A Combinatorial Look at Auctions

Dave Buchfuhrer Chris Umans

October 9, 2009

Outline

1. Introduction
   - Auctions
   - VCG Mechanisms
   - The Model
   - Background
2. Allocate All
   - Allocate All Items
   - Large Range
   - VC Dimension
3. Our Work
   - The Issues
   - Large Range
   - VC Dimension
   - Subset Sum
4. Conclusions

Outline

1. Introduction
   - Auctions
   - VCG Mechanisms
   - The Model
   - Background
2. Allocate All
   - Allocate All Items
   - Large Range
   - VC Dimension
3. Our Work
   - The Issues
   - Large Range
   - VC Dimension
   - Subset Sum
4. Conclusions

Auctions

We consider combinatorial auctions of m items to n bidders where we wish to maximize the social welfare.

- The VCG mechanism can be used for truthfulness
- An FPTAS can be used to approximate arbitrarily well
- Can we achieve efficiency and truthfulness *simultaneously?*

Outline

1. Introduction
   - Auctions
   - VCG Mechanisms
   - The Model
   - Background
2. Allocate All
   - Allocate All Items
   - Large Range
   - VC Dimension
3. Our Work
   - The Issues
   - Large Range
   - VC Dimension
   - Subset Sum
4. Conclusions

The VCG Mechanism

- By participating in the auction, each bidder harms the others

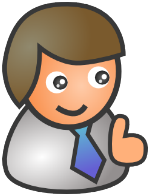

80

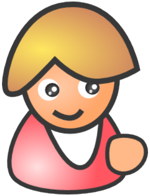

80

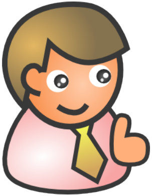

80

The VCG Mechanism

- By participating in the auction, each bidder harms the others

70

70

70

70

The VCG Mechanism

- By participating in the auction, each bidder harms the others

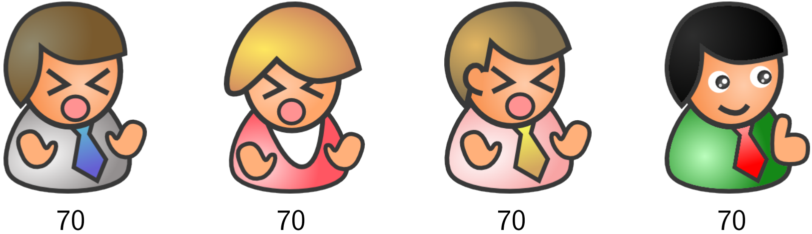

- To counter greed, each player is charged for this harm

The VCG Mechanism

- By participating in the auction, each bidder harms the others

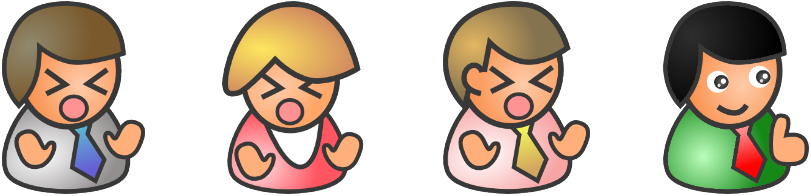

70 70 70 70

- To counter greed, each player is charged for this harm
- Intuitively, the player wants the social welfare maximized

Maximal-in-Range

The VCG mechanism allows for truth given exact solutions

- For VCG to work, simply maximize $W(a)$ over all allocations A

Maximal-in-Range

The VCG mechanism allows for truth given exact solutions

- For VCG to work, simply maximize $W(a)$ over all allocations A
- VCG works if we maximize $W(r)$ over any $R \subseteq A$

Maximal-in-Range

The VCG mechanism allows for truth given exact solutions

- For VCG to work, simply maximize $W(a)$ over all allocations A
- VCG works if we maximize $W(r)$ over any $R \subseteq A$
- These are exactly the types of algorithms for which VCG works

Maximal-in-Range

The VCG mechanism allows for truth given exact solutions

- For VCG to work, simply maximize $W(a)$ over all allocations A
- VCG works if we maximize $W(r)$ over any $R \subseteq A$
- These are exactly the types of algorithms for which VCG works

Example

Grouping all items into one lot, we can maximize over a range of size n. This yields a $1/n$ approximation.

MIR Example

- By giving all items to one player, we do well when welfare is concentrated

MIR Example

- By giving all items to one player, we do well when welfare is concentrated
- To do well when welfare is spread out, we can treat the auction as unit demand and solve exactly

MIR Example

- By giving all items to one player, we do well when welfare is concentrated
- To do well when welfare is spread out, we can treat the auction as unit demand and solve exactly
- One of these gets at least a $\min(n, 2\sqrt{m})$ approximation

Proof of Approximation Ratio

We know that it gets at least n, so let's see that we get $2\sqrt{m}$

- In an optimal allocation, bidders get $\leq \sqrt{m}$ or $> \sqrt{m}$ items

Proof of Approximation Ratio

We know that it gets at least n, so let's see that we get $2\sqrt{m}$

- In an optimal allocation, bidders get $\leq \sqrt{m}$ or $> \sqrt{m}$ items
- If most of the welfare goes to those with $\leq \sqrt{m}$ items, the unit allocation can get a $\sqrt{m}$ approximation on each of them

Proof of Approximation Ratio

We know that it gets at least n, so let's see that we get $2\sqrt{m}$

- In an optimal allocation, bidders get $\leq \sqrt{m}$ or $> \sqrt{m}$ items
- If most of the welfare goes to those with $\leq \sqrt{m}$ items, the unit allocation can get a $\sqrt{m}$ approximation on each of them
- If those with $> \sqrt{m}$ items get more welfare, giving all items to one bidder yields a $\sqrt{m}$ approximation for the group

Outline

1. Introduction
   - Auctions
   - VCG Mechanisms
   - The Model
   - Background
2. Allocate All
   - Allocate All Items
   - Large Range
   - VC Dimension
3. Our Work
   - The Issues
   - Large Range
   - VC Dimension
   - Subset Sum
4. Conclusions

The Model

- Each bidder has a valuation function v_i
- For each item j, bidder i has a value v_{ij}
- Each bidder i has a budget b_i
- For each subset $S \subseteq [m]$ of the items,

$$v_i(S) = \min\left(\sum_{j \in S} v_{ij}, b_i\right)$$

Example: Video Game Auction

Value: 40

Example: Video Game Auction

Value: 40

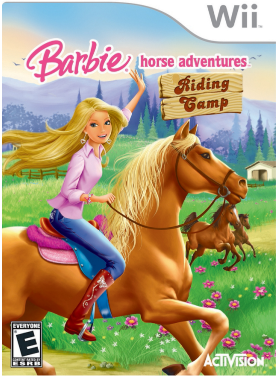

Value: 60

Example: Video Game Auction

Value: 40

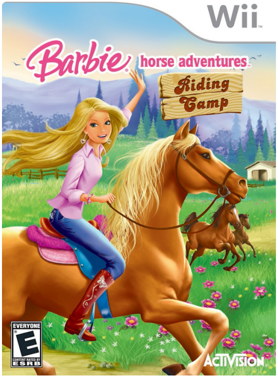

Value: 60

Value: 80

Outline

1. Introduction
   - Auctions
   - VCG Mechanisms
   - The Model
   - Background
2. Allocate All
   - Allocate All Items
   - Large Range
   - VC Dimension
3. Our Work
   - The Issues
   - Large Range
   - VC Dimension
   - Subset Sum
4. Conclusions

Previous Work

- Inapproximability for Combinatorial Public Projects (Schapira, Singer, 2008)
- n-bidder auctions can't approximate better than $(n+1)/2n$ (Mossel et al., 2009)

Previous Work

- Inapproximability for Combinatorial Public Projects (Schapira, Singer, 2008)
- n-bidder auctions can't approximate better than $(n+1)/2n$ (Mossel et al., 2009)
- We show that n-bidder auctions can't beat $\min(n, m^{1/2-\epsilon})$

Previous Work

- Inapproximability for Combinatorial Public Projects (Schapira, Singer, 2008)
- n-bidder auctions can't approximate better than $(n+1)/2n$ (Mossel et al., 2009)
- We show that n-bidder auctions can't beat $\min(n, m^{1/2-\epsilon})$
- The key to all of of these was VC dimension

VC Dimension

- Consider a subset $R \subseteq 2^{[m]}$
- By restricting to $S \subset [m]$, we get a new set R_S

Example

If $\{2, 3, 5\} \in R$ and $S = \{1, 2, 5\}$, then $\{2, 5\} \in R_S$.

- The VC dimension is size of the largest S such that $R_S = 2^S$
- For 2-bidder auctions, this is like allocating S in every way

Outline

1. Introduction
   - Auctions
   - VCG Mechanisms
   - The Model
   - Background
2. Allocate All
   - Allocate All Items
   - Large Range
   - VC Dimension
3. Our Work
   - The Issues
   - Large Range
   - VC Dimension
   - Subset Sum
4. Conclusions

Allocate All Items

Our work is based on a related proof for an easier case

- In auctions, items can be given to bidders or retained
- The social welfare is never harmed by giving out more items
- Doing so might result in not being maximal-in-range

Allocating All vs. Maximal-in-Range

Consider a 2 bidder, 2 item auction

Algorithm

- Let $\mathcal{M}$ maximize value with item 1, retain item 2
- Create $\mathcal{M}'$ by then giving item 2 to bidder 1

Allocating All vs. Maximal-in-Range

Consider a 2 bidder, 2 item auction

Algorithm

- Let $\mathcal{M}$ maximize value with item 1, retain item 2
- Create $\mathcal{M}'$ by then giving item 2 to bidder 1

Valuations

- Bidder 1 has value 2 for either item and budget 2
- Bidder 2 has value 1 for either item and budget 1

Allocating All vs. Maximal-in-Range

Consider a 2 bidder, 2 item auction

Algorithm

- Let $\mathcal{M}$ maximize value with item 1, retain item 2
- Create $\mathcal{M}'$ by then giving item 2 to bidder 1

Valuations

- Bidder 1 has value 2 for either item and budget 2
- Bidder 2 has value 1 for either item and budget 1

- $\mathcal{M}$ gives item 1 to bidder 1
- so $\mathcal{M}'$ gives both items to bidder 1

Allocating All vs. Maximal-in-Range

Consider a 2 bidder, 2 item auction

Algorithm

- Let $\mathcal{M}$ maximize value with item 1, retain item 2
- Create $\mathcal{M}'$ by then giving item 2 to bidder 1

Valuations

- Bidder 1 has value 2 for either item and budget 2
- Bidder 2 has value 1 for either item and budget 1

- $\mathcal{M}$ gives item 1 to bidder 1
- so $\mathcal{M}'$ gives both items to bidder 1
- but $\mathcal{M}'$ has a range that includes giving each bidder one item

Outline

1. Introduction
   - Auctions
   - VCG Mechanisms
   - The Model
   - Background
2. Allocate All
   - Allocate All Items
   - Large Range
   - VC Dimension
3. Our Work
   - The Issues
   - Large Range
   - VC Dimension
   - Subset Sum
4. Conclusions

Allocation Vectors

We start by looking at 2-bidder MIR allocate all mechanisms

- Associate a vector in $[2]^m$ with each allocation
- 1221 means bidder 1 gets 1 and 4, bidder 2 gets 2 and 3

Allocation Vectors

We start by looking at 2-bidder MIR allocate all mechanisms

- Associate a vector in $[2]^m$ with each allocation
- 1221 means bidder 1 gets 1 and 4, bidder 2 gets 2 and 3
- Associate a valuation function with each vector in $[2]^m$
- 1221 means bidder 1 values 1 and 4, bidder 2 values 2 and 3

Allocation Vectors

We start by looking at 2-bidder MIR allocate all mechanisms

- Associate a vector in $[2]^m$ with each allocation
- 1221 means bidder 1 gets 1 and 4, bidder 2 gets 2 and 3
- Associate a valuation function with each vector in $[2]^m$
- 1221 means bidder 1 values 1 and 4, bidder 2 values 2 and 3
- All values are 1 or 0, budgets are infinite

Allocation Vectors

We start by looking at 2-bidder MIR allocate all mechanisms

- Associate a vector in $[2]^m$ with each allocation
- 1221 means bidder 1 gets 1 and 4, bidder 2 gets 2 and 3
- Associate a valuation function with each vector in $[2]^m$
- 1221 means bidder 1 values 1 and 4, bidder 2 values 2 and 3
- All values are 1 or 0, budgets are infinite
- Social welfare is just how well the vectors match

Large Range

- Fix an allocation r in the range
- Pick a random value vector v

Large Range

- Fix an allocation r in the range
- Pick a random value vector v
- In expectation, r will achieve social welfare $m/2$

Large Range

- Fix an allocation r in the range
- Pick a random value vector v
- In expectation, r will achieve social welfare $m/2$
- By Chernoff bounds, $m(1/2 + \epsilon)$ is exponentially unlikely

Large Range

- Fix an allocation r in the range
- Pick a random value vector v
- In expectation, r will achieve social welfare $m/2$
- By Chernoff bounds, $m(1/2+\epsilon)$ is exponentially unlikely
- So it takes an exponentially large range to do well on all v

Outline

1. Introduction
   - Auctions
   - VCG Mechanisms
   - The Model
   - Background
2. Allocate All
   - Allocate All Items
   - Large Range
   - VC Dimension
3. Our Work
   - The Issues
   - Large Range
   - VC Dimension
   - Subset Sum
4. Conclusions

VC Dimension

- Since $|R| = 2^{\alpha m}$, R has VC dimension δm (Sauer's lemma)

VC Dimension

- Since $|R| = 2^{\alpha m}$, R has VC dimension δm (Sauer's lemma)
- So there is a subset of δm items on which we can solve exactly

VC Dimension

- Since $|R| = 2^{\alpha m}$, R has VC dimension δm (Sauer's lemma)
- So there is a subset of δm items on which we can solve exactly
- Using this subset as advice, we can solve welfare maximization

VC Dimension

- Since $|R| = 2^{\alpha m}$, R has VC dimension δm (Sauer's lemma)
- So there is a subset of δm items on which we can solve exactly
- Using this subset as advice, we can solve welfare maximization
- So approximating to $1/2 + \epsilon$ is impossible unless $NP \subseteq P/poly$

Outline

1. Introduction
   - Auctions
   - VCG Mechanisms
   - The Model
   - Background
2. Allocate All
   - Allocate All Items
   - Large Range
   - VC Dimension
3. Our Work
   - The Issues
   - Large Range
   - VC Dimension
   - Subset Sum
4. Conclusions

So what's the problem?

- We can't assume all items are allocated

So what's the problem?

- We can't assume all items are allocated
- So we focus in on some items where it's close to true

So what's the problem?

- We can't assume all items are allocated
- So we focus in on some items where it's close to true
- VC dimension doesn't generalize well to more than 2 bidders

So what's the problem?

- We can't assume all items are allocated
- So we focus in on some items where it's close to true
- VC dimension doesn't generalize well to more than 2 bidders
- So we form a meta-bidder out of all but one of the bidders

Outline

1. Introduction
   - Auctions
   - VCG Mechanisms
   - The Model
   - Background
2. Allocate All
   - Allocate All Items
   - Large Range
   - VC Dimension
3. Our Work
   - The Issues
   - Large Range
   - VC Dimension
   - Subset Sum
4. Conclusions

Coverings

- Suppose we have an approximation ratio of $1/n + \epsilon$
- For every $v \in [n]^m$, some $r \in R$ matches $(1/n + \epsilon)m$ indices

$$
\begin{aligned}
v &= 122221112212 \\
r &= 111221012210
\end{aligned}
$$

Coverings

- Suppose we have an approximation ratio of $1/n + \epsilon$
- For every $v \in [n]^m$, some $r \in R$ matches $(1/n + \epsilon)m$ indices

$$v = 122221112212$$
$$r = 111221012210$$

Coverings

- Suppose we have an approximation ratio of $1/n + \epsilon$
- For every $v \in [n]^m$, some $r \in R$ matches $(1/n + \epsilon)m$ indices

$$v = 122221112212$$
$$r = 111221012210$$

- For each S, T_S projects R to S

Coverings

- Suppose we have an approximation ratio of $1/n + \epsilon$
- For every $v \in [n]^m$, some $r \in R$ matches $(1/n + \epsilon)m$ indices

$$v = 12211221$$
$$r = 12211221$$

- For each S, T_S projects R to S

Coverings

- Suppose we have an approximation ratio of $1/n + \epsilon$
- For every $v \in [n]^m$, some $r \in R$ matches $(1/n + \epsilon)m$ indices

$$\begin{aligned} v &= 12211221 \\ r &= 12211221 \end{aligned}$$

- For each S, T_S projects R to S
- T_S filters out $r \in R$ such that any $s \in S$ is unassigned

Coverings

- Suppose we have an approximation ratio of $1/n + \epsilon$
- For every $v \in [n]^m$, some $r \in R$ matches $(1/n + \epsilon)m$ indices

$$\begin{aligned} v &= 12211221 \\ t &= 12211221 \end{aligned}$$

- For each S, T_S projects R to S
- T_S filters out $r \in R$ such that any $s \in S$ is unassigned
- $t \in T_S$ covers v if it is the projection of v to S

Coverings Continued

- If we fix $|S|$, each $v \in [n]^m$ is covered $\binom{(1/n+\epsilon)m}{|S|}$ times

$$
\begin{aligned}
v &= 122221112212 \\
r &= 111221012210
\end{aligned}
$$

Coverings Continued

- If we fix $|S|$, each $v \in [n]^m$ is covered $\binom{(1/n+\epsilon)m}{|S|}$ times

$$
\begin{aligned}
v &= 122221112212 \\
r &= 111221012210
\end{aligned}
$$

Coverings Continued

- If we fix $|S|$, each $v \in [n]^m$ is covered $\binom{(1/n+\epsilon)m}{|S|}$ times

$$
\begin{aligned}
v &= 122221112212 \\
r &= 111221012210
\end{aligned}
$$

Coverings Continued

- If we fix $|S|$, each $v \in [n]^m$ is covered $\binom{(1/n+\epsilon)m}{|S|}$ times

$$
\begin{aligned}
v &= 1222211112212 \\
r &= 1112210112210
\end{aligned}
$$

Coverings Continued

- If we fix $|S|$, each $v \in [n]^m$ is covered $\binom{(1/n+\epsilon)m}{|S|}$ times

$$\begin{aligned} v &= 122221112212 \\ r &= 111221012210 \end{aligned}$$

- Each $t \in T_S$ covers $n^{m-|S|}$ valuations

$$v = ***2***1****$$

Coverings Continued

- If we fix $|S|$, each $v \in [n]^m$ is covered $\binom{(1/n+\epsilon)m}{|S|}$ times

$$\begin{aligned} v &= 122221112212 \\ r &= 111221012210 \end{aligned}$$

- Each $t \in T_S$ covers $n^{m-|S|}$ valuations

$$v = ***2***1****$$

- So if

$$n^{cm}\binom{m}{|S|}n^{m-|S|} < n^m\binom{(1/n+\epsilon)m}{|S|},$$

there must be a T_S of size greater than n^{cm}

Coverings Continued

- If we fix $|S|$, each $v \in [n]^m$ is covered $\binom{(1/n+\epsilon)m}{|S|}$ times

$$\begin{aligned} v &= 122221112212 \\ r &= 111221012210 \end{aligned}$$

- Each $t \in T_S$ covers $n^{m-|S|}$ valuations

$$v = ***2***1****$$

- So if

$$n^{cm}\binom{m}{|S|}n^{m-|S|} < n^m\binom{(1/n+\epsilon)m}{|S|},$$

there must be a T_S of size greater than n^{cm}

Coverings Continued

- If we fix $|S|$, each $v \in [n]^m$ is covered $\binom{(1/n+\epsilon)m}{|S|}$ times

$$\begin{aligned} v &= 122221112212 \\ r &= 111221012210 \end{aligned}$$

- Each $t \in T_S$ covers $n^{m-|S|}$ valuations

$$v = ***2***1****$$

- So if
$$n^{cm}\binom{m}{|S|}n^{m-|S|} < n^m\binom{(1/n+\epsilon)m}{|S|},$$
there must be a T_S of size greater than n^{cm}

Coverings Continued

- If we fix $|S|$, each $v \in [n]^m$ is covered $\binom{(1/n+\epsilon)m}{|S|}$ times

$$\begin{aligned} v &= 122221112212 \\ r &= 111221012210 \end{aligned}$$

- Each $t \in T_S$ covers $n^{m-|S|}$ valuations

$$v = * * * 2 * * * 1 * * * *$$

- So if
$$n^{cm} \binom{m}{|S|} n^{m-|S|} < n^m \binom{(1/n+\epsilon)m}{|S|},$$
there must be a T_S of size greater than n^{cm}

Coverings Continued

- If we fix $|S|$, each $v \in [n]^m$ is covered $\binom{(1/n+\epsilon)m}{|S|}$ times

$$\begin{aligned} v &= 122221112212 \\ r &= 111221012210 \end{aligned}$$

- Each $t \in T_S$ covers $n^{m-|S|}$ valuations

$$v = ***2***1****$$

- So if
$$n^{cm}\binom{m}{|S|}n^{m-|S|} < n^m\binom{(1/n+\epsilon)m}{|S|},$$
there must be a T_S of size greater than n^{cm}

Coverings Continued

- If we fix $|S|$, each $v \in [n]^m$ is covered $\binom{(1/n+\epsilon)m}{|S|}$ times

$$
\begin{aligned}
v &= 122221112212 \\
r &= 111221012210
\end{aligned}
$$

- Each $t \in T_S$ covers $n^{m-|S|}$ valuations

$$v = * * * 2 * * * 1 * * * *$$

- So if

$$n^{cm}\binom{m}{|S|}n^{m-|S|} < n^m\binom{(1/n+\epsilon)m}{|S|},$$

there must be a T_S of size greater than n^{cm}

Where are we now?

So we not only have a large range, but by focusing in on S, we have a large range that allocates all items.

Next, we deal with the difficulty of using the VC dimension with more than two bidders.

Outline

1. Introduction
   - Auctions
   - VCG Mechanisms
   - The Model
   - Background
2. Allocate All
   - Allocate All Items
   - Large Range
   - VC Dimension
3. Our Work
   - The Issues
   - Large Range
   - VC Dimension
   - Subset Sum
4. Conclusions

VC Dimension

- Using Sauer's lemma requires an exponential subset of $[2]^m$
- We have an exponential subset of $[n]^m$
- This is a problem, as $[2]^m \subset [n]^m$ has exponential size but VC dimension 0

VC Dimension

- Using Sauer's lemma requires an exponential subset of $[2]^m$
- We have an exponential subset of $[n]^m$
- This is a problem, as $[2]^m \subset [n]^m$ has exponential size but VC dimension 0
- Solution: Map $[n]^m \rightarrow [2]^{nm}$

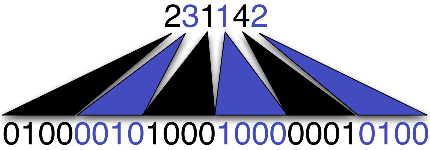

VC Dimension

- Using Sauer's lemma requires an exponential subset of $[2]^m$
- We have an exponential subset of $[n]^m$
- This is a problem, as $[2]^m \subset [n]^m$ has exponential size but VC dimension 0
- Solution: Map $[n]^m \rightarrow [2]^{nm}$

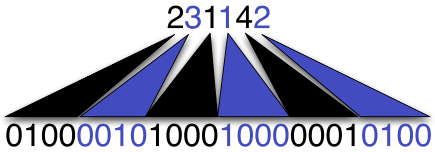

- 1 means i gets it, 0 means someone else does

VC Dimension

- Using Sauer's lemma requires an exponential subset of $[2]^m$
- We have an exponential subset of $[n]^m$
- This is a problem, as $[2]^m \subset [n]^m$ has exponential size but VC dimension 0
- Solution: Map $[n]^m \to [2]^{nm}$

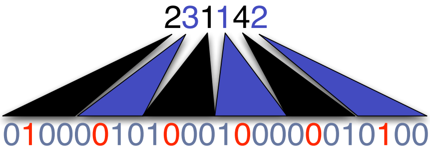

- 1 means i gets it, 0 means someone else does
- By sacrificing a factor of n, we can fix i

Now what do we know?

So we now see that the large range means that the range solves exactly over an auction with 2 bidders, one corresponding to a special bidder i and the rest forming a meta-bidder.

We do not know that this auction is hard yet, however, as the meta-bidder has a restricted class of valuations.

Outline

1. Introduction
   - Auctions
   - VCG Mechanisms
   - The Model
   - Background
2. Allocate All
   - Allocate All Items
   - Large Range
   - VC Dimension
3. Our Work
   - The Issues
   - Large Range
   - VC Dimension
   - Subset Sum
4. Conclusions

Embedding Subset Sum

- Let $a_1, \ldots, a_m$ be a subset sum instance with target τ

Embedding Subset Sum

- Let $a_1, \ldots, a_m$ be a subset sum instance with target τ
- The meta-bidder has $b = \infty$, $v_j = a_j$
- For bidder i, $b = 2\tau$, $v_j = 2a_j$

Embedding Subset Sum

- Let $a_1, \ldots, a_m$ be a subset sum instance with target τ
- The meta-bidder has $b = \infty$, $v_j = a_j$
- For bidder i, $b = 2\tau$, $v_j = 2a_j$
- A subset sums to τ iff we get welfare $\sum_j a_j + \tau$

Done

So if a maximal-in-range mechanism approximates the social welfare better than $\min(n, m^{1/2-\epsilon})$, subset sum has polynomial circuits.

Conclusions and Open Problems

- We showed that for any poly-bounded n, no poly-time MIR mechanism can beat $\min(n, m^{1/2-\epsilon})$
- This essentially solves the problem, as a $\min(n, 2\sqrt{m})$ approximation exists.
- The more general question of how well truthful mechanisms can perform is left open