

Distributed Resource Allocation Algorithms for Peer-to-peer Networks

Iordanis Koutsopoulos
Department of Computer and Communications
Engineering
University of Thessaly, Greece
jordan@uth.gr

George Iosifidis
Department of Computer and Communications
Engineering
University of Thessaly, Greece
giosifid@uth.gr

ABSTRACT

In peer-to-peer networks, each peer plays the role of client and server. As server, it receives content requests made by other peers and needs to decide on what basis and to what extent it will satisfy these requests by uploading content to others. As client, it addresses its own requests to appropriate peers to download desired content after resources are granted. We consider a network of peers in a star topology, where the bottleneck is the capacity of the access link connecting a peer to the backbone. Different peers have different utility functions which are private information and capture a peer's selfishness or desire for content. The objective is to maximize the sum of utilities of peers. We intend to answer the following questions in a peer-to-peer network: what portions of its link capacity does each peer allocate to upload flows from other peers and download flows for itself? How does a peer decide which portion of bandwidth will be allocated to each upload flow and download flow? How can these decisions be taken in a decentralized autonomous fashion? Although each peer directly obtains utility only from downloads, in the presence of an incentive protocol it would like to allow just enough capacity for uploads of others so that it is not punished by the protocol. The global link sharing problem of maximizing total utility is hard to solve in a distributed fashion because of coupled utility functions and constraints. That is, the utility of each peer depends on allocation decisions of others. By defining auxiliary variables and constraints, we transform the problem into one that is amenable to "distributization" by dual decomposition. The iterative algorithm involves solving separate optimization problems by each peer and updating Lagrange multipliers. Interestingly, the Lagrange multipliers corresponding to the newly added constraints are interpreted as reciprocals of pairwise reputation metrics. This leads us to a meaningful reputation-driven protocol with the desirable property that only the amounts of requested and granted bandwidth are circulated, and not reputations. The protocol is lightweight in terms of computational complexity and

overhead and converges to the globally optimal allocation.

Categories and Subject Descriptors

C.2 [Computer-Communication Networks]: Distributed networks; G.1 [Numerical analysis]: Optimization, Convex Programming

General Terms

Peer-to-peer networks, Algorithm, Protocol design

1. INTRODUCTION

The fundamental characteristic of the peer-to-peer network paradigm is that each peer acts simultaneously as client and as server. In peer-to-peer networks, peers form overlay structures. A link in the overlay network graph link connects two peers virtually, in the sense that it may correspond to one or more multi-hop paths in the physical network graph between the two peers. The overlay network offers an abstraction of the physical network graph that facilitates the study of issues such as content request routing and neighbor selection. Once the overlay is constructed, content traffic flows through the physical links. Peer-to-peer networks are most commonly used for file sharing and distribution (e.g. Gnutella, KaZaA, eDonkey, BitTorrent), and video and Television media distribution (e.g. Joost) VoIP (Skype). Other applications include distributed indexing and directory services, web caching, streaming media, grid-based computing. Recent applications include distributed storage [1] and social networking.

In peer-to-peer networks, a set of peers is formed based on declared interests of peers for content. For instance, in BitTorrent, each peer has to pre-register its content with a website, the tracker. A content request by a peer is followed by an advertisement by the tracker of a random set of peers that possess the content fully or partially. Each peer subsequently generates content requests which are submitted to other peers. The client problem faced by a peer is to determine the peers to which a request for download will be addressed. On the other hand, a peer as a server needs to adopt a bandwidth allocation policy to serve incoming requests with available upload link capacity. More fundamentally, each peer needs to decide the portion of its link bandwidth that will be allocated to uploads of other peers and downloads by that peer. These bandwidth sharing decisions need to be fully decentralized since they should be separate by each peer and should not imposed by a central authority such as a central server.

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.

ValueTools 2008, October 21 - 23, 2008, Athens, GREECE
Copyright 2008 ICST ISBN # 978-963-9799-31-8.

Each peer is characterized by a utility function of the allocated amount of resource, which quantifies the amount of satisfaction received by a peer. Given an amount of resource, a larger value of utility of a peer in a sense denotes the peer's willingness to obtain the resource. If the peer's tendency to take a selfish decision is driven by the aspired utility, then a higher utility might also be an indication of higher amount of selfishness. The network-wide objective is the enforcement of an operating point (namely, vector of peer utilities) with a desirable property such as fairness, Pareto optimality etc. Nevertheless, this is a non-trivial issue in peer-to-peer networks either due to limited amount of information in the network (e.g. about peer individual utilities, link capacities or availability of content), or because of the inherent inclination of peers to behave selfishly: since a peer obtains direct utility only from downloads and not for uploads, it is natural that a peer dedicate its bandwidth to downloads while neglecting or even shutting down other peers' uploads from it. This selfish behavior is widely known as free riding and it turns out to have significant impact on the operational point. Existing systems such as BitTorrent have incorporated mechanisms to provide incentives for resource provisioning to enhance peers' willingness to cooperate and improve overall system performance. In that context, the basic dilemma faced by peers is the following: peers want to act only as clients to maximize their derived benefit, yet they know that such behavior may lead to a subsequent reduction of their utility because of the activated incentive mechanism. In some contemporary peer-to-peer systems, this tradeoff is not dynamically controlled. That is, the ratio of download and upload bandwidth portions is either fixed or manually configured by each peer. Yet, such solutions are mainly based only on empirical rules that attempt to provide a viable solution to the free riding problem.

A peer can be connected to the backbone network either through an access link (cable modem or ADSL line), where upstream and downstream links are separated, or through a local hub with which it is connected with a wireless or non-wireless link with an Ethernet-like protocol [2]. In the latter case, both upload and download flows traverse the common access link of finite capacity. This case captures the interplay between the amount of resource dedicated to downloads from others and that granted to uploads to others. Increasing the amount of bandwidth for uploads directly affects the download bandwidth and thus the perceived utility. A typical scenario is when the underlying network infrastructure is wireless, and each peer is connected to an access point or gateway of the backbone network through a wireless link. Upload and download flows take place through the same wireless link between the peer and the access point. Essentially, it is the Carrier Sense Multiple Access (CSMA) protocol at the MAC layer that determines the portions of time where the peer will get access to upload its content through the backbone network to others, or the neighboring access point will send data to the peer. These time portions present themselves in the higher layers as portions of the link capacity dedicated to uploading or downloading. The high penetration of wireless devices and the proliferation of wireless community based networks such as FON [3] necessitate the adoption of methods that tackle problems arising in peer-to-peer networks in a wireless environment.

In this paper we take a network optimization approach and devise a resource allocation protocol that is specifically

applicable for the bandwidth sharing scenario above. We start from the system-wide goal of identifying the upload / download link shares for maximizing social welfare, namely the sum of peer utilities. The utility of each peer depends on its own upload decision but also on decisions of other peers since the latter will determine the amount of download for that peer. This optimization problem is difficult to solve in a distributed fashion due to coupled utility functions. By defining new auxiliary variables and introducing additional constraints, we transform the original problem to another one, which is amenable to distributed optimization through decomposition to separate optimization problems that are solved by each peer until convergence. Interestingly, it turns out by our formulation that the Lagrange multipliers that correspond to the newly introduced constraints can be interpreted as reciprocals (inverses) of pairwise reputation metrics and are updated with the gradient-like updates of the dual problem. Another important and desirable feature of our protocol is that only the amounts of requested and granted resource are circulated, and not the reputation metrics or utilities which have to be private information. The contributions of our work to the current literature are as follows: (i) we address the interplay between upload and download decisions in a peer-to-peer network through a network utility maximization approach; (ii) we show how to solve the problem of maximizing social welfare in a distributed fashion; (iii) we provide a meaningful interpretation of Lagrange multipliers and their updates as reciprocals of pairwise reputations which evolve according to the amounts of requested and contributed resources by peers; (iv) we utilize the framework above to design a lightweight protocol that provably reaches the social optimum, in which resource allocation decisions are driven by those reciprocal reputations. The rest of the paper is organized as follows. In section 2 we review the related work and in section 3 we present the system model and problem statement. In section 4 we present the problem analysis and the distributed algorithm. In section 5 we provide numerical results and in section 6 we conclude our study.

2. RELATED WORK

In incentive systems, a peer attempts to estimate the behavioral profile of others, based on the results of its transactions with them, and accordingly to impose certain resource allocation regimes. The authors in [4] present a system where peers use direct and second-hand information through weighted voting schemes to assess the behavior of other nodes. The work in [5] proposes a mechanism with the interesting property that the reputation of every peer converges to its inherent tendency to cooperate and reveals its true intentions. Once calculated, reputation values are used by servers in allocating resources among requesting peers, or by clients in order to select the most reputable among the servers. In [6], resource allocation and server selection policies are determined by local reputation values for different types of offered services. It is shown that peers tend to form coalitions and to cooperate by exchanging resources and services while free riding behavior is discouraged. Other sophisticated models have also been devised and studied, [10], [11], [2]. For example in [11], peers distribute their requests with the objective to minimize required download time and in [2] peers organize auctions in order to allocate their serving bandwidth. On a slightly different context,

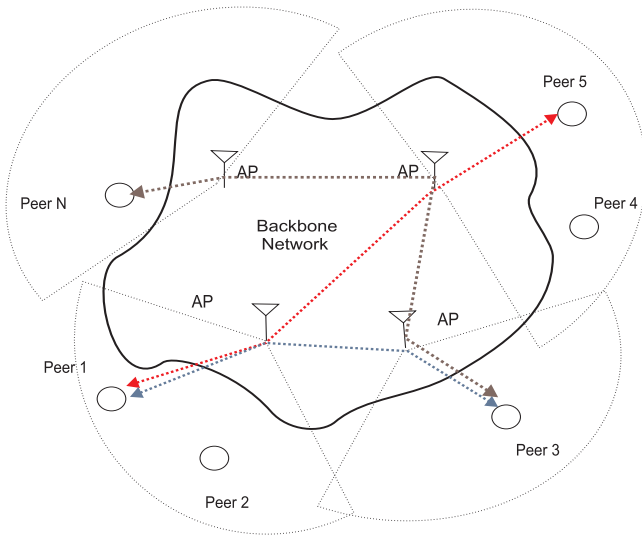


Figure 1: An instance of peer topology that is covered by our model: the backbone network consists of APs that communicate through wireline or wireless links. Each peer belongs to an AP and is connected to it through a wireless link.

the authors of [12] propose an algorithm for rate control in peer-to-peer overlays, thus extending the TCP congestion control mechanism. They formulate a network utility maximization (NUM) problem and derive a protocol based on its distributed implementation. Client and server peers consider link prices and make optimal routing and bandwidth allocation decisions. In [16], we model the hidden intentions of a peer to allow others' uploads with a single scalar parameter: the level of dissatisfaction the peer intends to cause by not fulfilling other peers' requests.

Incentive-based mechanisms have also been studied for wireless networks. For instance, in [7], the reputation metric quantifies the cooperation level of each node, i.e its propensity to provide packet routing. This is used when the node in turn requests service itself in terms of having its packets forwarded and routed. However, wireless channel impairments such as packet collisions, interference and channel quality variations may mislead peers and result in distorted reputation estimates. In [8] the authors argue that a tolerance threshold can be added to the pure Tit-for-Tat (TFT) strategy, such that a few defections (naturally encountered due to impairments) will not trigger any punishment. Another model is presented in [9], where a node which has defected is given a second opportunity to cooperate.

The framework of NUM has become very attractive in the years that followed the seminal work by Kelly *et.al* [13] both for wire-line and wireless networks for addressing and solving utility maximization problems in a distributed manner and designing novel protocols based on optimization theory. For wireless networks, decomposition is mapped to cross-layer design and placing emphasis on functionalities of a certain layer, while the circulated parameters stand for the information that needs to be exchanged among different layers [14], [15]. Our work is in line with this direction, and, to the best of our knowledge, this is the first work which applies network optimization to derive a protocol out of the

distributed solution of the problem.

3. SYSTEM MODEL AND PROBLEM STATEMENT

3.1 System model

We consider a set of N peers, each with some content at its disposal. The system is static for the period of peer interaction, namely we do not consider peers coming to or leaving the system. The case of dynamics of peer arrivals and departures may be incorporated in the formulation by allowing a certain interval for adjusting information about the system. This setting is similar to BitTorrent, where the peers form groups (*swarms*), based on their declared interest on a specific file. These groups are considered static for a certain time period, and the peers inside each group, interact by exchanging different fragments of the file. However, in this work, peers are not distinguished on the basis of specific content they possess and thus the algorithm for peer selection does not take requested content into account. We assume that the requested content by a peer is always available at other peers.

We consider peers in a star topology, where each peer i is connected to the backbone network through an access link of capacity C_i bits/sec. This can be wire-line or even wireless. In the latter case, a peer can be associated with a specific AP and is connected to the backbone network through this AP, as shown in Figure 1. The backbone network is assumed to be of high enough capacity such that it does not create in-network congestion. We do not consider any in-network related mechanisms such as overlay routing and request scheduling at intermediate backbone nodes. We assume that the last-mile wireless link between each peer i and its corresponding AP is the performance bottleneck with a capacity of C_i bits/sec that needs to be shared between upload and download flows. The allocation of capacity portions to these flows at the network layer is realized by regulation through the CSMA MAC protocol of the time portions when a peer gets access to the channel to upload content or the access point transmits such that the peer downloads. Link capacity C_i needs to be split to a portion related to peer downloads and one for peer uploads. This upload to download ratio is passed to the network layer through an application software running at each peer, such as eMule and BitTorrent clients. Then, it is communicated down to the network and MAC layers and determines wireless link usage.

A note about resource allocation is in place here. In the case of wire-line network, each peer is connected through a separate link of capacity C_i to the back-bone. In the case of wireless link, it is possible for two or more peers to be connected to the backbone network through the same AP. However in that case, a scheduling mechanism should be employed at the AP to specify the capacities dedicated to the uploads and downloads of different peers. In this work, we do not address this issue. Instead, we adopt the abstraction of a star topology with a fixed access link capacity for each peer, which needs to be split between uploads and downloads.

We assume a continuous backlog of arising content requests at each peer that need to be satisfied. This case arises for instance in continual requests for streaming con-

tent. Each peer as client addresses its content requests to other peers. Content requests can be directly viewed as amounts of requested bandwidth. Let us denote the requests of peer i by the resource request vector $\mathbf{x}_i = (x_{ij} : j = 1, \dots, N, j \neq i)$, where x_{ij} is the amount of bandwidth peer i requests from peer j . Each peer as server receives incoming requests and needs to serve them by granting them a certain amount of bandwidth. We denote the bandwidth granted by peer i by a resource provisioning vector $\mathbf{y}_i = (y_{ij} : j = 1, \dots, N, j \neq i)$, where y_{ij} is the amount of bandwidth that is granted from i to j . Denote by $\tilde{\mathbf{y}}_i = (y_{ji} : j = 1, \dots, N, j \neq i)$ the amounts of bandwidth provided to peer i by other peers $j \neq i$. We assume that a client always uses exactly the bandwidth it is granted. The global network decision is represented by resource request matrix $\mathbf{X} = (\mathbf{x}_i : i = 1, \dots, N)$ and resource provisioning matrix $\mathbf{Y} = (\mathbf{y}_i : i = 1, \dots, N)$. Clearly, \mathbf{y}_i and $\tilde{\mathbf{y}}_i$ are the i th row and i th column of \mathbf{Y} respectively. Although $\{x_{ij}\}$ are not present in the problem formulation, we choose to define them here, as they will naturally emerge later on as auxiliary variables in the solution of the problem.

3.2 Utility function

Each peer i is characterized by a differentiable utility function $U_i(\cdot)$, which quantifies the amount of satisfaction. In the current setup we relate it to the acquired bandwidth. In a different setup, it could also represent perceived peer delay from downloads.

Option A. A first alternative for the utility function $U_i(\cdot)$ of a peer i is that the received utility depends on the amounts of granted bandwidth y_{ji} by other peers $j \neq i$, (namely downloads of i), and on the amounts of uploads y_{ij} from peer i to others. Specifically, a peer receives utility from its downloads based on a rule of diminishing marginal returns, captured by a non-decreasing, differentiable concave function $u_i(\cdot)$ of the download bandwidth from each other peer. On the other hand, a peer experiences certain “dissatisfaction” due to the upload bandwidth of others: if that bandwidth were not devoted to uploads, the peer might have used it for its own downloads. Dissatisfaction is represented by a non-decreasing, differentiable convex cost function $w_i(\cdot)$ of the upload bandwidth of each other peer. The convexity accounts for increasing marginal dissatisfaction cost for larger amounts of uploading. A reasonable choice of utility function that captures the features above is:

$$U_i(\mathbf{y}_i, \tilde{\mathbf{y}}_i) = \sum_{j \neq i} [u_i(y_{ji}) - w_i(y_{ij})]. \quad (1)$$

Thus, $U_i(\cdot)$ is differentiable, non-increasing concave in \mathbf{y}_i and concave non-decreasing in $\tilde{\mathbf{y}}_i$. In a way, $U_i(\cdot)$ is the net utility of a peer. This type of utility function stems partially from the theory of public goods and has been adopted in similar forms in various works [17], [18], [2].

Option B. According to a second option, the utility function $U_i(\cdot)$ of a peer i depends only on the downloads of peer i , i.e.,

$$U_i(\tilde{\mathbf{y}}_i) = \sum_{j \neq i} u_i(y_{ji}). \quad (2)$$

While in option B, the utility of a peer depends on allocation decision of other peers which determine peer downloads, in option A a peer can control to a certain extent its own utility by regulating the amount of uploads of other peers.

A larger value of utility function of a peer in a sense denotes the peer’s willingness to obtain the resource. If the peer’s tendency to take a selfish decision is driven by the aspired utility, then a higher utility might also be an indication of higher amount of selfishness. The utility function is private information for each peer.

3.3 Problem Statement

Following the round of submission of bandwidth requests, each peer decides on the granted amounts of bandwidth in response to those requests. Each peer selects the amount it will upload to other peers based on maximizing its utility function. In the case of (1), where a peer i can control its derived utility, the obvious solution that maximizes $U_i(\cdot)$ is $\mathbf{y}_i = \mathbf{0}$, namely the peer should not allow uploads from other peers. Thus, in the absence of an incentive protocol, we are led to the phenomenon of free-riding, where ultimately no peer obtains utility.

In the presence of an incentive protocol, the resource allocation decisions of a peer become more complicated, given that both uploads and downloads are realized through a common link. On the one hand, the peer would like to assign a small portion of link bandwidth to other peers’ uploads so as to have spare bandwidth for its own downloads that provide direct utility. On the other hand, however, the peer would like to grant to other peers just enough bandwidth so that it does not get punished by the incentive protocol. Otherwise, in the presence of the incentive mechanism and due to reciprocity, other peers will not respond to this peer’s request in the future, or they will allocate small portion of their bandwidth, and the peer of interest will obtain small utility in the future. A peer hopes that, by serving with high bandwidth a peer that has served him well in the past, it is likely to ensure an even better treatment in the future. Addressing this tradeoff between instantaneous received utility and future anticipated utility is the objective of incentive mechanisms. In the context of a peer-to-peer environment and in the absence of centrally exercised control, fundamental questions are: how can the optimization problem above be solved in a distributed fashion that accounts for autonomous unsupervised operation of peers? How does the incentive protocol naturally come into stage in this optimization? In the sequel, we attempt to shed light in these two questions.

We start from the global, network-wide goal of maximizing social welfare, namely the sum of peer utilities, subject to link capacity constraints of each peer. The selection of this objective will allow the identification of a globally optimal operating point in terms of bandwidth sharing. For option A of the utility function, this problem can be formulated as:

$$\max_{\mathbf{y}_i, i=1, \dots, N} \sum_{i=1}^N U_i(\mathbf{y}_i, \tilde{\mathbf{y}}_i), \quad (3)$$

subject to:

$$\sum_{j \neq i} y_{ij} + \sum_{j \neq i} y_{ji} \leq C_i, \quad i = 1, \dots, N. \quad (4)$$

with $\mathbf{y}_i \geq \mathbf{0}$ and $\tilde{\mathbf{y}}_i \geq \mathbf{0}$. We refer to this problem as problem (P). In case of option B for the utility function, the objective in (3) is replaced by $U_i(\tilde{\mathbf{y}}_i)$.

4. DISTRIBUTED BANDWIDTH ALLOCATION

Our aim is to devise distributed algorithms to solve this problem. We consider dual decomposition methods as most appropriate for solving the problem for reasons that will become clear in the sequel. The distributed solution is attained by the circulation of implicit signals in the network about the behavior of peers, such that each peer solves separate optimization problems. By the formulation above, it can be observed first, that there exists coupling in the constraints. For each i , the function on the left-hand side of (6) depends on decisions $\{y_{ji}\}$ of peers other than i . This fact per se does not cause any difficulty in making the problem amenable to distributed solution. In fact, relaxing link capacity constraints and assigning a Lagrange multiplier to each one of them, allows for the decomposition of the original problem to optimization problems, each of which can be solved separately by every peer. Lagrange multipliers updates enable coordination among nodes [19], [13]. Note however that in such cases, the utility of a node depends only on decisions taken by that peer (e.g. in [19] is the source rate).

The challenge in our problem lies in the fact that, besides the constraints, the objective functions of peers are coupled as well. That is, the utility function of each peer i depends on allocation decisions $\tilde{\mathbf{y}}_i = (y_{ji} : j \neq i)$ of other peers, in addition to local decisions \mathbf{y}_i .

In order to decouple peer utilities, we adopt an alternative decomposition method in the spirit of the ones presented in, [20]. For each peer i , and for each external variable y_{ji} of i , we start by introducing auxiliary variables z_{ij} , whose role is to turn the external (for peer i) variables, y_{ji} , into local ones that can become part of the local optimization problem of peer i . Let $\mathbf{z}_i = \tilde{\mathbf{y}}_i$ denote the vector of auxiliary variables for peer i . The introduction of these variables adds new constraints in the formulation and leads to a new maximization problem (for utility option A), which we refer to as problem (Q):

$$\max_{\mathbf{y}_i, \mathbf{z}_i, i=1, \dots, N} \sum_{i=1}^N U_i(\mathbf{y}_i, \mathbf{z}_i), \quad (5)$$

subject to:

$$\sum_{j \neq i} y_{ij} + \sum_{j \neq i} y_{ji} \leq C_i, \quad i = 1, \dots, N. \quad (6)$$

$$z_{ij} = y_{ji}, \quad \forall i, j, \quad \text{with } i \neq j, \quad (7)$$

with $\mathbf{y}_i \geq \mathbf{0}$ and $\mathbf{z}_i \geq \mathbf{0}$.

For the problem above, let u^* denote its optimal value, namely the value of the objective function at the optimal solution. The utility of each peer becomes a function of only local variables, and the coupling has been transferred only to the constraints. Now the problem can be tackled by classical decomposition methods. To proceed, we relax only the problem's constraints. Consider the Lagrangian function for problem (Q),

$$\begin{aligned} L(\mathbf{Y}, \mathbf{Z}, \mathbf{R}, \boldsymbol{\lambda}) &= \sum_{i=1}^N U_i(\mathbf{y}_i, \mathbf{z}_i) + \sum_{i=1}^N \sum_{j=1}^N r_{ij} (y_{ji} - z_{ij}) \\ &+ \sum_{i=1}^N \lambda_i (C_i - \sum_{j \neq i} y_{ij} - \sum_{j \neq i} y_{ji}) \end{aligned} \quad (8)$$

where \mathbf{Z} denotes the matrix of auxiliary variables $\mathbf{Z} = (\mathbf{z}_i : i = 1, \dots, N)$, and $\boldsymbol{\lambda} \geq \mathbf{0}$ is the vector of Lagrange multipliers corresponding to capacity constraints. Also $\mathbf{R} = (r_{ij} : i, j = 1, \dots, N)$ is the $N \times N$ matrix of Lagrange multipliers (dual variables) r_{ij} corresponding to the newly introduced equality constraints. Note that $r_{ii} = 0$ for all $i = 1, \dots, N$. Let $\Omega = \{(\mathbf{Y}, \mathbf{Z}) : \mathbf{y}_i \geq \mathbf{0}, \mathbf{z}_i \geq \mathbf{0}, \forall i\}$, and let $\Omega_i = \{(\mathbf{y}_i, \mathbf{z}_i) \geq (\mathbf{0}, \mathbf{0})\}$.

The dual function is defined as,

$$g(\mathbf{R}, \boldsymbol{\lambda}) = \max_{(\mathbf{Y}, \mathbf{Z}) \in \Omega} L(\mathbf{Y}, \mathbf{Z}, \mathbf{R}, \boldsymbol{\lambda}) \quad (9)$$

and the dual problem corresponding to the primal problem (Q) is

$$\min_{\boldsymbol{\lambda} \geq \mathbf{0}, \mathbf{R}} g(\mathbf{R}, \boldsymbol{\lambda}). \quad (10)$$

Let d^* be the optimal value of the dual problem. Attempting to solve the original problem via its dual is a suitable approach; the dual problem is always a convex optimization problem since the objective to be minimized is always convex. From weak duality, it is $u^* \leq d^*$. When solving a problem through its dual, one is interested in exploring whether strong duality holds, i.e. $u^* = d^*$, so that by solving the dual, one can obtain the optimal solution. It turns out that, for the presented problem this is the case, since it falls within the context discussed in [21, pp.514-515], namely the constraints are linear functions of variables and the utility functions $U_i(\cdot)$ are concave over $\mathbb{R}^{2(N-1)}$ (the dimension of variables $(\mathbf{y}_i, \mathbf{z}_i)$).

4.1 Distributed algorithm

In dual decomposition, the basic idea is to come up with the primal-dual optimal points of the Lagrangian, $L(\cdot)$. By weak duality, we have

$$\max_{(\mathbf{Y}, \mathbf{Z}) \in \Omega} \min_{\boldsymbol{\lambda} \geq \mathbf{0}, \mathbf{R}} L(\mathbf{Y}, \mathbf{Z}, \mathbf{R}, \boldsymbol{\lambda}) \leq u^* \leq \min_{\boldsymbol{\lambda} \geq \mathbf{0}, \mathbf{R}} \max_{(\mathbf{Y}, \mathbf{Z}) \in \Omega} L(\mathbf{Y}, \mathbf{Z}, \mathbf{R}, \boldsymbol{\lambda}).$$

Under strong duality, the inequality becomes equality. A point $(\mathbf{Y}^*, \mathbf{Z}^*, \mathbf{R}^*, \boldsymbol{\lambda}^*)$ is primal-dual optimal for the problem if and only if it is a saddle point of $L(\mathbf{Y}, \mathbf{Z}, \mathbf{R}, \boldsymbol{\lambda})$, namely it satisfies

$$L(\mathbf{Y}, \mathbf{Z}, \mathbf{R}^*, \boldsymbol{\lambda}^*) \leq L(\mathbf{Y}^*, \mathbf{Z}^*, \mathbf{R}^*, \boldsymbol{\lambda}^*) \leq L(\mathbf{Y}^*, \mathbf{Z}^*, \mathbf{R}, \boldsymbol{\lambda}^*).$$

That is, $(\mathbf{Y}^*, \mathbf{Z}^*)$ maximizes $L(\mathbf{Y}, \mathbf{Z}, \mathbf{R}^*, \boldsymbol{\lambda}^*)$, while the pair $(\mathbf{R}^*, \boldsymbol{\lambda}^*)$ with $\boldsymbol{\lambda}^* \geq \mathbf{0}$ minimizes $L(\mathbf{Y}^*, \mathbf{Z}^*, \mathbf{R}, \boldsymbol{\lambda}^*)$.

4.1.1 Maximization of Lagrangian w.r.t. (\mathbf{Y}, \mathbf{Z})

The dual function can be equivalently written as:

$$\begin{aligned} g(\mathbf{R}, \boldsymbol{\lambda}) &= \sum_{i=1}^N \max_{(\mathbf{y}_i, \mathbf{z}_i) \in \Omega_i} \{U_i(\mathbf{y}_i, \mathbf{z}_i)\} \\ &+ \frac{1}{2} \sum_{j \neq i} [r_{ij}(y_{ji} - z_{ij}) + r_{ji}(y_{ij} - z_{ji})] \\ &- \lambda_i \sum_{j \neq i} y_{ij} - \sum_{j \neq i} \lambda_j y_{ij} + \sum_{i=1}^N \lambda_i C_i. \end{aligned} \quad (11)$$

Thus the maximization of the dual function can be decomposed into separate maximization problems, each of which can be solved by each peer i in terms of its strategy $(\mathbf{y}_i, \mathbf{z}_i)$. Observe also that only the i th row, $\mathbf{r}_i = (r_{ij} : j = 1, \dots, N)$ and the i th column, $\tilde{\mathbf{r}}_i = (r_{ji}, j = 1, \dots, N)$ of the Lagrange

multiplier matrix \mathbf{R} arise in the optimization problem encountered by each peer i . For given vectors \mathbf{r}_i and $\tilde{\mathbf{r}}_i$ and for given multiplier vector $\boldsymbol{\lambda}$, each peer i separately solves the optimization problem,

$$\begin{aligned} \max_{(\mathbf{y}_i, \mathbf{z}_i) \in \Omega_i} \{ & U_i(\mathbf{y}_i, \mathbf{z}_i) + \frac{1}{2} \sum_{j \neq i} [r_{ij}(y_{ji} - z_{ij}) \\ & + r_{ji}(y_{ij} - z_{ji})] - \lambda_i \sum_{j \neq i} y_{ij} - \sum_{j \neq i} \lambda_j y_{ij} \}. \end{aligned} \quad (12)$$

This is a convex optimization problem whose solution depends on the form of utility function $U_i(\cdot)$. Consider stage t of the algorithm, where vectors $\mathbf{r}_i^{(t-1)}$, $\tilde{\mathbf{r}}_i^{(t-1)}$ and $\boldsymbol{\lambda}^{(t-1)}$ of the previous stage are known to peer i . Peer i comes up with the optimal (for stage t) resource provisioning vector $\mathbf{y}_i^{(t)}$ and the optimal vector $\mathbf{z}_i^{(t)}$. Let us omit superscripts $(t-1)$. Each peer i derives the optimal solution (for stage t) by applying the KKT conditions to obtain the following equations:

$$\begin{aligned} \frac{du_i(z_{ij})}{dz_{ij}} - \frac{1}{2}r_{ij} &= 0, \\ -\frac{dw_i(y_{ij})}{dy_{ij}} + \frac{1}{2}r_{ji} - \lambda_i - \lambda_j &= 0, \end{aligned} \quad (13)$$

for $j = 1, \dots, N$. Peer i can then determine \mathbf{y}_i and \mathbf{z}_i . These vectors for $i = 1, \dots, N$ are then fed into the dual problem (10).

4.1.2 Minimization of Lagrangian w.r.t. $(\mathbf{R}, \boldsymbol{\lambda})$

Let us fix now the values of (\mathbf{Y}, \mathbf{Z}) . Next, the Lagrange multiplier matrix \mathbf{R} and multiplier vector $\boldsymbol{\lambda}$ that minimize $g(\mathbf{R}, \boldsymbol{\lambda})$ need to be computed. First, note that the minimization can be performed in two stages,

$$\min_{\lambda \geq 0, \mathbf{R}} L(\mathbf{Y}, \mathbf{Z}, \mathbf{R}, \boldsymbol{\lambda}) = \min_{\lambda \geq 0} \min_{\mathbf{R}} L(\mathbf{Y}, \mathbf{Z}, \mathbf{R}, \boldsymbol{\lambda}) \quad (14)$$

and

$$\min_{\lambda \geq 0, \mathbf{R}} g(\mathbf{R}, \boldsymbol{\lambda}) = \min_{\lambda \geq 0} \min_{\mathbf{R}} g(\mathbf{R}, \boldsymbol{\lambda}). \quad (15)$$

For given vectors $\mathbf{y}_i^{(t)}$ and $\mathbf{z}_i^{(t)}$, the dual problem can be seen to simplify as follows:

$$\begin{aligned} \sum_{i=1}^N \sum_{j \neq i} \{ & [\min_{r_{ij}} r_{ij}(y_{ji}^{(t)} - z_{ij}^{(t)})] + \min_{r_{ji}} [r_{ji}(y_{ij}^{(t)} - z_{ji}^{(t)})] \} \\ & + \sum_{i=1}^N \min_{\lambda_i} \lambda_i (C_i - \sum_{j \neq i} y_{ij}^{(t)} - \sum_{j \neq i} z_{ji}^{(t)}). \end{aligned} \quad (16)$$

Regarding multipliers \mathbf{R} , a separate problem can be solved for each pair of peers i, j . If the dual function $g(\cdot)$ is differentiable, a typical approach is to perform one step of the iteration of a gradient-based method for r_{ij} and r_{ji} :

$$\begin{aligned} r_{ij}^{(t)} &= r_{ij}^{(t-1)} - s_t(y_{ji}^{(t)} - z_{ij}^{(t)}), \\ r_{ji}^{(t)} &= r_{ji}^{(t-1)} - s_t(y_{ij}^{(t)} - z_{ji}^{(t)}), \end{aligned} \quad (17)$$

where $\partial g(\cdot)/\partial r_{ij} = y_{ji}^{(t)} - z_{ij}^{(t)}$, $\partial g(\cdot)/\partial r_{ji} = y_{ij}^{(t)} - z_{ji}^{(t)}$ at iteration t , and $\{s_t\}_{t \geq 0}$ is a sequence of positive steps for the gradient algorithm. Similarly, in order to update the multipliers $\{\lambda_i\}$, one can perform a similar type of gradient

update for each λ_i ,

$$\lambda_i^{(t)} = \left[\lambda_i^{(t-1)} - s_t \left(C_i - \sum_{j \neq i} y_{ij}^{(t)} - \sum_{j \neq i} z_{ji}^{(t)} \right) \right]^+ \quad (18)$$

where $x^+ = x$ when $x > 0$, otherwise it is 0, and the projection accounts for maintaining non-negativity of λ_i .

4.2 The algorithm and various interpretations

We now proceed to the distributed algorithm for the problem in hand. We first explain various interesting attributes of our approach. Consider the auxiliary variables z_{ij} of the primal problem. The vector $\mathbf{z}_i = (z_{ij} : j = 1, \dots, N, j \neq i)$ can be computed by each peer i separately from (13). Each variable z_{ij} can be interpreted as the amount of resource (bandwidth) that peer i requests from peer j , $j \neq i$. Thus at each round of the algorithm, and for given values of Lagrange multipliers, each peer specifies the amounts of requested bandwidth \mathbf{z}_i from other peers and the amounts of granted bandwidth to other peers \mathbf{y}_i that maximize its utility. Therefore, for the auxiliary variables we have that $\mathbf{z}_i \equiv \mathbf{x}_i$.

Now fix attention to the updates of multipliers $\{r_{ij}\}$ and $\{r_{ji}\}$ for all (i, j) with $i \neq j$ in (17). At each stage, each peer i updates the multipliers $\{r_{ij}\}_{j \neq i}$ based on the amount of bandwidth z_{ij} that i requested from j and the amount of bandwidth y_{ji} that j granted to i . If $y_{ji} < z_{ij}$, that is, peer j does not fully satisfy the request of i , then r_{ij} is increased. On the other hand, if $y_{ji} > z_{ij}$, i.e. there is overprovisioning of resources from j to i , then r_{ij} is decreased. Intuitively, r_{ij} can be interpreted as the *reciprocal of reputation* of peer j as perceived by peer i , or in other words, the inverse of the opinion that i forms for j . The inverse reputation metric r_{ij} here is viewed as the degree to which peer j conforms to the requests of peer i and in some sense it quantifies the tendency of each peer to cooperate through resource provisioning. The dynamic update of metric r_{ij} is straightforward and does not need any other information besides the amounts of requested and provisioned bandwidth, z_{ij} and y_{ji} . The former quantity is known and the latter can be readily available by measurements of peer i on its access link.

A note about the update of the reputation metric is in place here. In case $y_{ji} = z_{ij}$, i.e. peer j fully satisfies the requests of peer i , we have reputation $r_{ij}^{(t)} = r_{ij}^{(t-1)}$. i.e. the reputation remains unchanged. Nevertheless, one could argue that r_{ij} should be decreased (so that its inverse, the usual reputation, will increase), since j fully served i 's requests. This issue can be taken into consideration simply by changing the definition of auxiliary variables in (7) to $z_{ij} = y_{ji} + \Delta$, where $\Delta > 0$. Then, the update of r_{ij} will become $r_{ij}^{(t)} = r_{ij}^{(t-1)} - s_t(y_{ji}^{(t)} + \Delta - z_{ij}^{(t)})$ and when $y_{ji}^{(t)} = z_{ij}^{(t)}$, the reputation will still decrease by Δ , so that its inverse, which amounts to the usual notion of reputation, will increase.

In the same spirit with r_{ij} , multiplier r_{ji} is also updated by peer i . This metric can be interpreted as the reciprocal of reputation of peer i as perceived by j . In other words, a peer i updates its reciprocal reputation in the eyes of j based on its response y_{ij} to peer j 's request z_{ji} . Again, the update is straightforward and requires knowledge of only z_{ji} and y_{ij} . Clearly, r_{ij} depends on peer j 's response y_{ji} to peer i 's re-

quest z_{ij} . Peer j may not fully satisfy peer i 's requests due to various reasons such as different (high) reciprocal reputation metrics r_{ij} or r_{ji} , limited capacity, large amount of received requests or due to the fact that j needs to satisfy peers other than i . Naturally, peer i is agnostic of the reason it is not fully served and adjusts its reputation metrics appropriately, so that future allocations are more suitable.

Finally, the dual variables $\{\lambda_i\}_{i=1,\dots,N}$ admit the (standard in literature) interpretation of prices of peer i 's access link and serve as indicators of the congestion experienced in peer i 's link. If the total load on each link increases, λ_i will increase to denote that it is "expensive" to use that link. In our framework, these variables are used by the peers simply to regulate their own resource allocation decisions (both the amounts of requested and provisioned resources). They are also communicated in the network in order to signal potential congestion to other peers so that they appropriately control their own decisions. In particular, variables λ_i attempt to balance the demand by discouraging peers from requesting resources from peers that are very likely to be extremely popular and experience high request load. Through appropriate calibration of multipliers $\{\lambda_i\}$, peers are encouraged to address requests to less loaded peers. However, in a different setting with resource price advertisement and payments, these variables may represent actual link usage prices.

Each peer i has a utility function $U_i(\cdot)$ and access link bandwidth C_i . These are kept private for each peer, so that peers other than i do not need to know this information. The basic steps of the algorithm are as follows:

- **Step 0:** Initialization. Set $t = 0$. Set the multipliers to initial values $\lambda_i^{(0)} \geq 0$ and set the reciprocal reputation matrix $\mathbf{R}^{(0)} = \alpha \mathbf{I}$, where \mathbf{I} is the unit matrix and $\alpha > 0$.
- **Step 1:** Each peer i , independently of others, solves locally its separate optimization problem in (12) and computes the optimal (for stage j) vectors $\mathbf{y}_i^{(t)}$ and $\mathbf{z}_i^{(t)}$ with the help of (13). Subsequently, it communicates each $z_{ij}^{(t)}$ to peers $j : z_{ij}^{(t)} > 0$ and grants bandwidth $y_{ij}^{(t)}$ to other peers.
- **Step 2:** Each peer i measures $\{y_{ji}^{(t)}\}_{j \neq i}$ and updates the reciprocal reputation metrics $\{r_{ij}\}$ and $\{r_{ji}\}$ in (17). These are kept as private information.
- **Step 3:** Each peer i updates its multiplier λ_i according to (18) and broadcasts it to the network.
- **Step 4:** $t \leftarrow t + 1$. If a termination condition is satisfied, STOP. Else, go to Step 1.

While in usual distributed optimization methods the signals that are circulated in the network are the Lagrange multipliers, in our formulation, the multipliers $\{r_{ij}\}$ are not circulated. Instead, it is the auxiliary variables $\{z_{ij}\}$ that are naturally circulated, as they denote resource requests of peers. In addition, a peer j needs to know only the amounts of resource requests $\{z_{ij}\}_{i \neq j}$ addressed to itself by other peers and not the amounts $\{z_{ik}\}$, $i, k \neq j$ requested from other peers. As mentioned above, vectors $\mathbf{z}_i = (z_{ij} : j \neq i)$ and $\tilde{\mathbf{y}}_i = (y_{ji} : j \neq i)$ are readily available or easily measurable on the access link of a peer i . Interestingly, Lagrange multipliers $\{r_{ij}\}$ and $\{r_{ji}\}$ that correspond to the auxiliary

variable constraints can be maintained as private information stored locally at each peer and used for a peer's own resource allocation decisions. This attribute is highly desirable in our case, as these multipliers can be interpreted as pairwise reputation metrics that should not be revealed in a network.

If the utility functions $U_i(\cdot)$ are strictly concave, differentiable and bounded, and the step size s_t for multiplier updates is appropriately chosen, the algorithm above converges to the optimal solution of the original problem (P). The step sizes need to be diminishing with t and should satisfy:

$$\sum_{t=1}^{\infty} s_t = \infty, \quad \sum_{t=1}^{\infty} s_t^2 < \infty. \quad (19)$$

These requirements are satisfied for the sequence of steps $\{s_t\}_{t>0}$ with $s_t = 1/t$ or $s_t = (1+m)/(t+m)$, where $m > 0$ is a non-negative number.

If the utility functions $U_i(\cdot)$ are differentiable, the separate optimization problems (12) by each peer i in Step 1 can be solved by means of KKT conditions. Furthermore, the updates of the Lagrange multipliers in (17) and (18) in Steps 2 and 3 are performed with the gradient descent method outlined above since the gradient exists. If the utility functions are non-differentiable, the gradient does not exist in general, and the sub-gradient method [22], [15] needs to be employed for the optimization in Steps 1,2 and 3. For a function $f(\mathbf{x})$ a sub-gradient $\boldsymbol{\omega}$ at point \mathbf{x}_0 is a vector that satisfies $f(\mathbf{x}) \leq f(\mathbf{x}_0) + \boldsymbol{\omega}^T(\mathbf{x} - \mathbf{x}_0)$, for all \mathbf{x} .

For iteration t of Lagrange multipliers in Step 3, we start by computing a sub-gradient of $g(\cdot)$, $\boldsymbol{\omega}^{(t)}$ with respect to $\boldsymbol{\lambda}$ at point $\boldsymbol{\lambda}^{(t-1)}$. It can be easily shown that such a sub-gradient is

$$\boldsymbol{\omega}^{(t)} = \mathbf{C} - \boldsymbol{\sigma}^{(t)} - \tilde{\boldsymbol{\sigma}}^{(t)}, \quad (20)$$

where $\boldsymbol{\sigma}$, $\tilde{\boldsymbol{\sigma}}$ are vectors of dimension N whose i th component denotes the sum of elements of vector \mathbf{y}_i and the sum of elements of vector $\tilde{\mathbf{y}}_i$ respectively, and $\mathbf{C} = (C_1, \dots, C_N)$. Finally, an iteration similar to the one in (18) is obtained for each λ_i separately. Using similar arguments for the iteration of multipliers in Step 2, we can find a sub-gradient of $g(\cdot)$ with respect to \mathbf{R} at point $\mathbf{R}^{(t-1)}$ and come up with iterations similar to the ones in (17). For the separate optimization problem (12) of Step 1, we again compute sub-gradients of $g(\cdot)$ with respect to \mathbf{y}_i and \mathbf{z}_i at point $\mathbf{y}_i^{(t-1)}$ and $\mathbf{z}_i^{(t-1)}$. It turns out that a sub-gradient of $g(\cdot)$ with respect to \mathbf{y}_i is $\frac{1}{2}\tilde{\mathbf{r}}_i^{(t-1)} - \boldsymbol{\Lambda}^{(t-1)} + \boldsymbol{\omega}_u$, where $\boldsymbol{\Lambda}$ is the vector whose k th component is $\lambda_i + \lambda_k$, and $\boldsymbol{\omega}_u$ a sub-gradient of function $u(\cdot)$ that appears in the utility function (1). Also, a sub-gradient of $g(\cdot)$ with respect to \mathbf{z}_i is $-\frac{1}{2}\mathbf{r}_i^{(t-1)} + \boldsymbol{\omega}_w$ with $\boldsymbol{\omega}_w$ a sub-gradient of function $w(\cdot)$ that appears in the utility function (1). The distributed algorithm with sub-gradient updates also converges to the optimal solution of problem (P).

5. NUMERICAL RESULTS

We evaluate the performance of our approach in a simulated system of 10 peers. We consider two different scenarios for the utility $U(\cdot)$:

- Each peer differs from others in its function $u_i(\cdot)$. We consider $u_i(x) = G_i \log(1+x)$ where $G_i > 0$ is a private parameter that captures the valuation of bandwidth

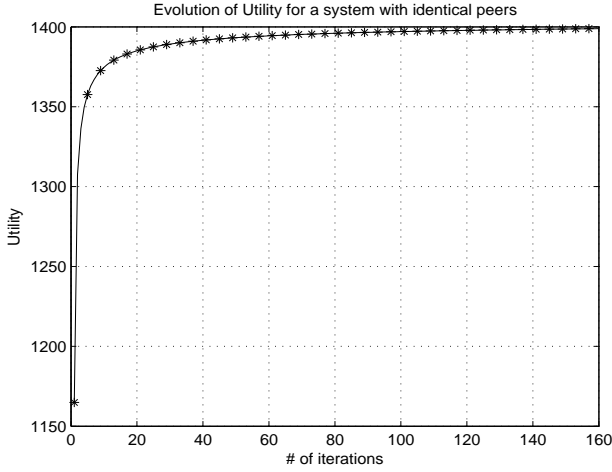


Figure 2: Utility vs. number of iterations in a system with identical peers

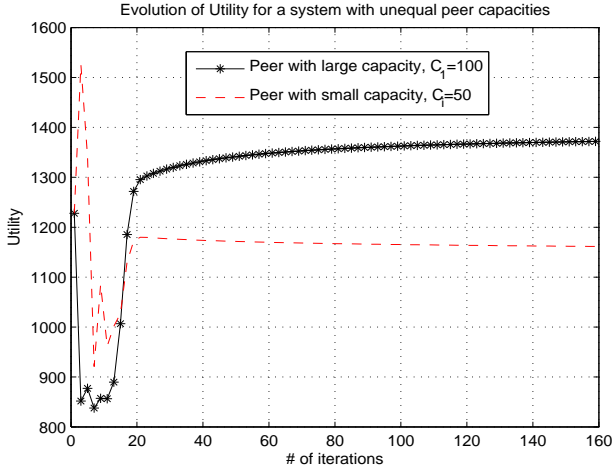


Figure 3: Utility vs. number of iterations in a system with different link capacities

by each peer i . This choice of utility function models peers with different bandwidth needs that places more emphasis on the amount of download.

- Each peer differs in its function $w_i(\cdot)$. We consider $w_i(y) = G_i y^2$, where $G_i > 0$. In a sense, this choice models selfish peers with different degree of “reluctance” to let uploads through their link.

For these choices of utility functions, we obtain from (13),

$$z_{ij} = \left(\frac{2G_i}{r_{ij}} - 1 \right)^+,$$

$$\text{and } y_{ij} = \frac{1}{2G_i} \left(\frac{1}{2} r_{ji} - \lambda_i - \lambda_j \right)^+. \quad (21)$$

We observe that a peer i aims at placing more requests to peers j with low r_{ij} . This accounts for an intuitive selection of peers as servers based on how much reputable they are in the eyes of peer i . Also note that, the higher the value

of G_i of a peer i , the larger the amount of requests of that peer. On the other hand, in deciding the amounts of granted bandwidth, a peer i wishes to do that so that it satisfies (and gains the attention of) peers which, according to its perception, do not consider peer i to be reputable. This can be seen from the fact that y_{ij} is proportional to r_{ji} , the reciprocal reputation that of i in the eyes of j . Note also, that a peer refrains from providing large amount of bandwidth if its link or peer j 's link are signaled as being congested, i.e the values of λ_i and λ_j are large.

The simulation model consists of $N = 10$ peers which interact in successive rounds according to the proposed protocol. We assume that the communication is synchronous. Thus, in each round, every peer's requests and allocations are successfully communicated to the intended receivers before the next round takes place. The basic attributes of each peer are the capacity C_i of the access link through which it is connected to the network and the parameter G_i which differentiates, either through $u_i(\cdot)$ or through $w_i(\cdot)$, its utility function. We consider the two different concave utility functions:

$$(A1) : U_i(\mathbf{z}_i, \mathbf{y}_i) = G_i \sum_{j \neq i} \log(1 + z_{ij}) - \alpha_1 \sum_{j \neq i} y_{ij}^2,$$

and

$$(A2) : U_i(\mathbf{z}_i, \mathbf{y}_i) = \alpha_2 \sum_{j \neq i} \log(1 + z_{ij}) - G_i \sum_{j \neq i} y_{ij}^2,$$

where α_1 and α_2 are scaling parameters and do not affect the conclusions.

First, we simulate a system in which all peers are described by utility functions of type (A1) and have identical profiles, i.e. they have equal link capacities, $C_i = 100$ units, and the same utility functions, i.e. parameter $G_i = 100$. In figure 2 the evolution of utility for every peer is shown. The system is gradually driven to a steady state while the utilization of each link is approximately 100%. Note that for this case, when all peers are identical, exactly the same results hold also for utility functions of type (A2). In figure 3 we simulate a system where a peer has link capacity equal to 50 units, lower than the capacity of the rest, which is 100 units. This peer submit smaller requests, grants less bandwidth to other peers, has worse reputation and eventually obtains less utility.

The case where different peers have different utility functions is studied in figures 4 and 5. In figure 4 we show the evolution of utility for a system where peers have utility functions of type (A1), equal link capacities, $C_i = 100$, but different values of G_i . Specifically, $G_1 = 80$, $G_2 = 120$, and $G_i = 100$ for every other peer. We observe that the utility is proportional to the value of the G parameter. The peer with the smallest G , submits the smallest requests, $x_{1j} = 4.68$ (compared to $x_{2j} = 5.54$), allocates the largest bandwidth, $y_{1j} = 5.90$ (compared to $y_{2j} = 5.22$), and has the best reputation as it is perceived by the other peers, $r_{j1} = 29.76$ (compared to $r_{j2} = 32.58$). Recall that r is the inverse reputation. It is interesting to notice that the peers submit the largest requests to the peer with the smallest G , $x_{j1} = 5.71$ while $x_{j2} = 5.13$, because they realize that this peer is more willing to serve. Similar results hold for the case where the peer utility functions are of type (A2). In this case, the peer with the largest value of G , $G_1 = 0.08$, submits the largest amount of requests and receives more bandwidth than other

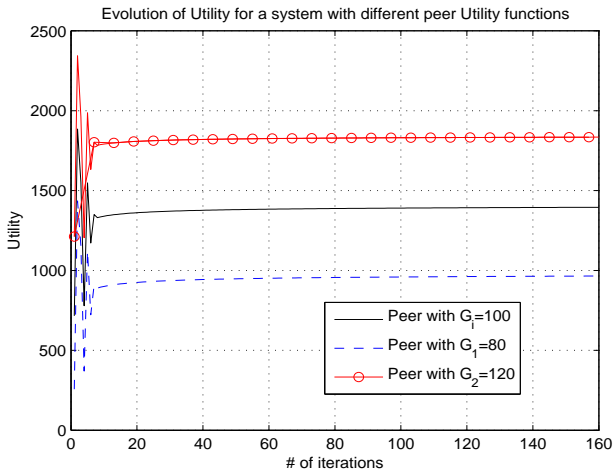


Figure 4: Utility vs. number of iterations in a system with type (A) utility functions and different G parameters.

peers. From the figures above, we observe that the peers with the largest utility functions, i.e. larger G for type (A1) or smaller G for type (A2), obtain more resources and finally higher perceived utility. This is expected for our problem formulation, (P), as social welfare maximization.

6. DISCUSSION

We presented a distributed network utility optimization framework for addressing the link bandwidth sharing problem that arises in peer-to-peer networks, where the finite-capacity access link from each peer to the backbone network is the bottleneck. Our algorithm allows dynamic and flexible control of the portions of bandwidth that are dedicated to upload and download flows of different peers. Our optimization framework leads to an insightful interpretation of Lagrange multipliers as reciprocal pairwise reputation metrics whose temporal evolution reflects the extent to which a peer responds to requests of others. This approach leads us to a meaningful protocol for autonomous bandwidth sharing among peers, which relies on readily available or measurable quantities. The protocol is lightweight in terms of computational complexity and overhead. Furthermore, it does not require pairwise reputation metrics to be broadcasted or circulated in the network; reputations as well as utilities are hidden information for each peer. The proposed protocol addresses the general case where uploads and downloads take place over the same link. The case where upload and download streams flow over two separate links can also be addressed by including in the formulation two separate link capacity constraints, one for the upload and one for the download traffic.

There exist several directions for future study. The selection of social welfare as our objective allows us to identify a socially optimal operating point in terms of bandwidth allocation and reach it by means of a distributed algorithm. As revealed by the numerical results, peers with larger utility function obtain more bandwidth; this is reasonable, due to the social welfare objective of the problem. In that sense, the reputation metrics do not have the impact of an incentive

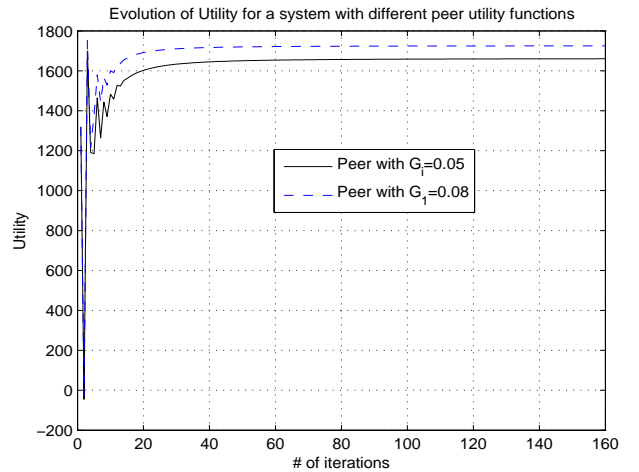


Figure 5: Utility vs. number of iterations in a system with type (B) utility functions and different G parameters.

mechanism that would reward cooperating peers and would punish selfish ones. The role and interpretation of reputations would actually be more insightful if we could alter the objective function and/or the constraints of the original problem such that some notion of fairness is enforced. For instance, an additional constraint could be that the obtained utility of a less collaborating or “selfish” peer at the optimal operating point be no more than that of a more collaborating peer. It would be interesting to see how can such a case become amenable to distributed implementation and how a reputation mechanism essentially guides the system to efficient operating points.

Our protocol does not distinguish among different content items that are possessed by peers. The approach could be extended to the more realistic situation of different items such as files residing (fully or partially) at different peers in the network. The obtained utility will then depend not only on the amount of provisioned bandwidth, but also on different items, in the sense that certain items may be valued more by certain peers than others. One way of addressing this issue is through several classes of items with e.g. different popularity. Pairwise reputation will then depend on individual exchanged items as well. This more general aspect gives rise to more generalized decision strategies. Now each peer has to decide on the item it will grant besides the bandwidth. In order to increase its reputation, it might be meaningful for a peer to serve requests less frequently or with less bandwidth, yet it can choose to provision very popular files. In that context, it would be interesting to study the long-term temporal behavior of utility and reputation, given that item popularity is itself time-varying, with very high peaks for a while after being posted in the network and decreasing later.

The model can be enhanced to account for a dynamic population where peers enter or leave the system. It could also include an average request load rate that each peer attempts to satisfy or it could incorporate recommendations from third parties in addition to pairwise experience. Moreover, it would be interesting to relax the synchronized (slotted) operation assumption and investigate the impact of de-

layed message passing in the overall protocol efficacy. In the case of wireless access links of peers, channel impairments such as packet collisions or interference will cause the amount of granted bandwidth as perceived by the grantee to differ from the amount of truly granted bandwidth. This will create discrepancies in the values of reciprocal pairwise reputations $\{r_{ij}\}$ and $\{r_{ji}\}$ that are computed by involved peers i and j that were not assumed in our approach. It would be interesting to analyze the impact of such a phenomenon on the convergence of the algorithm.

In the proposed algorithm, peers are rational, possibly selfish, but honest in the sense that they decide on the requested amounts based on (12). Thus, a peer is not prevented from placing false requests in an effort to disorientate the protocol. It would be interesting to devise rules by which peers would be able to determine whether another peer does not follow the protocol in terms of this aspect. Finally, our approach can give rise to a class of protocols with different definitions of auxiliary variables that in turn result in different constraints and different modes of reputation metric update. This was briefly mentioned in the paper in section 4.2 with parameter Δ . It would be interesting to study the dependence of the operating point to which the distributed algorithm converges from the definition of auxiliary variables.

7. ACKNOWLEDGEMENTS

The authors acknowledge support from the European Commission through the STREP project NET-REFOUND (FP6-IST-034413), STREP project NANDA (FP7-ICT-223850), and the NoE project NEWCOM++ (FP7-ICT-216715).

8. REFERENCES

- [1] N. Laoutaris, P. Rodriguez and L. Massoulie, "ECHOS: Edge Capacity Hosting Overlays of Nano data centers", *ACM SIGCOMM Comp. Comm. Review*, vol.38, no.1, pp.51-54, Jan. 2008.
- [2] M. Meo, and F. Milan, "A Rational Model for Service Rate Allocation in Peer-to-Peer Networks", *In Proc. of 8th IEEE Global Internet Symp., 2005* March 2005.
- [3] <http://www.fon.com>.
- [4] S. Marti, and H. Garcia-Molina, "Limited reputation sharing in P2P systems", *In Proc. of 5th ACM Conf. on Electr. Commerce*, 2004.
- [5] B. Mortazavi, and G. Kesidis, "Incentive-compatible cumulative reputation systems for peer-to-peer file-swapping", *In Proc. IEEE CISS*, 2006.
- [6] A. Satsiou, and L. Tassiulas, "A Trust-Based Exchange Framework for Multiple Services in P2P Systems", *In Proc. of 7th Intern. Conf. on P2P Comp.*, 2007.
- [7] S. Marti, T. J. Giuli, K. Lai, and M. Baker, "Mitigating Routing Misbehavior in Mobile Ad hoc Networks", *In Proc. of 6th MobiCom*, 2000.
- [8] F. Milan, J. J. Jaramilo, and R. Srikant, "Achieving cooperation in multihop wireless networks of selfish nodes", *In Workshop GameNets*, 2006.
- [9] J. J. Jaramilo, and R. Srikant, "DARWIN: Distributed and Adaptive Reputation Mechanism for Wireless Adhoc Networks", *Proc. of 13th MobiCom*, 2007.
- [10] R.T.B. Ma, S.C.M. Lee, J.C.S. Lui, and D.K.Y. Yau, "An Incentive Mechanism for P2P Networks", *In Proc. Intern. Confer. Distr. Comp. Systems (ICDCS)*, pp.516-523, 2004.
- [11] M. Adler, R. Kumar, K. Ross, D. Rubenstein, T. Suel and D.D. Yao, "Optimal peer selection for P2P downloading and streaming", *In Proc. INFOCOM*, 2005.
- [12] K. Eger, and U. Killat, "TCPeer: Rate Control in P2P over IP Networks", *LNCS 4516, In 20th ITC20*, 2007.
- [13] F.P. Kelly, A. Maulloo, and D. Tan, "Rate Control for communication networks: Shadow prices, proportional fairness and stability", *J. Oper. Res. Soc.*, vol.49, no.3, pp.237-252, 1998.
- [14] M. Chiang, S. H. Low, A. R. Calderbank, and J. C. Doyle, "Layering as optimization decomposition: A mathematical theory of network architectures", *Proc. IEEE*, vol. 95, no. 1, pp. 255-312, Jan. 2007.
- [15] B. Johansson, P. Soldati and M. Johansson, "Mathematical decomposition techniques for distributed cross-layer optimization of data networks", *IEEE J. Select. Areas Commun.*, vol. 24, no. 8, pp. 1535-1547, Aug. 2006.
- [16] G. Iosifidis and I. Koutsopoulos, "Reputation-assisted utility maximization algorithms for peer-to-peer networks", *In Proc. IEEE IWQoS*, 2008.
- [17] C. Courcoubetis, and R. Weber, "Incentives for large peer-to-peer systems", *IEEE J. on Select. Areas Commun.*, vol.24, no.5, pp.1034-1050, May 2006.
- [18] D. R. Figueiredo, J. K. Shapiro, and D. Towsley, "A public good model of availability in peer to peer systems", *Tech. Rep. 04-27, CSE Dept*, Michigan State University, 2004.
- [19] S. Low, D.E. Lapsley, "Optimization-based flow control I: Basic algorithms and convergence", *IEEE/ACM Trans. Networking*, vol.7, no.6, pp.861-874, Dec. 1999.
- [20] D. Palomar, and M. Chiang, "Alternative decompositions for distributed maximization of network utility: framework and applications", *In Proc. INFOCOM*, 2006.
- [21] D.P. Bertsekas, *Non-linear Programming*, Athena Scientific, 2003.
- [22] S. Boyd and L. Vandenberghe, *Convex Optimization*, Cambridge, U.K.: Cambridge Univ. Press, 2004.