On the Complexity of Equilibria *

Xiaotie Deng
Department of Computer Science, City University of Hong Kong

Christos Papadimitriou
Division of Computer Science, University of California
Berkeley, USA

Shmuel Safra
School of Computer Science and School of Mathematics
Tel Aviv University, Israel

ABSTRACT

We prove complexity, approximability, and inapproximability results for the problem of finding an exchange equilibrium in markets with indivisible (integer) goods, most notably a polynomial-time algorithm that approximates the market equilibrium arbitrarily closely when the number of goods is bounded and the utilities are linear. We also show a communication complexity lower bound, implying that the ideal informational economy of a market with unique individual optima is unattainable in general.

1. INTRODUCTION

The existence of equilibria in markets, a long-standing conjecture proved rigorously in the 1950's [1], is one of the most fundamental results in Mathematical Economics. Imagine n agents in a market with m kinds of goods (or commodities), each agent with an initial allocation, or endowment, $e_i \in \Re_+^m$ and her own utility function $u_i : \Re_+^m \mapsto \Re_+$, assumed to be strictly concave. The initial allocation may of course be suboptimal (this is precisely the reason why commerce exists), in that with certain mutually beneficial exchanges everybody is better off (some utilities are increased and none is decreased). An allocation from which this is impossible — there is no other allocation with some utilities better and none worse— is called a *Pareto allocation*. This is a broadly accepted concept of optimality in economics, a coveted goal. But it may take long sequences of complex multi-party exchanges, each carried out by complex, inefficient and deceitful bargaining, to arrive at such an allocation (by what we would call local search in our field). How do markets do it?

*Research partially supported by a grant from RGC of Hong Kong SAR, a grant of the City University of Hong Kong, an NSF ITR Grant, an IBM Faculty Development Award, and a grant from the US-Israel BSF Research Foundation.

STOC'02, May 19-21, 2002, Montreal, Quebec, Canada.

The answer proposed by the Arrow-Debreu theorem is, *by the price mechanism.* They prove that, in any such situation, there is a price vector $p \in \Re_+^m$, the *price equilibrium*, such that, if $x_i(p)$ denotes the allocation that optimizes $u_i(x)$ subject to $px \leq pe_i$ (unique by strict concavity), then $\sum_i x_i(p) = \sum_i e_i$. In other words, if each agent does the only reasonable thing —sell her endowment at market prices, and with the proceeds buy the best commodity combination she can afford— the market will clear exactly.

This is a beautiful and mathematically sophisticated result, whose proof uses Kakutani's theorem, a stronger version of Brouwer's fixpoint theorem. The obvious computational problem is, given the e_i's (the supply) and some representation for the u_i's (the demand), find the price equilibrium; this problem was first studied in the pioneering work of Scarf [14]. To date, there is no known polynomial-time algorithm for it, even though certain sophisticated algorithms perform quite well. NP-hardness is very unlikely, due to the guaranteed existence property (see [12] for an extensive discussion of the complexity classes appropriate for this and related problems).

It has been known since even before the 1950's that if the commodities are *indivisible* (that is, if we have the integer programming version of the equilibrium problem), then an equilibrium may well not exist. There is an extensive literature proposing remedies for this deficiency, some in fact reasonably close in spirit to the area of approximation algorithms. Most remedies are based on special structure of the market, or the fact that there are certain divisible commodities ("money") in sufficient supply. In such a situation, one may define the *deficiency* of a market as the smallest $\Delta \geq 0$ such that, for some price vector, the market clears as long as the agents are persuaded to obtain a basket that is within an (1-Δ) factor of their optimum utility. So, at equilibrium the deficiency is zero.

We prove that it is NP-hard to calculate — or even approximate within any factor better than 1/3— the deficiency of a market with indivisible goods; this is true even in the special case of *linear* utilities, even when there are only two agents (even though the inapproximability result does not apply there —and, as we shall see, for a reason). It is even

NP-hard to find equilibrium prices even if it is known that they exist.

We then provide polynomial-time algorithms for several aspects of the equilibrium problem. We give a polynomial-time algorithm for finding the price equilibrium if the utilities are linear, the goods *divisible* (that is, real-valued), and the number of commodities bounded (Theorem 2; there is no known polynomial algorithm for the unbounded case). Using some of the ideas and techniques of [13], we can find in polynomial time an ϵ-approximate Pareto allocation (one in which no exchange can increase all utilities by a factor of more than $(1+\epsilon)$) for any $\epsilon > 0$ (Theorem 3). Finally, our technically most demanding result is a polynomial-time algorithm for finding an approximate equilibrium in markets with a bounded number of indivisible commodities and linear utilities (Theorem 4).

One interesting way of looking at the Arrow-Debreu theorem on the existence of a price equilibrium is as a striking positive result in *communication complexity:* The prices, just $m-1$ reals, whose number is independent of n, summarize effectively the whole messy situation —the mn reals representing the supply (endowments), plus the arbitrarily complex utility functions (the demand)— and enable a protocol that produces a Pareto allocation by the simplest, communication-free distributed algorithm: Each agent will look at the prices, optimize locally, and help herself to the optimum basket. Since it is reasonable to think of the number of agents in a market as much larger than that of the number of commodities, this is a significant result. We have here of course a highly nonstandard communication model: Somebody (the market's invisible hand?) absorbs all information, computes the prices, and then announces them. *Only the announcement is costly.*

This ideally streamlined communication efficiency depends crucially on the assumption that utility functions are strictly concave —and hence the optimum basket for each agent is unique. If optimum baskets are not unique, heavy coordination may be required between agents to clear the markets. We show (Theorem 6) that, if the utilities are concave but not strictly so, then $\Omega(n\log(m+n))$ *bits will need to be announced in this communication model in order to achieve a Pareto allocation.* In the case of *integer-valued goods*, Pareto allocation with less communication *cannot even be achieved approximately, for any approximation ratio.* These bounds are valid even when $m \geq 2$ is constant.

Three remarks are in order regarding this result and its place in the literature. First, it was inspired by Nisan's recent results on communication complexity in economic contexts [9], see also [10]. Second, an apparent weakness of our result is that the comparison is between an upper bound of m reals in the strictly concave case and our lower bound of $n\log(m+n)$ *bits* in the general case —and obviously a real number can hide many bits. The significance of our result relies on the intuitive feeling one may have by inspecting the situation, that prices encode markets in a very direct and natural way, while our lower bound does not reflect in any way the arithmetic precision of the situation.

Thirdly, there is a parallel (technically incomparable) classical result in Economics[1][3, 5]. They consider "rational" (formal definition omitted) mechanisms (see [9, 11] for accessible introductions to similar concepts) in which agent i responds to a proposal K (a real vector) by a central authority by evaluating a differentiable function $g_i(K,\theta_i)$, where θ_i are private values (in our case, the endowments and the parameters of the utility function) and accepts the proposal if the value is zero. The result states that K must have dimension $mn-1$ —hence the price mechanism with baskets proposed by the government (with one basket implicit, and with normalized prices) is optimum. Notice that this result holds even for strictly concave utility functions, where $m-1$ is possible in our "self-service" model.

2. NP-HARDNESS

In a *linear market of indivisible goods*, n agents have linear utilities $u_i(x_1,\ldots,x_m)=\sum_{j=1}^m u_{ij}x_j$, for given nonnegative integers u_{ij} utility matrix, and n endowment vectors $e_i \in Z_+^m$. An *ϵ-approximate equilibrium* is a vector $p \in \Re_+^m$ of prices and an allocation of n vectors $x_i \in \tilde{Z}_+^m, i=1,\ldots,n$ such that

- The market approximately clears: $(1-\epsilon)\sum_i e_i \sum_i x_i \leq \sum_i e_i$
- For all i, the utility $\sum_j u_{ij}x_{ij}$ is at least $(1-\epsilon)$ times the optimum solution to the integer program "maximize over all integer vectors y the quantity $\sum_j u_{ij}y_j$ subject to $py \leq pe_I$."

A 0-approximate equilibrium is called simply an *equilibrium.* The *deficiency* of a market is the smallest $\epsilon \geq 0$ for which an ϵ-approximate equilibrium exists.

THEOREM 1. *The deficiency of a linear market with indivisible goods is NP-hard to compute, or even approximate to within* 1/3. *The problem is NP-hard (but not to approximate, see Theorem nteger) even if the number of agents is two.*

Sketch: By a reduction from EXACT COVER BY 3-SETS. We are given a family $S_1,\ldots,S_n$ of 3-subsets of $\{1,\ldots,3m\}$, and we are asked whether there is a subfamily of m sets covering all elements.

The idea in the reduction is this: There is a "shopper" for each element, a shopper for each set, plus a "broker." The endowment of all element-shoppers consists of a single unit of a commodity called an "egg," while set-shoppers have three eggs each. The broker has an endowment of commodities corresponding to the sets (three units each), plus m units of another commodity called "set substitute."

The utilities are as follows: The shoppers have no use for their eggs (the corresponding coefficient u_{ij} is zero), but each element-shopper wants the sets that contain it, while

[1]Many thanks to Tom Marschak for bringing that literature to our attention, and for helping us understand the connection.

each set-shopper wants either the same set or, three times more intensely, the set substitute (the u_{ij} coefficient for the substitute is three times that of the set itself). The broker is not interested in sets or set substitutes, but has a craving for eggs.

It is not very hard to prove that in any price equilibrium with zero deficiency each element-shopper will get her set, while each set-shopper will get either the set itself in triplicate, or one unit of set substitute —these latter set-shoppers must form a set cover. We then argue that, in the absence of an exact cover, any price equilibrium will need to have deficiency at least 1/3 —hence inapproximability.

For the case of a bounded number of agents, we reduce from the following version of the PARTITION problem (whose NP-hardness proof is quite nontrivial): Given a set of n integers, are there two different subsets with the same sum? □

3. ALGORITHMS FOR LINEAR MARKETS

Divisible Markets

That a linear market with divisible commodities has a price equilibrium follows from the general theory. However, we are not aware of a polynomial-time algorithm for finding such an equilibrium. Typically, a generalization of linear programming called *linear complementarity* [2, 12] is used to solve such problems; however, there is no polynomial-time algorithm for that problem either. The problem with the obvious linear programming formulation [4] is this: The prices (dual variables) of the program maximizing sum of utilities do not take supply (endowments) into account. What should be maximized instead is the weighted sum of the utilities *with appropriate weights* reflecting supply, but it is not clear how to find those weights.

However, the problem can be solved in polynomial time when the number of commodities, or agents, is bounded.

THEOREM 2. *There is a polynomial-time algorithm for finding an equilibrium in linear markets with a bounded number of divisible goods. Ditto for a polynomial number of agents.*

Sketch: We shall sketch the bounded commodity case only. Consider n agents and m (constant) commodities, and an integer utility matrix u_{ij}. Consider a price equilibrium p. Each agent i sorts her ratios u_{ij}/p_i ("bang for the buck" ratios) and buys as much as available going down this list. Each commodity will therefore belong, at equilibrium, to one three categories: It is either monopolized by i, or shared by i and others, or not wanted by i. Define the *type* of i to be a pair (A_i, B_i) of subsets of $[m]$, where A_i is the set of commodities monopolized by i and B_i those shared by i. The type of an allocation is the collection of all types of the agents in it. The result follows from these two observations: (a) If we know the type of an allocation, equilibrium prices are easy to find by linear programming (we compute the p_j's that arrange the ratios in the right order); and (b) There is a polynomial number of types for fixed u_{ij}'s (the A_i's are disjoint and their union is bounded by $[m]$, while each B_i represents $|B_i| - 1$ linear equations between the prices, for fixed u_{ij}, and therefore all B_i's can also be summarized by a family of disjoint sets). □

Approximate Equilibria

Suppose first that we have a linear market with indivisible commodities and a bounded number n of agents. Recall that a *Pareto allocation* is one whose n-dimensional utility vector is not dominated component-wise by another utility vector. An *ϵ-Pareto curve* is a set of allocations such that every other allocation is dominated component-wise by one of them *if multiplied by* $(1+\epsilon)$. We can prove the following:

THEOREM 3. *There is a polynomial-time algorithm for computing an ϵ-Pareto curve in linear markets with indivisible commodities and a fixed number of agents.*

Sketch: The main idea is very much along the lines of algorithms given in [13] for approximating the Pareto curve in multi-objective optimization problems. The algorithm is dynamic programming on the utilities, adding one commodity at a time, and restricting values to the set $\tilde{Z}_+$ (see the next proof). □

COROLLARY 1. *The deficiency of a market with linear utilities, indivisible commodities, and a fixed number of agents can be approximated within any given ϵ in polynomial time.*

The remainder of this section is dedicated to our main algorithmic result:

THEOREM 4. *If the number of goods is bounded, there is a polynomial-time algorithm which, for any linear indivisible market for which a price equilibrium exists, and for any $\epsilon > 0$, finds an ϵ-approximate equilibrium.*

Sketch: Let d be the bounded number of commodities, let E be the largest component of $\sum_i e_i$, and let δ be a small parameter depending on ϵ, d and n. Define the set of integers $\tilde{Z}_+ = \{0, 1, \ldots, \lceil 1/\delta \rceil\} \cup \{\lfloor (1+\delta)^i \rfloor : i \geq 0, (1+\delta)^i < E\}$. That is, $\tilde{Z}_+$ contains all small numbers, plus all rounded powers of $(1+\delta)$ up to E. Let M denote $|\tilde{Z}_+|$; notice that M is polynomial in the size of the instance. For an integer x, $\tilde{x}$ will denote the largest integer in $\tilde{Z}_+$ less than or equal to x.

It is immediate that, if x_i is the equilibrium allocation, then $\tilde{x_i}$ is an allocation that clears the markets within a factor of $(1 - n\delta)$, and is optimum for each agent within a factor of $(1 - \delta)$. Hence, we shall be looking for an allocation in $\tilde{Z}_+^n$.

Consider Agent i, and a price vector p. Define $A_i(p)$ to be the set of all vectors $x \in \tilde{Z}_+^d$ that maximize px subject to $px \leq e_i x$. Call two price vectors p and q *i-equivalent* if $A_i(p) = A_i(q)$. The equivalent classes of i-equivalence are polytopal subdivisions of the d-orthant by the hyperplanes normal to the vectors $\{e_i - z : z \in \tilde{Z}^d_+\}$.

Consider next all agents, and the common refinement of all i-equivalence classes, $i = 1, \ldots, n$. That is, two price vectors are *equivalent* if they are i-equivalent for all i; in other words, the set of optima in $\tilde{Z}^d_+$ for these vectors are the same for all agents. These classes are in turn polytopal subdivisions of the d-orthant by hyperplanes normal to the vectors $\{e_i - z : i = 1, \ldots, n; z \in \tilde{Z}^d_+\}$, nM hyperplanes in toto. Hence there are $O(n^d M^d)$ such subdivisions, and in fact they can be constructed in NC [6].

To find the allocation approximating the equilibrium allocation, we iterate over all (polynomially many) equivalence classes. For each class, we consider a representative vector p and the set of all vectors in $A_i(p), i = 1, \ldots, n$. We must find an allocation that combines one optimum from each agent. This is done by dynamic programming, starting from the zero allocation and adding the optima of the first agent, then the second agent, and so on. However, the values of the partial allocations are rounded down to $\tilde{Z}_+$. This results in a cummulative error of $O(n\delta)$. It follows that any feasible allocation discovered by this algorithm will be ϵ-approximate. $\square$

4. MARKET COMMUNICATION COMPLEXITY

If the utility functions are strictly concave, the Arrow-Debreu theorem enables a remarkable economy of information transfer, as the whole market can be summarized by $m-1$ real numbers —the prices, normalized. What if the utility functions are concave but not strictly so? A price equilibrium will still exist, but it is not enough for telling the agents exactly which one of the several possible optimum baskets they should choose so as to clear the market.

Markets suggest the following (we believe novel) variant of the one-way/nondeterministic communication complexity models, which we call *market communication*: Each of agents $1, \ldots, n$ has private information $x_i \in X_i$, for some domain X_i, and each must compute a (multivalued) function $f_i(x_1, \ldots, x_n)$. Agent 0 (the "invisible hand") knows all $x_i, i = 1, \ldots, n$. A protocol is a set of functions $g_0, g_1, \ldots, g_n$. Agent 0 computes $g_0 : X_1 \times \cdots \times X_n \mapsto X_0$, and then the other agents compute the required final results $g_i(g_0(x_1, \ldots, x_n), x_i) = f_i(x_1, \ldots, x_n), i = 1, \ldots, n$. The cost of this protocol on the present instance is the size of $g_0(x_1, \ldots, x_n)$ (the "announcement" of Agent 0).

The domains and the notion of size have been left open so far; the domains of interest here are of two kinds: Subsets of real vector spaces, or subsets of $\{0,1\}^{[p(m,n)]}$ for some polynomial p (we call this latter case *binary domains*). Size is in both cases length (of vector or string). The market communication complexity of $(f_1, \ldots, f_n)$ with domains $X_0, \ldots, X_n$ is the complexity of the protocol with the smallest worst-case complexity (a function of m and n) that computes these multivalued functions.

The Arrow-Debreu Theorem can be restated as follows:

THEOREM 5 (ARROW-DEBREU). *There is a protocol with market communication complexity $m-1$ for computing Pareto allocations of n agents and m goods, with domains (appropriate vector space representations of) strictly concave utility functions and endowments.*

In partial contrast, we can show the following:

THEOREM 6. *Any protocol with binary domains for computing Pareto allocations of n agents and m divisible commodities with concave utility functions (resp., ϵ-Pareto allocations for indivisible commodities, for any $\epsilon < 1$) must have market communication complexity $\Omega(n \log(m+n))$ bits.*

Proof: We shall explain the $\Omega(n)$ bound, the extension to the claimed one being trivial. Consider the following domains for the n agents:

Each odd-numbered agent $2i-1, i = 1, \ldots, n' = n/2$ has a domain for its utility function consisting of two possible concave utility functions: One that is $\min\{1, \lambda\}$ for each commodity vector of the form $\lambda v_{i1}, 0 \leq \lambda$, and zero everywhere else, where v_{i1} is a fixed integer vector particular to i, and a similar function for another vector, $v_i 2$. In other words, Agent $2i-1$ has one of these two utility functions, but only she knows which one. Agent $2i$ has a domain with only one possible concave utility function, one that is $\max\{1, \lambda_1 + \lambda_2\}$ for each commodity vector of the form $\lambda_1(K\mathbf{1} - v_{i1} + \lambda_2(K\mathbf{1} - v_{i2}), 0 \leq \lambda_1, \lambda_2$ where $\mathbf{1}$ is the vector of all ones and K is an integer to be determined; Agent $2i$'s utility is zero everywhere else. Each agent has an endowment $\frac{1}{2}K\mathbf{1}$ —that is to say, $\frac{1}{2}K$ from each good.

It is easy to see that there is a Pareto allocation with all utilities 1 —the maximum possible (hence this is the only Pareto allocation): Just allocate to Agent $2i-1$ vector $v_i 1$ or $v_i 2$ —whichever she prefers— and allocate $K\mathbf{1}$ minus that to Agent $2i$. However, for any protocol with function g_0 that achieves this allocation, an adversary argument shows that g_0 must have a different value for each sum of the form $\sum_{i=1}^n v_{ik_i}$, where $k_i \in \{1,2\}$ for each i.

It remains to show that we can choose the v_{ik}'s so that there are $2^{\Omega(n)}$ different such sums. This is done by randomization: If we choose the first $m/2$ entries of each vector independently to be a random number uniformly distributed between 1 an K, and arrange the remaining entries somehow so as to balance the sum to $1/over2Km$, then the probability that two given sums will have the same value is less than $K^{-m/2}$, and hence the number of different sums is at least $2^n(1 - 2^n K^{-m/2})$. Taking $K > 2^{2n/m}$ gives the $\Omega(n)$ result for divisible goods. The full result for divisible goods requires a polynomial number of v_{ik}'s per agent, instead of two. Finally, to see the inapproximability result for indivisible goods, just consider the utility u_{2i-1} that is the characteristic function of $\{v_{ik}\}$, for some k, and for u_{2i} the characteristic function of $\{K\mathbf{1} - v_{ik} : k = 1, 2, \ldots\}$. $\square$

5. REFERENCES

[1] K. K. Arrow, and G. Debreu, "Existence of an Equilibrium for a Competitive Economy", *Econometrica*, Vol. 22, pp.265-290, 1954.

[2] R. W. Cottle, G. B. Dantzig, "Complementary pivot theory of mathematical programming", *Linear Algebra and Appl.* 1 1968 no. 1, 103–125.

[3] J. Jordan, "The Competitive Mechanism is Dimensionally Efficient Uniquely," Journal of Economic Theory, 1982.

[4] D. Gale, "Equilibrium in a discrete exchange economy with money," *Internat. J. Game Theory* 13 (1984), no. 1, 61–64.

[5] L. Hurwicz, "On Information Decentralization and Efficiency in Resourse Allocation Mechanism", *Studies in Mathematical Economics* (S. Reiter, ed.), pp. 238-350, 1986.

[6] D. Kozen, C. Yap, "Algebraic Cell Decomposition in NC (Preliminary Version)," *Proc. FOCS 1985*, pp. 515-521.

[7] H. W. Lenstra, Jr. "Integer programming with a fixed number of variables," *Math. Oper. Res.* 8 (1983), no. 4, 538–548.

[8] A. K. Lenstra; H. W. Lenstra, Jr.; L. Lovasz "Factoring polynomials with rational coefficients", *Math. Ann.* 261 (1982), no. 4, 515–534.

[9] N. Nisan, "Communication Complexity of Combinatorial Auctions," manuscript, 2001.

[10] N. Nisan and I. Segal, "The Communication Complexity of Efficient Allocation Problems" presented in DIMACS workshop on Computational Issues in Game Theory and Mechanism Design, November 2001.

[11] C. H. Papadimitriou "Algorithms, games, and the internet," Proc. STOC 2001: pp. 749-753.

[12] C. H. Papadimitriou "On the Complexity of the Parity Argument and Other Inefficient Proofs of Existence," *JCSS 48*(3), pp. 498-532 (1994).

[13] C. H. Papadimitriou, M. Yannakakis "The complexity of tradeoffs, and optimal access of web sources", *Proc. FOCS 2000.*

[14] H. Scarf, *The Computation of Economic Equilibria* (with collaboration of T. Hansen), Cowles Foundation Monograph No. 24. New Haven: Yale University Press, 1973.