Cooperative Network Design: a Nash bargaining solution approach

Konstantin Avrachenkov[a], Jocelyne Elias[b], Fabio Martignon[1,c,*], Giovanni Neglia[a], Leon Petrosyan[d]

[a] *INRIA Sophia Antipolis, France.*
[b] *Paris Descartes University, France.*
[c] *IUF, Institut Universitaire de France*
[d] *St. Petersburg State University, Russia.*

Abstract

The Network Design problem has received increasing attention in recent years. Previous works have addressed this problem considering almost exclusively networks designed by selfish users, which can be consistently suboptimal. This paper addresses the network design issue using cooperative game theory, which permits to study ways to enforce and sustain cooperation among users. Both the Nash bargaining solution and the Shapley value are widely applicable concepts for solving these games. However, the Shapley value presents several drawbacks in this context.

For this reason, we solve the cooperative network design game using the Nash bargaining solution (NBS) concept. More specifically, we extend the NBS approach to the case of multiple players and give an explicit expression for users' cost allocations. We further provide a distributed algorithm for computing the Nash bargaining solution. Then, we compare the NBS to the Shapley value and the Nash equilibrium solution in several network scenarios, including real ISP topologies, showing its advantages and appealing properties in terms of cost allocation to users and computation time to obtain the solution.

Numerical results demonstrate that the proposed Nash bargaining solution approach permits to allocate costs fairly to users in a reasonable computation time, thus representing a very effective framework for the design of efficient and stable networks.

Index Terms: - Network Design, Cooperative Game Theory, Nash bargaining solution, Shapley value.

*Corresponding author, Tel: (+33) 01.69.15.68.16, Fax: (+33) 01.69.15.65.86
Email addresses: K.Avrachenkov@sophia.inria.fr (Konstantin Avrachenkov), jocelyne.elias@parisdescartes.fr (Jocelyne Elias), fabio.martignon@lri.fr (Fabio Martignon), Giovanni.Neglia@sophia.inria.fr (Giovanni Neglia), spbuoasis7@peterlink.ru (Leon Petrosyan)

1. Introduction

The Network Design (ND) problem has become increasingly important given the continued growth of computer networks such as the Internet. The design of such networks is generally carried out by a large number of self-interested actors (users, Internet Service Providers . . .), all of whom seek to optimize the quality and cost of their own operation. In general, for the ND problem we are given a directed graph, where each edge has a nonnegative cost, and a set of players. Each player is identified with a source-destination pair and wants to connect his source to the destination node with the minimum possible cost. Over the past years, the network design problem has been tackled almost exclusively from a non-cooperative point of view. Recent works [1, 2, 3, 4, 5, 6] have modeled how selfish agents can build or maintain a large network by paying for possible edges. Nash equilibria in such games, however, can be much more expensive than the optimal, centralized solution. This is mainly due to the lack of cooperation among network users, which leads to the design of costly networks.

The underlying assumption in all the above works is that agents are completely non-cooperative entities. However, this assumption could be not entirely realistic, for example when network design involves long-term decisions (e.g., in the case of Autonomous Systems peering relations). It is more natural that agents will discuss possible strategies and, as in other economic markets, form coalitions taking strategic actions that are beneficial to all members of the group. Moreover, incentives could be introduced by some external authority (e.g., the network administrator, government authority) in order to increase the users' cooperation level.

Preliminary works, like [7, 8], tried to overcome this limitation by incorporating a socially-aware component in the users' utility functions. This solution, though, can be insufficient to obtain cost-efficient networks in all scenarios. In fact, it has been demonstrated in [8] that, quite surprisingly, highly socially-aware users can form stable networks that are much more expensive than the networks designed by purely selfish users.

To address the above issues, in this paper we first formulate the network design problem as a *cooperative* game, where groups of players (named *coalitions*) coordinate their actions and pool their winnings; consequently, one of the problems is how to divide the cost savings among the members of the formed coalition.

Then, we propose a *Nash bargaining approach*[1] to solve the cooperative network design problem. The Nash bargaining solution (NBS) is a very effective tool to model interactions among negotiators, and is unique for bargaining games satisfying Pareto optimality, symmetry, scale independence, and independence of irrelevant alternatives [9, 10]. More specifically, as a key contribution, we extend the Nash bargaining solution for the cooperative network design problem to the case of multiple players with linear constraints, and give explicit expressions for users' cost allocations, assuming that the disagreement point corresponds to players' costs at Nash equilibrium (the cost for players to connect their source-destination nodes in a purely non-cooperative game). To the best of our knowledge, the derived explicit expressions are new. Our other major contribution, in fact, is the demonstration that our proposed cooperative game theory approach, based on the Nash Bargaining Solution, can be actually applied to large networks.

To complement our study, we further focus on the *Shapley value* concept, which is a widely applied solution for cooperative games, since it provides a unique and fair solution [11]. To compute the Shapley value of the cooperative ND game, we consider in this paper three different (natural) definitions for the characteristic function, which associates with every coalition (a subset of players) a real value respresenting the cost for the coalition. However, we show that the Shapley value presents several limitations in our context: (1) it is non-trivial to define meaningful characteristic functions, (2) the cost allocation determined by the Shapley value can be, in some cases and for some players, even costlier than that obtained at some Nash equilibrium, and (3) for our network design game, it cannot be determined in a reasonable computation time, even when approximation techniques (like that proposed in [12]) are applied.

Finally, we provide a distributed algorithm for computing the Nash bargaining solution. Furthermore, we perform a thorough comparison of the proposed

[1]The Nash bargaining approach studies situations where two or more agents need to select one of the many possible outcomes of a joint collaboration [9, 10]. Each party in the negotiation has the option of leaving the table, in which case the bargaining will result in a disagreement outcome.

Nash bargaining solution with other classic approaches like the Shapley value and the Nash equilibrium solutions, using different, large-scale network scenarios, including real Internet Service Provider (ISP) topologies. Both exact and approximate methods for computing the Shapley value are considered and compared to our approach.

Numerical results demonstrate that our Nash bargaining solution can provide efficient cost allocations in a short computing time, thus representing a very effective tool to plan efficient and stable networks.

The main contributions of this work can therefore be summarized as follows:

- **the formulation of the network design problem as a cooperative game, where players cooperate when connecting their source-destination pairs to reduce their costs.**
- **The proposition of a novel Nash bargaining solution for the n-person cooperative network design problem, which has appealing properties in terms of planning efficient networks and cost allocations in a short computation time.**
- **The proposition of three definitions for the characteristic function and computation of the Shapley value for the cooperative network design game, showing that this solution requires a long computation time for solving large-scale networks even when sampling-based approximation techniques are used.**
- **The construction of a distributed algorithm for computing the Nash bargaining solution.**
- **A thorough comparison of the proposed approach with classic solutions, viz. the Shapley value and the Nash equilibrium concepts, in several realistic and large-size network scenarios, including real ISP topologies.**

The paper is organized as follows: Section 2 discusses related work. Section 3 introduces the cooperative network design game, while Section 4 illustrates the proposed Nash bargaining solution along with a distributed approach we propose for its computation.

Application scenarios are discussed in Section 5. Section 6 presents numerical results that demonstrate the effectiveness of the NBS approach in several realistic network scenarios, including real ISP topologies. Finally, Section 7 concludes this paper.

2. Related Work

The network design problem has been addressed in several recent works, mainly in the context of non-cooperative games [1, 2, 8]. The works in [3, 5, 13, 14, 15] have further considered coordination issues among players.

The so-called *Shapley network design game* is proposed in [1]. In this non-cooperative network design game, each player chooses a path from its source to its destination, and the overall network cost is shared among the players in the following way: each player pays for each edge a proportional share $\frac{c_e}{x_e}$ of the edge cost c_e, where x_e is the number of players that choose such edge. In [8], the Shapley network design game is extended, adding a socially-aware component to users' utility functions.

The survey article in [13] presents the most notable works on network formation in *cooperative games*; furthermore, the existence of networks that are stable against changes in link choices by any coalition is studied in [16]. In [17], Andelman et al. analyze strong equilibria with respect to players' scheduling as well as a different class of network creation games in which links may be formed between any pair of agents. For these latter games, strong Nash equilibria (i.e., equilibria where no coalition can improve the cost of each of its members) achieve a constant Price of Anarchy, which is defined as the ratio between the cost of the worst Nash equilibrium and the social optimum. Strong Nash equilibria ensure stability against deviations by every conceivable coalition of agents. A similar problem is considered in [14], where nodes can collaborate and share the cost of creating any edge in the host graph.

The works in [3, 5] study the existence of strong Nash equilibria in network design games under different cost sharing mechanisms. More specifically, the authors in [3] show that there are graphs that do not admit strong Nash equilibria, and then give sufficient conditions for the existence of approximate strong Nash equilibria.

In [18], the authors investigate a Shapley value-based approach for profit distribution between Internet Service Providers (ISPs), in a network model with three classes of ISPs

(Content, Transit and Eyeball), and study its implications on the stability of prevalent bilateral settlements between ISPs and the pricing structure for differentiated services. However, in this paper, very specific structures of network scenarios are considered: (i) the topology between any two ISPs (content, transit or eyeball) is assumed to be a complete bipartite graph and (ii) transit ISPs form a fully meshed topology. Hence, it is feasible to decompose the initial network into subsystems, and then derive the Shapley revenue distribution for each of the decomposed subnetworks.

On the other hand, we show through extensive numerical analysis that the Shapley value approach has limited applicability in our scenario, since it is often hard to compute and even to approximate, while our approach based on the Nash bargaining solution exhibits low computation time.

Finally, it is worth noting that our proposed solution is complementary to the work in [18]. In fact, once the bilateral agreements and cost/revenue sharing mechanisms are well settled between Content, Transit and Eyeball ISPs in the substrate networks, our solution applies directly to any virtual/overlay network built upon such CTE model, and provides an effective tool to share costs among network users.

The idea of using the Nash bargaining solution in the context of telecommunication networks has been considered in different networking scenarios [19, 20, 21, 22, 23, 24]. Such approach was first presented for packet-switched data networks by Mazumdar et al. [19]. The concept of Nash bargaining solution is used by Yaiche et al. [20] to derive a price-based resource allocation scheme. In [21] the authors propose a scheme to allocate subcarrier, rate, and power for multiuser orthogonal frequency-division multiple-access systems. The approach considers a fairness criterion, which is a generalized proportional fairness based on Nash bargaining solutions and coalitions.

The work in [25] studies the application of cooperative game theory to the routing problem in a *parallel links* networking context, and focuses on the NBS as solution concept for cooperative networking games. The existence and uniqueness of a solution to the NBS is guaranteed under mild conditions. Finally, several performance measures (Price of Anarchy, Price of Stability, and Price of Heterogeneity (PoH)) are used to evaluate the performance of the NBS; PoH is introduced in the case of heterogeneous players with different objective functions.

The reader is referred to the next section, to the book by Muthoo [9] and the paper by Nash [10] for a general introduction to the Nash bargaining solution concept.

3. Cooperative Network Design Game: Definition and Shapley Value solution

This section illustrates the cooperative network design game considered in this work, and provides a review of the Shapley value approach for comparison reasons.

3.1. Network Model

We are given a directed graph $G = (V, E)$, where each edge e has a nonnegative cost c_e; each player $i \in \mathcal{I} = \{1, 2, \ldots n\}$ is identified with a source-destination pair (s_i, t_i), and wants to connect his source to the destination node with the minimum possible cost. Note that c_e represents the *total* edge cost, which is shared among the players according to the allocation algorithms we will describe in the following.

We consider a cooperative game in strategic form $G = \langle \mathcal{I}, A, \{J^i\} \rangle$, where $\mathcal{I}$ is the set of players, A_i is the set of actions for player i, $A = A_1 \times \ldots \times A_n$, and J^i is the objective (cost) function, which player i wishes to optimize (minimize).

In a cooperative game, players bargain with each other before the game is played. If an agreement is reached, players act according to such agreement, otherwise players act in a non-cooperative or antagonistic way. Note that the agreements must be binding, so players are not allowed to deviate from what is agreed upon.

3.2. The Shapley value solution

We now review the Shapley value solution approach, and discuss meaningful definitions for the characteristic function.

The Shapley value is a widely applied concept for solving cooperative games. It is a possible way to allocate the total costs (or "values") among the members of a coalition, taking into account their different importance for the coalition. The main advantage of the Shapley value is that it provides a solution that is both unique and fair: it is unique in the class of subadditive cooperative games (see definition below); it is fair in a sense that it satisfies a series of axioms intuitively associated with fairness (see [11]). However, while these are both desirable properties, the Shapley value has one major drawback:

for many coalition games, including our network design game, it cannot be determined in a reasonable time. We shall discuss computational aspects, along with approximation methods used to reduce such complexity, in more detail below.

A Shapley function ϕ is a function that assigns to each possible characteristic function v a vector of real numbers, i.e., $\phi(v) = [\phi_1(v), \ldots, \phi_i(v), \ldots, \phi_n(v)]$, where $\phi_i(v)$ represents the cost of player i in the game.

The characteristic function, v, is a real-valued function that associates with every non-empty subset $\mathcal{S}$ of $\mathcal{I}$ (i.e., a coalition) a real number $v(\mathcal{S})$, the cost of $\mathcal{S}$; $v(\mathcal{S})$ must satisfy the following properties[2]:

1. $v(\emptyset) = 0$.
2. (Subadditivity) if $\mathcal{S}$ and $\mathcal{T}$ are disjoint coalitions ($\mathcal{S} \cap \mathcal{T} = \emptyset$), then $v(\mathcal{S})+v(\mathcal{T}) \geq v(\mathcal{S} \cup \mathcal{T})$.

This latter property means that cooperation can only help but never hurt.

Note that defining the characteristic function is not straightforward for the cooperative network design game considered in this work, since a "natural" definition can violate the subadditivity property, as we will discuss in the following.

The three definitions reported hereafter "naturally" arise in our networking problem as *candidate* characteristic functions:

1. Players in $\mathcal{S}$ and players in $\mathcal{I} - \mathcal{S}$ form two separate coalitions. Each coalition tries to minimize the total cost for its members, taking into account the selfish behavior of the other coalition. A Nash equilibrium is reached, and $v(\mathcal{S})$ is defined as the total cost for members in $\mathcal{S}$ at this equilibrium[3].
2. The value of the coalition $\mathcal{S}$ is defined as its *security level*, i.e. as the minimum total cost that $\mathcal{S}$ can guarantee to itself when members in $\mathcal{I} - \mathcal{S}$ act collectively in order to maximize the cost for $\mathcal{S}$.
3. The value of coalition $\mathcal{S}$ is equal to the minimum cost that its members would incur if players in $\mathcal{I} - \mathcal{S}$ would be absent.

[2]The second one is required to guarantee the uniqueness of the Shapley value solution.

[3]If multiple such equilibria exist, we consider the one reached starting from the empty network using best response dynamics.

We note that, in our specific game, these three definitions give increasing value to a coalition $\mathcal{S}$. In fact, when players in $\mathcal{I}-\mathcal{S}$ minimize their own cost (first definition), their path choices cannot be as bad for $\mathcal{S}$ as when they try to maximize the cost for $\mathcal{S}$ (second definition). Still, when players in $\mathcal{I}-\mathcal{S}$ are present, they are obliged to select paths to connect their source-destination pairs, and some of these links may also be used by players in $\mathcal{S}$, so that $v(\mathcal{S})$ is smaller in the second definition than in the third.

To better illustrate the differences underlying these definitions, let us consider the hexagon network scenario of Figure 1, with 6 links and 3 players having the following source-destination pairs: (s_1, t_1), (s_2, t_2) and (s_3, t_3). All link costs are equal to 1, except for link $t_3 \rightarrow t_2$, which has a cost equal to $1-\epsilon$, ϵ being a very small constant.

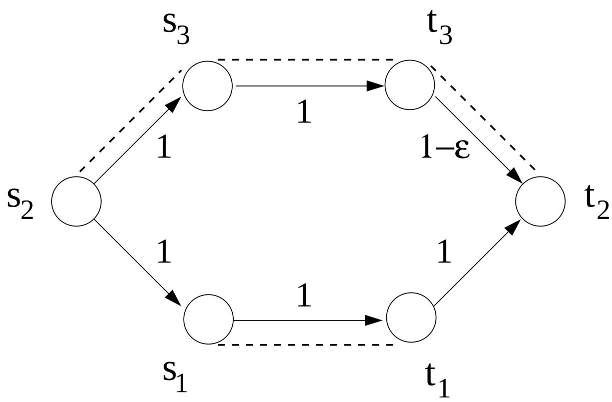

Figure 1: Hexagon network topology: the 3 players must connect their source-destination nodes (s_i, t_i). The optimal solution, which in this case coincides with both the Nash equilibrium point and the Nash bargaining solution, is illustrated with dashed lines.

Table 1: Hexagon network scenario: characteristic function values, $v(\mathcal{S})$, for definitions (1), (2) and (3).

	Characteristic Function value ($v(\mathcal{S})$)		
Coalition ($\mathcal{S}$)	Definition (1)	Definition (2)	Definition (3)
$\emptyset$	0	0	0
1	1	1	1
2	2.5-ϵ	2.5-ϵ	3-ϵ
3	0.5	1	1
12	3	3	3
13	1.5	1.5	2
23	3-ϵ	3-ϵ	3-ϵ
123	4-ϵ	4-ϵ	4-ϵ

Table I reports, for each of the three above definitions, the corresponding candidate characteristic function values. It can be easily checked that definition (1) does *not* lead to a characteristic function, since the subadditivity property is not satisfied (for example, $v(12) + v(3) < v(123)$), and therefore it cannot be used to compute Shapley values. Indeed, with such definition, cooperation among players can lead to costlier solutions. On the other hand, definitions (2) and (3) lead to characteristic functions.

Theorem 1. *In the Cooperative Network Design Game, the security level (definition 2) and the minimum cost of the coalition (definition 3) satisfy the axioms of a characteristic function.*

Proof: See the Appendix.

To calculate the Shapley value, suppose we form the grand coalition (the coalition containing all n players) by entering the players into this coalition one at a time. As each player enters the coalition, he is charged the cost by which his entry increases the cost of the coalition he has entered. The cost a player pays by this scheme depends on the order in which the players enter. The Shapley value is just the average cost charged to the players if they enter in a completely random order, i.e.,

$$\phi_i = \sum_{\mathcal{S} \subset \mathcal{I}, i \in \mathcal{S}} \frac{(|\mathcal{S}| - 1)!(n - |\mathcal{S}|)!}{n!} [v(\mathcal{S}) - v(\mathcal{S} - \{i\})]. \quad (1)$$

It can be proved that the problem of computing the Shapley value is an NP-complete problem. Polynomial methods, based on sampling theory, have been proposed in the literature for approximating the Shapley value [12, 26]; these estimations, though, are efficient only if the worth of any coalition $\mathcal{S}$ can be calculated in polynomial time, which is not the case for our problem.

In fact, we will show in Section 6.4 that even using the approximation methods proposed for example in [12], it is necessary to compute the worth of an extremely large number of coalitions, which is computationally very cumbersome, while as we see next, our proposed Nash bargaining solution needs only computing the worth of the grand coalition. Furthermore, we will demonstrate that the approximate solution obtained with the sampling technique is quite far from the optimal solution in several network scenarios.

4. Cooperative Network Design Game: Nash Bargaining Solution (NBS)

Since the computation time of the Shapley value can be extremely long in network scenarios with many players, in this paper we consider another approach to cooperative network design: Nash bargaining. We will show that the computation of the Nash bargaining solution is very light.

Let u_i denote the maximal acceptable cost that user i is willing to pay. In the present work we suggest the three following options:

1. the cost for user i to connect its source-destination nodes in a purely non-cooperative game (i.e., the Nash equilibrium solution);
2. the cost for user i to connect its source-destination nodes in a zero-sum game where all the other players are trying to maximize the cost of user i;
3. the cost for user i to connect its source-destination nodes when there is no other player.

The vector $u = \{u_1, u_2, \ldots u_n\}$ is also denoted as the *disagreement point* of the cooperative network design game (i.e., what will happen if players cannot come to an agreement). Clearly, the cost achieved by every player at any agreement point (every possible outcome of the bargaining game) has to be at most equal to the cost achieved at the disagreement point.

To determine the disagreement point u using the Nash equilibrium solution (option 1), we can consider the non-cooperative, potential network design game in [1, 27] and apply Best Response Dynamics, which is guaranteed to converge to a Nash Equilibrium Point. In the worst case, such dynamics can require an exponential time to converge. Indeed, it has been shown in [1] that the game converges to a Nash equilibrium in polynomial time for the case of two players, but that with n players, it can run for a time exponential in n. However, we prove hereafter that in our case the Best Response algorithm converges to an ϵ-Nash equilibrium in at most $n \cdot c_{max}|E|(\ln(n)+1)/\epsilon$ moves, where $|E|$ is the number of edges, $c_{max} = \max_{e \in E} c_e$, and c_e is the cost of edge e.

Let us first define the ϵ-Nash equilibrium. A strategy profile S corresponds to the ϵ-Nash equilibrium, if

$$J^i(S_i', S_{-i}) \leq J^i(S_i, S_{-i}) - \epsilon. \quad (2)$$

For our game, the potential function is given by ([1, 27]):

$$\Phi(S) = \sum_{e \in E} \sum_{x=1}^{x_e} \frac{c_e}{x}, \tag{3}$$

where x_e and c_e represent, respectively, the number of players' paths that go through edge e, and the cost of this latter.

At each improvement we decrease the potential function by at least ϵ. If no player can make a move decreasing the potential by at least ϵ, we stop and the reached profile corresponds to the ϵ-Nash equilibrium. Let us denote by $S^{(0)}$ the initial profile. Then, it will take no more than $n \cdot \Phi(S^{(0)})/\epsilon$ steps to achieve the ϵ-Nash equilibrium, where n is the number of players. We can suggest a bound for $\Phi(S^{(0)})$:

$$\Phi(S^{(0)}) \leq \sum_{e \in E} c_e(\ln(x_e) + 1) \leq c_{max}|E|(\ln(n) + 1),$$

with $c_{max} = \max_{e \in E} c_e$.

Therefore, we achieve the ϵ-Nash equilibrium in at most $n \cdot c_{max}|E|(\ln(n)+1)/\epsilon$ moves.

Furthermore, in our simulation campaign, we verified in practice that in all considered scenarios, Best Response Dynamics always converges to an equilibrium point in less than 10 iterations. Indeed, in the large majority of our considered topologies, only 5 iterations were necessary, in the worst case, to converge to a NEP starting from the empty network.

We now derive a Nash bargaining solution for allocating the total network cost to users. To this aim, we extend the well-known two-player NBS concept to the n-player network design game, considering transferable network costs, providing explicit expressions. This assumption means that the players or the system administrator can redistribute the total cost among the players.

Let u_{soc} denote the total network cost resulting from social optimization. This can be computed, for example, formulating the Generalized Steiner Tree (GST) problem [28] with an Integer Linear Program, using a mathematical programming model (like AMPL [29]), and solving it with a commercial solver (like CPLEX [30]). Furthermore, very efficient, polynomial-time approximation algorithms have been proposed [31, 32] to solve the GST problem in a reasonable computation time, even in a distributed, online fashion [33]. Solving such problem provides the least-cost network topology that connects all source-destination pairs.

Then, the Nash bargaining solution can be given in explicit form.

Theorem 2. *The Nash bargaining solution for player i, α_i is given by the following expression:*

$$\alpha_i = u_i - \frac{\sum_k u_k - u_{soc}}{m}, \tag{4}$$

where m coincides with the number of players n (i.e., $m \equiv n$) if we allow for negative costs (i.e., some α_i values are negative, which means that some players are actually paid to ensure their participation). Otherwise, if only non-negative costs are allowed (or equivalently, if no positive transfers are permitted), m is defined as the largest integer for which the following inequality is satisfied:

$$\frac{1}{m-1}(\sum_{i=1}^{m-1} u_i - u_{soc}) < u_m \tag{5}$$

having assumed, without loss of generality, that players are ordered such that $u_1 \geq u_2 \geq \ldots \geq u_n$.

Proof: See the Appendix.

We would like to emphasize that in the first case α values can be positive or negative, while in the second case α values are non-negative. In particular, m gives the number of non-zero α values, i.e., $\alpha_1, \alpha_2, \ldots \alpha_m$ are positive and given by expression (4), while $\alpha_{m+1}, \ldots \alpha_n$ are equal to zero.

4.1. Distributed algorithm for computing the NBS

We now outline a distributed algorithm for computing the Nash bargaining solution, which is detailed in Algorithm 1.

First of all, we would like to stress that, indeed, a possible architecture is the one that assumes a centralized implementation, where a single entity in the coalition collects all necessary information, performs the cost allocation and finally broadcasts it to all coalition members. In this section, however, we propose a fully distributed solution for the NBS approach, in order to obtain a more robust architecture without a single point of

Algorithm 1: Distributed Algorithm executed by each player i

Input : $\mathcal{I}$, $G = (V, E)$, $\{c_e, \forall e \in E\}$
Output: α_i

Compute u_i **as in Section 4;**
Broadcast u_i **to all players** $j \in \mathcal{I}, j \neq i$**;**
$u_{soc} = compute\_total\_network\_cost(\mathcal{I}, G, \{c_e\})$ **according to [28];**
if ***negative costs are allowed*** **then**
    $m = |\mathcal{I}|$**;**
else
    m **= largest integer such that:** $\frac{1}{m-1}(\sum_{i=1}^{m-1} u_i - u_{soc}) < u_m$**;**
end
$\alpha_i = u_i - \frac{\sum_k u_k - u_{soc}}{m}$**;**
Return α_i**;**

failure, and at the same time, reduce the memory size requirements and communication load/overhead.

The minimal requirements $u_i, \forall i \in \mathcal{I}$, defined in the previous section, can be easily computed. In the first two cases, it can be observed that the computation of the Nash equilibrium solution for a non-cooperative network design game can be performed in a distributed way, since the best-response dynamics is guaranteed to converge to a pure Nash equilibrium. If $|E|$ is the number of edges in the network, each player i, after choosing the best response path, must send to all other players a link utilization vector of size $|E|$ bits. This vector simply indicates, for each edge $e \in E$, if it belongs or not to the path chosen by player i to connect its source-destination pair. As for the third definition, the problem is simply to compute the least-cost path between all source-destination pairs, for which there exist several distributed algorithms, with known complexity [34].

In order to compute the Nash bargaining solution given by expression (4), it is necessary to find the cost of the socially optimal network (u_{soc}). To this aim, the technique proposed in [28] can be used. This is one of the first distributed algorithms proposed for the Generalized Steiner tree problem; it is a probabilistic algorithm with $O(\log n)$ expected approximation, based on a probabilistic tree embedding due to Fakcharoenphol et al. [35]. It has been demonstrated that the number of messages exchanged using such technique is upper-bounded by $O(|V|^{\frac{3}{2}} log(|V|))$, being V the set of nodes.

Then, each player i to calculate the NBS solution α_i needs only to know

the disagreement point u of the cooperative game; this can be simply achieved by each player broadcasting its minimal requirement u_i to all other players. If we consider the ϵ-Nash equilibrium concept discussed before, it is not hard to see that the overall signaling overhead is upper-bounded by $|V||V-1|n \cdot c_{max}|E|(\ln(n)+1)/\epsilon + O(|V|\sqrt{|V|}log(|V|))$, where n is the number of players, $|V|$ and $|E|$ are the number of nodes and edges, respectively, and $c_{max} = \max_{e \in E} c_e$ is the maximum cost among edge costs (c_e).

5. Application Scenarios

This section illustrates some notable application scenarios that can be envisaged for our proposed game. First of all, we underline that the application scenario we envisage more naturally for our proposed solution is that of *long term contracts* between network users and virtual network/service providers, since money transfers and long-term relationships between these two actors are involved. However, note that even if the underlying network dynamics is changing rapidly (due for example to topology changes, link failures, traffic changes and users' arrivals/departures, among others) the low complexity of our approach allows us to recompute on-the-fly the compensation for each player, and thus to adapt to such dynamics.

Our model directly applies to Virtual Networks/Service Overlay Networks [27, 36], as illustrated in Figure 2, which are application-layer networks built on top of the tradi-

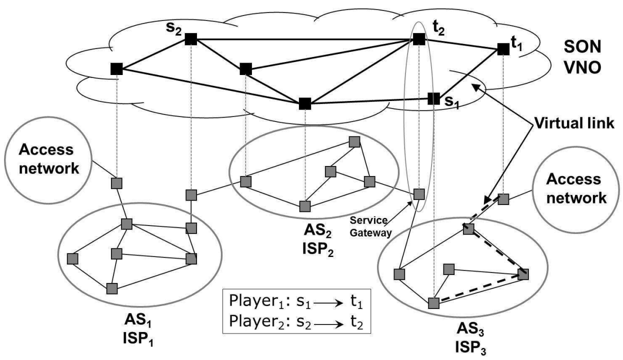

Figure 2: Application scenario: a Virtual Network (or Service Overlay Network), built on top of multiple Autonomous Systems (ASs) owned by different Internet Service Providers (ISPs), serves two players wishing to connect two terminals ($s_1 - t_1$, $s_2 - t_2$).

tional IP-layer network. In general, these networks are operated by a third-party Virtual Network operator, that owns a set of overlay nodes residing in the underlying ISPs' domains. The virtual operator must establish overlay links, purchasing them from the underlying ISPs. In these networks, virtual operators provide a service to users through the creation of an overlay.

A specific application, which will be detailed in the following, is selling connectivity with some specific guarantees. In such scenario the virtual (overlay) operator would buy bandwidth from the underlying ISPs and sell it to the users (e.g., companies) to connect their sites (for example, creating Virtual Private Networks). In another application scenario, users/players could further represent research institutes wishing to cooperate in sharing network resources in order to build a federation network to enable efficient distribution of large scientific data sets (e.g., in grid or cloud computing contexts). In all these cases, the users that make up such networks are basically cooperative (at an extreme case, they belong to the same administrative authority).

Then, the costs considered in our network model (c_e) are mainly the costs required for reserving some bandwidth from the underlying ISPs (plus the operating/management costs), increased by a given percentage to provide the overlay operator some revenues. We can expect that negotiating with the underlying ISPs would have some serious constraints in terms of (i) granularity of the purchased bandwidth [36] and (ii) the possibility to re-negotiate the contract on short time scales. Therefore, the overlay operator has interest to let the users group together, and share as much as possible the same overlay links in order to (i) reduce the number of contracts it has to manage with the underlying ISPs, and hence the management costs, and (ii) take advantage of the multiplexing gain. At the same time, inspired by the emerging group-buying services on the Internet, e.g., Groupon [37], network users have an incentive to voluntarily group together (in a *cooperative* way) to acquire and share the bandwidth/connectivity sold by virtual operators. In our work, we show that the advantage can indeed be consistent in all considered network scenarios. Hence, users could interact with each other in order to reduce their costs:

- One possibility is to let each user pay for each overlay link a share inversely proportional to the total number of users that pass through such link, or make users more altruistic and socially aware in order to improve the system's performance; each user pays for his connectivity service plus a given percentage of the total network cost. However, we showed in [27] that this solution can be insufficient to obtain cost-efficient networks in all scenarios.

- Another more effective possibility, proposed in this work, is to let users cooperate among themselves, using Bargaining-based cost sharing (or Shapley value) methods, thus further reducing their costs as well as the management complexity for the virtual operator.

Finally, another example of application scenario is represented by a physical infrastructure owned by a physical operator, and such infrastructure is shared by a set of virtual operators who aim at building their networks upon the physical infrastructure, minimizing their costs [38, 39]. These virtual operators can cooperate in order to reduce their costs, using our proposed NBS-based game. The same idea applies in a Cloud environment, where IaaS (Infrastructure as a Service) users' requests are mapped to the physical cloud infrastructure.

6. Numerical Results

This section reports the numerical results obtained applying our proposed Nash bargaining solution to cooperative network design games played in various network scenarios, including simple network instances and more general network topologies. To this end, we consider both randomly generated network instances and real ISP topologies mapped by the Rocketfuel tool [40].

The NBS, computed as suggested in the previous section, is compared both to the cost allocation provided by the Shapley value, as well as to a Nash equilibrium solution. This latter is determined in the non-cooperative network design framework proposed in [1], revised in Section 2, starting from the empty network and using a best response algorithm where each user greedily minimizes its path cost until an equilibrium is reached. Since it has been demonstrated in [1, 27] that such non-cooperative game is a potential game, Best Response Dynamics is guaranteed to converge to a Nash Equilibrium Point.

We assume that positive transfers are allowed. To compute the Shapley value, unless stated otherwise, we assume that the worth of a coalition $\mathcal{S}$ is the minimum cost that its members would incur if players in $\mathcal{I} - \mathcal{S}$ would be absent (definition 3). This allows us to consider the "worst case", i.e. the costlier definition for a coalition, as discussed before. We emphasize that definition 1 of the candidate characteristic function in Section 3.2 does not lead to a well-defined characteristic function which satisfies the axiom of subadditivity, and therefore cannot be used.

With the Nash Bargaining solution, on the contrary, any of the three definitions of the disagreement point in Section 4 can be used. We argue that the first definition, where the maximal acceptable cost is defined as a cost at the Nash equilibrium solution, is the most *natural* choice in the network design game setting. In fact, it is hard to imagine that in a large network with many players, all players will be able to agree to play against a single player or even to withdraw from the network, risking to lose altogether their revenues. Not following the recommendations and playing on their own appears to be a more likely outcome if the recommendations do not result in benefits for the players. Furthermore, as will be consistently demonstrated by all our examples, the Shapley value solution exhibits *instability* with respect to the Nash equilibrium solution. In all our examples there are players which have higher costs under the Shapley value solution than under the NBS with the disagreement point selected according to definition 1. This is the first reason why we suggest to use the NBS solution rather than the Shapley value. The second reason against the use of the Shapley value is the impracticality of its computation. We demonstrate that in the network design game the Shapley value can be computed in feasible time only for very small network sizes. In contrast, the NBS solution scales very well with network size and can be computed in a distributed fashion.

6.1. Simple Network Scenario

Let us first consider the simple network scenario already illustrated in Figure 1, with 6 links and 3 players. Table I already reported the characteristic function values $v(\mathcal{S})$, for all possible coalitions $\mathcal{S}$. The optimal network cost is here $u_{soc} = 4 - \epsilon$, and coincides with the cost of the network formed at the Nash Equilibrium Point (NEP). The Nash equilibrium and the Shapley value solutions for this scenario are reported in Table II, together with the Nash bargaining solution, which was computed using all the three proposed definitions for the disagreement point u_i.

Table 2: Hexagon network scenario with 3 players. The table reports the cost paid by each player at the Nash Equilibrium Point, the Shapley value and the Nash bargaining solution. The NBS is computed using all the three proposed definitions for the disagreement point u_i. The total network cost is equal to $4 - \epsilon$ for all allocation algorithms.

Algorithm	(s_1, t_1)	(s_2, t_2)	(s_3, t_3)
NEP	1	2.5 - ϵ	0.5
Shapley value	$\frac{5+\epsilon}{6} \approx 0.83$	$\frac{14-5\epsilon}{6} \approx 2.33$	$\frac{5-2\epsilon}{6} \approx 0.83$
NBS (def. 1)	1	2.5 - ϵ	0.5
NBS (def. 2)	$\frac{5}{6}$	$\frac{14}{6} - \epsilon \approx 2.33$	$\frac{5}{6}$
NBS (def. 3)	$\frac{2}{3}$	$\frac{8}{3} - \epsilon \approx 2.67$	$\frac{2}{3}$

We see that the solution given by the Shapley value for player 3 ($\frac{5-2\epsilon}{6} \approx 0.83$) is costlier than that of the Nash equilibrium, 0.5. We further observe that even defining the value of a coalition as its security level (definition 2, Section 3.2) leads to the same Shapley values reported in Table II. As a consequence, the Shapley value solution is somehow *unstable* for all the considered definitions of the characteristic function, since some players (i.e., player 3 in this scenario) can deviate to reduce their cost. This is surprising, because the Shapley value satisfies the *individual rationality* property, so that the Shapley value allocation is always preferable for each player than playing alone. The apparent paradox originates from the fact that the value of the single player coalition has been defined either as the cost incurred if all other players are absent (definition 3), or as its security value, considering that all the other players are trying to maximize its cost (definition 2). In reality, at the Nash equilibrium, the cost of player i is smaller than such values and, as we have shown, it can be even smaller than the Shapley value imputation.

At the same time, the NBS computed using as disagreement point both definition 3 (the cost for user i to connect its source-destination when there is no other player) and definition 2 (the cost for user i to connect its source-destination in a zero-sum game where all the other players try to maximize the cost for such user) exhibit the same *instability* discussed before, since in both cases at least one player has a costlier solution that at the Nash equilibrium, and can therefore deviate and reduce its cost.

On the other hand, if we assume that the disagreement point is the cost for user i to connect its source-destination nodes in a purely non-cooperative game (i.e., the Nash equilibrium solution), then the NBS coincides in this case with both the NEP and the optimal solution. Indeed, definition 1 appears to be the most natural in the context of network design games, since a selfish non-cooperation is a much more natural reaction for players than forming a coalition against a single player or withdrawing from the game.

6.2. Random Topologies

To study the Nash Bargaining solution behavior in more general topologies, we consider random network scenarios generated as follows: we randomly extract the position of N nodes, uniformly distributed on a square area with edge equal to 1000. As for the network links, which can be bought by players to connect their endpoints, we consider random geometric graphs, where links exist between any two nodes located within a range R. The link cost is set to its length.

Tables III, IV and V illustrate the results obtained in a random geometric graph scenario with 50 nodes, range $R = 500$ (which means approximately more than 1200 links) and, respectively, 5, 10 and 15 source-destination pairs (players). The tables report the costs for the players reached at the Nash equilibrium, the Shapley value as well as our proposed Nash bargaining solution. The total network cost is reported in the last column; note that such value corresponds, for the Shapley value and the Nash bargaining allocation algorithms, to the socially optimal solution (u_{soc} parameter), which can be obtained as explained in Section 3.

Table 3: Random geometric network scenario with $n = 5$ players. The table reports the cost paid by each player at the Nash Equilibrium Point, the Shapley value and the Nash bargaining solution. The total network cost is also reported.

Algorithm	P1	P2	P3	P4	P5	Total cost
NEP	299.1	149.4	400.0	824.4	580.3	2253.1
Shapley value	298.9	**167.5**	380.4	817.3	**589.0**	2253.1
NBS	299.1	149.4	400.0	824.4	580.3	2253.1

It can be observed that, in all scenarios, at least 2 players (marked in bold in the tables) have a Shapley value that is higher than the Nash equilibrium cost. However, the cost saving between the NEP and the optimal cost (which is approximately 700 and 1250 for the $n = 10$ and $n = 15$ scenarios, respectively) could be re-distributed, which is what the Nash bargaining solution does, increasing the appeal of the cost sharing solution.

Table 4: Random geometric network scenario with $n = 10$ players. The table reports the cost paid by each player at the Nash Equilibrium Point, the Shapley value and the Nash bargaining solution. The total network cost is also reported.

Algorithm	P1	P2	P3	P4	P5	P6	P7	P8	P9	P10	Total cost
NEP	283.0	149.4	235.3	824.4	714.8	450.5	674.0	195.6	186.0	266.9	3979.9
Shapley value	260.5	**170.6**	**253.9**	717.1	472.8	387.5	508.0	142.5	183.8	175.6	3272.2
NBS	212.3	78.6	164.6	753.6	644.0	379.7	603.2	124.8	115.3	196.2	3272.2

Table 5: Random geometric network scenario with $n = 15$ players. The table reports the cost paid by each player at the Nash Equilibrium Point, the Shapley value and the Nash bargaining solution. The total network cost is also reported.

Algorithm	P1	P2	P3	P4	P5	P6	P7	P8	P9	P10	P11	P12	P13	P14	P15	Total cost
NEP	283.0	149.4	235.3	824.3	626.6	417.7	566.6	195.6	186.0	133.5	554.8	140.2	94.6	191.5	476.8	5076.0
Shapley value	247.8	**160.2**	198.5	620.1	403.3	319.9	362.2	131.8	173.0	123.9	359.1	**160.5**	**100.2**	188.3	271.8	3820.7
NBS	199.3	65.7	151.6	740.7	542.9	334.0	482.9	111.9	102.3	49.8	471.1	56.5	10.9	107.9	393.2	3820.7

Obviously, since both the Shapley value and the NBS distribute the social cost (u_{soc}) among the players, there will be players whose allocation is costlier under the NBS than with the Shapley value allocation. This happens, in the numerical examples we considered, for players that have a large cost at the Nash equilibrium. However, every player is always better off under the NBS allocation than at the Nash equilibrium, since cost savings are redistributed.

Furthermore, we observe that computing the Shapley value for $n = 15$ players took several weeks of computation on the workstation used to obtain the numerical results reported in this paper, i.e., an Intel Pentium 4 (TM) processor with CPUs operating at 3 GHz and with 1024 Mbyte of RAM. Therefore, computing the Shapley value for a larger number of players is practically infeasible in such network scenario. On the other hand, our proposed n-person Nash bargaining solution is very simple to calculate, and could be computed within a few minutes in all considered network scenarios, thus representing a practical and efficient solution to the network design problem.

6.3. Real ISP topologies

We further consider three real ISP topologies mapped using Rocketfuel [40], listed in Table VI, with an increasing number of nodes and links. The link costs are those provided by Rocketfuel, which are equal to the delay experienced on each link. For each topology we performed a random selection of n source-destination pairs, with $n \in \{10, 15\}$.

In our scenario, and according to the real world applications we outlined in Section 5, this cost can be considered as the bandwidth plus management cost faced by players to create a virtual/overlay network between the desired endpoints. In this regard, the cost parameter c_e can be considered as a cost per time unit (hence it could be expressed, for example, as USD per month).

Tables VII and VIII show the results obtained in such topologies with 10 and 15 players, that is, the costs for the players reached at the Nash equilibrium, the Shapley value as well as our proposed Nash bargaining solution. The total network cost is also reported in the last column. Note that it was impossible to compute the Shapley value for the Abovenet topology (for $n = 15$ players), due to the large number of nodes and links.

It can be observed that in the case of Sprintlink and Abovenet topologies, respectively,

Table 6: Rocketfuel-inferred ISP topologies: number of network nodes and links.

Network	Location	Nodes	Links
Telstra	AU	108	306
Sprintlink	US	141	748
Abovenet	US	315	1944

Table 7: Rocketfuel-inferred ISP topologies with $n = 10$ players. The table reports the cost paid by each player at the Nash Equilibrium Point, the Shapley value and the Nash bargaining solution. The total network cost is also reported.

Algorithm	P1	P2	P3	P4	P5	P6	P7	P8	P9	P10	Total cost
Telstra											
NEP	7.4	3.5	30.4	8.2	13.4	8.0	8.9	11.4	17.2	6.4	115
Shapley value	7.4	3.5	30.4	6.9	13.4	7.4	8.9	10.8	16.9	6.4	112
NBS	7.1	3.2	30.1	7.9	13.1	7.7	8.6	11.1	16.9	6.1	112
Sprintlink											
NEP	13.4	24.3	21.3	31.0	3.5	11.2	56.0	4.0	46.5	37.8	249
Shapley value	**14.3**	22.0	20.3	27.4	3.4	10.1	52.9	4.0	46.0	35.6	236
NBS	12.1	23.0	10.0	29.7	2.2	9.9	54.7	2.7	45.2	36.5	236
Abovenet											
NEP	8.0	16.5	52.5	10.0	38.0	16.5	74.0	65.5	4.0	10.0	295
Shapley value	7.3	13.8	50.6	8.3	37.5	14.4	70.5	**66.1**	3.6	7.0	279
NBS	6.4	14.9	50.9	8.4	36.4	14.9	72.4	63.9	2.4	8.4	279

Table 8: Rocketfuel-inferred ISP topologies with $n = 15$ players. The table reports the cost paid by each player at the Nash Equilibrium Point, the Shapley value and the Nash bargaining solution. The total network cost is also reported.

Algorithm	P1	P2	P3	P4	P5	P6	P7	P8	P9	P10	P11	P12	P13	P14	P15	Total cost
Telstra																
NEP	6.2	2.2	30.0	8.2	13.0	6.5	6.8	9.7	10.8	5.2	22.0	16.7	9.5	11.0	10.2	168
Shapley value	6.2	2.2	**30.1**	6.3	**13.1**	6.3	6.8	8.6	10.7	5.2	22.0	16.4	9.5	11.0	8.7	163
NBS	5.8	1.8	29.7	7.9	12.7	6.2	6.5	9.3	10.5	4.8	21.7	16.4	9.2	10.7	9.8	163
Sprintlink																
NEP	9.8	15.2	16.4	30.7	3.0	9.1	28.0	4.0	24.4	34.0	18.4	14.2	31.5	17.1	24.0	280
Shapley value	**10.7**	13.9	15.5	25.4	2.7	7.9	26.6	3.9	24.4	30.8	15.5	12.8	29.7	8.5	**24.5**	253
NBS	8.0	13.4	14.6	28.9	1.2	7.3	26.2	2.2	22.6	32.2	16.6	12.4	29.7	15.3	22.2	253
Abovenet																
NEP	8.0	8.8	50.8	9.3	23.0	12.8	69.3	65.5	4.0	10.0	23.0	59.5	10.8	4.0	39.0	398
Shapley value	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
NBS	5.9	6.7	48.7	7.2	20.9	10.7	67.2	63.4	1.9	7.9	20.9	57.4	8.7	1.9	36.9	366

two players have a Shapley value that is higher than the Nash equilibrium cost. On the other hand, in all three topologies, the NBS redistributes fairly the gap between the total cost under the Nash solution and the social optimum among the players.

As argued previously, the NBS can be computed very rapidly with respect to the Shapley value, which may take several weeks of computation to determine such values. This is due to the time necessary to determine the worth of each coalition, which considerably increases with the network size.

6.4. Polynomial approximation of the Shapley Value

In order to reduce the computation time for determining the Shapley value, several techniques have been proposed in the literature. To test and compare such techniques in our context, we further implemented the polynomial method proposed in [12], based on sampling theory, to estimate the Shapley value in the same network scenarios considered before. This method consists in the following steps:

1. considering the set of all possible orders of n players,
2. randomly and uniformly extracting a subset of q orders (called sampling unit),
3. computing the marginal contribution of each player observed in each sampling unit,
4. calculating the mean of such marginal contribution over the subset in step (2).

More specifically, we considered an increasing number of samples q, from 20 to 200, and for each q value we performed 100 random extractions. Table IX illustrates the corresponding results for the random network topology with $n = 15$ players, in terms of the *average* and *maximum* error with respect to the exact Shapley value (which was already illustrated in Table V). Table X further reports the same performance figures for all three ISP topologies, where $n = 10$ players are involved.

It can be observed that when a small number of samples is used ($q = 20$), the approximation is not satisfying, and significant errors can be observed with respect to exact Shapley values (i.e., approximately 14.4% for random topologies and 22% for ISP networks, in average, up to 61.8% in the worst case). The average and maximum errors decrease for increasing q values, but they are still quite high (up to 20% and 45%, in the worst case, for random and ISP topologies, respectively), and even increasing the number of samples does not permit to improve the precision of the estimation procedure.

Furthermore, we observe that such method (as several sampling methods) assumes that the worth of any coalition can be calculated in polynomial time, which is not the case for our game, since this involves finding the minimum cost Generalized Steiner tree that connects all source-destination pairs [28]. In a cooperative network design context,

Table 9: Sampling-based approximation of the Shapley value. Random geometric network scenario with $n = 15$ players: average and maximum error (percentage) with respect to the exact Shapley values illustrated in Table V, for different numbers of samples q.

q	Average Error (%)	Maximum Error (%)
20	14.40	45.53
50	9.41	28.32
100	8.08	23.66
150	6.95	20.63
200	6.84	20.22

Table 10: Sampling-based approximation of the Shapley value. Rocketfuel-inferred ISP topologies with $n = 10$ players: average and maximum error (percentage) with respect to the exact Shapley values illustrated in Table VII, for different numbers of samples q.

	Telstra		Sprintlink		Abovenet	
q	Avg. Error (%)	Max. Error (%)	Avg. Error (%)	Max. Error (%)	Avg. Error (%)	Max. Error (%)
20	13.62	36.06	21.62	61.76	6.04	19.19
50	10.71	26.28	18.39	51.54	4.61	14.25
100	9.69	22.70	17.48	47.57	3.93	12.50
150	8.80	20.43	16.93	45.87	3.63	11.45
200	8.61	20.22	16.53	45.23	3.42	10.94

even for relative small q values it is necessary to compute the worth of hundreds of coalitions, which means solving each time several NP-hard problems. As stated before, solving each of these problems, to determine just the worth of a single coalition, may require several days of computing time if networks formed by a large number of players are considered. For even larger topologies, the computing time grows exponentially, while our proposed NBS requires much less computation and can be obtained in any realistic network scenario within a small time.

In summary, the approximation methods based on sampling, like that in [12], are of limited interest in our problem since (1) for small q values they exhibit low precision in the approximation and (2) they are inapplicable for larger topologies and number of players, since they do not permit to reduce the computing time necessary to find an approximate solution.

7. Conclusion

In this paper we proposed a novel and efficient Nash bargaining solution for the cooperative network design problem with n players. Our solution has very appealing properties in terms of planning efficient networks and determining cost allocations in a very short computation time, even when compared to approximation techniques for Shapley value estimation, which obtain suboptimal solutions and still require an exponential time to be computed.

In all scenarios, our proposed solution permits to re-distribute cost savings in an efficient, fair and distributed way, increasing the appeal of the cost sharing solution. On the other hand, several players exhibit a Shapley value that is higher than the Nash equilibrium cost.

We compared our proposed solution to classic approaches, like the Shapley value and the Nash equilibrium concepts, in simple and large-size network topologies (including *random* and *real ISP* networks), with an increasing number of players. In particular, it appears that the Shapley value solution can be unstable and very difficult to compute.

In fact, we found that our proposed solution can be computed in few minutes (at most) in all considered, real-life scenarios. At the same time, we showed that the approximation methods for computing the Shapley value,

based on sampling techniques, are of limited interest in our problem since (1) when a small number of samples is used, they exhibit low precision in the approximation and (2) they are inapplicable for larger topologies and number of players, since they do not permit to sufficiently reduce the computing time necessary to find an approximate solution. For example, when 20 samples are used, the approximation is not satisfying, and significant errors can be observed with respect to exact Shapley values (i.e., approximately 14.4% for random topologies and 22% for ISP networks, in average, up to 61.8% in the worst case).

To summarize, numerical results demonstrate that our approach permits to achieve very effective cost allocations in a short computing time, thus representing an efficient and promising framework for the planning of stable networks.

Acknowledgments

This work was partially supported by ANR in the framework of the ANR Green-Dyspan project.

References

[1] E. Anshelevich, A. Dasgupta, J. Kleinberg, E. Tardos, T. Wexler, and T. Roughgarden. The Price of Stability for Network Design with Fair Cost Allocation. In *SIAM Journal on Computing, Vol. 38, no. 4*, pages 1602–1623, 2008.

[2] H.L. Chen and T. Roughgarden. Network Design with Weighted Players. In *Theory of Computing Systems, Vol. 45, no. 2*, pages 302–324, August 2009.

[3] S. Albers. On the value of coordination in network design. In *SIAM Journal on Computing, Vol. 38, no. 6*, pages 2273–2302, 2009.

[4] G. Smaragdakis, N. Laoutais, P. Michiardi, and A. Bestavros. Distributed Network Formation for n-Way Broadcast Applications. *IEEE Transactions on Parallel and Distributed Systems*, pages 1427–1441, vol. 21, no. 10, October 2010.

[5] A. Epstein, M. Feldman, and Y. Mansour. Strong equilibrium in cost sharing connection games. In *Games and Economic Behavior, Vol. 67, no. 1*, pages 51–68, September 2009.

[6] H.L. Chen, T. Roughgarden, and G. Valiant. Designing Network Protocols for Good Equilibria. In *SIAM Journal on Computing, Vol. 39, no. 5*, pages 1799–1832, 2010.

[7] A.P. Azad, E. Altman, and R. El-Azouzi. From Altruism to Non-Cooperation in Routing Games. In *Proceedings of International Workshop on Wireless Networks: Communication, Cooperation and Competition (WNC3 2010)*, Avignon, France, May 2010.

[8] J. Elias, F. Martignon, K. Avrachenkov, and G. Neglia. Socially-Aware Network Design Games. In *Proceedings of INFOCOM 2010*, San Diego, CA, USA, Mars 2010.

[9] A. Muthoo. *Bargaining theory with applications*. Cambridge Univ. Press, 1999.

[10] J.F. Nash Jr. The bargaining problem. *Econometrica*, 18(2):155–162, 1950.

[11] R.J. Aumann, R.B. Myerson, and A. Roth. The Shapley Value. *Game-Theoretic Methods in General Equilibrium Analysis*, pages 61–66, 1994.

[12] J. Castro, D. Gómez, and J. Tejada. Polynomial calculation of the Shapley value based on sampling. *Computers and Operations Research*, pages 1726–1730, vol. 36, no. 5, May, 2009.

[13] A. van den Nouweland. *Models of network formation in cooperative games*. Cambridge Univ. Press, 2005.

[14] E.D. Demaine, M. Hajiaghayi, H. Mahini, and M. Zadimoghaddam. The Price of Anarchy in Cooperative Network Creation Games. In *Proceedings of STACS 2009*, pages 301–312, February 26-28, 2009, Freiburg, Germany.

[15] W. Saad, Z. Han, T. Basar, M. Debbah, and A. Hjorungnes. Network formation games among relay stations in next generation wireless networks. *IEEE Transactions on Communications*, 59(9):2528–2542, September 2011.

[16] M. Jackson and A. van den Nouweland. Strongly stable networks. *Games and Economic Behavior*, pages 420–444, vol. 51, 2005.

[17] N. Andelman, M. Feldman, and Y. Mansour. Strong price of anarchy. In *Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2007)*, January 7-9, 2007, New Orleans, Louisiana.

[18] R.T.B. Ma, D.M. Chiu, J.C.S. Lui, V. Misra, and D. Rubenstein. On cooperative settlement between content, transit, and eyeball internet service providers. *IEEE/ACM Transactions on Networking*, 19(3):802–815, 2011.

[19] R. Mazumdar, L.G. Mason, and C. Douligeris. Fairness in network optimal flow control: Optimality of product forms. *IEEE Transactions on Communications*, 39(5):775–782, 1991.

[20] H. Yaiche, R.R. Mazumdar, and C. Rosenberg. A game theoretic framework for bandwidth allocation and pricing in broadband networks. *IEEE/ACM Transactions on Networking*, 8(5):667–678, 2000.

[21] Zhu Han, Zhu Ji, and K. J. Ray Liu. Fair Multiuser Channel Allocation for OFDMA Networks Using Nash Bargaining Solutions and Coalitions. *IEEE Transactions on Communications*, pages 1366–1376, vol. 53(8), August 2005.

[22] K. Avrachenkov, J. Elias, F. Martignon, G. Neglia, and L. Petrosyan. A Nash bargaining solution for Cooperative Network Formation Games. In *Proceedings of Networking 2011*, Valencia, Spain, May 2011.

[23] Y. Wu and W.-Z. Song. Cooperative Resource Sharing and Pricing for Proactive Dynamic Spectrum

Access via Nash Bargaining Solution. *IEEE Transactions on Parallel and Distributed Systems*, Preprint, November 2013.

[24] S.U. Khan, N.D. Fargo, and I. Ahmad. A Cooperative Game Theoretical Technique for Joint Optimization of Energy Consumption and Response Time in Computational Grids. *IEEE Transactions on Parallel and Distributed Systems*, pages 346–360, vol. 20, no. 3, March 2009.

[25] G. Blocq and A. Orda. How good is bargained routing. In *Proceedings of the IEEE Infocom 2012*, pages 2453–2461, 25-30 March 2012,Orlando, FL, USA.

[26] R. Stanojevic, N. Laoutaris, and P. Rodriguez. On economic heavy hitters: Shapley value analysis of 95th-percentile pricing. *Proceedings of the 10th annual conference on Internet measurement (IMC '10)*, pages 75–80, 2010.

[27] J. Elias, F. Martignon, K. Avrachenkov, and G. Neglia. A Game Theoretic Analysis of Network Design with Socially-Aware Users. *Elsevier Computer Networks*, 55(1):106–118, January 2011.

[28] M. Khan, F. Kuhn, D. Malkhi, G. Pandurangan, and K. Talwar. Efficient distributed approximation algorithms via probabilistic tree embeddings. In *Proceedings of the 27th symposium on Principles of Distributed Computing*, pages 263–272, 2008.

[29] *AMPL: A Modeling Language for Mathematical Programming.* Available at http://www.ampl.com.

[30] ILOG Optimization Products. IBM ILOG CPLEX optimizer. http://www.ibm.com/software/integration/optimization/cplex-optimizer/.

[31] A. Agrawal, P. Klein, and R. Ravi. A general approximation technique for constrained forest problems. *SIAM Journal on Computing*, pages 440–456, vol. 24, no. 3, 1995.

[32] M. X. Goemans and D. P. Williamson. A general approximation technique for constrained forest problems. *SIAM Journal on Computing*, pages 296–317, vol. 24, no. 2, 1995.

[33] B. Awerbuch, Y. Azar, and Y. Bartal. On-line generalized Steiner problem. In *Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 1996)*, January 1996, San Francisco, CA.

[34] S. Haldar. An all pair shortest paths distributed algorithm using $2n^2$ messages. *Journal of Algorithms*, pages 20–36, vol. 24, July 1997.

[35] J. Fakcharoenphol, S. Rao, and K. Talwar. A tight bound on approximating arbitrary metrics by tree metrics. In *Journal of Computer and System Sciences*, pages 485–497, vol. 69(3), November 2004.

[36] H.T. Tran and T. Ziegler. A design framework towards the profitable operation of service overlay networks. *Computer Networks*, pages 94–113, vol. 51, 2007.

[37] P. Lin, X. Feng, Q. Zhang, and M. Hamdi. Groupon in the Air: A Three-stage Auction Framework for Spectrum Group-buying. In *Proceedings of IEEE INFOCOM 2013*, Turin, Italy, April 14-19, 2013.

[38] ETSI. Network Functions Virtualisation - Introductory White Paper. `http://portal.etsi.org/NFV/NFV_White_Paper.pdf`.

[39] N.M. Chowdhury and R. Boutaba. A survey of network virtualization. *Computer Networks*, 54(5):862–876, 2010.

[40] N. Spring, R. Mahajan, and D. Wetherall. Measuring ISP topologies with Rocketfuel. In *Proceedings of ACM SIGCOMM 2002*, Pittsburgh, PA, USA, August 2002.

Appendix

This appendix provides detailed proofs for Theorems 1 and 2.

7.1. Proof of Theorem 1

Theorem 1: In the Cooperative Network Design Game, the security level (definition 2) and the minimum cost of the coalition (definition 3) satisfy the axioms of characteristic function.

Proof: We need to prove that definitions (2) and (3) lead to subadditive characteristic functions.

Let us first consider the characteristic function defined as in (2). We prove the result directly for pure strategy games, i.e. when players cannot use more paths according to a probability distribution.

Let μ_S be a strategy available to players in $\mathcal{S}$, i.e. $\mu_{\mathcal{S}} \in \Pi_{i\in\mathcal{S}} A_i$. Observe that, given a coalition $\mathcal{S} \cup \mathcal{T}$ with a strategy $\mu_{\mathcal{S}\cup\mathcal{T}}$, this latter can be expressed as product of a strategy of coalition $\mathcal{S}$ and a strategy of coalition $\mathcal{T}$. We denote by $C_{\mathcal{S}}$ the total cost paid by users in $\mathcal{S}$ for a given outcome of the game (i.e., for a given set of strategies of the players). Then, we can express definition (2) as follows:

$$v(\mathcal{S}) = \min_{\mu_{\mathcal{S}}} \max_{\mu_{\mathcal{I}-\mathcal{S}}} C_{\mathcal{S}}(\mu_{\mathcal{S}}, \mu_{\mathcal{I}-\mathcal{S}}).$$

Let us denote by $\mathcal{V}$ the set $\mathcal{I} - (\mathcal{S} \cup \mathcal{T})$. The value of the coalition $\mathcal{S} \cup \mathcal{T}$ can therefore be expressed as reported in equation (4) at the top of next page, showing that $v(\mathcal{S} \cup \mathcal{T}) \leq v(\mathcal{S}) + v(\mathcal{T})$.

If we consider also *mixed strategies*, then the result follows immediately from the following theorem in [38].

Theorem in [38]: Denote by $C_{\mathcal{S}}(\mu_{\mathcal{S}}, \nu_{\mathcal{I}-\mathcal{S}})$ the cost of the coalition $\mathcal{S}$ in the zero-sum game between coalition $\mathcal{S}$ and $\mathcal{I} - \mathcal{S}$, where $\mu_{\mathcal{S}}$ is the strategy of coalition $\mathcal{S}$ and $\nu_{\mathcal{I}-\mathcal{S}}$ is

$$v(\mathcal{S} \cup \mathcal{T}) = \min_{\mu_{\mathcal{S}\cup\mathcal{T}}} \max_{\mu_{\mathcal{V}}} C_{\mathcal{S}\cup\mathcal{T}}(\mu_{\mathcal{S}\cup\mathcal{T}}, \mu_{\mathcal{V}}) = \min_{\mu_{\mathcal{S}},\mu_{\mathcal{T}}} \max_{\mu_{\mathcal{V}}} (C_{\mathcal{S}}(\mu_{\mathcal{S}}, \mu_{\mathcal{T}}, \mu_{\mathcal{V}}) + C_{\mathcal{T}}(\mu_{\mathcal{S}}, \mu_{\mathcal{T}}, \mu_{V})) \leq$$
$$\leq \min_{\mu_{\mathcal{S}},\mu_{\mathcal{T}}} \left(\max_{\mu_{\mathcal{V}},\tilde{\mu}_{\mathcal{T}}} C_{\mathcal{S}}(\mu_{\mathcal{S}}, \tilde{\mu}_{\mathcal{T}}, \mu_{\mathcal{V}}) + \max_{\mu_{\mathcal{V}},\tilde{\mu}_{\mathcal{S}}} C_{\mathcal{T}}(\tilde{\mu}_{\mathcal{S}}, \mu_{\mathcal{T}}, \mu_{\mathcal{V}}) \right) =$$
$$= \min_{\mu_{\mathcal{S}}} \max_{\mu_{\mathcal{V}},\tilde{\mu}_{\mathcal{T}}} C_{\mathcal{S}}(\mu_{\mathcal{S}}, \tilde{\mu}_{\mathcal{T}}, \mu_{\mathcal{V}}) + \min_{\mu_{\mathcal{T}}} \max_{\mu_{\mathcal{V}},\tilde{\mu}_{\mathcal{S}}} C_{\mathcal{T}}(\tilde{\mu}_{\mathcal{S}}, \mu_{\mathcal{T}}, \mu_{\mathcal{V}}) =$$
$$= \min_{\mu_{\mathcal{S}}} \max_{\mu_{\mathcal{I}-\mathcal{S}}} C_{\mathcal{S}}(\mu_{\mathcal{S}}, \mu_{\mathcal{I}-\mathcal{S}}) + \min_{\mu_{\mathcal{T}}} \max_{\mu_{\mathcal{I}-\mathcal{T}}} C_{\mathcal{T}}(\mu_{\mathcal{T}}, \mu_{\mathcal{I}-\mathcal{T}}) = v(\mathcal{S}) + v(\mathcal{T}) \quad (4)$$

---

the strategy of coalition $\mathcal{I} - \mathcal{S}$. Then the function $v(\mathcal{S}) = \inf_{\mu_{\mathcal{S}}} \sup_{\nu_{\mathcal{I}-\mathcal{S}}} C_{\mathcal{S}}(\mu_{\mathcal{S}}, \nu_{\mathcal{I}-\mathcal{S}})$ is subadditive.

The fact that the function defined by (2) is subadditive in the case of mixed strategies follows immediately from the above theorem.

In order to prove that the function defined by (3) is also subadditive, we modify the network so that the characteristic function according to (3) in the original game coincides with the characteristic function defined as in (2) in the modified network. In particular, we introduce auxiliary links with infinite cost connecting source and destination nodes of each player. Now, the best strategy for the coalition $\mathcal{I} - \mathcal{S}$ in the setting (2) is just to choose these auxiliary links, which is equivalent to remove the players in $\mathcal{I} - \mathcal{S}$ from the game (as required by definition (3)). Thus, subadditivity of the characteristic function (3) follows from the above results for function (2).

7.2. Proof of Theorem 2

Theorem 2: The Nash bargaining solution for player i, α_i is given by the following expression:

$$\alpha_i = u_i - \frac{\sum_k u_k - u_{soc}}{m}, \quad (5)$$

where m coincides with the number of players n (i.e., $m \equiv n$) if we allow for negative costs (i.e., some α_i values are negative, which means that some players are actually paid to ensure their participation). Otherwise, if only non-negative costs are allowed (or equivalently, if no positive transfers are permitted), m is defined as the largest integer for which the following inequality is satisfied:

$$\frac{1}{m-1}(\sum_{i=1}^{m-1} u_i - u_{soc}) < u_m \tag{6}$$

having assumed, without loss of generality, that players are ordered such that $u_1 \geq u_2 \geq \ldots \geq u_n$.

Proof: We consider a Nash bargaining solution for the n-players cooperative game with transferable cost. The assumption about the transferable cost means that the players, or the system administrator, can redistribute the total cost among the players.

Then, the Nash bargaining solution is given by the following optimization problem:

$$\max_{\alpha_i} \prod_{i=1}^{n} (u_i - \alpha_i), \tag{7}$$

subject to

$$\sum_{i=1}^{n} \alpha_i = u_{soc}. \tag{8}$$

Below we consider two cases: (a) individual costs α_i can be negative (this can be interpreted as the coalition pays some members to ensure their participation) and (b) costs cannot be negative, i.e., no positive transfers are allowed; this precludes paying players to participate to the network.

Both scenarios make practical sense.

7.2.1. The case without the requirement on the positivity of costs

If there are no positivity constraints on α_i, the Lagrangian is given by

$$L_{NP} = \prod_{i=1}^{n} (u_i - \alpha_i) + \mu(\sum_{i=1}^{n} \alpha_i - u_{soc}),$$

and the Karush-Kuhn-Tucker condition takes the form:

$$\frac{\partial L_{NP}}{\partial \alpha_i} = -\prod_{j \neq i} (u_j - \alpha_j) + \mu = 0, \tag{9}$$

plus constraint (8). Multiplying (9) by $(u_i - \alpha_i)$ and dividing by μ, we obtain

$$u_i - \alpha_i = \frac{1}{\mu}\prod_{j=1}^{n}(u_j - \alpha_j).$$

Hence, the difference $u_i - \alpha_i$ for the optimal solution does not depend on the index i and we can denote its value by δ. Thus, we have

$$\alpha_i = u_i - \delta, \tag{10}$$

where the value of δ can be found from condition (8):

$$\delta = \frac{1}{n}(\sum_{i=1}^{n} u_i - u_{soc}).$$

It is interesting to observe that in this case every player gains an equal share of the difference between the total cost of the Nash equilibrium and the total socially optimal cost. Some players might actually be reimbursed.

7.2.2. The case with the requirement on the positivity of costs

In this case, in addition to the equality constraint (8), we have n inequality constraints

$$\alpha_i \geq 0, \quad i = 1, ..., n. \tag{11}$$

This formulation corresponds to the following Lagrangian

$$L_P = \prod_{i=1}^{n}(u_i - \alpha_i) + \sum_{i=1}^{n}\lambda_i\alpha_i + \mu(\sum_{i=1}^{n}\alpha_i - u_{soc}).$$

The Karush-Kuhn-Tucker condition takes the form:

$$\frac{\partial L_P}{\partial \alpha_i} = -\prod_{j\neq i}(u_j - \alpha_j) + \lambda_i + \mu = 0, \tag{12}$$

$$\lambda_i \geq 0, \quad \lambda_i\alpha_i = 0, \quad i = 1, ..., n, \tag{13}$$

plus conditions (8) and (11).

Without loss of generality, we have ordered the players such that $u_1 \geq u_2 \geq ... \geq u_n$. It follows that $\alpha_1 \geq \alpha_2 \geq ... \geq \alpha_n$. In fact, let us assume that $u_i > u_j$ (then $i < j$) and $\alpha_j > \alpha_i$: the vector of costs $(\alpha_1, \cdots, \alpha_{i-1}, \alpha_i, \cdots, \alpha_{j-1}, \alpha_j, \cdots, \alpha_n)$ cannot be a solution of the problem (7), since the vector $(\alpha_1, \cdots, \alpha_{i-1}, \alpha_j, \cdots \alpha_{j-1}, \alpha_i, \cdots, \alpha_n)$ corresponds to a higher value of the optimization function, being that $(u_i - \alpha_j)(u_j - \alpha_i) > (u_i - \alpha_i)(u_j - \alpha_j)$.

Let us denote by m the number of non-zero α's. In particular, it may happen that $m = n$. We shall now illustrate how to determine m.

If $\alpha_1, ..., \alpha_m > 0$, by the complementarity slackness condition $\lambda_i \alpha_i = 0$, we have $\lambda_i = 0$ for $i = 1, ..., m$. Similarly to the first case, multiplying (12) by $(u_i - \alpha_i)$ and dividing by μ, we conclude that

$$u_i - \alpha_i = \hat{\delta}, i = 1, ..., m.$$

In addition, from (12) we have:

$$\mu = (u_1 - \alpha_1)(u_2 - \alpha_2)...(u_m - \alpha_m)u_{m+1}...u_n = \hat{\delta}^{m-1} u_{m+1}...u_n.$$

Then, for k such that $m + 1 \leq k \leq n$, the equation (12) gives

$$\frac{\partial L_P}{\partial \alpha_k} = -\prod_{j \neq k} (u_j - \alpha_j) + \lambda_k + \mu = 0,$$

$$-(u_1 - \alpha_1)...(u_m - \alpha_m)u_{m+1}...u_{k-1}u_{k+1}...u_n + \lambda_k + \mu = 0,$$

$$\lambda_k = \hat{\delta}^m u_{m+1}...u_{k-1}u_{k+1}...u_n - \hat{\delta}^{m-1} u_{m+1}...u_n,$$

$$\lambda_k = [\hat{\delta} - u_k]\hat{\delta}^{m-1} u_{m+1}...u_{k-1}u_{k+1}...u_n.$$

Since according to the Karush-Kuhn-Tucker condition λ_i must be non-negative, we obtain the following condition

$$\hat{\delta} \geq u_k, \quad k = m + 1, ..., n.$$

This condition allows us to determine m. Before proceeding towards this goal, we determine $\hat{\delta}$ from condition (8), which gives us

$$\hat{\delta} = \frac{1}{m}(\sum_{i=1}^{m} u_i - u_{soc}).$$

Then, we have the following algorithm to determine m: first, check if

$$\frac{1}{n-1}(\sum_{i=1}^{n-1} u_i - u_{soc}) < u_n. \tag{14}$$

If it is the case, then all α_i are positive. We note that $\frac{1}{n}(\sum_{i=1}^{n} u_i - u_{soc}) < u_n$ is an equivalent condition. If condition (14) is not satisfied, find the largest m for which the following condition holds

$$\frac{1}{m-1}(\sum_{i=1}^{m-1} u_i - u_{soc}) < u_m. \tag{15}$$

This m gives the number of non-zero α's.