Multi-player and Multi-round Auctions with Severely Bounded Communication

Liad Blumrosen[1], Noam Nisan[1], and Ilya Segal[2]

[1] School of Engineering and Computer Science.
The Hebrew University of Jerusalem, Jerusalem, Israel.
liad,noam@cs.huji.ac.il

[2] Department of Economics,
Stanford University, Stanford, CA 94305
ilya.segal@stanford.edu

Abstract. We study auctions in which bidders have severe constraints on the size of messages they are allowed to send to the auctioneer. In such auctions, each bidder has a set of k possible bids (i.e. he can send up to $t = log(k)$ bits to the mechanism). This paper studies the loss of economic efficiency and revenue in such mechanisms, compared with the case of unconstrained communication. For any number of players, we present auctions that incur an efficiency loss and a revenue loss of $O(\frac{1}{k^2})$, and we show that this upper bound is tight. When we allow the players to send their bits sequentially, we can construct even more efficient mechanisms, but only up to a factor of 2 in the amount of communication needed. We also show that when the players' valuations for the item are not independently distributed, we cannot do much better than a trivial mechanism.

1 Introduction

Computers on the Internet are owned by different parties with individual preferences. Trying to impose protocols and algorithms on them in the traditional computer-science way is doomed to fail, since each party might act for its own selfish benefit. Thus, designing protocols for Internet-like environments requires the usage of tools from other disciplines, especially microeconomic theory and game theory. This intersection between computer science theory and economic theory raises many interesting questions. Indeed, much theoretical attention was given in recent years to problems with both game theoretic and algorithmic aspects (see e.g. the surveys [10, 18, 5]). Many of the algorithms for such distributed environments are closely related to the theory of *mechanism design* and in particular to *auction theory* (see [6] for comprehensive survey about auctions). An auction is actually an algorithm, that allocates some resources among a set of players. The messages (bids) that the players send to the auctioneer are the input for this algorithm, and it outputs an allocation of the resources and payments for the players. The main challenge in designing auctions is related to the incomplete information that the designer has about the players' secret data (for example, how much they are willing to pay for a certain resource). The auction mechanism must somehow elicit this information from the selfish participants, in order to achieve global or "social" goals (e.g. maximize the seller's revenue). Recent results show that auctions are hard to implement in practice. The reasons might be computational (see e.g. [12, 8]), communication-related ([13]), uncertainty about timing or participants ([4, 7],) and many more. This, and the growing usage of auctions in e-commerce (e.g. [14, 21, 15]) and in various computing systems (see e.g. [11, 19, 20]) led researchers to take computational effects into consideration when designing auctions.

Much interest was given in the economic literature to the design of *optimal auctions* and *efficient auctions*. Optimal auctions are auctions that maximize the seller's revenue. Efficient auctions maximize the social welfare, i.e. they allocate the resources to the players that want them the most. A positive correlation usually exist between the two measures: a player is willing to pay more for an item that is worth a higher value to her. Nevertheless, efficient auctions are not necessarily optimal, and vice versa. In our model, each player has a private valuation for a single item (i.e. she knows how much she values the item, but this value is a private information for herself). The goal of the auction's designer (in the Bayesian framework) is, given distributions on the players' valuations, to find auctions that maximize the expected revenue or the expected welfare, when the players act selfishly. For the single item case, these problems are in fact solved: the Vickrey auction (or the 2nd-price auction, see [17]) is efficient; Myerson, in a classic paper ([9]), fully characterize optimal auctions when the players' valuations are independently distributed. In the same paper, Myerson also shows that Vickrey's auction (with some reservation price) is also optimal (i.e. revenue maximizing), when the distribution functions hold some regularity property. Optimal auctions and efficient auctions were studied lately also by computer scientists (e.g. [4, 16]). Recently, Blumrosen and Nisan ([1]) initiated the study of auctions with severely bounded communication, i.e. settings where each player can send a message of up to t bits to the mechanism. In other words, each bidder can choose a bid out of a set of $k = 2^t$ possible bids. The players' valuations, however, can be any real numbers in the range $[0, 1]$.

Here, we generalize the main results from [1] for multi-player games. We also study the effect of relaxing some of the assumptions made in [1], namely the simultaneous bidding and the independence of the valuations.

Severe constraints on the communication are expected in settings where we need to design quick, and cheap auctions that should be performed frequently. For example, if a route for a packet over the Internet is auctioned, we can dedicate for this purpose only a small number of bits. Otherwise, the network will be congested very quickly. For example, we might want to use some unused bits in existing networking protocols (e.g. IP or TCP) to transfer the bidding information. This is opposed to the traditional economic approach that views the information sent by the players as real numbers (representing these can take infinite number of bits!). Low communication also serves as a proxy for other desirable properties: with low communication the interface for the auction is simpler (the players have a small number of possible bids to choose from), the information revelation is smaller and only a small number of discrete prices is used. In addition, understanding the tradeoffs between communication and auctions' optimality (or efficiency) might help us find feasible solutions for settings which are currently computationally impossible (combinatorial auctions' design is the most prominent example).

Under severe communication restrictions, [1] characterizes optimal and efficient auctions among *two players*. They prove that the welfare loss and the revenue loss in mechanisms with t-bits messages is mild: for example, with only *one* bit allowed for each player (i.e. $t = 1$) we can have 97 percent of the efficiency achieved by auctions that allow the players to send infinite number of bits (with uniform distributions)! Asymptotically, they show that the loss (for both measures) diminishes exponentially in t (specifically $O(\frac{1}{2^{2t}})$ or $O(\frac{1}{k^2})$ where $k = 2^t$). These upper bounds are tight: for particular distribution functions, the expected welfare loss and the expected revenue loss in *any* mechanism are $\Omega(\frac{1}{k^2})$.

In this work, we show *n-player* mechanisms that, despite using very low communication, are nearly optimal (or nearly efficient). These mechanisms are an extension of the *"priority-games"* and *"modified priority-games"* concepts described in [1], and they achieve the asymptotically-optimal results with dominant strategies equilibrium and with individual-rationality constraints (see formal definitions in the body of the paper). For both measures, we characterize mechanisms that incur a loss of $O(\frac{1}{k^2})$, and we show that for some distribution functions (e.g. the uniform distribution) this bound is tight.

We also extend the framework to the following settings:

- **multi-round auctions:** By allowing the bidders to send the bits of their messages one bit at a time, in alternating order, we can strictly increase the efficiency of auctions with bounded communication. In such auctions, each player knows what bits where sent by all players up to each stage. However, we show that the same extra gain can be achieved in simultaneous auctions that use less than double amount of communication.
- **Joint distributions:** When the players' valuations are statistically dependent, we show that we cannot do better (asymptotically) than a trivial mechanism that achieves an efficiency loss of $O(\frac{1}{k})$. Specifically, we show that for some joint distribution functions, *every* mechanism with k possible bids incurs a revenue loss of at least $\Omega(\frac{1}{k})$.
- **Bounded distribution functions:** We know ([1]) that we cannot construct one mechanism that incurs a welfare loss of $O(\frac{1}{k^2})$ for **all** distribution functions. Nevertheless, if we assume that the density functions are bounded from above or from below, a trivial mechanism achieves results which are asymptotically optimal.

The organization of the paper is as follows: section 2 describes the formal model of auctions with bounded communication. Section 3 gives tight upper bounds for the optimal welfare loss and revenue loss in n-player mechanisms. Section 4 studies the case of bounded density functions and joint distributions.

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 $\frac{2}{3}$

Fig. 1. A matrix representation for a mechanism with two possible bids. E.g., when Alice bids $''1''$ and Bob bids $''0''$, Alice wins the item and pays $\frac{1}{3}$.

Finally, section 5 discusses multi round auctions. All the omitted proofs can be found in the full version ([2]).

2 The Model

We consider single item, sealed bid auctions among n risk-neutral players. Player i has a private data (*valuation*) $v_i \in [0,1]$ that represents the maximal payment he is willing to pay for the item. For every player i, v_i is independently drawn from a density function f_i $\left(\int_0^1 f_i(v)dv = 1\right)$ which is commonly known for all participants. The cumulative distribution for player i is F_i. Throughout the paper we assume that the distribution functions are continuous and always positive. We also assume a normalized model, i.e. players' valuations for not having the item are zero. The seller's valuation for the item is zero, and the players' valuations depend only on whether they win the item or not (no externalities).

Players aim to maximize their utilities, which are *quasi-linear*, i.e. the utility of player i from the item is $v_i - p_i$ when p_i is his payment.

The unique assumption in our model, is that each player can send a message of no more than $t = \lg(k)$ **bits** to the mechanism, i.e. players can choose one of k possible **bids** (or messages). Denote the possible set of bids for the players as $\beta = \{0, 1, 2, ..., k-1\}$. In each auction, player i chooses a bid $b_i \in \beta$. A mechanism determines the allocation and payments given a vector of bids $b = (b_1, ..., b_n)$:

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

- *$a : (\beta \times ... \times \beta) \rightarrow [0,1]^n$ is the allocation scheme. 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 $\forall b \; \sum_{i=1}^{n} a_i(b) \leq 1$.*
- *$p : (\beta \times ... \times \beta) \rightarrow \Re^n$ is the payment scheme. $p_i(b)$ is player i's payment given a bids' vector b (paid only upon winning).*

Definition 2 *In a mechanism with k-possible bids, $|\beta| = k$. Denote the set of all mechanisms with k-possible bids among n players by $\mathbf{G_{n,k}}$.*

Figure 1 describes the matrix representation of a 2-player mechanism with two possible bids ("0" or "1").

All the results in this paper are achieved with *ex-post Individually-Rational* (IR) mechanisms, i.e. mechanisms in which players can always ensure themselves not to pay more than their valuations for the item (or 0 when they lose). (We equivalently use the term: mechanisms with ex-post individual rationality.)

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

Denote $\varphi_k = \{s \,|\, s : [0,1] \rightarrow \{0, 1, ..., k-1\}\}$ (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 and $v = (v_1, ..., v_n)$. 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 $s = (s_i, s_{-i})$.

2.1 Optimality Measures

The players in our model choose strategies that maximize their utilities. We are interested in games with stable behaviour for all players, i.e. such that these strategies form an equilibrium.

Definition 6 *Let $u_i(g, s)$ be the expected utility of player i from game g when bidders use the 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 *The strategies $s = (s_1, ..., s_n)$ form a Bayesian-Nash equilibrium in a mechanism $g \in G_{n,k}$, if for every player i, s_i is the best response for the strategies s_{-i} of the other players, i.e.* $\forall i \quad \forall \widetilde{s_i} \in \varphi_k \quad u_i(g, (s_i, s_{-i})) \geq u_i(g, (\widetilde{s_i}, s_{-i}))$

Definition 8 *A strategy s_i for player i is dominant in mechanism $g \in G_{n,k}$ if regardless of the other players' strategies s_{-i}, i cannot gain a higher utility by changing his strategy, i.e.* $\forall \widetilde{s_i} \in \varphi_k \quad \forall s_{-i} \quad u_i(g, (s_i, s_{-i})) \geq u_i(g, (\widetilde{s_i}, s_{-i}))$

We say that a mechanism g has *a dominant strategies equilibrium* if for every player i there exists a strategy s_i which is dominant. Clearly, a dominant strategies equilibrium is also a Bayesian-Nash equilibrium.

Each bidder aims to maximize her expected utility. As mechanisms' designers, we aim 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 valuation of the winning players (if any).

Definition 9 *Let $w(g, s)$ denote the expected welfare in the n-player game g when bidders' strategies are s, i.e.* $w(g,s) = E_{v \in [0,1]^n} \left(\sum_{i=1}^{n} a_i\left(s(v)\right) \cdot v_i\right)$

Definition 10 *Let $r(g,s)$ denote the expected revenue in the n-player game g when bidders' strategies are s, i.e. $r(g,s) = E_{v\in[0,1]^n}\left(\sum_{i=1}^{n} a_i\left(s(v)\right)\cdot p_i\left(s(v)\right)\right)$*

Definition 11 *We say that a mechanism $g \in G_{n,k}$ achieves an expected welfare **(revenue)** of α if g has a Bayesian-Nash equilibrium s for which the expected welfare **(revenue)** is α, i.e. $w(g,s) = \alpha$ ($\mathbf{r(g,s)} = \alpha$).*

Definition 12 *We say that a mechanism $g \in G_{n,k}$ incurs a welfare loss of c, if there is a Bayesian-Nash equilibrium s in g such that the difference between $w(g,s)$ and the maximal welfare with unbounded communication is c.*

We say that g incurs a revenue loss of c, if there is an individually-rational Bayesian-Nash equilibrium s in g, such that the difference between $r(g,s)$ and the optimal revenue, achieved in an individually-rational mechanism with Bayesian-Nash equilibrium in the unbounded communication case, is c.

Recall that an equilibrium is individually rational, if the expected utility of each player, given his own valuation, is non negative. The mechanism described in Fig. 1 has a dominant strategy equilibrium that achieves an expected welfare of $\frac{35}{54}$ (with uniform distributions). Alice's dominant strategy is the threshold strategy based on $\frac{1}{3}$, i.e. she bids "0" when her valuation is below $\frac{1}{3}$, and "1" otherwise. The threshold strategy based on $\frac{2}{3}$ is dominant for Bob. We know ([17]) that the optimal welfare from a 2-player auction with unconstrained communication is $\frac{2}{3}$. Thus, the welfare loss incurred by this mechanism is $\frac{2}{3} - \frac{35}{54} = \frac{1}{54}$.

3 Multi-player Mechanisms

In this section, we construct n-player mechanisms with bounded communication which are asymptotically optimal (or efficient). We prove that they incur losses of welfare and revenue of $O(\frac{1}{k^2})$, and that these upper bounds are tight.

It was shown in [1] that "priority-games" (PG) and "modified priority-games" (MPG) are efficient and optimal (respectively) among all the 2-player mechanisms with bounded communications. For the n-player case, the characterization of the welfare maximizing and the revenue maximizing mechanisms remains an open question. We conjecture that PG's (and MPG's) with optimally chosen payments are efficient (optimal). We show that PG's and MPG's achieve **asymptotically**-optimal welfare and revenue (respectively). Note, that even though our model allows lotteries, our analysis presents only deterministic mechanisms. Indeed, [1] shows that optimal results are achieved by deterministic mechanisms.

Definition 13 *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.*

For example, Fig. 1 describes a priority game: the player with the highest bid wins, and ties are always broken in favour of Bob.

Definition 14 *A game is called a* ***modified priority-game*** *if it has an allocation as in priority-games, but no allocation is done when all players bid* 0.

Definition 15 *An* n*-player priority-game based on a profile of threshold values' vectors* $\overrightarrow{t} = (t^1, ..., t^n) \in \times_{i=1}^{n} \Re^{k+1}$ *(where for every* i, $t_0^i \leq t_1^i \leq ... \leq t_k^i$*) is a mechanism that its allocation is as in a priority game and its payment scheme is as follows: when player* j *wins the item for the bids vector* b *she pays the smallest valuation she might have and still win the item, given that she uses the threshold strategy* s_j *based on* t^j*. I.e.* $p_j(b) = \min\{v_j | a_j\,(s_j(v_j), b_{-j}) = 1\}$*. We denote this mechanism as* $PG_k(\overrightarrow{t})$*. A modified priority game with a similar payment rule is called a modified priority-game based on a profile of threshold values' vectors, and is denoted by* $MPG_k(\overrightarrow{t})$*.*

For example, Fig. 1 describes a priority game based on the threshold values $(0, \frac{1}{3}, 1)$ and $(0, \frac{2}{3}, 1)$. When Bob bids 0, the minimal valuation of Alice for which she still wins is $\frac{1}{3}$, thus this is her payment upon winning, and so on. We first show that these mechanisms have dominant-strategies and ex-post IR:

Proposition 1 *For every profile of identical threshold values' vectors* $\overrightarrow{t} = (x, x, ..., x)$, $x \in R^{k+1}$ *and* $x_0 \leq x_1 \leq ... \leq x_k$*, the threshold-strategies based on these threshold values are dominant in* $PG_k(\overrightarrow{t})$*, and this mechanism is ex-post IR.*

3.1 Asymptotically Efficient Mechanisms

Now, we show that given any set of n distribution functions of the players, we can construct a mechanism that incurs a welfare loss of $O(\frac{1}{k^2})$. In [1], a similar upper bound was given for the case of 2-player mechanisms:

Theorem 1 *[1] For every set of distribution functions on the players' valuations, the 2 player mechanism* $PG_k(x, y)$ *incurs an expected welfare loss of* $O(\frac{1}{k^2})$ *(for some threshold values vectors* x, y*). Moreover, when all valuations are distributed uniformly, the expected welfare loss is at least* $\Omega(\frac{1}{k^2})$ *in any mechanism.*

Here, we prove that n-player priority games are *asymptotically* efficient:

Theorem 2 *For any number of players* n*, and for any set of distribution functions of the players' valuations, the mechanism* $PG_k(\overrightarrow{t})$ *incurs a welfare loss of* $O(\frac{1}{k^2})$*, for some threshold values vector* $\overrightarrow{t} \in \times_{i=1}^{n} \Re^{k+1}$*. This mechanism has a dominant-strategies equilibrium with ex-post IR.*

In the following theorem we show that for uniform distributions, the welfare loss is proportional to $\frac{1}{k^2}$:

Theorem 3 *When valuations are distributed uniformly, and for any (fixed) number of players* n*, any mechanism* $g \in G_{n,k}$ *incurs a welfare loss of* $\Omega(\frac{1}{k^2})$*.*

Proof. Consider only the case where players 1 and 2 have valuations greater than $\frac{1}{2}$, and the rest of the players have valuations below $\frac{1}{2}$. This occurs with the constant probability of $\frac{1}{2^n}$ (n is fixed). For maximal efficiency, a mechanism with k possible bids always allocates the item to player 1 or 2. But due to theorem 1, a welfare loss of $\Omega(\frac{1}{k^2})$ will still be incurred (the fact that in theorem 1 the valuations' range is $[0,1]$ and here it is $[\frac{1}{2},1]$ only changes the constant c). Thus, any mechanism will incur a welfare loss which is $\Omega(\frac{1}{k^2})$.

3.2 Asymptotically Optimal Mechanisms

Now, we present mechanisms that achieve asymptotically optimal expected *revenue*. We show how to construct such mechanisms and give tight upper bounds for the revenue loss they incur.

Most results in the economic literature on revenue-maximizing auctions, assume that the distribution functions of the players' valuations holds a *regularity* property (as defined by Myerson [9], see below). For example, only when the valuations of all players are distributed with the same regular distribution-function, it is known that Vickrey's 2nd-price auction, with an appropriately chosen reservation price, is revenue-optimal ([17, 9, 3]).

Definition 16 *([9]) Let f be a density function, and let F be its cumulative function. We say that f is regular, if the function $\widetilde{\mathbf{v}}(\mathbf{v}) = \mathbf{v} - \frac{\mathbf{1-F(v)}}{\mathbf{f(v)}}$ is monotone, strictly increasing function of v. We call $\widetilde{v}$ the virtual utility.*

We define the virtual utility of all the players, except the winner, as zero. The seller's virtual utility is equal to his valuation for the item (zero in our model). Myerson ([9]) observed that in equilibrium, the expected revenue equals the expected virtual-utility (i.e. the average virtual utility of the winning players):

Theorem 4 *([9]) Consider a model with unbounded communication, in which losing players pay zero. Let h be a direct-revelation mechanism, which is incentive compatible (i.e. truth telling by all players forms Nash equilibrium) and individually rational. Then in h, the expected revenue equals the expected virtual utility.*

Simple arguments show (see [1]) that Myerson's observation also holds for auctions with bounded communication:

Proposition 2 *([1]) Let $g \in G_{n,k}$ be a mechanism with Bayesian Nash equilibrium $s = (s_1, ..., s_n)$ and ex-post individual rationality. Then, the expected revenue of s in g is equal to the expected virtual-utility in g.*

Using this property, the revenue optimization problem can be reduced to a welfare optimization problem, which was solved for the n-player case in theorems 2 and 3. We extend the techniques used in [1] for the n-player case: we optimize the expected *welfare* in settings where the players consider their virtual utility as their valuations (see [2] for the proof). We show that for a fixed n, and for every regular distribution, there is a mechanism that incurs a revenue loss of $O(\frac{1}{k^2})$. Again, this bound is tight: for uniform distributions the optimal revenue loss is proportional to $\frac{1}{k^2}$.

Theorem 5 *Assume that all valuations are distributed with the same regular distribution function. Then, for any number of players n, $MPG_k(\overrightarrow{t})$ incurs a revenue loss of $O(\frac{1}{k^2})$, for some threshold values vector $\overrightarrow{t} \in \times_{i=1}^{n}\Re^{k+1}$. This mechanism has dominant strategies equilibrium with ex-post IR.*

Theorem 6 *Assume that the players' valuations are distributed uniformly. Then, for any (fixed) number of players n, any mechanism $g \in G_{n,k}$ incurs a revenue loss of $\Omega(\frac{1}{k^2})$.*

4 Bounded Distributions and Joint Distributions

In previous theorems, we showed how to construct mechanisms with asymptotically optimal welfare and revenue, given a set of distribution functions. Can we design a particular mechanism that achieve similar results for **all** distribution functions? Due to [1], the answer in general is no. The simple mechanism $PG_k(x,x)$ where $x = (0, \frac{1}{k}, \frac{2}{k}, ..., \frac{k-1}{k}, 1)$ incurs a welfare loss of $O(\frac{1}{k})$ and no better upper bound can be achieved. Nevertheless, we show that if the distribution functions are bounded from above or from below, this trivial mechanism for two players achieves an expected welfare which is asymptotically optimal.

Definition 17 *We say that a density function f is bounded from above **(below)** if for every x in its domain, $f(x) \leq c$ $(\mathbf{f(x) \geq c})$, for some constant c.*

Proposition 3 *For every pair of distribution functions of the players' valuations which are bounded from above, 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 of $O(\frac{1}{k^2})$.*
*For every pair of distribution functions which are bounded from below, **every** mechanism incurs an expected welfare loss of $\Omega(\frac{1}{k^2})$.*

So far, we assumed that the players' valuations are drawn from statistically independent distributions. Now, we relax this assumption and deal with general joint distributions of the valuations. For this case, we show that a trivial mechanism is actually the best we can do (asymptotically). Particularly, it derives a tight upper bound of $O(\frac{1}{k})$ for the efficiency loss in 2-player games.

Theorem 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 joint distribution ϕ on the players' valuations. Moreover, for every k there is a joint distribution function ϕ_k such that **every** mechanism $g \in G_{2,k}$ incurs a welfare loss $\geq c \cdot \frac{1}{k}$ (where c is some positive constant independent of k).*

B \ A	0	1
0	$A, 0$	$B, \frac{1}{4}$
1	$A, \frac{1}{3}$	$B, \frac{3}{4}$

Fig. 2. (h_1) This sequential game (when A bids first, then B) achieves higher expected welfare than any simultaneous mechanism with the same communication complexity (2 bits). The welfare is achieved with Bayesian-Nash equilibrium.

5 Multi-round Auctions

In previous sections, we analyzed auctions with bounded communication in which players simultaneously send their bids to the mechanism. Can we get better results with *multi-round* (or *sequential*) mechanisms? I.e. mechanisms in which players send their bids one bit at a time, in alternating order. In this section, we show that sequential mechanisms can achieve better results. However, the additional gain (in the amount of communication) is up to a factor of 2.

5.1 Sequential Mechanisms Can Do Better

The definitions in this section are similar in spirit to the model described in section 2. For simplicity, we present this model less formally.

Definition 18 *A sequential (or multi-round) mechanism is a mechanism in which players send their bids one bit at a time, in alternating order. In each stage, each player knows the bits the other players sent so far. Only after all the bits were transmitted, the mechanism determines the allocation and payments.*

Definition 19 *The communication complexity of a mechanism is the total amount of bits which are sent by the players.*

Definition 20 *A strategy for a player in a sequential mechanism is the way she determines the bits she transmits, at every stage, given her valuation and given the other players' bits up to this stage.*
A strategy for a player in a sequential mechanism is called a threshold strategy if in each stage i of the game, the player determines the bit she sends according to some threshold value x_i; I.e. if her valuation is smaller than this threshold she bids 0, or bids 1 otherwise.

Denote the following sequential mechanism by h_1 (see Fig. 2): Alice sends one bit to the mechanism first. Bob, knowing Alice's bid, also sends one bit. When Alice bids 0: Bob wins if he bids 1 and pays $\frac{1}{4}$; If he bids zero Alice wins and pays zero. When Alice bids 1: Bob also wins when he bids 1, but now he pays $\frac{3}{4}$; If he bids zero, Alice wins again, but now she pays $\frac{1}{3}$.

The communication complexity of this mechanism is 2 (each player sends one bit to the mechanism). When players' valuations are distributed uniformly, this mechanism achieves an expected welfare which is greater than the optimal welfare from simultaneous mechanisms with the same communication complexity:

Proposition 4 *When valuations are distributed uniformly, the mechanism h_1 above has a Bayesian-Nash equilibrium and an expected welfare of 0.653.*

Proof. Consider the following strategies: Alice uses a threshold strategy based on the threshold value $\frac{1}{2}$, and Bob uses the threshold $\frac{1}{4}$ when Alice bids "0" and the threshold $\frac{3}{4}$ when Alice bids 1. It is easy to see that these strategies form a Bayesian-Nash equilibrium, with expected welfare of 0.653.

The communication complexity of the mechanism h_1 above is 2 bits (each player sends one bit). The efficient simultaneous mechanism, with 2 bits' complexity, achieves an expected welfare of 0.648 ([1]). Thus, we can gain more efficiency with sequential mechanisms. Note that this expected welfare is achieved in h_1 with Bayesian-Nash equilibrium, as opposed to dominant strategies equilibria in all previous results.

5.2 The Extra Gain from Sequential Mechanisms is Limited

How significant is the extra gain from sequential mechanisms? The following theorem states that for every sequential mechanism there exists a simultaneous mechanism that achieves at least the same welfare with less than double amount of communication. Note that in sequential mechanisms the players must be informed about the bits the other players sent (we do not take this into account in our analysis), so the total gain in communication can be very mild. We start by proving that optimal welfare can be achieved with threshold-strategies.

Lemma 1 *Given a sequential mechanism h and a profile of strategies $s = (s_1, ..., s_n)$ of the players, there exists a profile of **threshold** strategies $\overline{s} = (\overline{s_1}, ..., \overline{s_n})$ that achieves at least the same welfare with h as s does.*

Theorem 8 *Let h be a 2-player sequential mechanism with communication complexity m. Then, there exists a **simultaneous** mechanism g that achieves at least the same expected welfare as h, with communication complexity of $2m - 1$.*

Proof. Consider a 2-player, sequential mechanism h with a Bayesian-Nash equilibrium, and with communication complexity m (we assume m is even, i.e. each player sends $\frac{m}{2}$ bits). Due to lemma 1, there exists a profile $s = (s_1, s_2)$ of threshold-strategies that achieves at least the same expected welfare on h as the equilibrium welfare. Now, we will count the number of different thresholds of player A: at stage 1, she uses a single threshold. After B sends his first bit, A also uses a threshold, but she might have a different one for each history, i.e. $2^2 = 4$ thresholds. This way, it is easy to see that the number of thresholds for A is : $\alpha_A(m) = 2^0 + 2^2 + ... + 2^{m-2}$, and for player B is $\alpha_B(m) = 2^1 + 2^3 + ... + 2^{m-1}$. Next, we construct a simultaneous mechanism g that achieves at least the same expected welfare with a communication complexity smaller than $2m - 1$. In g, each player simply "tells" the mechanism within which 2 of the threshold values his valuations is. The number of bits the two players need for transmitting this information is:

$$log(\alpha_A(m)+1) + log(\alpha_A(m)+1) < log(2^{m-1}) + log(2^m) = 2m - 1$$
In the full paper ([2]) we show that the new strategies forms an equilibrium.

Acknowledgments. The work of the first two authors was supported by a grant from the Israeli Academy of Sciences. The third author was supported by the National Science Foundation.

References

1. Liad Blumrosen and Noam Nisan. Auctions with severely bounded communications. In *43th FOCS*, 2002.
2. Liad Blumrosen, Noam Nisan, and Ilya Segal. Multi-player and multi-round auctions with severely bounded communications, 2003. Full version. available from http://www.cs.huji.ac.il/~liad.
3. Riley J. G. and Samuelson W. F. Optimal auctions. *American Economic Review*, pages 381–392, 1981.
4. Andrew V. Goldberg, Jason D. Hartline, and Andrew Wright. Competitive auctions and digital goods. In *Symposium on Discrete Algorithms*, pages 735–744, 2001.
5. Papadimitriou C. H. Algorithms, games, and the internet. *ACM Symposium on Theory of Computing*, 2001.
6. Paul Klemperer. The economic theory of auctions. *Edward Elgar Publishing*, 2000.
7. Ron Lavi and Noam Nisan. Competitive analysis of incentive compatible on-line auctions. In *ACM Conference on Electronic Commerce*, pages 233–241, 2000.
8. 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.
9. R. B. Myerson. Optimal auction design. *Mathematical of operational research*, pages 58–73, 1981.
10. Noam Nisan. Algorithms for selfish agents. In *STACS*, 1999.
11. Noam Nisan, Shmulik London, Ori Regev, and Noam Camiel. Globally distributed computation over the internet – the popcorn project. In *ICDCS*, 1998.
12. Noam Nisan and Amir Ronen. Algorithmic mechanism design. In *STOC*, 1999.
13. Noam Nisan and Ilya Segal. Communication requirements of efficiency and supporting lindahl prices, 2003. working paper available from http://www.cs.huji.ac.il/~noam/mkts.html.
14. Web page. ebay. http://www.ebay.com.
15. Web page. vertical-net. http://www.verticalnet.com.
16. Amir Ronen and Amin Saberi. Optimal auctions are hard. In *43th FOCS*, 2002.
17. W. Vickrey. Counterspeculation, auctions and competitive sealed tenders. *Journal of Finance*, pages 8–37, 1961.
18. Rakesh Vohra and Sven de Vries. Combinatorial auctions: A survey, 2000. Availailabe from www.kellogg.nwu.edu/faculty/vohra/htm/res.htm.
19. 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.
20. 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.
21. Web-page. commerce-one. http://www.commerceone.com.