

A Separation Principle for Optimal IaaS Cloud Computing Distribution

Felix Kottmann, Saverio Bolognani, Florian Dörfler
Automatic Control Laboratory, ETH Zürich
Zürich, Switzerland
Email: {felixk, bsaverio, dorfler}@ethz.ch

Abstract—Due to the raising importance of cloud computing and infrastructure as a service (IaaS), the first markets for the exchange of computational power over the internet are being implemented. As of today, bandwidth constraints are not explicitly embedded in these market mechanisms. In this paper, the problem of optimal allocation of the computing power and of the corresponding data flows, according to bandwidth and computing capacity constraints, is modeled as a bilevel optimization program. It is shown that this program, which is generally non convex and hard to solve, has the same optimal solution of its convex relaxation. This allows to state a fundamental separation result, showing how the congestion control protocols employed in the network do not affect the optimal allocation problem, and allows to compute the shadow prices of the available computational resources.

Index Terms—Cloud computing, network control, congestion control, bilevel optimization, IaaS.

I. INTRODUCTION

Infrastructure as a Service (IaaS) is a form of cloud computing that provides virtualized computing resources over the Internet (e.g., Amazon Web Services). In the recent years, this paradigm has proved to be an effective way to exploit the economy of scale of big providers of computational services, while allowing smaller clients to avoid upfront infrastructure costs for their specific computational tasks.

The growing number of participating users (including consumers, providers, and users with mixed behavior) has recently led to the commoditization of cloud resources, and thus to the standardization of the provided service performance, and to the creation of exchange markets, e.g., the Deutsche Börse Cloud Exchange [1] and the Universal Compute Xchange [2].

This invites the fundamental question of what is the economic efficiency of the allocations returned by this market mechanisms, considering the complexity of the problem of exchanging computational services over a shared communication network. While the scheduling and distribution of computational tasks across different computing units is made possible by some specialized middleware (e.g. Openstack), the control and management of the corresponding data flow over the Internet can be a much more challenging task. In fact, the bandwidth available for the exchange of necessary data is affected by the data flows of other users, by the general

Internet traffic, and also by the congestion control protocols that are employed in the data network.

In this paper, we model this cross-layer interaction using the tools of bilevel optimization, a generalization of Stackelberg games. By showing that the social optimum problem is an exact convexification of the bilevel problem, we obtain a fundamental separation result: the optimal allocation is provably independent on the congestion control mechanisms employed by the data network. This result allows to shed some light on the economic efficiency of an allocation, via the identification of shadow prices for the computational resources that reflect the effect of both service costs and limited bandwidth.

II. NETWORK ARCHITECTURE AND CONGESTION CONTROL

We consider a data network operated according to the OSI-layer model, and in particular to the TCP/IP network protocol stack, where TCP implements the transportation layer while the IP implements the internet layer [3].

The network is described as an undirected graph $\mathcal{G} = (\mathcal{N}, \mathcal{L})$ with $|\mathcal{N}|$ nodes (vertices) and $|\mathcal{L}|$ links (edges).

We assume that routing is fixed for the time scale of interest, and we therefore do not model dynamic routing mechanisms.

Definition 1 (Routing matrix).

The routing matrix $R \in \{0, 1\}^{|\mathcal{L}| \times |\mathcal{N}|^2}$ is defined as

$$R_{\ell, ij} = \begin{cases} 1 & \text{if } \ell \in \mathcal{P}_{ij} \\ 0 & \text{otherwise,} \end{cases}$$

where $\mathcal{P}_{ij} \subset \mathcal{L}$ is the set of links that are traversed by the data flow from node i to node j .

A. Congestion Control

To avoid congestion in a TCP network, a distributed mechanism is used to find a fair equilibrium on the connections. In practice, these distributed mechanisms control the data flow on a path \mathcal{P}_{ij} based on a feedback measurement performed on the path itself, typically the packet loss rate or the round-trip delay. In this work, we consider the abstraction proposed in [4], [5] where the congestion control protocol is interpreted

as a distributed algorithm for the solution of the following Network Utility Maximization program,

$$\begin{aligned} & \max_{b_{ij} \geq 0} \sum_{i,j \in \mathcal{N}} W_{ij}(b_{ij}) \\ & \text{subject to } Rb \leq B \end{aligned} \quad (1)$$

where b_{ij} is to be interpreted as the data rate of the communication from i to j , the functions W_{ij} are concave and strictly increasing, and $B \in \mathbb{R}^{|\mathcal{L}|}$ is the vector containing the bandwidth capacity B_ℓ of each link. Here and in the rest of the paper, we use the notation b (without indices) to represent the vector obtained by stacking a set of scalars $\{b_{ij}\}$.

Following this abstraction, the congestion control mechanisms can be interpreted as *distributed primal-dual algorithms* that solve (1). For the analysis presented in this paper, it suffices to say that different TCP protocols correspond to different choice of physical signals as Lagrange multipliers (e.g. packet loss in TCP Reno and queuing delay in TCP Vegas) and different choices of the cost function W_{ij} (e.g., \tan^{-1} in TCP Reno and \log in TCP Vegas).

III. PROBLEM FORMULATION

In this section, we formulate the problem of optimal allocation of computing resources in a data network, subject to bandwidth constraints. We first present a model for the distribution of computational jobs, and for the corresponding requirements in terms of bandwidth. We then formulate the resulting social utility maximization problem, whose solution is the optimal allocation that we are seeking. This is of immediate interest in those scenarios where users act as cooperative players (e.g., Worldwide LHC Computing Grid by CERN, D-Grid Initiative by German academic institutions). In other cases, users act as competitive players. In this latter context, the social utility maximization problem serves as a benchmark: we will show in Section VI that the solution of its dualized version returns the *shadow prices* under which the social optimum and the competitive equilibrium correspond.

A. Computational load

Let any computational job have the size D , counted in operations or instructions, or any other customized metric. We define the computational power $d(t)$ as a service rate that amounts to the job size,

$$\int_{t_0}^{t_c} d(t) dt = D, \quad (2)$$

where the job starts at t_0 and ends at t_c .

Each node i of the graph \mathcal{G} corresponds to a user.¹ We define by $d_{ij} \geq 0$ the computational load that user i executes on the computational resources of agent j . In order to do that, user i needs to reserve a computational capacity $p_{ij} \geq 0$ at node j (corresponding for example to the reservation limits of

¹In the model we do not differentiate between users that need (and therefore buy) computational power, those that sell it, and those that do not participate to any exchange. Such a scenario is a special case of our model.

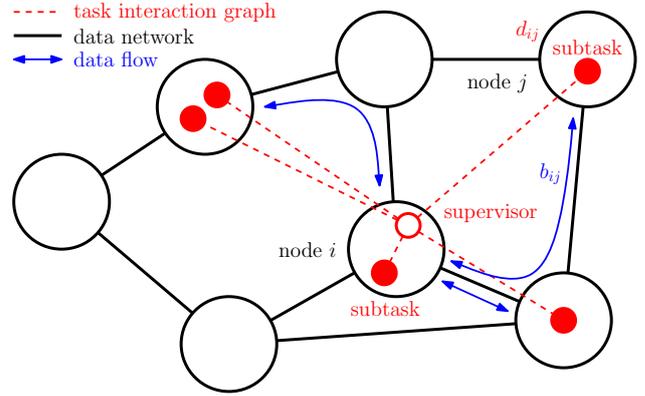


Fig. 1. An example of distribution of the 5 subtasks of a job across the nodes of a data network. In this case, the *task interaction graph* is a star.

the corresponding virtual machine). Each node $j \in \mathcal{N}$ has a maximal available computational capacity P_j . Clearly, it holds

$$0 \leq d_{ij} \leq p_{ij} \quad \text{and} \quad \sum_{i \in \mathcal{N}} p_{ij} \leq P_j.$$

B. Distribution of a computational job

A job is assumed to consist of subtasks, which are programmer defined and are considered to be indivisible, and that can be executed on separate computational resources available in the network. For the distribution of the subtasks to different nodes, the following assumptions are considered [6].

- User i 's job is managed at node i , where input data is available and output data needs to be collected.
- Maximal concurrency is achieved, therefore tasks can be executed in parallel.
- Subtasks need to exchange data during execution, according to a *task interaction graph* (e.g. Figure 1).
- All edges of the task interaction graph correspond to an identical requirement in terms of exchanged data.
- The communication between two subtasks allocated at different nodes i and j happens via the data network according to the routing matrix R , i.e. using the links belonging to the path \mathcal{P}_{ij} . Communication between two sub tasks on the same node does not use the network.

In this work, we consider the case where each job is decomposed according to a Data Parallel Model or a Master-Slave Model, which corresponds to the typical goal of promoting locality or reducing subtask interaction costs. Under these models, the task interaction graphs is a star graph, whose root is located at the owner node i .

By assuming a fine granularity and homogeneity in the subtask decomposition, we interpret d_{ij} as the number of subtasks that user i executes on node j . This corresponds to a bandwidth requirement which is linear in the assigned computational load, i.e.

$$b_{ij} = \alpha_i d_{ij}$$

for some α_i that depends on the specific job i .

In the rest of the analysis, we consider the generalization given by the following assumption.

Assumption 1 (Intra-task interaction). *Let d_{ij} be the computational load that user i executes on the computational resources of node j . Then the communication between node i and j amounts to*

$$b_{ij} = f_i(d_{ij}) \geq 0 \quad (3)$$

where f is a quasi-convex increasing function, with $f(0) = 0$.

C. Social optimum problem

To define the optimal distribution of subtasks over the available networked computational resources, a social utility maximization is formulated. In this problem, the sum of all users' utility minus the sum of all computational costs is maximized, subject to the network bandwidth constraints.

We assume that the cost of computation is an increasing convex cost of the reserved computational resources, either because of energy consumption costs or because of the opportunity costs connected to the reservation. For each node j , we then have the cost

$$g_j \left(\sum_{i \in \mathcal{N}} p_{ij} \right).$$

On the other hand, the utility provided to user i is an increasing concave function of the form

$$U_i \left(\sum_{j \in \mathcal{N}} d_{ij} \right).$$

The concavity assumption is satisfied in the very natural scenario in which users are interested in minimizing their completion time t_c as defined in (2). In fact, given the inversely proportional dependence of t_c on $\sum_j d_{ij}$, the marginal utility U'_i results to be positive and decreasing.

We therefore have the following convex program in the nonnegative decision variables $\{p_{ij}, d_{ij}, b_{ij}\}$, $i, j \in \mathcal{N}$.

$$\begin{aligned} \min_{\{p_{ij}, d_{ij}, b_{ij}\}} \quad & \sum_{j \in \mathcal{N}} g_j \left(\sum_{i \in \mathcal{N}} p_{ij} \right) - \sum_{i \in \mathcal{N}} U_i \left(\sum_{j \in \mathcal{N}} d_{ij} \right) \\ \text{subject to} \quad & \begin{cases} \sum_{i \in \mathcal{N}} p_{ij} \leq P_j & \forall j \in \mathcal{N} \\ d_{ij} \leq p_{ij} & \forall i, j \in \mathcal{N} \\ b_{ij} = f_i(d_{ij}) & \forall i, j \in \mathcal{N} \\ Rb \leq B \end{cases} \end{aligned} \quad (4)$$

We have the following technical result.

Lemma 1. *The optimization program (4) has a unique solution $\{p_{ij}^*, d_{ij}^*, b_{ij}^*\}$, which satisfies $p_{ij}^* = d_{ij}^*$ for all $i, j \in \mathcal{N}$.*

Proof: Existence and uniqueness of the solution descends from the convexity of the cost function, Assumption 1, and the fact that the feasible set is not empty (as $p_{ij} = d_{ij} = b_{ij} = 0$ belongs to it). Finally, assume $d_{ij}^* < p_{ij}^*$ for some i, j . Then a smaller $p_{ij} = d_{ij}^*$ would also be feasible, and would yield a smaller cost, leading to a contradiction. ■

IV. BILEVEL FORMULATION

The optimization problem (4) defines the optimal allocation of networked computational resources, but cannot be implemented in a real data network. In fact, it is not possible for the operator to directly decide on the amount of computation d_{ij} that user i can execute on node j , as this is also determined by the bandwidth limitations of the network. In this section, we model that by describing the internet congestion control as a nested optimization problem, that responds to the decided computing resource allocation (the decision variables p_{ij}). This leads to the following *bilevel optimization* problem [7], [8], a generalization of a Stackelberg game, which is in general non-convex.

$$\begin{aligned} \min_{\{p_{ij}, d_{ij}, b_{ij}\}} \quad & \sum_{j \in \mathcal{N}} g_j \left(\sum_{i \in \mathcal{N}} p_{ij} \right) - \sum_{i \in \mathcal{N}} U_i \left(\sum_{j \in \mathcal{N}} d_{ij} \right) \\ \text{subject to} \quad & \begin{cases} \sum_{i \in \mathcal{N}} p_{ij} \leq P_j & \forall j \in \mathcal{N} \\ b_{ij} = f_i(d_{ij}) & \forall i, j \in \mathcal{N} \\ b \in \Psi(p) \end{cases} \end{aligned} \quad (5)$$

where

$$\Psi(p) = \arg \min_{\{b_{ij}\}} \left\{ \sum_{(i,j) \in \mathcal{N}} -W_i(b_{ij}) \mid \begin{bmatrix} R \\ I \end{bmatrix} b \leq \begin{bmatrix} B \\ f(p) \end{bmatrix} \right\} \quad (6)$$

Problem (5) is the *upper level (leader) problem* and represents the allocation layer. Problem (6) is the *lower level (follower) problem* and models the internet congestion control. The design of a solver for the allocation layer is challenging without the explicit knowledge of the induced set $\Psi(p)$, which is determined by the response of the underlying data network.

Notice that the constraint $d_{ij} \leq p_{ij}$ is now embedded in the lower level problem, as it affects the behavior and the steady state solution of the network congestion protocol. In fact, also in the absence of bandwidth constraints on a path \mathcal{P}_{ij} between two nodes, the bandwidth b_{ij} will be limited by the maximum data flow generated at the application level (in our case, by the algorithm subtasks executed on the available computational resources p_{ij} , thus the limit $b_{ij} \leq f_i(p_{ij})$). Notice also that, in order to maintain generality, the self-loop terms $\{b_{ii}\}$ are also included in the lower level optimization layer. In practice, no bandwidth limit applies to these terms, and therefore $b_{ii} = f_i(p_{ii})$, yielding also $d_{ii} = p_{ii}$.

V. MAIN RESULT

We begin this section by discussing the relation between the bilevel optimization problem (5)-(6) and the social utility maximization problem (4).

Lemma 2 (Convex relaxation). *The social utility maximization problem (4) is a convex relaxation of the bilevel optimization problem (5)-(6).*

Proof: The statement follows from the fact that (4) is a convex program, and that each feasible point for (5)-(6) is also feasible for (4) (and not vice-versa). ■

VI. SHADOW PRICES

In the following theorem, we prove that such a convex relaxation is tight, i.e. the solution of the relaxed convex program corresponds to the solution of the generally non-convex bilevel program.

Theorem 3 (Tightness of the convex relaxation). *Let (p^*, d^*, b^*) be the optimal solution of the convex social utility maximization program (4). Then (p^*, d^*, b^*) is also the global optimal solution of the bilevel problem (5)-(6).*

Proof: From Lemma 2 we know that (4) is a convex relaxation of (5)-(6). We therefore need to prove that the solution (p^*, d^*, b^*) belongs to the non convex feasible set of (5)-(6), i.e. that $b^* \in \Psi(p^*)$.

Let us define $\mathcal{F}(p)$ as the feasible set of the lower level optimization program introduced in (6), namely

$$\mathcal{F}(p) := \left\{ b \mid \begin{bmatrix} R \\ I \end{bmatrix} b \leq \begin{bmatrix} B \\ f(p) \end{bmatrix} \right\}$$

Notice that the constraint $Rb^* \leq B$ is explicitly included in the social utility maximization program (4). From Lemma 1 we know that $d^* = p^*$, and therefore the constraint $b_{ij} = f_i(d_{ij})$ in (4) yields $b^* = f(p^*)$. The constraint $b^* \leq f(p^*)$ is therefore trivially satisfied, and thus b^* belongs to $\mathcal{F}(p^*)$.

We have then to prove that b^* is the minimizer of $\sum_{(i,j) \in \mathcal{N}} -W_i(b_{ij})$ over the set $\mathcal{F}(p^*)$. For this, we use again the fact that $b^* = f(p^*)$. Based on that, any other element $b' \in \mathcal{F}(p^*)$ will necessarily satisfy $b'_{ij} \leq b^*_{ij}$. Using the fact that all the terms $W_i(b_{ij})$ are strictly increasing, we conclude that $W_i(b'_{ij}) \leq W_i(b^*_{ij})$ for any $b' \in \mathcal{F}(p^*)$, and therefore $b^* \in \Psi(p^*)$. ■

Theorem 3 is a fundamental assessment of the tractability of the problem of optimal allocation of computational resources in a data network. It shows that the problem, although formulated as a bilevel optimization program which is general is non-convex, is instead tractable. A further important implication follows, as stated in the following corollary.

Corollary 4 (Congestion control transparency). *The optimal solution of the problem of optimal allocation of computational resources in a data network does not depend on the specific congestion control protocol employed.*

Proof: The result follows directly from the fact that, via Theorem 3, the bilevel program (5)-(6) has the same optimal solution of (4), which is independent on the specific choice of the cost functions $\{W_i\}$ employed in the congestion control protocol. ■

Corollary 4 shows how it is then possible to design solutions and distributed mechanisms for the application that we are considering, without taking the congestion control layer into account. This *separation principle* is a fundamental feature of the layered architecture of data networks, and it is encouraging to see that it holds in this complex application. We point the interested reader to the discussion in [9], where a similar result is obtained for a different network scenario (separation of congestion control and packet scheduling).

In order to gain more insight regarding the economical interpretation of these results, we identify two subsets of \mathcal{N} : the set $\mathcal{B} \subset \mathcal{N}$ of buyers of computational power, and the set $\mathcal{S} \subset \mathcal{N}$ of sellers.

Notice that problem (4), in general, cannot be solved by the end users $i \in \mathcal{B}$ and $j \in \mathcal{S}$, as the routing matrix R and the bandwidth capacity of the “inner” links of the network are not known. We therefore introduce the following assumption, which describes quite accurately the typical structure of the Internet [10] and of large-scale data networks, where data traffic is routed over a large-capacity backbone [11].

Assumption 2. *Each user i taking part to the problem of optimal allocation of computational resources is connected to the network via one link (gateway) of bandwidth B_i , while the rest of the network is assumed to be infinite-capacity.*

Based on this assumption, the social utility maximization problem (4) can be rewritten in the equivalent form, using also the result of Lemma 1

$$\begin{aligned} \min_{\{p_{ij}, d_{ij}, b_{ij}\}} & \sum_{j \in \mathcal{S}} g_j \left(\sum_{i \in \mathcal{B}} p_{ij} \right) - \sum_{i \in \mathcal{B}} U_i \left(\sum_{j \in \mathcal{S}} d_{ij} \right) \\ \text{subject to} & \begin{cases} \sum_{i \in \mathcal{B}} p_{ij} \leq P_j & \forall j \in \mathcal{S} \\ \sum_{i \in \mathcal{B}} b_{ij} \leq B_j & \forall j \in \mathcal{S} \\ \sum_{j \in \mathcal{S}} f_i(d_{ij}) \leq B_i & \forall i \in \mathcal{B} \\ d_{ij} = p_{ij} & \forall i \in \mathcal{B}, j \in \mathcal{S} \\ b_{ij} = f_i(d_{ij}) & \forall i \in \mathcal{B}, j \in \mathcal{S} \end{cases} \end{aligned} \quad (7)$$

Notice that all the parameters of program (7) are now known, either by the sellers or by the buyers.

Let us construct the partial Lagrangian for the program (7), where only the two equality constraints are dualized. We obtain the partial Lagrangian

$$\begin{aligned} \mathcal{L}(p, d, b, \lambda, \mu) &= \sum_{j \in \mathcal{S}} g_j \left(\sum_{i \in \mathcal{B}} p_{ij} \right) - \sum_{i \in \mathcal{B}} U_i \left(\sum_{j \in \mathcal{S}} d_{ij} \right) \\ &+ \sum_{i \in \mathcal{B}, j \in \mathcal{S}} \lambda_{ij} (d_{ij} - p_{ij}) \\ &+ \sum_{i \in \mathcal{B}, j \in \mathcal{S}} \mu_{ij} (f_i(d_{ij}) - b_{ij}). \end{aligned}$$

This formulation directly suggests the following interpretation. Consider the primal problem in the variables p_{iJ} and b_{iJ} for a fixed seller $J \in \mathcal{S}$, evaluated at the optimal primal-dual solution $(p^*, d^*, b^*, \lambda^*, \mu^*)$. We have

$$\begin{aligned} \{p^*_{iJ}, b^*_{iJ}\}_{j=J} &= \arg \min_{\{p_{ij}, b_{ij}\}_{j=J}} g_J \left(\sum_{i \in \mathcal{B}} p_{iJ} \right) - \sum_{i \in \mathcal{B}} \lambda^*_{iJ} p^*_{iJ} \\ &\quad - \sum_{i \in \mathcal{B}} \mu^*_{iJ} b^*_{iJ} \\ \text{subject to} & \begin{cases} \sum_{i \in \mathcal{B}} p^*_{iJ} \leq P_J \\ \sum_{i \in \mathcal{B}} b^*_{iJ} \leq B_J \end{cases} \end{aligned} \quad (8)$$

Similarly, the primal optimality condition for the variables d_{Ij} for a fixed buyer $I \in \mathcal{B}$ is

$$\begin{aligned} \{d_{ij}^*\}_{i=I} = \arg \max_{\{d_{ij}\}_{i=I}} & U_I \left(\sum_{j \in \mathcal{S}} d_{Ij} \right) - \sum_{j \in \mathcal{S}} \lambda_{Ij}^* d_{Ij} \\ & - \sum_{j \in \mathcal{S}} \mu_{Ij}^* f_I(d_{Ij}) \quad (9) \\ \text{subject to} & \sum_{j \in \mathcal{S}} f_I(d_{Ij}) \leq B_I \end{aligned}$$

The conditions given by (8) and (9) show that the optimal values of the dual variables $\{\lambda_{ij}, \mu_{ij}\}$ can be interpreted as *shadow prices* for the computational resources and for the bandwidth resources offered by the sellers to the buyers. As shown in the following example, this analysis reveals the complexity of the pricing mechanisms that are necessary in these exchange markets.

A. Example

Consider a simple network of four agents, two buyers and two sellers. The two buyers (node 1 and 2) have the same logarithmic utility function. The gateway of node 1 has a limited capacity, while the gateway of node 2 has unlimited bandwidth. Moreover, the computational job of node 1 has high bandwidth requirements while the job of node 2 is less demanding.

The first seller (node 3) offers cheap computational power but is bandwidth-constrained, while the second seller (node 4) is more expensive, but has unlimited bandwidth.

The problem parameters and the resulting allocations p are reported in Table I, together with the *marginal prices* q obtained from the shadow prices λ, μ as

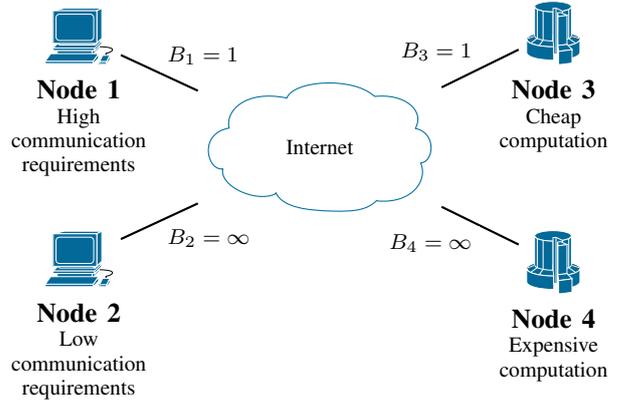
$$q_{ij} := \lambda_{ij}^* + f_i'(d_{ij}^*) \mu_{ij}^*$$

where f_i' is the derivative of the scalar function f_i .

Some observations can be made on this simple case. For example, the cheap but bandwidth-constrained computational resources of node 3 are allocated to the buyer with less communication requirements. More importantly, the proposed shadow price analysis shows how the same computational resource needs to be priced differently for different buyers.

VII. CONCLUSION

In this work, we have investigated the optimal allocation problem of computational capacity over a bandwidth-constrained network controlled by congestion control mechanisms. We showed how the resulting optimization problem is a bilevel program, a class of non-convex problems that is in general hard to solve. However, we showed that the corresponding social utility maximization problem is a convex relaxation of the same problem, and we proved that the relaxation is tight, i.e. the two problems have the same optimal solution. Remarkably, this also shows that the optimal allocation does not depend on the specific congestion control protocol employed in the network, yielding an important separation result for this application. The re-formulation of



	Seller (3)	Seller (4)
	$g_3(z) = 0.1(z + z^2)$	$g_4(z) = 0.2(z + z^2)$
	$P_3 = 1$	$P_4 = 1$
Buyer (1)	$p_{13} = 0$	$p_{14} = 0.25$
$U_1(z) = (1 + \log z)$	$q_{13} = 1.3$	$q_{14} = 0.8$
$f_1(z) = 4z$		
Buyer (2)	$p_{23} = 1$	$p_{24} = 0.33$
$U_2(z) = (1 + \log z)$	$q_{23} = 0.43$	$q_{24} = 0.43$
$f_2(z) = z$		

TABLE I
EXAMPLE OF OPTIMAL ALLOCATION AND MARGINAL PRICES.

the problem as a social utility maximization problem allows to identify and compute the equilibrium prices that characterize this complex market environment. As a future step, we plan to investigate the existence of auction mechanisms (such as the one proposed in [12]) that can be used to compute these prices in a distributed, privacy preserving, and efficient way.

REFERENCES

- [1] "Deutsche börse cloud exchange." [Online]. Available: cloud.exchange
- [2] "Universal Compute Xchange." [Online]. Available: www.ucxchange.com
- [3] D. P. Bertsekas, R. G. Gallager, and P. Humblet, *Data networks*. Prentice-Hall International New Jersey, 1992, vol. 2.
- [4] S. Low, F. Paganini, and J. Doyle, "Internet congestion control," *Control Systems, IEEE*, vol. 22, no. 1, pp. 28–43, Feb 2002.
- [5] R. Srikant, *The Mathematics of Internet Congestion Control (Systems and Control: Foundations and Applications)*. SpringerVerlag, 2004.
- [6] A. Grama, A. Gupta, G. Karypis, and V. Kumar, *Introduction to Parallel Computing*. Addison Wesley, 2003.
- [7] J. F. Bard, *Practical Bilevel Optimization: Algorithms and Applications (Nonconvex Optimization and Its Applications)*. Secaucus, NJ, USA: Springer-Verlag New York, Inc., 2006.
- [8] S. Dempe, *Foundations of bilevel programming*. Kluwer Academic, 2002, no. 61.
- [9] J. Ghaderi, T. Ji, and R. Srikant, "Flow-level stability of wireless networks: Separation of congestion control and packet scheduling," *arXiv:1209.5464 [cs.NI]*, 2012.
- [10] C. Berge, *Hypergraphs*, ser. North-Holland Mathematical Library. North-Holland, 1989, vol. 45, combinatorics of Finite Sets.
- [11] C. Leiserson, "Fat-trees: Universal networks for hardware-efficient supercomputing," *IEEE Transactions on Computers*, vol. C-34, no. 10, pp. 892–901, Oct 1985.
- [12] G. Iosifidis and I. Koutsopoulos, "Double auction mechanisms for resource allocation in autonomous networks," *IEEE Journal on Selected Areas in Communications*, vol. 28, no. 1, 2010.