Auctions with Severely Bounded Communication

Liad Blumrosen
School of Engineering and Computer Science
The Hebrew University of Jerusalem
Jerusalem, Israel
liad@cs.huji.ac.il

Noam Nisan
School of Engineering and Computer Science
The Hebrew University of Jerusalem
Jerusalem, Israel
noam@cs.huji.ac.il

ABSTRACT

We study auctions with severe bounds on the communication allowed: each bidder may only transmit t bits of information to the auctioneer. We consider both welfare-maximizing and revenue-maximizing auctions under this communication restriction. For both measures, we determine the optimal auction and show that the loss incurred relative to unconstrained auctions is mild. We prove non-surprising properties of these kinds of auctions, e.g. that discrete prices are informationally efficient, as well as some surprising properties, e.g. that asymmetric auctions are better than symmetric ones.

1. INTRODUCTION

Recent years have seen the emergence of the Internet as a central platform of interaction between computers, humans, and firms. The different parties that interact on the Internet are in various levels and modes of cooperation and competition with each other. This happens in all levels of the interaction, from the lowest technical level of computer communication, routing, storage, and computing, and reaching to the highest level of electronic commerce in its many forms. Studying such types of computer systems that are distributed in terms of the participants' goals and incentives, naturally requires using a combination of techniques from economics, game theory and computer science. Indeed much recent work has been done on this borderline, see e.g. the surveys [23, 12].

Studying these types of distributed-incentive-computer-systems naturally leads to many new problems in each of the participating research fields – questions that arise due to considerations from the other fields. In particular, many new questions in economics arise due to the necessity of taking computational questions into account. This paper deals with such a question: how to design efficient auctions that are restricted to using a very small amount of communication.

Auctions have been suggested many times as an efficient mechanism for resource allocation in computer systems. See e.g. [23, 9, 28, 33] and the many references therein. The basic argument goes along these lines: when some computational resource needs to be allocated in a distributed system, we would like the system to allocate it in the most beneficial way. If we auction the resource among all the conflicting uses in an economically efficient way then we will do just that. Thus, for example, congestion over some communication link in a network can be handled by auctioning the bandwidth of the link. This type of idea has indeed been applied in various forms both for low level resources like network bandwidth [17, 32, 30, 13] or computing resources [36, 38, 37, 14, 25, 27], and for many e-commerce systems [2, 3, 1, 5, 4].

Several researchers have considered the effect of various computational considerations on the design of auctions: on-line behavior [16, 6], unbounded supply [8, 11, 6], computational complexity in combinatorial auctions [35, 18, 24, 31, 40], timing uncertainty [29], and more. This paper studies the effect of severely restricting the amount of communication allowed in an auction of a single item. Each bidder is only allowed to send a single t-bit message to the auctioneer, who must then allocate the item and determine the price according to the messages received. I.e. each bidder has a set of k possible messages it can send (where $k = 2^t$). The simplest case is $t = 1$, and thus $k = 2$, i.e. each bidder sends a single bit of information, and the auctioneer must determine an allocation and pricing according to the bits received from the bidders. This is in stark contrast to the usual treatment in economics where we assume the communication is in terms of real numbers. The only treatment of similar issues we know of in the economic literature is [19], who considered similar questions in cases of restricting bid levels in oral auctions to discrete levels, and [39, 20] that analyze the inefficiency caused by discrete priority classes of customers.

The reader may ask why we bother studying such severe restrictions on communication, as a single real number does not seem like an excessive amount of information to transfer in any computer system. Well, there are several motivations for studying auctions with such severe restrictions on the communication. First, if auctions are to be used for allocation of low level resources in computer systems, then only a very small amount of computational effort can be spent on them. Thus, e.g., an auction for routing a single packet will need to require very little communication overhead, certainly not a whole real number. One would ideally like to "waste" only a bit or two on the bidding information.

Preferably, the bidding information can be piggy-backed on some unused bits in the packet header of existing networking protocols (such as IP, TCP, or those used for QoS). Second, restrictions on communication can sometimes function as a proxy for other simplifications in the auction: low communication means low information revelation; low communication reduces the amount of required human input (and thus simpler user interface); low communication means a small number of payment amounts and thus may simplify handling them electronically, etc. Since we will design auctions that are very efficient despite using very low communication, we get all such properties as a bonus. Third, while single item auctions require only a single real number in communication, efficient combinatorial auctions require an exponential amount of communication [26] and are thus impossible computationally. Understanding the trade-offs between communication and allocative efficiency will thus allow enlarging the envelope of auctions that are efficient both in the economic sense and in the computational sense.

We consider both the question of optimizing total social welfare and the question of maximizing seller revenue (under individual rationality constraints) under the restriction of bounded communication. We completely characterize the optimal auctions in the case of 2-bidders, a characterization that holds for all notions of equilibria. We describe two families of auctions called "priority games" and "modified priority games" each having a dominant strategy equilibrium. Each of these families is parametrized by certain threshold vectors. We derive the optimal values for these parameters, for which we prove:

THEOREM 1.1. *For any pair of distributions on the valuations of two bidders, and any bound k on the number of possible messages allowed to each bidder, a priority game with the derived parameters achieves maximal social welfare. The loss in social welfare compared to auctions that are unconstrained in communication is $O(\frac{1}{k^2})$*

THEOREM 1.2. *For any regular distribution on the valuations of two symmetric bidders, and any bound k on the number of possible messages allowed to each bidder, a modified priority-game with the derived parameters achieves maximal seller revenue (under individual rationality and Bayesian-Nash equilibrium constraints). The loss in seller revenue compared to auctions that are unconstrained in communication is $O(\frac{1}{k^2})$.*

The $O(\frac{1}{k^2})$ bound on the loss incurred is tight for some distributions, as we show by analysing the case of valuations uniformly distributed in $[0, 1]$. In this case the loss of social welfare is exactly $\frac{1}{6\cdot(2k-1)^2}$ and the loss of seller revenue is $\Omega(\frac{1}{k^2})$.

Our analysis implies some expected as well as some unexpected results:

- **Low welfare and revenue loss:** Even severe bounds on communication result in a mild loss of efficiency. E.g. for the case of two bidders whose valuations are uniformly distributed in $[0, 1]$, we obtain a 1-bit auction with expected welfare 0.648, compared to 0.667 which is what can be reached without any restriction on communication.
- **Asymmetry helps:** Asymmetric auctions may be more efficient than symmetric ones with the same communication bounds. E.g. for the case of two bidders whose valuations are uniformly distributed in $[0, 1]$, symmetric 1-bit auctions can only achieve expected welfare of 0.625, compared to 0.648 for asymmetric ones. We also show that asymmetry helps achieving higher revenue.
- **Discrete Prices are Informationally Efficient:** We show that in the optimal auctions with k messages, bidders simply partition the valuation range to k continuous price ranges and bid their price range.
- **Dominant Strategy equilibrium incurs no additional cost:** The efficient auction we design has a dominant strategies equilibrium and yet is optimal among all auctions regardless of their definition of equilibrium.

We start by presenting, in section 2, a self-contained treatment of the simplest case: 2 bidders with uniformly distributed valuations, each allowed a single bit of communication. We continue with the general case: section 3 provides the model definition and introduces our notations, section 4 analyzes welfare and revenue maximizing in auctions among two players. Finally, section 5 discusses the generalization to an arbitrary number of bidders. All proofs appear in the full version of this paper ([7]). Here, we give the proof's outline for some of the main theorems.

2. 2-PLAYERS, 2 POSSIBLE BIDS

We start with a description of the simplest case: auctions among 2 players where every player can send only a single bit to the auctioneer (or mechanism). With this simple case, we demonstrate the properties of the optimal solutions in the general cases, when we allow any number of possible bids or any number of players.

2.1 The simple model

The players in our model are risk-neutral, have independent private values for the item, and quasi-linear utilities. The valuation of player i is distributed in the range $[0, 1]$ with a commonly-known distribution function f_i. In this section, we assume players' valuations are distributed uniformly. Throughout the paper, we deal with *ex-post Individually Rational* (IR) mechanisms, i.e. games where a winning player never pays more then her valuation for the item, and a losing player always pays zero.

Our unique assumption is that every player has only two possible bids to choose from. Such mechanisms can be described with a 2x2 matrix, where the 1st player (Alice) chooses a row, and the 2nd (Bob) chooses a column. Each entry of the matrix specifies the allocation and payments given a bids' combination. The mechanism can toss coins to determine the allocations. Figures 1 and 2 depicts examples for *2-players 1-bit* mechanisms.

A Strategy s_i for player i is a function $s_i : [0, 1] \to \{0, 1\}$. A strategy determines the bid of player i according to his valuation v_i.

Each selfish bidder wants to maximize her expected utility. As the mechanism's designers, we will try to optimize "social" criteria such as expected *welfare* and *revenue*. The *expected welfare* (or efficiency) achieved by a mechanism is the expected valuation of the player that wins the item (if

B / A	0	1
0	B wins and pays 0	B wins and pays 0
1	A wins and pays $\frac{1}{3}$	B wins and pays $\frac{2}{3}$

Figure 1: (g_1) A 2-player 1-bit game that achieves maximal expected welfare (efficiency)

B / A	0	1
0	No allocation	B wins and pays $\frac{5}{8}$
1	A wins and pays $\frac{1}{2}$	B wins and pays $\frac{5}{8}$

Figure 2: (g_2) A 2-player 1-bit game that achieves maximal expected revenue

any). The *expected revenue* from a mechanism is the expected sum of bidders' payments.

2.2 Welfare and revenue in simple mechanisms

In the case of unlimited communication between the bidder and the mechanism, when valuations are distributed uniformly, we know that the optimal expected welfare is $\frac{2}{3}$ (2nd-price auction, see [34]) and the optimal expected revenue is $\frac{5}{12} = 0.417$ (2nd-price auction with a reservation price of $\frac{1}{2}$, see [10, 22]). In this section, we study how close can 2-player games, with 2 possible bids for each player, get to these values, with dominant strategies equilibrium and ex-post individual rationality. Note that with no communication at all, the mechanism can achieve an expected welfare of $\frac{1}{2}$ (e.g. when A always wins the object).

Let g_1 denote the mechanism described in figure 1 and g_2 denote the mechanism in figure 2.

THEOREM 2.1. *The mechanism g_1 has a dominant strategies equilibrium, with ex-post individual rationality and it achieves expected welfare of $\frac{35}{54} = \frac{2}{3} - \frac{1}{54} = 0.648$.*

The mechanism g_2 has dominant strategies equilibrium, with ex-post individual rationality, and it achieves expected revenue of $\frac{25}{64} = 0.39$

PROOF. (sketch) Consider the following strategy: "*bid 1 if your valuation is greater than $\frac{1}{3}$, else bid 0*". Clearly, this strategy is dominant for player A in g_1: when her valuation is smaller than $\frac{1}{3}$ she will gain a negative utility if she bids "1". When her valuation is greater than $\frac{1}{3}$, bidding "0" gives her zero utility, but she can get positive utility by bidding "1". We call this kind of strategies *threshold-strategies.* Similarly, a threshold-strategy with $\frac{2}{3}$ is dominant for player B in g_1. The threshold-strategies with the values $\frac{1}{2}, \frac{5}{8}$ are dominant for A, B, respectively, in g_2.

Bidding "0", guarantees a zero payment for the players in both games. Thus, both games are ex-post individually-rational.

Next, we calculate the expected welfare in g_1 when both players play their dominant strategies described above. When a player uses a threshold strategy with x as a threshold, he bids "0" with probability x, and "1" with probability $1-x$ (uniform distribution). The expected valuation of this player, given that he bids "0", is $\frac{x}{2}$, and it equals $\frac{1+x}{2}$ given that he bids "1". The expected welfare, is therefore:

$$\frac{1}{3}\frac{2}{3}\frac{\left(\frac{2}{3}\right)}{2} + \frac{1}{3}(1-\frac{2}{3})\frac{(1+\frac{2}{3})}{2} + (1-\frac{1}{3})\frac{2}{3}\frac{(1+\frac{1}{3})}{2}$$

$$+(1-\frac{1}{3})(1-\frac{2}{3})\frac{(1+\frac{2}{3})}{2} = \frac{35}{54}$$

Similarly, we can calculate the expected revenue in g_2, when both players play their dominant strategies:

$$0 + \frac{1}{2}(1-\frac{5}{8})\frac{5}{8} + (1-\frac{1}{2})\frac{5}{8}\frac{1}{2} + (1-\frac{1}{2})(1-\frac{5}{8})\frac{5}{8} = \frac{25}{64}$$

□

Recall that with no communication limitations, the optimal welfare is $\frac{2}{3} = \frac{36}{54}$. We surprisingly see that despite severely limiting the communication from infinitely many bits to a single bit, the welfare loss is mild (only $\frac{1}{54}$).

2.3 Optimal mechanisms

Next, we claim that the mechanisms g_1 and g_2 (described in figures 1 and 2) achieve optimal expected welfare and revenue, respectively.

THEOREM 2.2. *No 2-player 1-bit mechanism achieves strictly greater expected welfare, than the mechanism g_1 from figure 1 (even without restriction to mechanisms with dominant strategies equilibrium).*

PROOF. (sketch) This theorem is proved in 3 steps:

Step 1: We show that every mechanism g can achieve its optimal welfare with a pair of threshold strategies. I.e. there exists a pair of threshold strategies such that no other strategies achieve strictly greater expected welfare in g.

Step 2: Consider mechanisms in which the item is allocated to the player with the highest bid, and in case of equal bids, the item is always allocated using a pre-defined order on the players. We call this family of mechanisms *priority-games.* (For example, g_1 is a priority-game in which we always break ties in favor of B.) We show that for each priority-game there exists a pair of strategies (not necessarily in equilibrium) that achieves the maximal welfare among all 2-player 1-bit mechanisms with any pair of strategies.

Step 3: From the previous steps we know that optimal welfare can be achieved in priority-games with threshold-strategies. Next, we express the expected welfare in a priority-game as a function of the threshold-values that the players use:

$$w(g,x,y) = xy\frac{x}{2} + x(1-y)\frac{1+y}{2} +$$

$$(1-x)y\frac{1+x}{2} + (1-x)(1-y)\frac{1+y}{2}$$

This function achieves unique maximum ($x, y \in [0, 1]$) when $(x, y) = (\frac{1}{3}, \frac{2}{3})$. □

THEOREM 2.3. *No ex-post individually-rational mechanism achieves strictly greater revenue than g_2 (see figure 2) .*

Theorems 4.5 and 4.9 gives a generalization for theorems 2.2 and 2.3 respectively.

Observe the following properties of g_1 and g_2, which demonstrate the properties of the optimal mechanisms in the general case:

	0	1
0	w.p. $\frac{1}{2}$ A wins, pays 0 w.p. $\frac{1}{2}$ B wins, pays 0	B wins and pays $\frac{1}{4}$
1	A wins and pays $\frac{1}{4}$	w.p. $\frac{1}{2}$ A wins, pays $\frac{1}{2}$ w.p. $\frac{1}{2}$ B wins, pays $\frac{1}{2}$

Figure 3: (m_1) 2-players 1-bit symmetric mechanism that achieves optimal welfare

	0	1
0	No allocation	B wins and pays $\frac{1}{\sqrt{3}}$
1	A wins, pays $\frac{1}{\sqrt{3}}$	w.p. $\frac{1}{2}$ A wins, pays $\frac{1}{\sqrt{3}}$ w.p. $\frac{1}{2}$ B wins, pays $\frac{1}{\sqrt{3}}$

Figure 4: (m_2) 2-players 1-bit symmetric mechanism that achieves optimal revenue

- Both optimal welfare and optimal revenue are achieved when the players use threshold-strategies.

- The mechanism g_1 achieves the maximal welfare achievable by any 1-bit mechanism and any pair of strategies, without restrictions to any kind of equilibria. Nevertheless, we found a game with a dominant strategies equilibrium that achieves this welfare. g_2 achieves maximal revenue among all the ex-post IR mechanisms with Bayesian-Nash equilibrium. Yet, g_2 achieves this optimal revenue with dominant strategies equilibrium.

- The welfare maximizing threshold values $(x, y) = (\frac{1}{3}, \frac{2}{3})$ are what we call *mutually-centered*. I.e. x is the expected valuation of B given that B bids 0 (i.e. $x = E(v_B|0 \leq v_B \leq y) = \frac{0+y}{2}$), and y is the expected valuation of A given that A bids 1 (i.e. $y = E(v_A|x \leq v_A \leq 1) = \frac{x+1}{2}$).

- The optimal mechanisms are asymmetric (priority games are asymmetric by definition). Actually we show in section 2.4 that symmetric mechanisms must achieve strictly smaller expected welfare.

- When optimizing revenue, the ex-post IR assumption plays an important role, although our results also hold for interim IR. We could alternatively assume ex-ante IR, where the players participate only if their *expected* utility is non-negative. In this case the mechanism "extracts" the whole welfare of the players (thus the optimal revenue is $\frac{35}{54}$).

2.4 Optimal symmetric games

The optimal mechanisms we have presented so far in this paper are asymmetric. Can we find symmetric mechanisms that achieve these optimal results? In this section we show that the answer is no: symmetric games achieve smaller welfare and revenue, though the differences are small. All proofs appear in the full paper ([7]).

Let m_1 be the mechanism described in figure 3. m_1 achieves maximal welfare among all the symmetric mechanisms, and this welfare is smaller than the optimal welfare in asymmetric mechanisms (0.648).

PROPOSITION 2.4. *The mechanism m_1 achieves welfare of $\frac{5}{8}$ with dominant strategies equilibrium and ex-post IR. No symmetric 2-players 1-bit mechanism achieves a strictly higher welfare (with any strategies, not necessarily dominant).*

Let m_2 be the mechanism described in figure 4. m_2 achieves maximal revenue (for symmetric mechanisms), which is smaller than the optimal expected revenue in asymmetric mechanisms (0.39).

PROPOSITION 2.5. *The mechanism m_2 achieves revenue of* 0.385 *with dominant strategies equilibrium and ex-post IR. No symmetric mechanism with ex-post IR, achieves strictly greater revenue than m_2.*

3. THE GENERAL MODEL

3.1 The players and the mechanism

We consider single item, sealed bid auctions among n risk-neutral players. Player i has a private valuation for the object $v_i \in [0, 1]$. The valuations are independently drawn from a distribution function f_i $\left(\forall_{v \in [0,1]}\ f_i(v) \geq 0, \int_0^1 f_i(v)dv = 1\right)$ and the cumulative function is F_i.

Throughout the paper we assume that the distribution functions are continuous and always positive. Players want to maximize their utilities, which are *quasi-linear.* We assume a normalized model, i.e. players' valuations for not having the item are zero. The seller's valuation for the item is zero. We also assume players utilities depend only on whether they win the item or not (no externalities).

In our model, each player i can send a message of $t_i = \lg(k_i)$ **bits** to the mechanism, i.e. player i can choose one of possible k_i **bids** (or actions). In most parts of the paper, we assume that all players have the same number of possible bids, k. Denote the possible set of bids for player i as $\beta_i = \{0, 1, 2, ..., k_i - 1\}$. In each auction, player i chooses a bid $b_i \in \beta_i$. Let $b = \{b_1, ..., b_n\}$ be a vector of bids. A mechanism should determine the allocation and payments given a vector of bids b:

DEFINITION 1. *A mechanism g is composed of a pair (a, p) where:*

- *$a : (\beta_1 \times ... \times \beta_n) \to [0, 1]^n$ is the allocation scheme (not necessarily deterministic). We denote the i'th coordinate of $a(b)$ by $a_i(b)$, which is player i's probability for winning the item when the bidders bid b. Clearly, $\forall i\ \forall b\ a_i(b) \geq 0$ and $\forall b\ \sum_{i=1}^n a_i(b) \leq 1$.*

- *$p : (\beta_1 \times ... \times \beta_n) \to \Re^n$ is the payment scheme. $p_i(b)$ is player i's payment given a bids' vector b, which she pays only if she wins the item.*

Note that we allow non-deterministic allocations, but we ignore non-deterministic payments (since we are interested in expected values, using lottery for the payments has no effect on our results).

DEFINITION 2. *In a mechanism with k-possible bids, for every player i, $|\beta_i| = k$. We denote the set of all the mechanisms with k-possible bids among n players by $G_{n,k}$. We denote the set of all the n-player mechanisms in which $|\beta_i| = k_i$ for each player i, by $G_{n,(k_1,...,k_n)}$.*

Throughout the paper, we deal with *ex-post Individually Rational* (IR) mechanisms, i.e. mechanisms in which the player that wins the item will never pay more then her valuation, and a losing player will always pay zero. In other words, in ex-post IR mechanisms, the utility of zero is assured for each player (including the seller). (We sometimes equivalently use the term: mechanisms with ex-post individual rationality.) Therefore, in our model the payment for each player is always in the range $[0, 1]$.

Next, we define the notion of a strategy for a player, and show how players choose their strategies.

DEFINITION 3. *A Strategy s_i for player i in a game $g \in G_{n,k}$ describes how a player determines his bid according to his valuation, i.e. it is a function $s_i : [0, 1] \to \{0, 1, ..., k-1\}$.*

Denote $\varphi_k = \{s \,|\, s : [0, 1] \to \{0, 1, ..., k-1\}\}$ (i.e. the set of all strategies for players with k possible bids).

DEFINITION 4. *A real vector $c = (c_0, c_1, ..., c_k)$ is a vector of threshold-values if $c_0 \leq c_1 \leq ... \leq c_k$.*

DEFINITION 5. *A strategy $s_i \in \varphi_k$ is a threshold-strategy based on a vector of threshold-values $c = (c_0, c_1, ..., c_k)$, if $c_0 = 0$ and $c_k = 1$ and for every $c_i \leq v_i < c_{i+1}$ we have $s_i(v_i) = i$. We say that s_i is a threshold strategy, if there exists a vector c of threshold values such that s_i is a threshold strategy based on c.*

We use the notations: $s(v) = (s_1(v_1), ..., s_n(v_n))$, when s_i is a strategy for bidder i, $s = (s_1, ..., s_n)$ and $v = (v_1, ..., v_n)$. Note that $b = s(v)$ is a vector of bids. When $s = (s_1, ..., s_n)$ is a vector of strategies, and s_i a strategy for player i, let s_{-i} denote the strategies of the players except i, i.e. $s_{-i} = (s_1, ..., s_{i-1}, s_{i+1}, ..., s_n)$. We sometimes use the notation $s = (s_i, s_{-i})$.

3.2 Optimality criteria

The players in our model choose strategies that maximize their utilities. We are interested in games where such strategies forms equilibria.

DEFINITION 6. *Let $u_i(g, s)$ be the expected utility of player i from game g when bidders use strategies s, i.e.*

$$u_i(g, s) = E_{v\in[0,1]^n} \left(a_i\left(s(v)\right) \cdot \left(v_i - p_i\left(s(v)\right)\right)\right)$$

DEFINITION 7. *A strategy s_i for player i is dominant in mechanism $g \in G_{n,(k_1,...,k_n)}$ if regardless of the other players' strategies s_{-i}, i cannot gain a higher utility by changing his strategy, i.e.*

$$\forall \widetilde{s_i} \in \varphi_{k_i} \;\; \forall s_{-i} \;\; u_i(g, (s_i, s_{-i})) \geq u_i(g, (\widetilde{s_i}, s_{-i}))$$

We say that a mechanism g has *a dominant strategies equilibrium* if for every player i there exists a strategy s_i which is a dominant.

DEFINITION 8. *Strategies $s = (s_1, ..., s_n)$ forms a Bayesian-Nash equilibrium in a mechanism $g \in G_{n,(k_1,...,k_n)}$, if for every player i, s_i is the best response for the strategies s_{-i} of the other players, i.e.*

$$\forall i \;\; \forall \widetilde{s_i} \;\; u_i(g, (s_i, s_{-i})) \geq u_i(g, (\widetilde{s_i}, s_{-i}))$$

Our goal is to find optimal communication-bounded mechanisms. Each selfish bidder wants to maximize her expected utility. As the mechanism designers, we will try to optimize "social" criteria such as *welfare* (efficiency) and *revenue*.

The *expected welfare* from a mechanism g, when bidders use strategies s, is the expected sum of bidders valuations. Because the item is indivisible, the expected welfare is actually the expected valuation of the player that wins the item (if any). Note that the expected welfare does not directly depend on the payments in the mechanism.

DEFINITION 9. *Let $w(g, s)$ denote the expected welfare in the n-player game g when bidders' strategies are s, i.e.*

$$w(g, s) = E_{v\in[0,1]^n} \left(\sum_{i=1}^{n} a_i\left(s(v)\right) \cdot v_i\right)$$

and let $w^{opt}_{n,k}$ denote the maximal possible expected welfare from any n-player game where each player has k possible bids, with any vector of strategies allowed, i.e.

$$w^{opt}_{n,k} = \max_{g\in G_{n,k},\, s\in\varphi_1\times...\times\varphi_n} w(g, s)$$

DEFINITION 10. *Let $r(g, s)$ denote the expected revenue in the n-player game g when bidders' strategies are s, i.e.*

$$r(g, s) = E_{v\in[0,1]^n} \left(\sum_{i=1}^{n} a_i\left(s(v)\right) \cdot p_i\left(s(v)\right)\right)$$

and let $r^{opt}_{n,k}$ denote the maximal possible expected revenue from any ex-post individually-rational, n-players, k-possible-bids game and strategies s that forms a Bayesian-Nash equilibrium:

$$r^{opt}_{n,k} = \max_{\substack{g \in G_{n,k} \text{ is } ex-post\ IR \\ s \in \times_{i=1}^{n}\varphi_i \text{ in B. Nash equilibrium}}} r(g, s)$$

Note that we define the optimal welfare as the maximal welfare among all mechanisms and strategies, not necessarily in equilibria, and we define the optimal revenue as the maximal revenue among all mechanisms with Bayesian-Nash equilibrium and ex-post IR. Yet, the optimal mechanisms (for both measures) that we present in this paper achieve these optimal values with dominant strategies equilibria.

DEFINITION 11. *We say that a mechanism $g \in G_{n,k}$ achieves expected welfare (revenue) of α if there is a vector s of dominant strategies for which the expected welfare (revenue) from g and s is α, i.e. $w(g, s) = \alpha$ ($r(g, s) = \alpha$). Mechanism $g \in G_{n,k}$ achieves optimal excepted welfare (revenue) if it achieves $w^{opt}_{n,k}$ ($r^{opt}_{n,k}$).*

4. OPTIMAL MECHANISMS FOR 2 PLAYERS

In this section we present bounded communication mechanisms that achieve optimal welfare and revenue.

DEFINITION 12. *A game is called a priority-game if it allocates the item to the player i that bids the highest bid (i.e. when $b_i > b_j$ for all $j \neq i$, the allocation is $a_i(b) = 1$ and $a_j(b) = 0$ for $j \neq i$), with ties consistently broken according to a pre-defined order on the players.*

	0	1	2	..	k-2	k-1
0	$\mathbf{B},0$	$\mathbf{B},0$	$\mathbf{B},0$	...	$\mathbf{B},0$	$\mathbf{B},0$
1	A,x_1	$\mathbf{B},y_1$	$\mathbf{B},y_1$	...	$\mathbf{B},y_1$	$\mathbf{B},y_1$
2	A,x_1	A,x_2	$\mathbf{B},y_2$	...	$\mathbf{B},y_2$	$\mathbf{B},y_2$
...	...	...	...	...	...	...
k-2	A,x_1	A,x_2	A,x_3	...	$\mathbf{B},y_{k-2}$	$\mathbf{B},y_{k-2}$
k-1	A,x_1	A,x_2	A,x_3	...	A,x_{k-1}	$\mathbf{B},y_{k-1}$

Figure 5: A priority-game based on the threshold-values x,y. In each entry, the left argument denotes the winning player, and the right argument is the price she pays. When x,y are mutually-centered, this mechanism achieves optimal welfare, among all the mechanisms and all possible-strategies.

	0	1	2	..	k-2	k-1
0	ϕ	$\mathbf{B},y_1$	$\mathbf{B},y_1$	...	$\mathbf{B},y_1$	$\mathbf{B},y_1$
1	A,x_1	$\mathbf{B},y_1$	$\mathbf{B},y_1$	...	$\mathbf{B},y_1$	$\mathbf{B},y_1$
2	A,x_1	A,x_2	$\mathbf{B},y_2$	...	$\mathbf{B},y_2$	$\mathbf{B},y_2$
...	...	...	...	...	...	...
k-2	A,x_1	A,x_2	A,x_3	...	$\mathbf{B},y_{k-2}$	$\mathbf{B},y_{k-2}$
k-1	A,x_1	A,x_2	A,x_3	...	A,x_{k-1}	$\mathbf{B},y_{k-1}$

Figure 6: A modified priority-game based on the threshold strategies x,y. In each entry, the left argument denotes the winning player, and the right argument is the price she pays. For optimally chosen values of x,y, this mechanism achieves optimal revenue.

DEFINITION 13. *A game is called a modified priority-game if it has an allocation as of priority-games, except no allocation is done when all players bid* 0.

We will assume (w.l.o.g) throughout this paper that in 2-players priority-games $B \succ A$, i.e. the mechanism allocates the item to A if she bids a higher bid than B, and otherwise to B.

DEFINITION 14. *A 2-players priority-game based on the threshold-values $x = (x_0 \le ... \le x_k)$ and $y = (y_0 \le ... \le y_k)$ is a mechanism that its allocation is as in a priority-game and given a pair of bids (i,j) of A,B respectively, A pays x_{j+1} whenever she wins, and B pays y_i when he wins. We denote this mechanism by $PG_k(x,y)$. The mechanism $PG_k(x,y)$ is presented in figure 5.*

DEFINITION 15. *A 2-players modified priority-game based on the threshold-values $x = (x_0 \le ... \le x_k)$ and $y = (y_0 \le ... \le y_k)$ is a mechanism that its allocation is as in a modified priority-game and given a pair of bids (i,j) of A,B respectively, A pays x_{j+1} whenever she wins and when B wins he pays y_i when $i > 0$ and y_1 when $i = 0$. We denote this mechanism by $MPG_k(x,y)$. The mechanism $MPG_k(x,y)$ is presented in figure 6.*

It turns out that priority-games achieve optimal welfare, and modified priority-games achieve optimal revenue.

PROPOSITION 4.1. *For every pair of threshold-values x,y, the threshold-strategies based on these threshold-values are dominant in both $PG_k(x,y)$ and $MPG_k(x,y)$, and these mechanisms are ex-post individually-rational.*

Actually, we know that every monotone mechanism has a unique pricing scheme that achieves these properties (see [15, 21]). In this pricing scheme, the price must equal to the minimum valuation under which the buyer is supposed to get the object, given the other buyers' valuations. Priority games and modified priority games are monotone: if the winner raised his bid, he would still win the object.

DEFINITION 16. *The threshold-values*

$$x = (x_0, x_1, ..., x_{k-1}, x_k = 1)$$

$$y = (y_0, y_1, ..., y_{k-1}, y_k = 1)$$

are mutually-centered, if the following constraints hold:

$$\forall 1 \le i \le k-1 \quad x_i = E(v_B \,|y_{i-1} \le v_B \le y_i) = \frac{\int_{y_{i-1}}^{y_i} f_B(v_B) \cdot v_B dv_B}{F_B(y_i) - F_B(y_{i-1})}$$

$$\forall 1 \le i \le k-1 \quad y_i = E(v_A \,|x_i \le v_A \le x_{i+1}) = \frac{\int_{x_i}^{x_{i+1}} f_A(v_A) \cdot v_A dv_A}{F_A(x_{i+1}) - F_A(x_i)}$$

LEMMA 4.2. *For any pair of distribution functions on the players' valuations, and for any values of x_0 and y_0, there exist a unique pair of mutually-centered threshold-values x,y.*

4.1 Welfare-optimal 2 players mechanisms with k possible bids

In this subsection, we fully characterize efficient mechanisms for 2 players, given any pair of distribution functions on the players' valuations. We also give an asymptotically tight upper bound on the welfare loss they incur.

Before presenting the welfare-optimal mechanisms, we prove two lemmas. The first lemma enables us to focus our search only on threshold-strategies. The second lemma states that priority-games achieve optimal-welfare.

LEMMA 4.3. *Given any mechanism $g \in G_{n,(k_1,...,k_n)}$, there exists a vector of threshold-strategies $s \in \times_{i=1}^{n} \varphi_{k_i}$ that achieves the optimal welfare in g among all possible strategies, i.e.*

$$w(g,s) = \max_{\tilde{s} \in \times_{i=1}^{n} \varphi_k} w(g,\tilde{s})$$

PROOF. (sketch) Consider a strategies vector s that achieves optimal welfare in g, and assume that the strategy of player i is not a threshold strategy. Since s_i is not a threshold-strategy, there must exist valuations $0 \le a < b < c \le 1$ such that player i bids the same bid when her valuation is a or c but bids differently when her valuation is b (i.e. $s_i(a) = s_i(c) \ne s_i(b)$.) Notice that given a fixed valuation v_i of player i, the expected welfare is linear in v_i. Also observe that b is a convex combination of a and c. Thus, given that player i's valuation is b, bidding $s_i(a)$ achieves maximal expected welfare (compared with the other possible bids for player i). Therefore, we can modify s_i such that player i bids $s_i(a)$ when her valuation is b, and the expected welfare will not decrease. We can repeat this process until s_i becomes a threshold strategy. □

LEMMA 4.4. *There exists a priority-game $g \in G_{2,k}$ that achieves, with some pair of strategies, the optimal welfare (i.e. $w_{2,k}^{opt}$).*

PROOF. We prove this lemma using the following three claims:

CLAIM 1. *There exists a mechanism that achieves optimal welfare (i.e. $w_{2,k}^{opt}$) which is deterministic (i.e. the winner is fixed for each combination of bids) and in which the item must be sold (for every bids' combination).*

PROOF. (sketch) This claim is easy. For each bids' combination b we can always allocate the item to a player i with the highest expected valuation, given that his bid is b_i (i.e. $i \in argmax_j(E(v_j|s_j(v_j) = b_j))$). □

CLAIM 2. *Every mechanism can be modified to be monotone and the expected welfare it achieves with a given pair of strategies will not decrease. (In monotone mechanisms, if player i wins the item for some bids' vector b, she would win the item also if she sent a higher bid, i.e. for $(\widetilde{b_i}, b_{-i})$ where $\widetilde{b_i} > b_i$.)*

PROOF. (sketch) We can assume, w.l.o.g, that for each bidder the thresholds are ordered from lowest to highest (i.e. if a bidder bids "m" for all $v \in [c_t, c_{t+1}]$, she will bid "$m+1$" for all $v \in [c_{t+1}, c_{t+2}]$). Changing the mechanism as done in the previous claim, will modify the mechanism to be not only deterministic, but also monotone. Given a bids vector (b_i, b_{-i}) for which player j has maximal expected valuation over all other players, if he increases his bid to $\widetilde{b_i} > b_i$, he will still have the maximal expected valuation. □

CLAIM 3. *Consider the matrix representation of a 2-players game with k possible bids. In a deterministic, monotone mechanism in which the item must be sold, that achieves optimal welfare, no two rows (or columns) can have an identical allocation scheme.*

PROOF. (sketch) Consider such an optimal mechanism $g \in G_{2,k}$ with two identical rows. g's monotonicity derives that these rows are adjacent (still assuming w.l.o.g increasing thresholds). Thus, there is a mechanism $\widetilde{g} \in G_{2,(k-1,k)}$ that achieves exactly the same expected welfare as g (when the identical rows are united to one). We prove (see [7]) that the optimal welfare from a game where both players have k possible bids cannot be achieved when one of the players has only $k-1$ possible bids (i.e. $w_{2,k}^{opt}$ ¿ $w_{2,(k-1,k)}^{opt}$). □

Due to claims 1 and 2, there is a deterministic, monotone game in which the item must be sold that achieves $w_{2,k}^{opt}$. In such games, the allocation scheme in row i looks like $[A, ..., A, B...B]$. Due to claim 3, in the matrix representation of this optimal game, there are no two rows with the same allocation scheme. Thus there are $k+1$ possible rows for the game matrix. However, our mechanism has only k rows. Similarly, we have k (of possible $k+1$) columns in the mechanism. Assume that the row $[B, B, ..., B]$ is in g. Then, the column $[A, A..., A]$ is not in g. Therefore, our game matrix consists of all the columns except $[A, A, ..., A]$, which compose the priority game where $B \succ A$. If the row $[B, B, ..., B]$ is not in g, then g is the priority-game where $A \succ B$. This concludes the proof of the lemma.

Let $x^w = (x_0^w, ..., x_{k-1}^w, x_k^w = 1)$ and $y^w = (y_0^w, ..., y_{k-1}^w, y_k^w = 1)$ be mutually-centered threshold values, where $x_0^w = y_0^w = 0$.

THEOREM 4.5. *For any pair of distribution functions on the players' valuations, the mechanism $PG_k(x^w, y^w)$ achieves optimal welfare (i.e. $w_{2,k}^{opt}$) with dominant strategies equilibrium and ex-post IR. The welfare loss it incurs, compared with the efficient auction with no communication bounds, is $O(\frac{1}{k^2})$.*

PROOF. (sketch) Due to lemmas 4.3 and 4.4, optimal welfare can be achieved in priority-games with threshold-strategies. Thus, we can express the expected welfare in a priority game as a function of these threshold-values:

$$\sum_{i=1}^{k} F_A(x_i) \cdot \int_{y_{i-1}}^{y_i} f_B(v_B)v_B dv_B$$

$$+\sum_{i=2}^{k} F_B(y_{i-1}) \cdot \int_{x_{i-1}}^{x_i} f_A(v_A)v_A dv_A$$

First order conditions on the $2k-2$ variables derives exactly the conditions for mutually-centered threshold values. Thus, the given mechanism is welfare-optimal.

For showing the $O(\frac{1}{k^2})$ upper bound, we construct a priority game in which both players have the same dominant threshold strategy, such that the probability that every player bids each bid is smaller than $\frac{2}{k}$. This is done by dividing the density functions of both players to $\frac{k}{2}$ intervals with equal mass, then combining these thresholds to a vector of k threshold-values $(0, x_1, ..., x_{k-1}, 1)$. Because the players use the same threshold-strategy, a welfare loss is possible only when the bids are equal (i.e. on the diagonal). Both players will bid i with probability which is smaller than $\frac{4}{k^2}$, and the maximal welfare-loss possible in this case is $x_{i+1} - x_i$. Thus, the expected welfare loss is smaller than

$$\sum_{i=1}^{k} \frac{4}{k^2} \cdot (x_{i+1} - x_i) = \frac{4}{k^2} \cdot \sum_{i=1}^{k} (x_{i+1} - x_i) = \frac{4}{k^2}$$

The optimal mechanism can only do better. □

We give an explicit solution for the case of uniformly-distributed valuations in $[0, 1]$. We show that the asymptotic upper bound on the efficiency loss is tight, i.e. the efficiency loss in the worst case, is $\Theta(\frac{1}{k^2})$.

THEOREM 4.6. *When players' valuations are uniformly distributed, the mechanism $PG_k(\overrightarrow{x}, \overrightarrow{y})$ achieves optimal welfare where*

$$x = (0, \frac{1}{2k-1}, \frac{3}{2k-1}, ..., \frac{2k-3}{2k-1}, 1)$$

$$y = (0, \frac{2}{2k-1}, \frac{4}{2k-1}, ..., \frac{2k-2}{2k-1}, 1)$$

and the welfare loss it incurs is exactly $\frac{1}{6 \cdot (2k-1)^2}$. The optimal welfare is achieved with dominant strategies equilibrium and ex-post IR.

Note that the threshold values from theorem 4.6 are mutually centered.

Due to theorem 4.5, for any pair of distribution functions we can construct a mechanisms that incurs a welfare loss of $O(\frac{1}{k^2})$. But can we design a mechanism that regardless of the distribution functions, will always incur a low welfare loss? The following theorem presents a mechanism with k-possible bids that incurs a welfare loss of $O(\frac{1}{k})$ regardless of the players' distribution functions, and we also show that no mechanism can do better.

PROPOSITION 4.7. *The mechanism $PG_k(x,x)$, where $x = (0, \frac{1}{k}, \frac{2}{k}, ..., \frac{k-1}{k}, 1)$, incurs an expected welfare loss $\leq \frac{1}{k}$ for any pair of distribution functions of the players' valuations. Moreover, for any mechanism there exists a pair of distribution functions for which the expected welfare loss is $\Omega(\frac{1}{k})$.*

4.2 Revenue-optimal 2-players mechanisms with k possible bids

Most results in the literature on revenue-maximizing auctions, assume that the distribution functions of the players' valuations are *regular* (as defined below). When the valuations of all players are distributed with the same regular distribution-function, it is well known that Vickrey's 2nd-price auction, with an appropriately chosen reservation price, is revenue-optimal ([34, 22, 10]).

DEFINITION 17. *([22]) Let f be a distribution function on a finite range, and let F be its cumulative function. We say that f is regular, if the function*

$$\widetilde{v}(v) = v - \frac{1-F(v)}{f(v)}$$

is monotone, strictly increasing function of v. We call $\widetilde{v}$ virtual utility or virtual surplus.

The key observation of Myerson ([22]), which we also use, is that in equilibrium, the expected revenue equals the expected virtual-utility. We use this property to reduce the revenue optimization problem to the welfare optimization problem, for which we have already given a full characterization. We first show that this observation also holds for auctions with bounded communication:

PROPOSITION 4.8. *Let $g \in G_{n,k}$ be a mechanism with Bayesian Nash equilibrium and ex-post individual rationality. Then, the expected revenue in g is equal to the expected virtual-utility in g.*

PROOF. (sketch) Let $g \in G_{n,k}$ be a mechanism with dominant strategies equilibrium $(s_1, ..., s_n)$. Consider the following direct-revelation mechanism g_d: each player i bids her true valuation v_i. The mechanism calculates $s_i(v_i)$ for every i, and determines the allocation and payments according to g. An easy observation is that g_d is incentive-compatible (i.e. truthful bidding is dominant strategy for the players) and ex-post individually rational. According to Myerson's observation, in g_d (which is individually-rational with incentive-compatible Bayesian-Nash equilibrium), the expected revenue is equal to the expected virtual-utility. However, for every combination of bids, the allocation and payments in both mechanisms are identical. Thus, both the expected revenue and the expected virtual-utility is equal in both mechanisms. □

Now, we can characterize revenue-optimal mechanisms. For optimally chosen $\widetilde{x_0}$ and $\widetilde{y_0}$ (see [7]), let $\widetilde{x} = (\widetilde{x_0}, ..., \widetilde{x_{k-2}}, 1)$ and $\widetilde{y} = (\widetilde{y_0}, ..., \widetilde{y_{k-2}}, 1)$ be mutually centered threshold values. Let $x^r = (x^r_0, ..., x^r_k)$ and $y^r = (y^r_0, ..., y^r_k)$ be threshold values that satisfy $x^r_0 = y^r_0 = 0$, $x^r_k = y^r_k = 1$ and for every $1 \leq i \leq k-1$ $\widetilde{x_{i-1}} = \widetilde{v}(x^r_i)$.

THEOREM 4.9. *When both players' valuations are distributed with the same regular distribution function, the mechanism $MPG_k(x^r, y^r)$ achieves optimal expected revenue (i.e. $r^{opt}_{2,k}$). It incurs a revenue loss, compared with the optimal auction with no communication limitations, of $O(\frac{1}{k^2})$.*

PROOF. (sketch, see [7] for full proof): Using proposition 4.8, we can reduce the revenue optimization problem to a welfare optimization problem which we have already solved (in theorem 4.5): we can design a mechanism $\widetilde{g}$ that achieves maximal expected welfare, in a model where players consider their virtual utilities as their own valuation (with dominant strategies equilibrium and ex-post IR). Consider the mechanism g with the same allocation scheme as in $\widetilde{g}$, but with the payment scheme transformed by $\widetilde{v}^{-1}$ (i.e. each payment $\widetilde{p}$ is replaced with payment p such that $\widetilde{p} = \widetilde{v}(p)$). In this mechanism, players behave exactly as in $\widetilde{g}$, and still with dominant strategies (here we use the regularity of the distribution function), and thus has the same expected virtual-utility as in $\widetilde{g}$. Therefore, g achieves maximal expected virtual-utility and thus maximal revenue. Some adjustments have to be done for situations when the seller has a higher virtual-utility than the other players (when the players' virtual utilities are negative or equivalently when their valuations are below the reservation-price). These adjustments derive that g is a modified priority-game (and not a standard priority game). □

Now we show that this upper bound is asymptotically tight. As in the case of welfare optimization, we give an explicit solution for the case of uniform distribution functions. This optimal mechanism incurs a revenue loss of $\Omega(\frac{1}{k^2})$, which is (due to theorem 4.9) asymptotically the worst case.

THEOREM 4.10. *When players' valuations are distributed uniformly, the modified priority-game $MPG_k(x,y)$ achieves optimal expected revenue among all the individually-rational mechanisms in $G_{2,k}$, where*

$$x = (0, \frac{1}{2}, t + \frac{1\cdot(1-t)}{2k-3}, ..., t + \frac{(2k-5)\cdot(1-t)}{2k-3}, 1)$$

$$y = (0, t, t + \frac{2\cdot(1-t)}{2k-3}, ..., t + \frac{(2k-4)\cdot(1-t)}{2k-3}, 1)$$

and $t = \frac{-2\alpha+\sqrt{1+3\alpha}}{2(1-\alpha)}$ for $\alpha = \frac{1}{(2k-3)^2}$. This mechanism incurs revenue loss of $\Omega(\frac{1}{k^2})$.

Note that when the players' valuations are distributed uniformly, the transformation $\widetilde{v}^{-1}$ is linear, and thus the threshold values x and y from theorem 4.10, without the first zero element, are mutually-centered.

5. RESULTS FOR ANY NUMBER OF PLAYERS

We consider games among n players where each player has 2 possible bids (i.e. she can send only 1 bit to the mechanism). In the full paper ([7]) we give a characterization of the efficient mechanism for general distribution functions. Here, we give a characterization of the welfare-optimal and the revenue-optimal mechanisms, when players' valuations are distributed uniformly. We give an $O(\frac{1}{n})$ upper bound for both the welfare loss and the revenue loss incurred by these mechanisms. According to simulations we made, these upper bounds are tight, but we haven't proved it yet. All proofs appear in the full paper ([7]).

It is easy to see, that when all the players use the threshold strategy with $1 - \frac{ln(n)}{n}$ as a threshold, both the revenue loss and the welfare loss are $O(\frac{logn}{n})$ (see [7] for proof). For an upper bound of $O(\frac{1}{n})$, we construct less trivial mechanisms.

Let $(x_1, ..., x_n) \in [0, 1]^n$ be threshold values for players $1, ..., n$ respectively. Consider the following recursive constraints:

$$x_1 = \frac{x_n}{2} \quad (1)$$

$$x_{m+1} = \frac{1}{2} + \frac{x_m^2}{2} \quad \forall m \in \{1, ..., n-2\} \quad (2)$$

$$x_n = \frac{\sum_{i=1}^{n-1} \left(\prod_{j=i+1}^{n-1} x_j \right) (1 - x_i^2)}{2 \left(1 - \prod_{i=1}^{n-1} x_i\right)} \quad (3)$$

In the spirit of definitions 14 and 15 we say that an n-players 1-bit priority-game (or modified priority-game) is based on the threshold values $x_1, ..., x_n$ $(1 \prec ... \prec n)$, if when player i bids 1 and wins, she will pay x_i. (when all bids are zero all players pay zero.)

THEOREM 5.1. *Consider the threshold values $x_1, ..., x_n$ for which equations 1, 2, 3 hold. When the players' valuations are distributed uniformly, a modified priority-game g which is based on $x_1, ..., x_n$ achieves optimal welfare. The welfare loss it incurs is $O(\frac{1}{k})$.*

THEOREM 5.2. *Consider the threshold values $x_1, ..., x_n$ for which equation 2 holds for every $m \in \{1, ..., n-1\}$ and $x_1 = \frac{1}{2}$. When the players' valuations are distributed uniformly, a modified priority-game g which is based on the threshold values $x_1, ..., x_n$ achieves optimal revenue. The revenue loss it incurs is $O(\frac{1}{k})$.*

Acknowledgment. We gratefully thank Ilya Segal for suggesting the use of "virtual utilities" for proving bounds for the revenue loss. We also thank Motty Perry, Daniel Lehmann, Abraham Neyman, Ron Lavi and Ahuva Mualem for helpful discussions.

6. REFERENCES

[1] commerce-one. Web Page: http://www.commerceone.com.

[2] ebay. Web Page: http://www.ebay.com.

[3] enron. Web Page: http://www.enron.com.

[4] orbis-online. Web Page: http://www.orbisonline.com.

[5] vertical-net. Web Page: http://www.verticalnet.com.

[6] Z. Bar-Yossef, K. Hildrum, and F. Wu. Incentive-compatible online auctions for digital goods. In *13th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2002.

[7] Liad Blumrosen and Noam Nisan. Auctions with severely bounded communications. Web Page: http://www.cs.huji.ac.il/~liad/Bounded.ps, 2002.

[8] J. Feigenbaum, C. H. Papadimitriou, and S. Shenker. Sharing the cost of muliticast transmissions. In *ACM Symposium on Theory of Computing*, pages 218–227, 2000.

[9] Donald F. Ferguson, Christos Nikolaou, and Yechiam Yemini. Economic models for allocating resources in computer systems. In Scott Clearwater, editor, *Market-Based Control: A Paradigm for Distributed Resource Allocation.* World Scientific, 1995.

[10] Riley J. G. and Samuelson W. F. Optimal auctions. *American Economic Review*, pages 381–392, 1981.

[11] Andrew V. Goldberg, Jason D. Hartline, and Andrew Wright. Competitive auctions and digital goods. In *Symposium on Discrete Algorithms*, pages 735–744, 2001.

[12] Papadimitriou C. H. Algorithms, games, and the internet. 2001.

[13] Y.A Korilis, A. A. Lazar, and A. Orda. Architecting noncooperative networks. *IEEE Journal on Selected Areas in Communication (Special Issue on Advances in the Fundamentals of Networking)*, 13(7):1241–1251, September 1995.

[14] Levy L., Blumrosen L., and Nisan N. Online markets for distributed market services: the majic system. In *USITS 01*, 2001.

[15] J. Laffont and E. Maskin. A differentiable approach to dominant strategy mechanisms. *Econometrica*, 48:1506–1520, 1980.

[16] Ron Lavi and Noam Nisan. Competitive analysis of incentive compatible on-line auctions. In *ACM Conference on Electronic Commerce*, pages 233–241, 2000.

[17] A.A. Lazar and N. Semret. The progressive second price auction mechanism for network resource sharing. In *8th International Symposium on Dynamic Games*, Maastricht, The Netherlands, July 1998.

[18] Daniel Lehmann, Liadan Ita O'Callaghan, and Yoav Shoham. Truth revelation in rapid, approximately efficient combinatorial auctions. In *1st ACM conference on electronic commerce*, 1999.

[19] Harstad R. M. and Rothkopf M. H. On the role of discrete bid levels in oral auctions. *Europian Journal of Operations Research*, pages 572–581, 1994.

[20] P. McAfee. Coarse matching. Web Page: http://www.eco.utexas.edu/~mcafee/Papers/PDF/Categories.pdf, 2002.

[21] P. Milgrom and I. Segal. Envelope theorems for arbitrary choice sets. *Econometrica*, 70(2):583–601, 2002.

[22] R. B. Myerson. Optimal auction design. *Mathematical of operational research*, pages 58–73, 1981.

[23] Noam Nisan. Algorithms for selfish agents. In *STACS*, 1999.

[24] Noam Nisan. Bidding and allocation in combinatorial auctions. In *ACM Conference on Electronic Commerce*, 2000.

[25] Noam Nisan, Shmulik London, Ori Regev, and Noam Camiel. Globally distributed computation over the

internet – the popcorn project. In *ICDCS*, 1998.

[26] Noam Nisan and Ilya Segal. The communication complexity of efficient allocation problems, 2001. Preliminary version available from http://www.cs.huji.ac.il/˜noam/mkts.html.

[27] Ori Regev and Noam Nisan. The popcorn market - an online market for computational resources. In *ICE*, 1998.

[28] Jeffrey S. Rosenschein and Gilad Zlotkin. *Rules of Encounter: Designing Conventions for Automated Negotiation Among Computers*. MIT Press, 1994.

[29] A. E. Roth and A. Ockenfels. Last minute bidding and the rules for ending second-price auctions: Theory and evidence from a natural experiment on the internet. Technical Report NBER Working Papers 7729, National Bureau of Economic Research, Inc., 2000.

[30] Shenker S. Making greed work in networks: A game-theoretic analysis of switch service disciplines. In *Proc. of the ACM SIGCOMM*, 1994.

[31] T. Sandholm, S. Suri, A. Gilpin, and Levine D. Cabob: A fast optimal algorithm for combinatorial auctions. In *IJCAI*, 2001.

[32] S. Shenkar, Clark D. E., and Hertzog S. Pricing in computer networks: Reshaping the research agenda. *ACM Computational Comm. Review*, pages 19–43, 1996.

[33] Hal R. Varian. Economic mechanism design for computerized agents. In *Proceedings of the First Usenix Conference on Electronic Commerce*, New York, July 1995.

[34] W. Vickrey. Counterspeculation, auctions and competitive sealed tenders. *Journal of Finance*, pages 8–37, 1961.

[35] Rakesh Vohra and Sven de Vries. Combinatorial auctions: A survey, 2000. Availailabe from www.kellogg.nwu.edu/faculty/vohra/htm/res.htm.

[36] C. A. Waldspurger, T. Hogg, B. A. Huberman, J. O. Kephart, and W. S. Stornetta. Spawn: A distributed computational economy. *IEEE Transactions on Software Engineering*, 18(2), 1992.

[37] W.E. Walsh and M.P. Wellman. A market protocol for decentralized task allocation. In *The Proceedings of the Third International Conference on Multi-Agent Systems (ICMAS-98)*, 1998.

[38] W.E. Walsh, M.P. Wellman, P.R. Wurman, and J.K. MacKie-Mason. Auction protocols for decentralized scheduling. In *Proceedings of The Eighteenth International Conference on Distributed Computing Systems (ICDCS-98)*, 1998.

[39] R. Wilson. Efficient and competitive rationing. *Econometrica*, 57:1–40, 1989.

[40] Edo Zurel and Noam Nisan. An efficient approximate allocation algorithm for combinatorial auctions. In *ACM conference on electronic commerce*, 2001.