Manipulability of the Price Mechanism for Data Centers

Greg Bodwin[1], Eric Friedman[2,3,4], and Scott Shenker[3,4]

[1] Department of Computer Science, Tufts University, Medford, Massachusetts 02155
[2] School of ORIE, Cornell University, Ithaca, NY 14853
[3] International Computer Science Institute, Berkeley, CA 94720
[4] Department of Computer Science, University of California, Berkeley, CA 94720

Abstract. We consider the manipulability of the price mechanism on Leontief economies, a natural abstraction that arises in the design and optimization of modern data centers. Although there are well known, and commonly used, (non-price based) strategyproof mechanisms in this setting, we show that pricing is not strategyproof. However, we prove that the maximum gain from manipulation is bounded by a factor of $2 - 1/n$ when there are n players and in many natural settings the maximum gain is significantly less. For example, in homogeneous settings or those with a dominant good there are no profitable manipulations, while in extremely heterogeneous situations, the gain from strategic manipulations appears to be small and vanish in the limit of large n. The gains from collusive manipulations are proportionally smaller with a bound of $2 - k/n$ when there are k colluding players out of n.

Keywords: Cloud Computing, Pricing, Strategyproofness, Approximation

1 Introduction

In this paper we study price mechanisms for economies with Leontief utility functions. These economies naturally arise in the design of modern data centers and cloud computing systems [1] where players submit tasks (computer programs), which require a fixed ratio of resources, such as CPU, memory, data transfer, etc., and their utility is simply the number of completed tasks. This is the model behind the development of the Dominant Resource Fairness algorithm [5, 4] which is used by the Mesos platform for sharing commodity clusters between multiple diverse cluster computing frameworks. Mesos (and its underlying DRF algorithm) are in use in several production and research clusters.

The design of the DRF algorithm was based on both fairness and strategic considerations, including the requirement that the system be immune to manipulation [5]; indeed, DRF (and other related "water-filling" algorithms) can be characterized by basic axioms combined with strategyproofness [4]. However, DRF is a centralized algorithm (albeit simple to implement) and among those operating datacenters there is much interest in decentralized algorithms with a particular focus on price based algorithms. In this paper we study whether one can design simple price based mechanisms for these systems that achieve many of the properties of DRF, focusing on strategyproofness.

Our first result is negative; decentralized pricing is not strategyproof and there are no pricing mechanisms that can implement DRF. However, our next results are more nuanced and less negative. We show that in the worst case the gain from manipulation is at most a multiplicative factor of $2 - 1/n$ for the manipulator when there are n players while in many natural economies the gain from manipulation is much smaller. For example, in economies with homogeneous preferences or those with a dominant good there are no possible gains from manipulation, while at the other extreme, in very heterogeneous economies manipulation can be profitable, but the gains are typically less than 7% with two players, 1% with 6 players and 1/2% for 10. We then extend the analysis to group manipulations and show that the gains from group deviations lead to smaller multiplicative gains (but greater absolute gains) which are bounded by $2 - k/n$ in the worst case when k out of n players collude.

2 Model and Definitions

Our model is that of a Leontief economy.[5] We will assume that there are n players and 1 unit of each of m goods. Let x_{ij} be the amount of good j allocated to agent i subject to the constraint $\sum_i x_{ij} \leq 1$ for all goods j. The utility of an agent, or number of tasks, is given by $U(x_i; a_i) = \min_j x_{ij}/a_{ij}$ where each player's preferences are parameterized by a preference vector $a_i \in \Re^n_+$.

We consider mechanisms M which define the allocation given a set of preferences, i.e. $x = M(a)$. We assume that mechanisms are symmetric under permutations of the players and goods. The two main requirements for mechanisms are efficiency and strategyproofness. Efficiency implies that the vector of utilities induced by M, $u(M(a); a)$, is Pareto efficient, there is no allocation that strictly dominates it. Strategyproofness implies that no player can misrepresent their utility to improve their outcome. Formally, this can be written as

$$u(M(a); a_i) \geq u(M(a_{-i}, a'_i); a_i)$$

for all $a'_i \in \Re^n_+$. Equivalently, for Leontief utilities, this requires that there does not exist any a'_i such that $M(a)_i > M(a_{-i}, a'_i)_i$.

In general the construction of efficient and strategyproof mechanisms for economies is challenging and for many classes of economies such mechanisms do not exist [6, 9]; however, for Leontief economies there do exist many efficient and strategyproof mechanisms [8, 5, 4, 7]. Our main example is a class of water filling mechanisms which are efficient and strategyproof [4].[6] The idea behind these mechanisms is to normalize the preference vectors $a_i \to a_i/N(a_i)$, where

[5] Note that we are considering a steady state model for an online dynamic process. See [5] for more details.

[6] In fact they are even group strategyproof.

N is some vector norm and then to find the largest utility that all players can attain, i.e. find the largest $\gamma > 0$ such that the allocation where $x_i = \gamma a_i / N(a_i)$ is feasible. The DRF mechanism is simply the water filling mechanism for the L^∞ norm, i.e. $N(a_i) = \max_j a_{ij}$. Note that unlike many analyses of economies, we do not require that all goods are fully allocated. In a data center, as in many other systems, there is no need to allocate goods that are not required by any player. (See [7] for a further discussion of the implications of this point.)

3 The Price Mechanism

We now consider price mechanisms. A price mechanism is a set of prices $p \in \Re^m_+$, one for each good, where the players choose the bundle of goods that maximizes their utility while not exceeding their budget which we define to be 1, i.e. $\sum_j x_{ij} p_j \leq 1$ for all players i. We will assume that players are non-wasteful and do not purchase more of any good than they require, so x_i will be proportional to a_i.

The main advantage of a price mechanism is its decentralized implementation via the tâtonnement [11]. This is the mechanism by which prices for each good are adjusted independently according to a simple procedure – lower the price if supply exceeds demand and raise it if demand exceeds supply. Although the tâtonnement need not converge in general economies [10] our numerical results suggest that it does in Leontief economies.

It is important to note that if the tâtonnement converges, it will converge to a set of prices where the price of any slack good must be 0. We call these decentralized price mechanisms and consider them in the remainder of the paper as the value of a price mechanism is in its implementation through the tâtonnement.

In a Leontief economy the decentralized price mechanism is unique as shown in [2], which also provided a polynomial algorithm for their computation. Thus, we will call these prices and the associated purchases by the players, *the* decentralized price mechanism, denoted by M^P.

One can also consider price mechanisms which do not satisfy the dual-slackness conditions arising from the tâtonnement. In this case there can be many different price mechanisms, but it can be shown (under certain regularity conditions) that even these non-decentralizable mechanisms can not be strategyproof.

For example, consider the following preferences:

$$a_1 = (1, 0),\ a_2 = (0, 1),\ a_3 = (1, \epsilon),\ a_4 = (1, 2\epsilon)$$

for small $\epsilon > 0$. In the DRF mechanism, both players 3 and 4 attain the same utility, but this is not possible in any pricing mechanism. Since both goods must have non-zero prices due to the single-mindedness of players 1 and 2, the cost (for the same utility) will be strictly greater for player 4 than for player 3, but they both have the same amount of money to spend.

4 Manipulability

It is easy to see that the decentralized price mechanism is not strategyproof. For example, consider:

$$a_1 = (1, 0),\ a_2 = (1, 3/2),\ a_3 = (1, 3/2).$$

If all users purchase truthfully, the first resource will be priced and the second will be free. This means user 1 will get 1/3 of the available supply of the first good, resulting in $u_1 = 1/3$. If user 1 instead purchases according to a false preference vector of $a'_1 = (1, 3/5)$, the first good will be free and the second will be priced. This means user 1 will get 1/3 of the available supply of the second good (and throw it away), so he will also take 5/9 the available supply of the first good, resulting in $u'_1 = 5/9$. This is a proportional improvement of $\frac{1/3}{5/9} = 5/3$. In fact, this example is the worst possible case for 3 users.

Theorem 1. *For Leontief economy with n players and preferences given by a for any any a'_i,*

$$u_i(M^P(a_{-i}, a'_i); a_i) \leq (2 - 1/n)u_i(M^P(a), a_i).$$

Proof: Let A be the matrix of preferences normalized so that $A_{i1} = 1$ and good 1 is chosen so that the best deviation by player 1 is to increase a_{1j} for some set of $j \neq 1$. In addition, we can assume that the columns of A are linearly independent by small perturbations and then removing slack goods. Define t to be the row vector such that t_i is the utility of player i. Feasibility implies that $tA = e$ where $e_j = 1$ for all goods j. Multiplying both sides by $A^\dagger$ yields $tD = s$ where

$$D_{ij} = \sum_k a_{ik}a_{jk}$$

and $s_j = \sum_k a_{jk}$.

When the preferences a and a'_1 are chosen to maximize proportional gain for user 1, before (resp. after) user i changes her demand, user 1's utility must be in local minimum (resp. maximum) so $\partial t_i/\partial a_{1k} = 0$. Some algebra yields $D^{-1}_{1k} = sd$ where

$$d_i = (D^{-1}(\partial D/\partial a_{ik})D^{-1})_i.$$

Then note that $\partial D/\partial a_{ik}$ is a square matrix with i^{th} row and column as $a_{\cdot k}$ and 0 elsewhere. From this one sees that swapping rows or columns (except the first ones) in A does not affect any of the terms in this equation, which implies that the original preferences were of the form

$$a_1 = (1, \alpha, \alpha, \ldots, \alpha)$$

and for $j > 1$,

$$a_j = (1, \beta, \beta, \ldots, \beta).$$

Solving this yields $t_1 = (\beta - 1)/(\beta - \alpha)$. Choosing $\alpha = 0$, $\beta = n/(n-1)$ and the deviation $\alpha' = n/(2n-1)$ yields an initial utility of $t_1 = 1/n$ and a final utility of $t'_1 = (2 - 1/n)/n$ proving the result. $\square$

In addition, in many common cases, the gain from strategic manipulation is smaller. First consider an economy with a unique dominant good, i.e., the largest element, j such that $a_{ij} = \max_j a_{ij}$, of every preference vector is the same.

Theorem 2. *For Leontief economy with n players and preferences given by a, where there is a unique dominant good, there are no profitable deviation, i.e., for any any a'_i,*

$$u_i(M^P(a_{-i}, a'_i); a_i) \leq u_i(M^P(a), a_i).$$

Proof: To simplify the presentation assume that there are 2 goods and the dominant good is strict for player 1. The extension to the general case follows readily. Thus, we can assume that $a_{i1} = 1$ and $a_{i,2} < 1$ for all i. First note that the equilibrium must allocate the same amount of good 1 and set $p_2 = 0$. In order to influence p_1 player 1 must increase a_{12} past the point where the demand for good 2 becomes tight under the original prices which implies that $a'_{12} > 1$. Now, for any set of equilibrium prices p_1, p_2 player 1 will receive the smallest share of good 1, and thus $x'_{11} \leq 1/n$, with no gain in utility. $\square$

Note that this covers the homogeneous case where all players have the same preferences. Next consider very heterogeneous preferences. In this case, where the a_{ij}'s are chosen i.i.d. uniformly on $[0, 1]$ we find that the worst manipulation is small with high probability. The size of this bound appears to decrease with the number of players, with gains of less than 7% for 2 players, decreasing to gains of less than 1% with 6 players and gains of less than 1/2% with 10 players. Thus, while the worst case bound increases with n we conjecture that the average case bound decreases (eventually to 0) with n.

Lastly, we consider the case in which multiple players collude and coordinate their misrepresentations of their preferences and show that this does not change the analysis significantly.

Theorem 3. *For Leontief economy with n players and preferences given by a for any subset S of the players and any a'_S,*

$$u_i(M^P(a_{-S}, a'_S); a_i) \leq (2 - |S|/n) u_i(M^P(a), a_i)$$

for some $i \in S$.

Proof: This proof is similar to the one for theorem 1, and reduces to the case of 2 goods where all colluding agents share a preference vector as do all the non-colluding players. $\square$

5 Open Problems

In the data center setting, which is best modeled resource allocation with Leontief utilities, our results give both bad news and good news about price mechanisms. While they cannot replicate DRF nor even be strategy proof, their manipulability appears to be rather limited. Thus, price mechanisms in such settings are neither a panacea nor a plague for resource allocation in the data center setting, and more work is needed to completely characterize their properties. For instance, it is not clear whether the the tâtonnement process converges in Leontief (and other special) economies. In addition, we have only begun to characterize the degree of manipulability, and need more precise manipulability bounds under various assumptions for the demand profiles. Lastly, are there normative characterizations for the allocations pricing mechanisms produce, similar in spirit to the characterizations for DRF [4] and the Nash Bargaining solution [3]?

References

1. M. Armbrust, A. Fox, R. Griffith, A.D. Joseph, R.H. Katz, A. Konwinski, G. Lee, D.A. Patterson, A. Rabkin, I. Stoica, et al. Above the clouds: A berkeley view of cloud computing. *EECS Department, University of California, Berkeley, Tech. Rep. UCB/EECS-2009-28*, 2009.
2. B. Codenotti and K. Varadarajan. Efficient computation of equilibrium prices for markets with leontief utilities. *Automata, Languages and Programming*, pages 257–287, 2004.
3. E. Friedman, A. Ghodsi, S. Shenker, and I. Stoica. Nash bargaining without scale invariance. Technical Report, 2011.
4. E. Friedman, A. Ghodsi, S. Shenker, and I. Stoica. Strategyproofness, Leontief economies and the Kalai-Smorodinsky solution. Technical Report, 2011.
5. A. Ghodsi, M. Zaharia, B. Hindman, A. Konwinski, S. Shenker, and I. Stoica. Dominant resource fairness: fair allocation of multiple resource types. In *Proceedings of the 8th USENIX conference on Networked systems design and implementation*, pages 24–24. USENIX Association, 2011.
6. L. Hurwicz. On 1nformationally decentralized systems. *Decision and organization: A volume in honor of Jacob Marschak*, 12:297, 1972.
7. J. Li and J. Xue. Egalitarian fair division under leontief preferences. mimeo, 2011.
8. A. Nicolo. Efficiency and truthfulness with Leontief preferences. A note on two-agent, two-good economies. *Review of Economic Design*, 8(4):373–382, 2004.
9. T.R. Palfrey and S. Srivastava. On bayesian implementable allocations. *The Review of Economic Studies*, 54(2):193, 1987.
10. H. Scarf. Some examples of global instability of the competitive equilibria. *International Economic Review*, 1:157–72, 1960.
11. L. Walras. *Elements of Pure Economics Or The Theory of Social Wealth.* George Allen & Unwin Ltd., 1954.