CS 684 Algorithmic Game Theory	Scribe: Christina Tavoularis
Instructor: Eva Tardos	Monday April 26th, 2004

1 Gibbard-Satterthwaite Theorem

We will begin by recalling an important result that was introduced in the last lecture. For an election with at least 3 candidates and where each voter has an ordered preference for the candidates, consider two claims:

- *Claim 1: Unanimity Preserving* If $x > y$ and candidate x is preferred by all voters then y is not selected.
- *Claim 2: Independent Irrelevant Alternatives* If x is selected on some preferences, and the preference of alternative z is changed, then that cannot cause $y \neq x, z$ to be selected, although z may be selected as a result.

Arrow's Theorem states that when players are truthful[1], claims 1 and 2 are only satisfied in a dictatorial system. Recall that in a dictatorial system, the outcome is based on one players highest preference. Now we can present a stronger result, courtesy of Gibbard and Satterthwaite. Assume we have a set X of candidates and the goal is to select exactly 1. Voters have utilities $u_i(x) \geq 0$ for each candidate $x \in X$ such that the utility function encodes users' preferences: $u_i(x) > u_i(y) \sim x > y$. Furthermore, we will assume that side payments (bribes) are not allowed.

When there are only 2 candidates, it is clear that there is no harm to admitting one's preferences. In this case, the truthful election is simply majority voting. A voter against the majority can only change her strategy by voting with the majority which has no benefit. The following theorem tells us that we cannot make the same inference for more than 3 candidates.

Theorem 1 (Gibbard-Satterthwaite) *No truthful election system can be achieved when there are 3 or more candidates and where candidates can be elected by user preferences.*

Example Assume we have three candidates x, y, z. A truthful system can be achieved if it is known in advance that candidate z cannot get elected. In this system, those whose first preference is z can use their second preference.

Before formally proving the theorem, we will comment on the monotonicity of any individual utility function. Assume candidate x is selected. Monotonicity implies that for some user i if $u_i(x)$ increases, or $u_i(y)$ decreases for $y \neq x$, then the outcome does not change.

Claim Truthful $\Rightarrow$ monotone.

Proof. Suppose $u_i(x)$ increasing to $u_i'(x)$ causes the outcome to change from x to y. If $u_i(y) > u_i(x)$ then player i is happy and should be claiming $u_i'(x)$ even when $u_i(x)$ is the truth. If $u_i(y) < u_i(x)$ then he should stay with his original claim $u_i(x)$ even when $u_i'(x)$ is true. Thus, being truthful implies monotonicity. The same argument can be used for $u_i(y)$ decreasing. ■

[1] Truthfulness requires that every voter (i.e. player) cannot express untruthful preferences.

Proof. (of Gibbard-Satterthwaite) Assume candidate x has been selected. First we wil prove the independence of irrelevant alternatives. For some i, assume i decreases its utility for candidate z. By monotonicity, x will remain the outcome. But what happens if i increases its preference for z from $u(z)$ to $u'(z)$? Suppose that the winning candidate is z when player i uses preference u'. Let $u''_i(z) = min(u_i(z), u'_i(z))$. Under the u preference scheme x is elected, and by monotonicity $\Rightarrow$ x is still elected in u''. However, under the u' system y is elected, and by monotonicity y is elected in u'', contradiction. Thus, both system must elect the same candidate.

Now we will show that this scheme is unanimity preserving. Suppose not. Then, $u_i(x) > u_i(y)$, $\forall i \Rightarrow$ y is elected. For simplicity, let $u_i(z) = 0\ \forall i$ and $\forall\ z \neq x, y$ (by monotonicity we can show that we are solving the same problem). This simplifies the game to a two person system. There exists a preference $\hat{u}_i$ that elects x (again we will assume that $\hat{u}_i(z) = 0\ \forall i$ and $\forall\ z \neq x, y$). Consider switching players one by one from using u to $\hat{u}$, forming an ordered list. After some finite number of switches, the outcome changes. In the first case, it changes to $z \neq x, y$. This outcome is undesirable for user i because $\hat{u}_i(z) = 0$ and thus user i would never admit such a change. Otherwise, the new outcome is x, in which case the subset of players are better off claiming $\hat{u}$ because x is in fact preferred by all players. ■

2 Roberts Theorem and Related Results

Arrow's and the Gibbard-Satterthwaite theorem are election results. We can think of systems such as combinatorial auctions and network routing as elections. For any algorithm, the candidates are the possible outcomes.[2] The Gibbard-Satterthwaite theorem holds for no side payments and arbitrary valuations of outcomes. We would also like to consider algorithms with no externalities such that players only care about their personal gain and there is no correlation between players preferences. A more general result presented by Roberts allows for arbitrary preferences, and side payments where preferences are linear in the payment amount. As usual, we will think of the problem as an election. Assume that x is elected and user i has preference $u_i(x) + p_i$ where p_i is a payment.

Recall the VCG auction which maximizes the social happiness $\max\limits_x \sum\limits_i u_i(x)$ according to a payment system where player i pays $p_i = \sum_{j \neq i} u_j(x)$. The weighted VCG auction selects x according to $\max\limits_x \sum\limits_i w_i u_i(x) + h(x)$ where $h(x)$ is a penalty. The weighted VCG auction can be implemented truthfully with side payments. User i should receive payment $p_i = 1/w_i[\sum_{j \neq i} w_j u_j(x) + h(x)] - h_i(\text{other users})$. We require however that $h_i(\cdot)$ is such that $p_i \geq 0$.

Theorem 2 (Roberts) *The only truthful outcome with payments is a weighted VCG.*

Proof. The proof of Roberts theorem requires both Arrow's and Gibbard-Satterthwaite theorems. The proof is omitted.

Lavi, Mu'alem and Nisan present an analog of Roberts theorem for combinatorial auction with no externalities in preferences. They provide three axioms:

1. Truthful auction $\Rightarrow$ monotone

[2] Note that in this case there could be exponentially many candidates, however this theorem is not concerned with the complexity of the problem.

2. monotone + independence of irrelevant alternatives $\Rightarrow$ weighted VCG.

3. If preference $u \to a$ and $u' \to b$ then there exists a user i such that $u_i(a) - u_i(b) > u'_i(a) - u'_i(b)$.

Axiom 3 is a stronger result that encodes the linearity of preferences, whereas Roberts only codes preferences according to payoff. However, the general implications are much like those of Gibbard-Satterthwaite.

References

[1] K. Roberts: The characterization of implementable choice rules. In Aggregation and Revelation of Preferences, Laffont, J.-J. (Ed.), North-Holland.

[2] Ron Lavi, Ahuva Mu'alem, Noam Nisan: Towards a Characterization of Truthful Combinatorial Auctions. FOCS 2003: 574-583

[3] Ahuva Mu'alem, Noam Nisan: Truthful Approximation Mechanisms for Restricted Combinatorial Auctions. AAAI/IAAI 2002: 379-384