CSE 6319 Notes 3: Mechanism Design (Part 3)

(Last updated 4/2/24 12:21 PM)

3.I. VCG AND SCORING RULES (KP 16; N 9; R 7)

Social Surplus Maximization and the General VCG Mechanism (KP 16.2)

Example 16.2.4 - Roads for three cities

Example 16.2.5 - Employee housing

Example 16.2.8 - Spectrum auctions

Scoring Rules (KP 16.3 - SKIP)

3.J. COMBINATORIAL AUCTIONS (N 11; R 8)

Introduction (N 11.1)

m indivisible items, n bidders

Definition 11.1 A valuation v is a real-valued function that for each subset S of items, $v(S)$ is the value that bidder i obtains if he receives this bundle of items. A valuation must have "free disposal," i.e., be monotone: for $S \subseteq T$ we have that $v(S) \leq v(T)$, and it should be "normalized": $v(\emptyset) = 0$.

Sets S and T with $S \cap T = \varnothing$:

Complements: $v(S \cup T) > v(S) + v(T)$

Substitutes: $v(S \cup T) < v(S) + v(T)$

Definition 11.2 An *allocation* of the items among the bidders is $S_1, \ldots, S_n$ where $S_i \cap S_j = \emptyset$ for every $i \neq j$. The *social welfare* obtained by an allocation is $\sum_i v_i(S_i)$. A socially efficient allocation (among bidders with valuations $v_1, \ldots, v_n$) is an allocation with maximum social welfare among all allocations.

Issues

Computational complexity

Representation and communication

Strategic behavior

Applications: `bichler.pdf` `newman.pdf` `parkes_iBundle.pdf`

Single-Minded Case (N 11.2)

Definition 11.3 A valuation v is called *single minded* if there exists a bundle of items S^* and a value $v^* \in \Re^+$ such that $v(S) = v^*$ for all $S \supseteq S^*$, and $v(S) = 0$ for all other S. A single-minded bid is the pair (S^*, v^*).

Definition 11.4 The allocation problem among single-minded bidders is the following:
INPUT: (S_i^*, v_i^*) for each bidder $i = 1, \ldots, n$.
OUTPUT: A subset of winning bids $W \subseteq \{1, \ldots, n\}$ such that for every $i \neq j \in W$, $S_i^* \cap S_j^* = \emptyset$ (i.e., the winners are compatible with each other) with maximum social welfare $\sum_{i \in W} v_i^*$.

Intractability

Proposition 11.5 *The allocation problem among single-minded bidders is NP-hard. More precisely, the decision problem of whether the optimal allocation has social welfare of at least k (where k is an additional part of the input) is NP-complete.*

(Proof is by reduction from Independent-Set)

Proposition 11.6 *Approximating the optimal allocation among single-minded bidders to within a factor better than $m^{1/2-\epsilon}$ is NP-hard.*

Incentive-Compatible Approximation

Definition 11.7 Let V_{sm} denote the set of all single-minded bids on m items, and let A be the set of all allocations of the m items between n players. A mechanism for single-minded bidders is composed of an allocation mechanism $f : (V_{sm})^n \to A$ and payment functions $p_i : (V_{sm})^n \to \Re$ for $i = 1, \ldots, n$. The mechanism is computationally efficient if f and all p_i can be computed in polynomial time. The mechanism is incentive compatible (in dominant strategies) if for every i, and every $v_1, \ldots, v_n, v_i' \in V_{sm}$, we have that $v_i(a) - p_i(v_i, v_{-i}) \geq v_i(a') - p_i(v_i', v_{-i})$, where $a = f(v_i, v_{-i})$, $a' = f(v_i', v_{-i})$ and $v_i(a) = v_i$ if i wins in a and zero otherwise.

Issue with VCG - loses incentive compatibility

Theorem 11.8 *The greedy mechanism is efficiently computable, incentive compatible, and produces a $\sqrt{m}$ approximation of the optimal social welfare.*

The Greedy Mechanism for Single-Minded Bidders:

Initialization:

- Reorder the bids such that $v_1^*/\sqrt{|S_1^*|} \geq v_2^*/\sqrt{|S_2^*|} \geq \ldots \geq v_n^*/\sqrt{|S_n^*|}$.
- $W \leftarrow \emptyset$.

For i = 1...n do: if $S_i^* \cap \left(\bigcup_{j \in W} S_j^*\right) = \emptyset$ then $W \leftarrow W \cup \{i\}$.

Output:

Allocation: The set of winners is W.

Payments: For each $i \in W$, $p_i = v_j^*/\sqrt{|S_j^*|/|S_i^*|}$, where j is the smallest index such that $S_i^* \cap S_j^* \neq \emptyset$, and for all $k < j, k \neq i$, $S_k^* \cap S_j^* = \emptyset$ (if no such j exists then $p_i = 0$).

Figure 11.1. The mechanism achieves a $\sqrt{m}$ approximation for combinatorial auctions with single-minded bidders.

(For Payments, S_j is the first bundle after S_i with an element in common with S_i. Thus, S_j is the first bundle "disqualified" from W by S_i.)

Lemma 11.9 *A mechanism for single-minded bidders in which losers pay 0 is incentive compatible if and only if it satisfies the following two conditions:*

(i) **Monotonicity:** *A bidder who wins with bid (S_i^*, v_i^*) keeps winning for any $v_i' > v_i^*$ and for any $S_i' \subset S_i^*$ (for any fixed settings of the other bids).*

(ii) **Critical Payment:** *A bidder who wins pays the minimum value needed for winning: the infimum of all values v_i' such that (S_i^*, v_i') still wins.*

Lemma 11.10 *Let OPT be an allocation (i.e., set of winners) with maximum value of $\sum_{i \in OPT} v_i^*$, and let W be the output of the algorithm, then $\sum_{i \in OPT} v_i^* \leq \sqrt{m} \sum_{i \in W} v_i^*$.*

Walrasian Equilibrium and the LP Relaxation (N 11.3)

Winner Determination Problem = Determine the Allocation

May be stated as a (integer/fractional) linear program (N p. 276)

Dual LP Relaxation also includes prices and utilities

Definition 11.11 For a given bidder valuation v_i and given item prices $p_1, \ldots, p_m$, a bundle T is called a *demand* of bidder i if for every other bundle $S \subseteq M$ we have that $v_i(S) - \sum_{j \in S} p_j \leq v_i(T) - \sum_{j \in T} p_j$.

Definition 11.12 A set of nonnegative prices $p_1^*, \ldots, p_m^*$ and an allocation $S_1^*, \ldots, S_m^*$ of the items is a *Walrasian equilibrium* if for every player i, S_i^* is a demand of bidder i at prices $p_1^*, \ldots, p_m^*$ and for any item j that is not allocated (i.e., $j \notin \cup_{i=1}^n S_i^*$) we have $p_j^* = 0$.

Theorem 11.13 **(The First Welfare Theorem)** *Let* $p_1^*, \ldots, p_m^*$ *and* $S_1^*, \ldots, S_n^*$ *be a Walrasian equilibrium, then the allocation* $S_1^*, \ldots, S_n^*$ *maximizes social welfare. Moreover, it even maximizes social welfare over all fractional allocations, i.e., let* $\{X_{i,S}^*\}_{i,S}$ *be a feasible solution to the linear programming relaxation. Then,* $\sum_{i=1}^n v_i(S_i^*) \geq \sum_{i \in N,\ S \subseteq M} X_{i,S}^* v_i(S)$.

Theorem 11.15 **(The Second Welfare Theorem)** *If an integral optimal solution exists for LPR, then a Walrasian equilibrium whose allocation is the given solution also exists.*

Corollary 11.16 *A Walrasian equilibrium exists in a combinatorial-auction environment if and only if the corresponding linear programming relaxation admits an integral optimal solution.*

Bidding Languages (N 11.4)

Atom: (S, p) - price p for a bundle S of items

({TV, DVD player}, \$100)

OR: any subset of the atoms may be satisfied, but an item may be matched only once

({TV}, \$200) OR ({PC}, \$700)

XOR: only one of the atoms may be satisfied

({TV}, \$200) XOR ({PC}, \$700)

Maximization:

More formally, both OR and XOR bids are composed of a collection of pairs (S_i, p_i), where each S_i is a subset of the items, and p_i is the maximum price that he is willing to pay for that subset. For the valuation $v = (S_1, p_1)\, XOR, \ldots, XOR\, (S_k, p_k)$, the value of $v(S)$ is defined to be $max_{i|S_i \subseteq S}\, p_i$. For the valuation $v = (S_1, p_1)\, OR, \ldots, OR\, (S_k, p_k)$, one must be a little careful and the value of $v(S)$ is defined to be the maximum over all possible "valid collections" W, of the value of $\sum_{i \in W} p_i$, where W is a valid collection of pairs if for all $i \neq j \in W$, $S_i \cap S_j = \emptyset$.

Combinations of OR and XOR

Definition 11.18 Let v and u be valuations, then $(v\ XOR\ u)$ and $(v\ OR\ u)$ are valuations and are defined as follows:

- $(v\ XOR\ u)(S) = \max(v(S), u(S))$.
- $(v\ OR\ u)(S) = \max_{R,T \subseteq S,\ R \cap T = \emptyset}\ v(R) + u(T)$

Negative results on "compactly representing" downward sloping valuations . . .

Dummy Items

Representing XORs as ORs using dummy items:

(S_1, p_1) XOR (S_2, p_2) becomes $(S_1 \cup \{d\}, p_1)$ OR $(S_2 \cup \{d\}, p_2)$

OR* - Implicitly augments each set of items with the same dummy item

Formally, we let each bidder i have its own set of dummy items D_i, which only he can bid on. An OR* bid by bidder i is an OR bid on the augmented set of items $M \cup D_i$. The value that an OR* bid gives to a bundle $S \subseteq M$ is the value given by the OR bid to $S \cup D_i$. Thus, for example, for the set of items $M = \{a, b, c\}$, the OR* bid $(\{a, d\}, 1)\ OR\ (\{b, d\}, 1)\ OR\ (\{c\}, 1)$, where d is a dummy item, is equivalent to $((\{a\}, 1)\ XOR\ (\{b\}, 1))\ OR\ (\{c\}, 1)$.

An equivalent but more appealing "user interface" is to let bidders report a set of atomic bids together with "constraints" that signify which bids are mutually exclusive. Each constraint can then be converted into a dummy item that is added to the conflicting atomic bids. Despite its apparent simplicity, this language can simulate general OR/XOR formulae.

Theorem 11.21 *Any valuation that can be represented by OR/XOR formula of size s can be represented by OR* bids of size s, using at most s^2 dummy items.*

Iterative Auctions: The Query Model

Concept: Develop valuation over time rather than expecting complete elicitation upfront.

Value query: *The auctioneer presents a bundle S, the bidder reports his value $v(S)$ for this bundle.*

Demand query *(with item prices[2]): The auctioneer presents a vector of item prices $p_1, \ldots, p_m$; the bidder reports a demand bundle under these prices, i.e., some set S that maximizes $v(S) - \sum_{i \in S} p_i$.*

Relationship:

Lemma 11.22 *A value query may be simulated by mt demand queries, where t is the number of bits of precision in the representation of a bundle's value.*

Lemma 11.23 *An exponential number of value queries may be required for simulating a single demand query.*

Linear programming for demand queries . . .

N p. 286 (classes of CA solvers and quality of approximation) and 287 (classes of valuations)

Communication Complexity (N 11.6)

https://amturing.acm.org/award_winners/yao_1611524.cfm

Theorem 11.27 *For every $\epsilon > 0$, approximating the social welfare in a combinatorial auction to within a factor strictly smaller than $\min\{n, m^{1/2-\epsilon}\}$ requires exponential communication.*

Ascending Auctions (N 11.7)

Ascending Item-Price Auctions

Definition 11.28 A valuation v_i satisfies the *substitutes (or gross-substitutes)* property if for every pair of item-price vectors $\vec{q} \geq \vec{p}$ (coordinate-wise comparison), we have that the demand at prices q contains all items in the demand at prices p whose price remained constant. Formally, for every $A \in argmax_S\{v(S) - \sum_{j \in S} p_j\}$, there exists $D \in argmax_S\{v(S) - \sum_{j \in S} q_j\}$, such that $D \supseteq \{j \in A | p_j = q_j\}$.

(Goods may be substitutes or independent, but not complements.)

Also implies submodularity, for every two bundles S and T,
$v(S) + v(T) \geq v(S \cup T) + v(S \cap T)$

An item-price ascending auction for substitutes valuations:

Initialization:
For every item $j \in M$, set $p_j \leftarrow 0$.
For every bidder i let $S_i \leftarrow \emptyset$.

Repeat
For each i, let D_i be the demand of i at the following prices:
p_j for $j \in S_i$ and $p_j + \epsilon$ for $j \notin S_i$.
If for all i $S_i = D_i$, exit the loop;
Find a bidder i with $S_i \neq D_i$ and update:
- For every item $j \in D_i \setminus S_i$, set $p_j \leftarrow p_j + \epsilon$
- $S_i \leftarrow D_i$
- For every bidder $k \neq i$, $S_k \leftarrow S_k \setminus D_i$

Finally: Output the allocation $S_1, ..., S_n$.

Figure 11.3. An item-price ascending auction that ends up with a nearly optimal allocation when bidders' valuations have the (gross) substitutes property.

Definition 11.29 An allocation $S_1, \ldots, S_n$ and a prices $p_1, \ldots, p_m$ are an ϵ-Walrasian equilibrium if $\bigcup_i S_i \supseteq \{j | p_j > 0\}$ and for each i, S_i is a demand of i at prices p_j for $j \in S_i$ and $p_j + \epsilon$ for $j \notin S_i$.

Theorem 11.30 *For bidders with substitutes valuations, the auction described in Figure 11.3 ends with an ϵ-Walrasian equilibrium. In particular, the allocation achieves welfare that is within $n\epsilon$ from the optimal social welfare.*

Claim 11.31 *At every stage of the auction, for every bidder i, $S_i \subseteq D_i$.*

$m \bullet v_{max} / \in$ stages (iterations of **Repeat**)

Similar to the Uniform-Price Multi-Unit Auction for Budgeted Bidders, demand reduction to improve utility (payoff) is possible (N Example 11.32)

Ascending Bundle-Price Auction

$p_i(S)$ - personalized bundle price on bundle S for bidder i

Demand for bidder i are the bundles that maximize $v_i(S)$ - $p_i(S)$

A bundle price auction:

Initialization: For every player i and bundle S, let $p_i(S) \leftarrow 0$.

Repeat

- Find an allocation $T_1, ..., T_n$ that maximizes revenue at current prices, i.e., $\sum_{i=1}^{n} p_i(T_i) \geq \sum_{i=1}^{n} p_i(Y_i)$ for any other allocation $Y_1, ..., Y_n$. (Bundles with zero prices will not be allocated, i.e., $p_i(T_i) > 0$ for eve ry i.)
- Let L be the set of losing bidders, i.e., $L = \{i | T_i = \emptyset\}$.
- For every $i \in L$ let D_i be a demand bundle of i under the prices $\vec{p_i}$.
- If for all $i \in L$, $D_i = \emptyset$ then terminate.
- For all $i \in L$ with $D_i \neq \emptyset$, let $p_i(D_i) \leftarrow p_i(D_i) + \epsilon$.

Figure 11.4. A bundle price auction which terminates with the socially efficient allocation for any profile of bidders.

Definition 11.33 Personalized bundle prices $\vec{p} = \{p_i(S)\}$ and an allocation $S = (S_1, \ldots, S_n)$ are called a *competitive equilibrium* if:

- For every bidder i, S_i is a demand bundle, i.e., for any other bundle $T_i \subseteq M$, $v_i(S_i) - p_i(S_i) \geq v_i(T_i) - p_i(T_i)$.
- The allocation S maximizes *seller's revenue* under the current prices, i.e., for any other allocation $(T_1, \ldots, T_n)$, $\sum_{i=1}^{n} p_i(S_i) \geq \sum_{i=1}^{n} p_i(T_i)$.

Definition 11.35 A bundle S is an ϵ-*demand* for a player i under the bundle prices $\vec{p_i}$ if for any other bundle T, $v_i(S) - p_i(S) \geq v_i(T) - p_i(T) - \epsilon$. An ϵ-*competitive equilibrium* is similar to a competitive equilibrium (Definition 11.33), except each bidder receives an ϵ-demand under the equilibrium prices.

Theorem 11.36 *For any profile of valuations, the bundle-price auction described in Figure 11.4 terminates with an ϵ-competitive equilibrium. In particular, the welfare obtained is within $n\epsilon$ from the optimal social welfare.*

Finding each allocation is NP-hard.

(2016 FCC) Spectrum Auctions

R 8 and appendix to first chapter of Milgrom book `https://www.amazon.com/dp/023117599X`

`https://www.fcc.gov/auctions`

Goal: Reallocate 500 MHz from TV to wireless internet (and reduce US national debt)

Concepts:

Forward auction to allocate bandwidth (upload, download, interference)

Vendors need channels per "partial economic area"

Figure 1. Interference graph visualizing the FCC's constraint data[9] (2 990 stations; 2 575 466 channel-specific interference constraints).

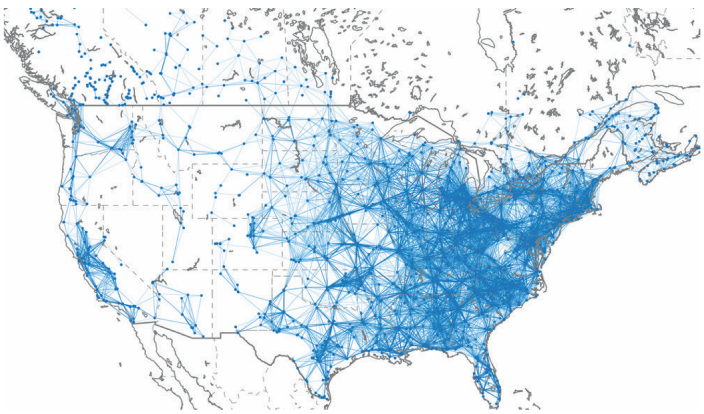

(`newman.pdf`)

Reverse auction to acquire bandwidth

UHF stations get money and possibly VHF channel assignment

VHF stations get money and go out of business

Value index = (population served • degree of interference)$^{0.5}$

Opening bid total $120B with goal of decreasing to $86B

Use of SAT solver to check feasibility of "repacking"
(`potassco.org` Knuth: `www.amazon.com/dp/0134397606`)

Forward Auction Features:

Multiple Round Simultaneous Clock Auctions

Rule: If the price isn't increasing, can't decrease demand

Rule: May not increase overall activity from round to round

Rule: Up-front cash deposit to cover activity

Mandatory bid increments to avoid "signaling"

Bidders must avoid "exposure problem"

3.K. MATCHING MARKETS

Maximum Weighted Matching (Assignment Problem) (KP 17.1)

Notes 1, p. 5 (KP 3.2) introduced:

Maximum matching (bipartite)

Minimum vertex cover

Hall's marriage theorem

Konig's Lemma: | maximum matching | = | minimum vertex cover |

Hide and Seek game

Matching market problem:

Input: valuations for n buyers on n items (one seller)

Find price vector p^* and maximum matching M to maximize the social surplus

$$\sum_i (v_{iM(i)} - p^*_{M(i)}) + \sum_j p^*_j = \sum_i v_{iM(i)}$$

Map this need to generalized König's Lemma:

THEOREM 17.1.1. *Given a nonnegative matrix* $V = (v_{ij})_{n\times n}$, *let*

$$K := \{(\mathbf{u},\mathbf{p}) \in \mathbb{R}^n \times \mathbb{R}^n : \quad u_i, p_j \geq 0 \ \text{and} \ u_i + p_j \geq v_{ij} \ \forall i, j\}.$$

Then

$$\min_{(\mathbf{u},\mathbf{p})\in K} \Big\{\sum_i u_i + \sum_j p_j\Big\} = \max_{\text{matchings } M} \Big\{\sum_i v_{i,M(i)}\Big\}.$$

(**u**, **p**) is a minimum (fractional) cover. M is a maximum weight matching.

Observation: u_i and p_j cannot exceed the largest value in V.

Classic Algorithms for Matching:

Minimization instead of maximization . . .

Integers instead of floating-point . . .

Trivial: O(n^4)

Start with trivial matching

Iteratively find negative cycles to improve (Floyd-Warshall)

Hungarian method: O(n^3)

Papadimitriou & Steiglitz `https://www.amazon.com/dp/0486402584/`

Knuth `https://www-cs-faculty.stanford.edu/~knuth/sgb.html`

Envy-Free Prices (KP 17.2)

Preferred item(s) - Based on price vector **p** and buyer *i*, item *j* such that

$$\forall k \quad v_{ij} - p_j \geq v_{ik} - p_k \text{ and } v_{ij} \geq p_j$$

Demand graph *D*(**p**)

Bipartite graph connecting buyers to their preferred items

p is *envy-free* if *D*(**p**) is a perfect matching

LEMMA 17.2.2. *Let* $V = (v_{ij})_{n \times n}$, $\mathbf{u}, \mathbf{p} \in \mathbb{R}^n$, *all nonnegative, and let* M *be a perfect matching from* $[n]$ *to* $[n]$. *The following are equivalent:*

(i) $(\mathbf{u}, \mathbf{p})$ *is a minimum cover of* V *and* M *is a maximum weight matching for* V.

(ii) *The prices* $\mathbf{p}$ *are envy-free prices,* M *is contained in the demand graph* $D(\mathbf{p})$, *and* $u_i = v_{iM(i)} - p_{M(i)}$.

COROLLARY 17.2.3. *Let* $\mathbf{p}$ *be an envy-free pricing for* V *and let* M *be a perfect matching of buyers to items. Then* M *is a maximum weight matching for* V *if and only if it is contained in* $D(\mathbf{p})$.

LEMMA 17.2.4. *The envy-free price vectors for* $V = (v_{ij})_{n \times n}$ *form a* ***lattice:*** *Let* $\mathbf{p}$ *and* $\mathbf{q}$ *be two vectors of envy-free prices. Then, defining*

$$a \wedge b := \min(a, b) \quad \textit{and} \quad a \vee b := \max(a, b),$$

the two price vectors

$$\mathbf{p} \wedge \mathbf{q} = (p_1 \wedge q_1, \ldots, p_n \wedge q_n) \quad \textit{and} \quad \mathbf{p} \vee \mathbf{q} = (p_1 \vee q_1, \ldots, p_n \vee q_n)$$

are also envy-free.

COROLLARY 17.2.5. *Let* $\mathbf{p}$ *minimize* $\sum_j p_j$ *among all envy-free price vectors for* V. *Then:*

(i) *Every envy-free price vector* $\mathbf{q}$ *satisfies* $p_i \leq q_i$ *for all* i.

(ii) $\min_j p_j = 0$.

THEOREM 17.2.6. *Given an $n \times n$ nonnegative valuation matrix V, let M^V be a maximum weight matching and let $\|M^V\|$ be its weight; that is, $\|M^V\| = \sum_i v_{iM^V(i)}$. Write V_{-i} for the matrix obtained by replacing row i of V by $\mathbf{0}$. Then the lowest envy-free price vector $\mathbf{p}$ for V and the corresponding utility vector $\mathbf{u}$ are given by*

$$M^V(i) = j \implies p_j = \|M^{V_{-i}}\| - (\|M^V\| - v_{ij}), \quad (17.5)$$

$$u_i = \|M^V\| - \|M^{V_{-i}}\| \quad \forall i. \quad (17.6)$$

COROLLARY 17.2.9 gives symmetric details for the highest envy-free price vector

Introducing seller value s_j (i.e. reserve price) (LP 17.2.2)

Replace each v_{ij} by max(v_{ij} - s_j, 0)

Envy-Free Division of Rent (KP 17.3; `cake.sun.pdf` `cake.su.pdf`)

Previous use of Sperner's lemma for cake division may be adapted to rent problem (indivisible rooms). P. 940 of `cake.su.pdf` alludes to this.

`https://www.nytimes.com/interactive/2014/science/rent-division-calculator.html`

Assuming that at least one matching has weights whose sum is no less than the sum for the lowest envy-free rent vector (THEOREM 17.2.6) and no more than the sum for the highest envy-free rent vector (COROLLARY 17.2.9), envy-free rent division may be achieved. (KP p. 305).

`https://ranger.uta.edu/~weems/NOTES6319/AUCTION/fairRent.c`

Maximum Matching by Ascending Auctions (KP 17.4)

V is a non-negative integer matrix

Much like "item-price ascending auction for substitutes valuations"

- Fix the minimum bid increment $\delta = 1/(n+1)$.
- Initialize the prices $\mathbf{p}$ of all items to 0 and set the matching M of bidders to items to be empty.
- As long as M is not perfect:
  - one unmatched bidder i selects an item j in his demand set

    $$D_i(\mathbf{p}) := \{j \mid v_{ij} - p_j \geq v_{ik} - p_k \quad \forall k \quad \text{and } v_{ij} \geq p_j\}$$

    and bids $p_j + \delta$ on it.
    (We will see that the demand set $D_i(\mathbf{p})$ is nonempty.)
  - If j is unmatched, then $M(i) := j$; otherwise, say $M(\ell) = j$, remove (ℓ, j) from the matching and add (i, j), so that $M(i) := j$.
  - Increase p_j by δ.

THEOREM 17.4.1. *Suppose that the elements of the valuation matrix $V = (v_{ij})$ are integers. Then the above auction terminates with a maximum weight matching M, and the final prices* $\mathbf{p}$ *satisfy*

$$M(i) = j \quad \Longrightarrow \quad v_{ij} - p_j \geq v_{ik} - p_k - \delta \quad \forall k. \tag{17.11}$$

Matching Buyers and Sellers (Assignment Games) (KP 17.5)

n buyers, *n* sellers, v_{ij} is the value *i* assigns to house *j* (value to owner is 0)

j selling to *i* at price p_j gives utility of $u_i = v_{ij} - p_j$

DEFINITION 17.5.1. An **outcome** $(M, \mathbf{u}, \mathbf{p})$ of the assignment game is a matching M between buyers and sellers and a partition (u_i, p_j) of the value v_{ij} on every matched edge; i.e., $u_i + p_j = v_{ij}$, where $u_i, p_j \geq 0$ for all i, j. If buyer i is unmatched, we set $u_i = 0$. Similarly, $p_j = 0$ if seller j is unmatched.

We say the outcome is **stable**[2] if $u_i + p_j \geq v_{ij}$ for all i, j.

PROPOSITION 17.5.2. *An outcome $(M, \mathbf{u}, \mathbf{p})$ is stable if and only if M is a maximum weight matching for V and $(\mathbf{u}, \mathbf{p})$ is a minimum cover for V. In particular, every maximum weight matching supports a stable outcome.*

DEFINITION 17.5.3. Let $(M, \mathbf{u}, \mathbf{p})$ be an outcome of the assignment game. Define the **excess** β_i of buyer i to be the difference between his utility and his **best outside option**[3]; i.e., (denoting $x_+ := \max(x, 0)$),

$$\beta_i := u_i - \max_k \{(v_{ik} - p_k)_+ \ : \ (i, k) \notin M\}.$$

Similarly, the **excess** s_j of seller j is

$$s_j := p_j - \max_\ell \{(v_{\ell,j} - u_\ell)_+ \ : \ (\ell, j) \notin M\}.$$

The outcome is **balanced** if it is stable and, for every matched edge (i, j), we have $\beta_i = s_j$.

THEOREM 17.5.5. *Every assignment game has a balanced outcome. Moreover, the following process converges to a balanced outcome: Start with the minimum cover $(\mathbf{u}, \mathbf{p})$ where $\mathbf{p}$ is the vector of lowest envy-free prices and a maximum weight matching M. Repeatedly pick an edge in M to balance, ensuring that every edge in M is picked infinitely often.*

LEMMA 17.5.6. *Let $(M, \mathbf{u}, \mathbf{p})$ be a stable outcome with $\beta_i \geq s_j \geq 0$ for every $(i, j) \in M$. Pick a pair $(i, j) \in M$ with $\beta_i > s_j$ and balance the pair by performing the update*

$$u_i' := u_i - \frac{\beta_i - s_j}{2} \quad \text{and} \quad p_j' := p_j + \frac{\beta_i - s_j}{2},$$

leaving all other utilities and profits unchanged. Then the new outcome is stable and has excesses $\beta_i' = s_j'$ and $\beta_k' \geq s_\ell' \geq 0$ for all $(k, \ell) \in M$.

(LP 17.5.1 Positive seller values)

Application to Weighted Hide-and-Seek Games (KP 17.6)

Instead of 0/1 weights (section 3.2), general payoffs h_{ij} are used.

Theorem 17.1.1 is applied to obtain minimax result for zero-sum game.

Example 17.6.2