utility function affects the properties of our game and the values $\hat{\varphi}_i$ and $\hat{\pi}_i$. In Fig. 5, $\hat{\varphi}_i$ and $\hat{\pi}_i$ are depicted as d increases while the threshold value is kept fixed at $\underline{l} = 600$.

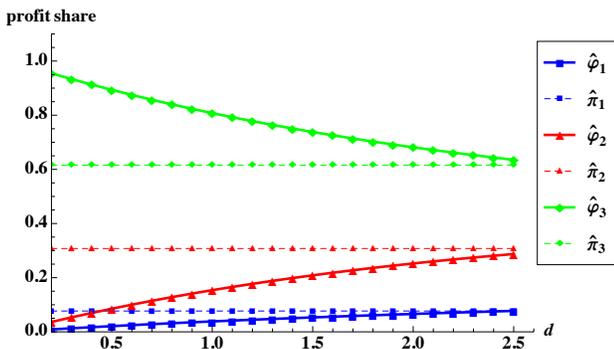


Figure 5: Profit share with respect to d

We see that, as d increases, the Shapley values get closer to the proportional share, since the smaller coalitions lose their importance compared to the larger ones due to the convexity of the utility function.

4.3 The role of demand

4.3.1 Resource requirements

Until now, we considered a very simple setup of our game in which the amount of resources per location does not matter. Consider now the same scenario but with the following resource values per facility: $R_1 = 80$, $R_2 = 20$, $R_3 = 10$. Notice that with these values all $L_i * R_i$ are equal. Let there also be a number of experiments k of the same type, with utility function (1) and $r_{kl} = t_{kl} = 1, \forall k, l$: enough in number to fill the system's capacity.

Fig. 6 depicts the Shapley value as \underline{l} increases while we take into account the resources available per location. We observe that facilities offering exactly the same amount of total resources can have very different contributions to the total value generated by a fed-

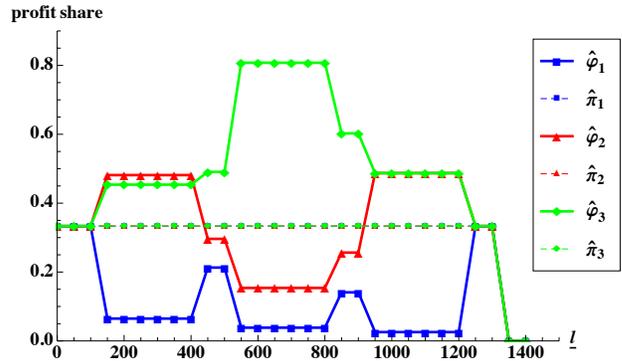


Figure 6: Profit share with respect to \underline{l}

eration, depending on the diversity provided and the corresponding minimum user requirements.

This result is due to the fact that the dominant part of the system's value comes from diversity. Resources are just a requirement for diversity to be exploited. Similarly, on the Internet it is the access to the customers that is valued and not the resources used for the communication, which ideally should be minimized (e.g., the smaller the number of hops, the better).

Notice also that the normalized Shapley value is the same at the extremes:⁵ when experiments are easy to multiplex or when *all* facilities are required to cooperate for any experiment to be meaningful. The latter is the case for the Internet (where full connectivity is the minimum requirement).

4.3.2 Different types of experiments

Another level of complexity is added when we consider different types of experiments. Fig. 7 depicts φ_i and π_i for $0 \leq \sigma \leq 1$, where σ is the ratio between two type of experiments with $\underline{l}_1 = 0$, $\underline{l}_2 = 700$. Now, $R_1 = 80$, $R_2 = 50$, and $R_3 = 30$.

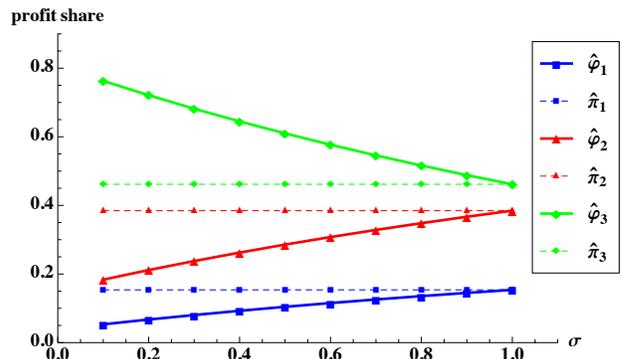


Figure 7: Profit share with respect to σ

⁵This is so, because $L_i * R_i$ is equal for all i . Otherwise we would get the proportional fair solution $\hat{\pi}_i$ for $\underline{l} < L_1$ and the equal share solution for $\underline{l} \geq L_2 + L_3$.

Clearly, the mixture of experiments also affects the suitable profit sharing function, and, as expected, the more diversity-sensitive experiments (in our simple example type 2) the more the Shapley value departs from standard proportional sharing. It is thus important to be able to classify experiments into a few meaningful categories and, based on the expected mixture, adjust the federation policies implemented in practice.

For the case of PlanetLab, we have extensively analyzed user behavior based on available measurement data [23]. This analysis suffered from the opacity of users' utility functions. We have now initiated changes to the PlanetLab interface⁶ to allow users to explicitly select resources on the basis of their properties (geographic location, reliability, etc.) This will allow us to construct more realistic utility functions in the future.

4.3.3 Volume of demand

Now we study how the Shapley value and the proportional share change as a function of the demand volume, expressed as the number of experiments K requesting access to the infrastructure for $\underline{l} = 250$. Fig. 8 plots $\hat{\varphi}_i$, $\hat{\pi}_i$, and $\hat{\rho}_i$ as demand increases for $R_1 = 80$, $R_2 = 60$, and $R_3 = 20$. Notice that now that demand does not necessarily exceed capacity, the proportionally fair share based on consumed resources, $\hat{\rho}$, as defined in (7), is different than $\hat{\pi}$, and thus we include it in our numerical comparison.

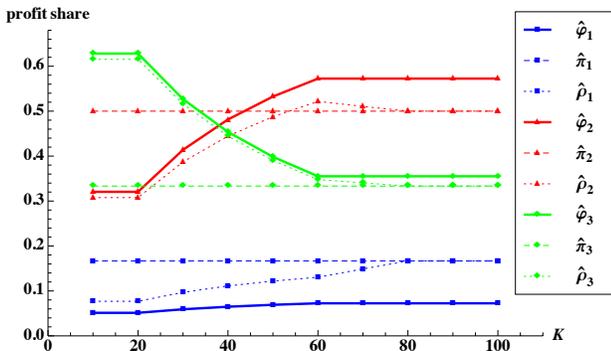


Figure 8: Profit share with respect to K

Notice that while the proportional share $\hat{\pi}_i$ does not depend on the volume of demand, both $\hat{\rho}_i$ and $\hat{\varphi}_i$ do. This means that including this dimension in the profit sharing scheme, as market-based schemes do by construction, can influence significantly the incentives of facilities for participation and resource provision.

To this end, it is important to note that in systems with low demand, evaluating contribution based on consumption, instead of resource availability, might be sub-optimal in terms of incentives for resource provision.

⁶See our MySlice project at <https://trac.myslice.info/>.

The reason is that when demand is low, certain value generators, such as diversity, should be treated as non-rivalrous goods. Understanding in more depth how the outcome of our global game is affected by the inclusion of actual demand in policy design will shed more light on the debate between markets versus rule-based schemes for resource management in P2P systems [16].

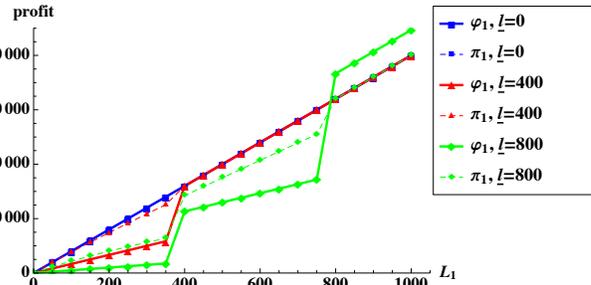


Figure 9: Profit of facility 1 with respect to L_1

4.4 Incentives for resource provision

Using the same configuration, but assuming that demand exceeds capacity, Fig. 9 depicts how the individual profit of facility 1 changes as it increases the number of locations that it covers. This figures shows, as expected, that facilities will have different incentives to upgrade their contribution to the federation depending on the type of demand and on the profit sharing mechanism employed.

More specifically, the Shapley value mechanism creates powerful incentives for resource provision around the threshold points when diversity is important. This is a potential weakness of using the Shapley value since it could cause instability in the system.

However, as already mentioned, we do expect the Shapley value to be used more as an input to the complicated process of policy design rather than an absolute policy parameter calculation tool. For example, one could compute off-line the values $\hat{\varphi}_i$ using some more detailed versions of the examples provided above based on the system characteristics in terms of expected demand. Then these values could be used as generic weights for sharing profits and resource allocation decisions instead of the proportionally fair ones.

5. RELATED WORK

The federation concept applies to a wide set of environments: ISP interconnection, peer-to-peer systems, the Grid, and cloud computing, are all cases in which independent administrative domains can share their local resources and services for their mutual benefit. Numerous studies in the literature try to capture the most important dimensions of these resource exchange economies.

Work on ISP interconnection has focused largely on the economics of peering and transit agreements (e.g., [24, 25]). There are also more theoretical studies on incentive mechanisms for optimal routing (e.g., [26]), and the effects of interconnection policies on the Internet graph (e.g., [27]).

Most relevant to our work is Ma et al.’s (including an author of the current paper) [7] proposal to use the Shapley value to guide sharing of the profit generated by users’ communication with content providers via the federation of backbone and local ISPs. The key characteristic that differentiates ISPs in this context is the redundancy of available paths connecting consumers to producers of content. Implementing such a scheme in practice would require centralized computation of the profit sharing algorithm and coordination of profit distribution. Although this is unrealistic for the Internet as a whole, it is conceivable for federation of today’s experimental facilities.

Aram et al. [8] formulate a similar coalitional game to study the benefits of sharing wireless access between independent providers in a commercial scenario. They prove the existence of the core for a variety of versions of their game including customer ownership as an important contribution of a coalition member. Their model does not address the value of diversity in a way that is relevant for our scenario and different profit sharing schemes are evaluated only in terms of their solution lying in the core of the game when it exists.

On the demand side, the economic models most relevant to our work are Bellagio [4] and GridEcon [5]. Bellagio is a combinatorial auction system designed for allocating PlanetLab resources on the basis of a virtual currency. While combinatorial auctions provide a general approach to meeting diverse requirements, they lead to complex solutions and they do not provide the means to share profits among a set of independent providers whose resources are part of a certain bid.

GridEcon is a spot market based on double auctions, which trades virtual machines for a specific duration. In terms of profit sharing, it has similarities with the parallel synchronized markets proposed for bandwidth-on-demand services [21, 20, 22]. In all these cases, profit between independent organizations is shared implicitly through the market ignoring the possible complementarities in the valuation of the users.⁷ Finally, there are numerous approaches that build on various market-based resource allocation mechanisms [3] and typically trade in CPU cycles.

Market mechanisms are also in general complex to implement and generate inconvenience and uncertainty for the users. Moreover, they are not well suited for the P2P scenario because of the standard issues of vir-

tual currencies (e.g., inflation or deflation) and the fact that when demand is lower than supply they could lead to under-provisioning of resources. For this scenario, Courcoubetis and Weber [6] propose a simple model mostly appropriate for resource hungry experiments and argues that a proportional sharing system based on contribution would lead the system close to the theoretically optimal level. We have showed that when considering the diversity dimension, the normalized Shapley value is a more accurate measure of contribution than a standard proportional sharing scheme.

6. DISCUSSION AND FUTURE WORK

In our work so far we have only scratched the surface of this problem. Our simple models and examples demonstrate that simple contribution- or consumption-based valuation schemes do not properly reflect the value of a player in the federation. We also showed that a Shapley value-based approach leads to a much more fair valuation of resources. The approach leads to quantification of intuitive notions of the value and bargaining power of the players depending on the demand of the users. Diversity was one simple demand parameter that yielded Shapley values that departed significantly from proportional valuations. We believe our approach extends easily to other forms of demand patterns and can lead to economically sound, incentive-compatible profit sharing schemes that can have practical use in the ongoing federation efforts around the PlanetLab facility. They can also be used as a basis to study the global federation game both for the commercial and the P2P cases, and other future Internet scenarios.

To achieve this objective we made a first step toward assessing the importance of various important system parameters to the relative contribution of the members of a federation. Our numerical analysis provides a deeper understanding of how different profit sharing schemes perform under various utility functions, expected minimum requirements of users in terms of diversity, and volumes of demand.

There are many interesting directions in which to extend the analysis and gain valuable insights. For instance, we can develop more sophisticated demand models, use a loss networks formulation, and compute the Shapley value in a manner similar to Paschalidis and Liu [28]. Such an approach can also lead to a more careful analysis of the tight coupling between demand and the value of a resource, as opposed to a simple proportional valuation based on resource contribution or consumption. Integrating complex topologies, hierarchical federations, the formal comparison with market-based approaches for sharing value inside a federation, and the effect of competition in the form of cloud facilities like Amazon EC2 are also avenues for further research.

⁷This is actually the price one has to pay for avoiding the complexity of the combinatorial auctions.

7. ACKNOWLEDGEMENT

The research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement n°224263-OneLab2.

8. REFERENCES

- [1] A. Bavier, M. Bowman, B. Chun, D. Culler, S. Karlin, S. Muir, L. Peterson, T. Roscoe, T. Spalink, and M. Wawrzoniak, "Operating system support for planetary-scale network services," in *Proc. USENIX NSDI*, 2004.
- [2] S. Fdida, T. Friedman, and T. Parmentelat, "OneLab: An open federated facility for experimentally driven future internet research," in *New Network Architectures*, ser. SCI, T. Tronco, Ed. Springer, 2010, vol. 297, pp. 141–152.
- [3] J. Broberg, S. Venugopal, and R. Buyya, "Market-oriented grids and utility computing: The state-of-the-art and future directions," *J. Grid Computing*, vol. 6, no. 3, 2008.
- [4] A. A. Young, B. N. Chun, A. C. Snoeren, and A. Vahdat, "Resource allocation in federated distributed computing infrastructures," in *Proc. OASIS Workshop*, 2004.
- [5] J. Altmann, C. Courcoubetis, G. D. Stamoulis, M. Dramitinos, T. Rayna, M. Risch, and C. Bannink, "GridEcon: A market place for computing resources," in *Proc. GECON Workshop*, 2008.
- [6] C. Courcoubetis and R. Weber, "Economic issues in shared infrastructures," in *Proc. ACM SIGCOMM VISA Workshop*, 2009.
- [7] R. T. Ma, D. Chiu, J. C. Lui, V. Misra, and D. Rubenstein, "On cooperative settlement between content, transit and eyeball Internet service providers," in *Proc. ACM CoNEXT*, 2008.
- [8] A. Aram, C. Singh, S. Sarkar, and A. Kumar, "Cooperative profit sharing in coalition based resource allocation in wireless networks," in *Proc. IEEE INFOCOM*, 2009.
- [9] W. Saad, Z. Han, M. Debbah, A. Hjørungnes, and T. Başar, "Coalitional games for distributed collaborative spectrum sensing in cognitive radio networks," in *Proc. IEEE INFOCOM*, 2009.
- [10] V. Misra, S. Ioannidis, A. Chaintreau, and L. Massoulié, "Incentivizing peer-assisted services: A fluid Shapley value approach," in *Proc. ACM SIGMETRICS*, 2010.
- [11] M. Madiman, "Cores of cooperative games in information theory," *EURASIP Journal on Wireless Communications and Networking*, 2008.
- [12] W. Saad, Z. Han, M. Debbah, A. Hjørungnes, and T. Başar, "Coalitional game theory for communication networks: A tutorial," *IEEE Signal Processing Magazine*, vol. 26, no. 5, 2009.
- [13] L. Shapley, "A value for n-person games," in *Contributions to the Theory of Games II*, H. W. Kuhn and A. W. Tucker, Eds. Princeton Univ. Press, 1953, pp. 307–317.
- [14] C. Elliott, "GENI: Global environment for network innovations," in *Proc. IEEE LCN*, 2008.
- [15] A. Gavras, A. Karila, S. Fdida, M. May, and M. Potts, "Future Internet research and experimentation: The FIRE initiative," *ACM SIGCOMM CCR*, vol. 37, no. 3, 2007.
- [16] P. Antoniadis, T. Friedman, and X. Cuvellier, "Resource provision and allocation in shared network testbed infrastructures," in *Proc. ROADS Workshop*, 2007.
- [17] R. B. Myerson, *Game Theory, Analysis of Conflict*. Harvard Univ. Press, 1991.
- [18] A. Mas-Colell, M. D. Whinston, and J. R. Green, *Microeconomic Theory*. Oxford Univ. Press, 1995.
- [19] L. Peterson, R. Ricci, A. Falk, and J. Chase, "Slice-based Federation Architecture," 2010, working draft V2.0.
- [20] R. Jain and J. Walrand, "An efficient Nash-implementation mechanism for network resource allocation," *Automatica*, vol. 26, no. 8, 2010.
- [21] M. Dramitinos, G. D. Stamoulis, and C. Courcoubetis, "An auction mechanism for allocating the bandwidth of networks to their users," *Computer Networks*, vol. 51, no. 18, 2007.
- [22] L. He and J. Walrand, "Pricing and revenue sharing strategies for Internet service providers," in *Proc. IEEE INFOCOM*, 2005.
- [23] K. Park and V. Pai, "CoMon: A mostly-scalable monitoring system for planetlab," *ACM SIGOPS OSR*, vol. 40, no. 1, 2006.
- [24] S. Shakkottai and R. Srikant, "Economics of network pricing with multiple ISPs," *IEEE/ACM ToN*, vol. 14, no. 6, 2006.
- [25] N. Economides, "The economics of the Internet backbone," in *Handbook of Telecommunications Economics*, S. Majumdar, I. Vogelsang, and M. Cave, Eds. North Holland, 2006.
- [26] J. Feigenbaum, C. Papadimitriou, R. Sami, and S. Shenker, "A BGP-based mechanism for lowest-cost routing," in *Proc. ACM PODC*, 2002.
- [27] J. Corbo, S. Jain, M. Mitzenmacher, and D. C. Parkes, "An economically principled generative model of Internet interdomain connectivity," in *Proc. IEEE INFOCOM*, 2009.
- [28] I. Ch. Paschalidis and Y. Liu, "Pricing in multiservice loss networks: static pricing, asymptotic optimality, and demand substitution effects," *IEEE/ACM ToN*, vol. 10, no. 3, 2002.