The Communication Complexity of Combinatorial Auctions in Graphs

George Christodoulou[1,2] Elias Koutsoupias[3]
Annamária Kovács[4] **Ioannis Vlachos**[5,2]

[1]Aristotle University of Thessaloniki, Greece
[2]Archimedes, Athena Research Center, Greece
[3]University of Oxford, UK
[4]Goethe University, Frankfurt M., Germany
[5]Athens University of Economics and Business, Greece

STACS 2026

Combinatorial Auctions

- n agents
- m items
- Valuation function v_i of agent i
  - $v_i : 2^E \mapsto \mathbb{R}^+$
- Output allocation $a = (X_1, \ldots, X_n)$
  - $X_i \cap X_j = \emptyset$ for $i \neq j$
- **Objective**: Maximize social welfare

$$SW(a, v) = \sum_{i \in N} v_i(X_i)$$

- ρ-approximate protocol: $\frac{OPT}{ALG} \leq \rho$

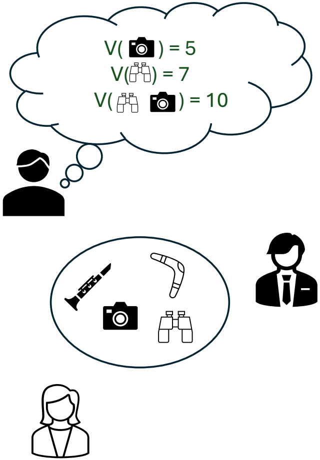

Valuation Classes

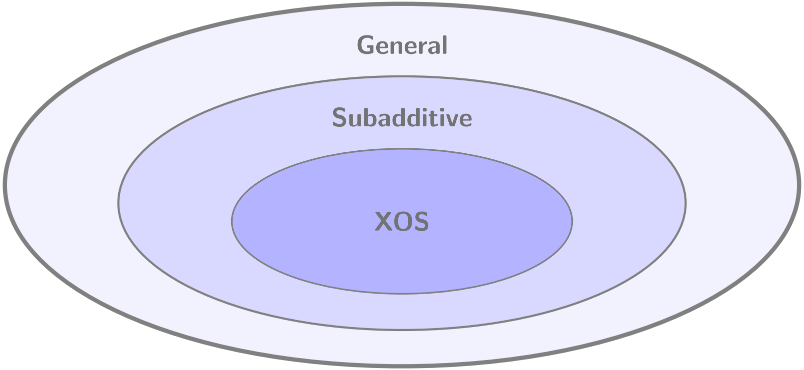

Communication in Auctions

Model:

- Bidders' values are **private**.
- Agents exchange info with the auctioneer to achieve high social welfare.
- Agents have unbounded computational power.

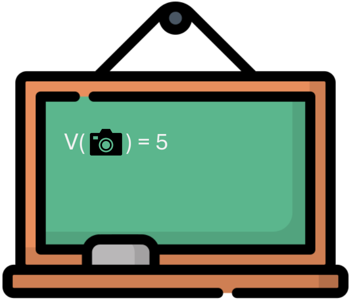

Protocols

This Problem is Well-Understood:

- Under poly(n,m)-communication, tight approximation bounds across different valuation classes.

Protocols

This Problem is Well-Understood:

- Under poly(n,m)-communication, tight approximation bounds across different valuation classes.

Non-Truthful (Tight Bounds)

- **General**: $\min\{n, \Theta(\sqrt{m})\}$

  [Nisan, Segal '06; Blumrosen, Nisan '05]

- **Subadditive**: 2

  [Dobzinski, Nisan, Schapira '05; Feige '06]

- **XOS**: ≈ 1.58

  [Dobzinski, Nisan, Schapira '05; Feige '06]

What if we require Truthfulness?

Non-Truthful (Tight Bounds)

- **General**: $\min\{n, \Theta(\sqrt{m})\}$

  [Nisan, Segal '06; Blumrosen, Nisan '05]

- **Subadditive**: 2

  [Dobzinski, Nisan, Schapira '05; Feige '06]

- **XOS**: ≈ 1.58

  [Dobzinski, Nisan, Schapira '05; Feige '06]

What if we require Truthfulness?

Incentive Compatibility:

- Bidders are strategic and seek to maximize their own utility: $u_i = v_i(X_i) - p_i$.
- A mechanism is **truthful** if reporting the true valuation v_i is always a dominant strategy.

Non-Truthful (Tight Bounds)

- **General**: $\min\{n, \Theta(\sqrt{m})\}$

  [Nisan, Segal '06; Blumrosen, Nisan '05]

- **Subadditive**: 2

  [Dobzinski, Nisan, Schapira '05; Feige '06]

- **XOS**: ≈ 1.58

  [Dobzinski, Nisan, Schapira '05; Feige '06]

What if we require Truthfulness?

Incentive Compatibility:

- Bidders are strategic and seek to maximize their own utility: $u_i = v_i(X_i) - p_i$.
- A mechanism is **truthful** if reporting the true valuation v_i is always a dominant strategy.

Truthfulness vs Non-Truthfulness:

- There is a gap of $\tilde{O}(\sqrt{m})$.

Non-Truthful (Tight Bounds)

- **General**: $\min\{n, \Theta(\sqrt{m})\}$
  [Nisan, Segal '06; Blumrosen, Nisan '05]
- **Subadditive**: 2
  [Dobzinski, Nisan, Schapira '05; Feige '06]
- **XOS**: ≈ 1.58
  [Dobzinski, Nisan, Schapira '05; Feige '06]

Truthful Deterministic (Best Known)

- **General**: $O(m/\log m)$
  [Qiu, Weinberg 2024]
- **Subadditive**: $O(\sqrt{m/\log m})$
  [Qiu, Weinberg 2024]
- **XOS**: $O(\sqrt{m/\log m})$
  [Qiu, Weinberg 2024]

Graph Setting

Can we bridge the gap between truthful and non-truthful mechanisms by exploiting a "structured" setting?

Can we achieve better approximations in restricted domains?

Graph Setting

Can we bridge the gap between truthful and non-truthful mechanisms by exploiting a "structured" setting?

Can we achieve better approximations in restricted domains?

Model agents and items using **Graphs**.

Graph Setting

- Bidders are represented by vertices.
- Each item is an edge connecting potential recipients.
- An item e can only be assigned to an agent i if i is incident to e.

Simple Graph

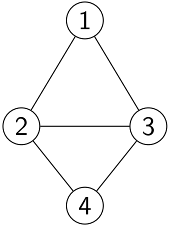

Multigraph

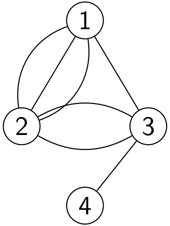

Hypergraph

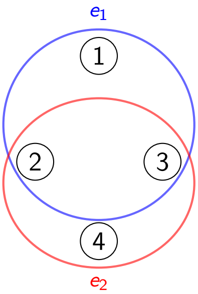

Graph Setting

- Bidders are represented by vertices.
- Each item is an edge connecting potential recipients.
- An item e can only be assigned to an agent i if i is incident to e.

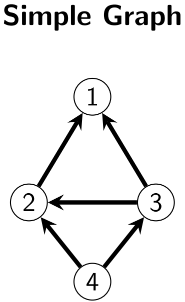

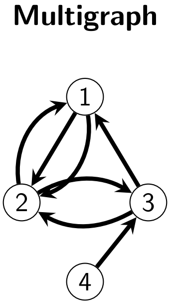

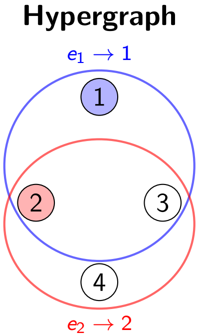

Graph Setting

- **Realistic**: Captures scenarios where items are only relevant to a fixed set of agents (e.g., neighbors in a network).
- **Literature**:
  - Crucial in settling the **Nisan-Ronen conjecture** [Christodoulou, Koutsoupias, Kovacs '23].
  - Extensively studied in **Fair Division**
    - e.g., EFX allocations [Amanatidis, Ratsikas, Sgouritsa '24; Zhou, Wei, M Li, B Li '24]
  - Used in **Graph Balancing** [Verschae, Wiese '14; Ebenlendr, Krčál, Sgall '14]

Contributions: General Valuations

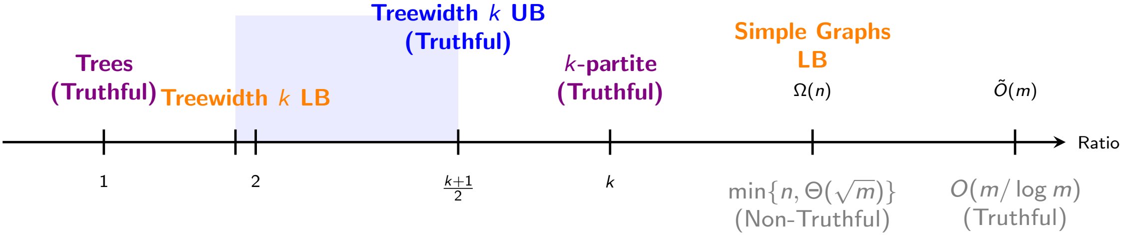

Upper Bound **Lower Bound**

Tight Bound **Prior Work**

VCG on Trees

Trees

For general valuations on trees, there exists a polynomial-communication truthful mechanism that finds the optimal solution.

VCG is implementable in poly-communication.

VCG on Trees

- Process leaves up to the root.

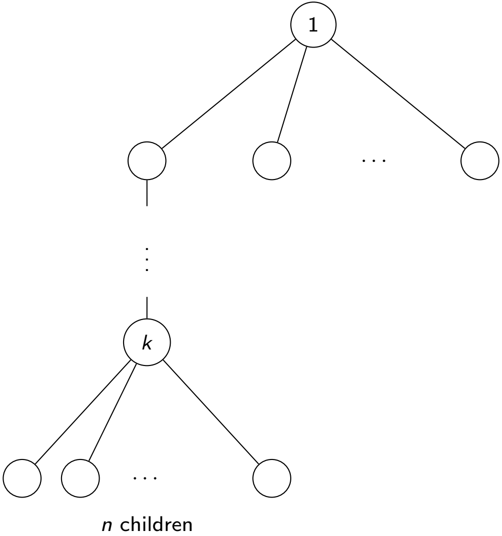

VCG on Trees

- Process leaves up to the root.
- The leaf nodes send their valuations to their parent.

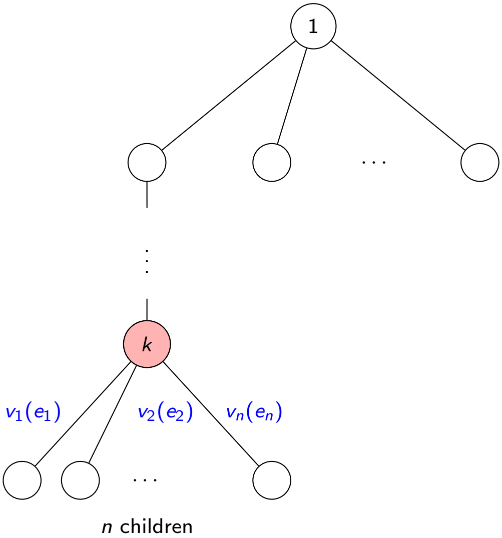

VCG on Trees

- Process leaves up to the root.
- The leaf nodes send their valuations to their parent.
- Vertex k calculates the optimal allocation for its subtree. This may take **exponential computation**.

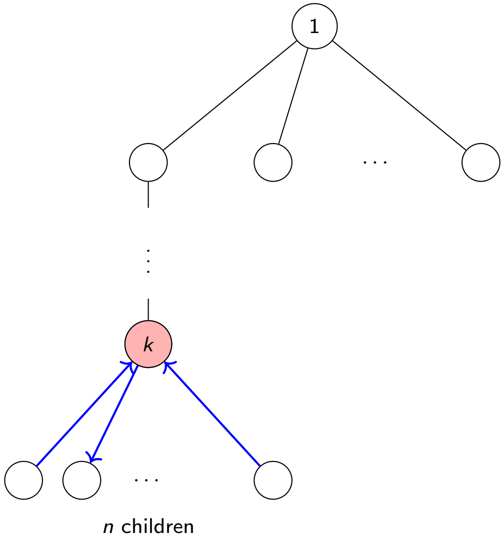

VCG on Trees

- Process leaves up to the root.
- The leaf nodes send their valuations to their parent.
- Vertex k calculates the optimal allocation for its subtree. This may take **exponential computation**.
- Repeat upwards.

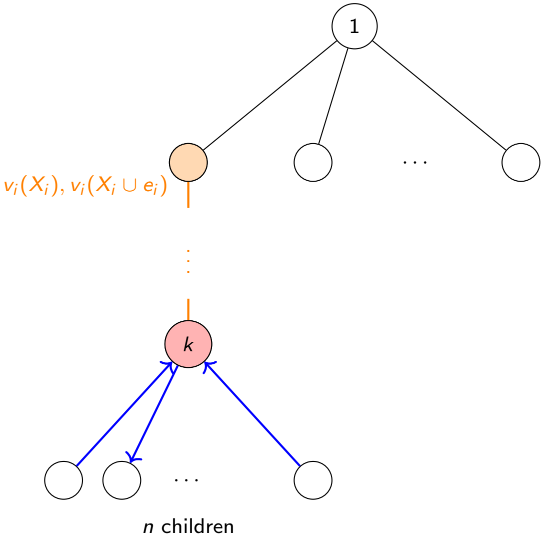

VCG on Trees

- Process leaves up to the root.
- The leaf nodes send their valuations to their parent.
- Vertex k calculates the optimal allocation for its subtree. This may take **exponential computation**.
- Repeat upwards.
- The root calculates the final optimal allocation.

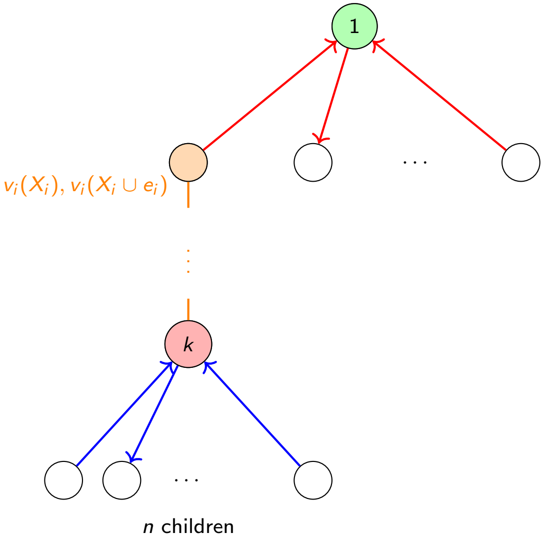

VCG on Trees

- Process leaves up to the root.
- The leaf nodes send their valuations to their parent.
- Vertex k calculates the optimal allocation for its subtree. This may take **exponential computation**.
- Repeat upwards.
- The root calculates the final optimal allocation.
- Apply payments to ensure **truthfulness**.

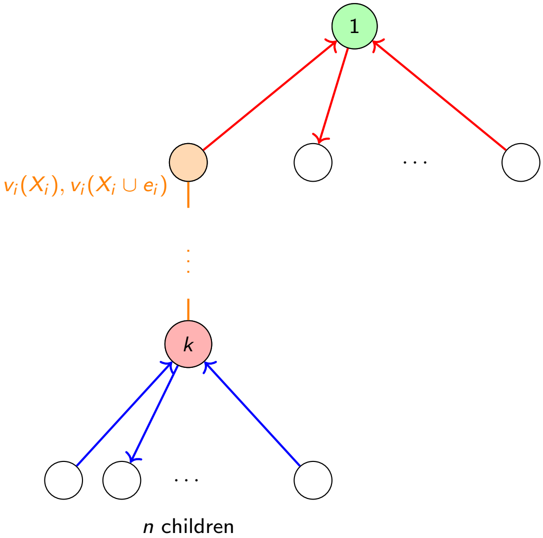

k-Degenerate Graphs

Definition

A graph is k-**degenerate** if there exists a linear ordering of its vertices such that every vertex has **at most k neighbors** appearing after it in the ordering.

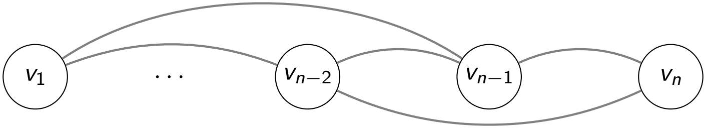

Each vertex has at most $k = 2$ edges pointing to the right $\longrightarrow$

k-Degenerate Graphs

$(k+1)$-Colorability

Because each vertex looks "forward" to at most k neighbors, a simple greedy protocol passing backwards through the ordering guarantees the graph is $(k+1)$-**colorable**.

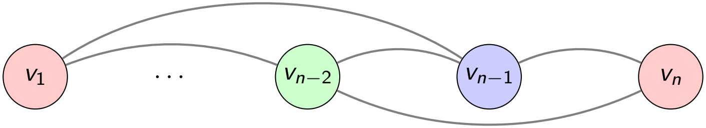

Each vertex has at most $k = 2$ edges pointing to the right $\longrightarrow$

k-Degenerate Graphs

MIR (Maximal-in-Range)

- The mechanism evaluates $k+1$ possible allocations by assigning all incident items (edges) to the agents of a color class.
- Selecting the best is truthful (as MIR) and $(k+1)$-approximate.

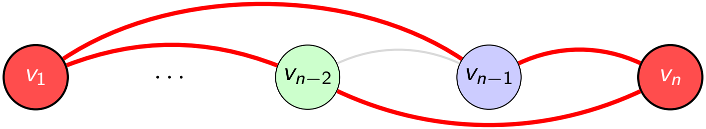

Each vertex has at most $k = 2$ edges pointing to the right $\longrightarrow$

k-Degenerate Graphs

MIR (Maximal-in-Range)

- The mechanism evaluates $k+1$ possible allocations by assigning all incident items (edges) to the agents of a color class.
- Selecting the best is truthful (as MIR) and $(k+1)$-approximate.

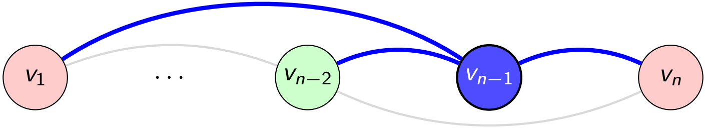

Each vertex has at most $k = 2$ edges pointing to the right $\longrightarrow$

k-Degenerate Graphs

MIR (Maximal-in-Range)

- The mechanism evaluates $k+1$ possible allocations by assigning all incident items (edges) to the agents of a color class.
- Selecting the best is truthful (as MIR) and $(k+1)$-approximate.

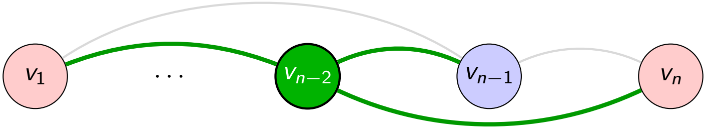

Each vertex has at most $k = 2$ edges pointing to the right $\longrightarrow$

Graphs of Treewidth k

k-Tree

Each vertex has **exactly** k forward neighbors these **form a clique**.

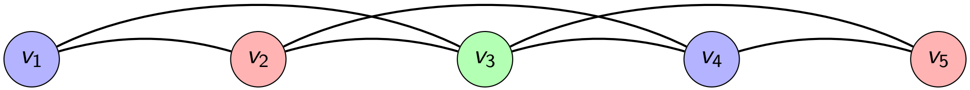

Graphs of Treewidth k

Mechanism

Because of the clique property, the subgraph induced by any **pair** of color classes is always a Forest!

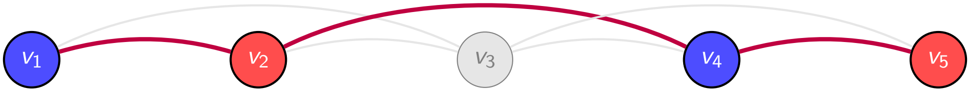

Graphs of Treewidth k

Mechanism

Because of the clique property, the subgraph induced by any **pair** of color classes is always a Forest!

1. Run the VCG tree-protocol on the forests induced by all pairs of color classes.

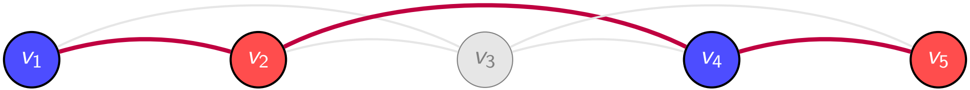

Graphs of Treewidth k

Mechanism

Because of the clique property, the subgraph induced by any **pair** of color classes is always a Forest!

1. Run the VCG tree-protocol on the forests induced by all pairs of color classes.
2. This pairwise tree-solving guarantees at least a $\frac{2}{k+1}$ fraction of OPT.

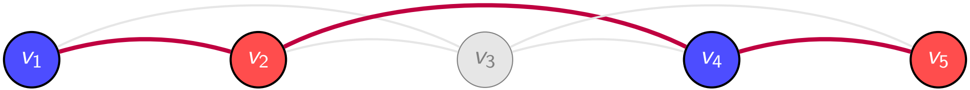

Graphs of Treewidth k

Mechanism

Because of the clique property, the subgraph induced by any **pair** of color classes is always a Forest!

1. Run the VCG tree-protocol on the forests induced by all pairs of color classes.
2. This pairwise tree-solving guarantees at least a $\frac{2}{k+1}$ fraction of OPT.
3. Run the MIR mechanism $\Rightarrow \frac{k+1}{2}$-approximation.

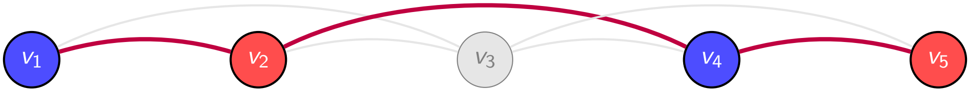

The n Lower Bound: Simple Graphs

Simple Graphs

- The problem of approximating optimal welfare is *at least as hard* as in the general case.
- Better than n-approximation requires exponential communication.
  Based on techniques from Nisan & Segal 2002.

Approximate Disjointness: n players each hold a t-bit string (representing sets B_i).

Case 1: Shared Index

$\exists s \in \cap B_i$.

	1	2	…	s^*	…	t
Agent 1	0	1	…	**1**	…	0
Agent 2	1	0	…	**1**	…	1
⋮	⋮	⋮	⋱	⋮	⋱	⋮
Agent n	0	0	…	**1**	…	0

Case 2: Disjoint Sets

All B_i are pairwise disjoint.

	1	2	…	s	…	t
Agent 1	**1**	0	…	0	…	0
Agent 2	0	**1**	…	0	…	0
⋮	⋮	⋮	⋱	⋮	⋱	⋮
Agent n	0	0	…	0	…	**1**

The n Lower Bound: Simple Graphs

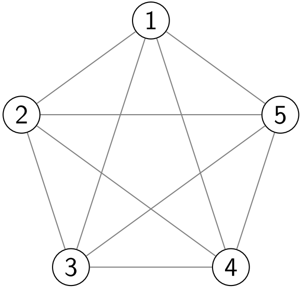

Balanced Orientation (P_s)

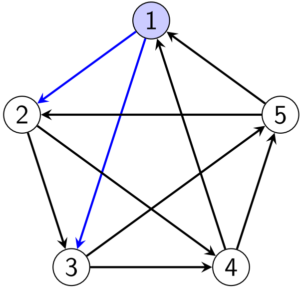

Every agent i receives exactly $(n-1)/2$ incident items.

The n Lower Bound: Simple Graphs

Overlap Property

There exists a set $\mathcal{F}$ of balanced partitions, exponential in size, where any $r = n^{\epsilon}$ distinct partitions intersect.

	P_1	P_2	$\dots$	P_k	$\dots$	P_t
v_1	P_1^1	P_2^1	$\dots$	P_k^1	$\dots$	P_t^1
v_2	P_1^2	P_2^2	$\dots$	P_k^2	$\dots$	P_t^2
v_3	P_1^3	P_2^3	$\dots$	P_k^3	$\dots$	P_t^3
$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$	$\ddots$	$\vdots$
v_n	P_1^n	P_2^n	$\dots$	P_k^n	$\dots$	P_t^n

The n Lower Bound: Simple Graphs

Approximate-Disjointness

	1	2	...	s	...	t
Agent 1	1	0	...	0	...	0
Agent 2	0	1	...	0	...	0
$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$	$\ddots$	$\vdots$
Agent n	0	0	...	0	...	1

+

Balanced Partitions

$\mathcal{F}$	P_1	P_2	...	P_{s^*}	...	P_t
	P_1^1	P_2^1	...	$P_{s^*}^1$	...	P_t^1
	P_1^2	P_2^2	...	$P_{s^*}^2$	...	P_t^2
$\vdots$	$\vdots$	$\vdots$	$\ddots$	$\vdots$	$\ddots$	$\vdots$
	P_1^n	P_2^n	...	$P_{s^*}^n$	...	P_t^n

Valuation v_i of agent i

- Agent i is "interested" in bundles corresponding to their set B_i.
- $v_i(X) = 1$ if $\exists s \in B_i$ such that $P_s^i \subseteq X$.
- $v_i(X) = 0$ otherwise.

The n Lower Bound: Simple Graphs

Let an $n^{1-\epsilon}$-approximate protocol. We will prove it requires exponential communication.

Case 1

There exists a partition giving everyone a valued bundle.

$$
\begin{array}{cccccc}
P_1 & P_2 & \dots & P_{s^*} & \dots & P_t \\
\hline
P_1^1 & P_2^1 & \dots & P_{s^*}^1 & \dots & P_t^1 \\
P_1^2 & P_2^2 & \dots & P_{s^*}^2 & \dots & P_t^2 \\
\vdots & \vdots & \ddots & \vdots & \ddots & \vdots \\
P_1^n & P_2^n & \dots & P_{s^*}^n & \dots & P_t^n
\end{array}
$$

Case 2

No single column contains more than one valued bundle.

$$
\begin{array}{cccccc}
P_1 & P_2 & \dots & P_s & \dots & P_t \\
\hline
P_1^1 & P_2^1 & \dots & P_s^1 & \dots & P_t^1 \\
P_1^2 & P_2^2 & \dots & P_s^2 & \dots & P_t^2 \\
\vdots & \vdots & \ddots & \vdots & \ddots & \vdots \\
P_1^n & P_2^n & \dots & P_s^n & \dots & P_t^n
\end{array}
$$

The n Lower Bound: Simple Graphs

Let an $n^{1-\epsilon}$-approximate protocol. We will prove it requires exponential communication.

Case 1

There exists a partition giving everyone a valued bundle.

$$\begin{array}{cccccc} P_1 & P_2 & \dots & P_{s*} & \dots & P_t \\ \hline P_1^1 & P_2^1 & \dots & P_{s*}^1 & \dots & P_t^1 \\ P_1^2 & P_2^2 & \dots & P_{s*}^2 & \dots & P_t^2 \\ \vdots & \vdots & \ddots & \vdots & \ddots & \vdots \\ P_1^n & P_2^n & \dots & P_{s*}^n & \dots & P_t^n \end{array}$$

- Optimal social welfare is n.

Case 2

No single column contains more than one valued bundle.

$$\begin{array}{cccccc} P_1 & P_2 & \dots & P_s & \dots & P_t \\ \hline P_1^1 & P_2^1 & \dots & P_s^1 & \dots & P_t^1 \\ P_1^2 & P_2^2 & \dots & P_s^2 & \dots & P_t^2 \\ \vdots & \vdots & \ddots & \vdots & \ddots & \vdots \\ P_1^n & P_2^n & \dots & P_s^n & \dots & P_t^n \end{array}$$

The n Lower Bound: Simple Graphs

Let an $n^{1-\epsilon}$-approximate protocol. We will prove it requires exponential communication.

Case 1

There exists a partition giving everyone a valued bundle.

$$\begin{array}{cccccc} P_1 & P_2 & \dots & P_{s*} & \dots & P_t \\ \hline P_1^1 & P_2^1 & \dots & P_{s*}^1 & \dots & P_t^1 \\ P_1^2 & P_2^2 & \dots & P_{s*}^2 & \dots & P_t^2 \\ \vdots & \vdots & \ddots & \vdots & \ddots & \vdots \\ P_1^n & P_2^n & \dots & P_{s*}^n & \dots & P_t^n \end{array}$$

- Optimal social welfare is n.
- Due to approximation, the protocol outputs welfare **at least** n^{ϵ}.

Case 2

No single column contains more than one valued bundle.

$$\begin{array}{cccccc} P_1 & P_2 & \dots & P_s & \dots & P_t \\ \hline P_1^1 & P_2^1 & \dots & P_s^1 & \dots & P_t^1 \\ P_1^2 & P_2^2 & \dots & P_s^2 & \dots & P_t^2 \\ \vdots & \vdots & \ddots & \vdots & \ddots & \vdots \\ P_1^n & P_2^n & \dots & P_s^n & \dots & P_t^n \end{array}$$

The n Lower Bound: Simple Graphs

Let an $n^{1-\epsilon}$-approximate protocol. We will prove it requires exponential communication.

Case 1

There exists a partition giving everyone a valued bundle.

$$\begin{array}{cccccc} P_1 & P_2 & \dots & P_{s^*} & \dots & P_t \\ \hline P_1^1 & P_2^1 & \dots & P_{s^*}^1 & \dots & P_t^1 \\ P_1^2 & P_2^2 & \dots & P_{s^*}^2 & \dots & P_t^2 \\ \vdots & \vdots & \ddots & \vdots & \ddots & \vdots \\ P_1^n & P_2^n & \dots & P_{s^*}^n & \dots & P_t^n \end{array}$$

- Optimal social welfare is n.
- Due to approximation, the protocol outputs welfare **at least** n^{ϵ}.

Case 2

No single column contains more than one valued bundle.

$$\begin{array}{cccccc} P_1 & P_2 & \dots & P_s & \dots & P_t \\ \hline P_1^1 & P_2^1 & \dots & P_s^1 & \dots & P_t^1 \\ P_1^2 & P_2^2 & \dots & P_s^2 & \dots & P_t^2 \\ \vdots & \vdots & \ddots & \vdots & \ddots & \vdots \\ P_1^n & P_2^n & \dots & P_s^n & \dots & P_t^n \end{array}$$

- Because of the overlap property, the welfare is **at most** n^{ϵ}.

The n Lower Bound: Simple Graphs

The protocol can distinguish between the two cases of the **Approximate Disjointness** problem.

The n Lower Bound: Simple Graphs

The protocol can distinguish between the two cases of the **Approximate Disjointness** problem.

$$\Downarrow$$

It requires $\Omega(t/n^4)$ communication [Alon, Matias, Szegedy '96], which is exponential.

Contributions: Subadditive Valuations

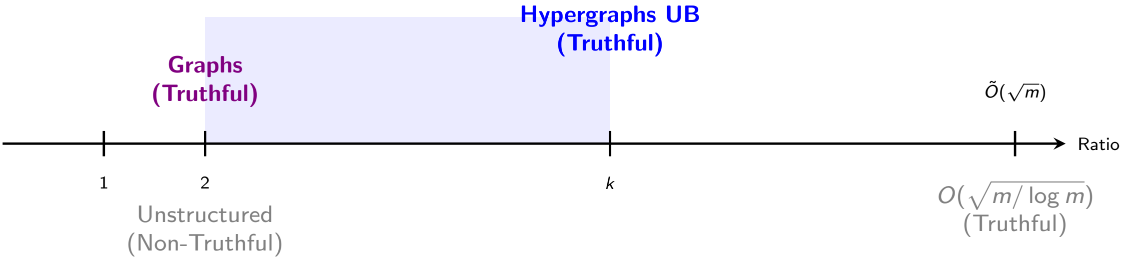

Upper Bound **Lower Bound**

Tight Bound **Prior Work**

Contributions: XOS Valuations

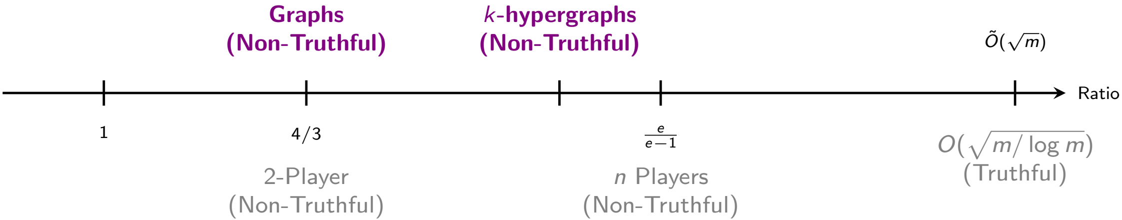

Upper Bound **Lower Bound**

Tight Bound **Prior Work**

Future Directions

- **Graph Setting**:
  - Treewidth-k graphs with general valuations have a gap between **LB** ≈ 2 and **UB** $\frac{k+1}{2}$ in general valuations.
  - k-hypergraphs with subadditive valuations have a gap between **LB** 2 and **UB** k.
  - Truthful graphs with XOS valuations have a gap between **LB** $4/3$ and **UB** 2.

- **General Setting**:
  - Truthful XOS mechanisms remain open even for 2 agents.

Thank You!