Computational Complexity Theory, Fall 2008 14 December 2009

Lecture 13: Equilibrium refinements

Lecturer: Peter Bro Miltersen Scribe: Andreas Sand

1 Practical information

Exams will take place on January 7th and January 20th in Turing-030, and there will be QA sessions in Turing-030 at 1pm on January 6th and January 19th. Students who have not yet sent an email to Peter (bromille@cs.au.dk) with their preferences regarding exam dates should do this as soon as possible. An exam plan will be posted on the course web page before christmas.

2 Course summary

The first part of this lecture gave an overview of the game representations, solution concepts and algorithms for finding solutions, we have seen so far in the course. In the second part of the lecture we studied certain weaknesses of the solution concepts and saw how to refine the solution concepts to remove these weaknesses and how to modify the algorithms to compute the refined solutions.

The main solution concepts studied in this course are *maximin* (or minimax) and *Nash equilibrium*. Maximin (minimax) is a prescriptive solution concept, which is best thought of as an advice we can give a player of the game to guarantee some outcome, while Nash equilibrium is a descriptive solution concept stating a stable way of playing a game, in which no player has anything to gain by deviating from the given strategy. For two-player zero-sum games we have seen that the set of Nash Equilibria is identical to the cartesian product of the sets of maximin and minimax strategies.

We have particularly studied three types of games: Perfect Information Games, Matrix Games and Extensive Form Games.

Perfect Information Games are games in which all players know all moves in the game. They are represented by *game trees* in which each node belongs to a single player or to nature, and payoffs are annotated at the leafs. Pure maximin/minimax strategies are found using backwards induction: a greedy algorithm working its way from the leafs to the root.

A Matrix Game is represented by a matrix in which each row corresponds to a pure strategy of player one and each column corresponds to a pure strategy of player two. The value in the (i, j)th entry of the matrix is the payoff to player one when player one's ith strategy is played against player 2's jth strategy. Maximin/minimax strategies for Matrix Games are found using linear programming, and *von Neuman's Minimax Theorem* states that for zero-sum games the guaranteed payoff gained by player one when playing by a maximin strategy is equal to the guaranteed payoff lost by player two when playing by a minimax strategy.

Extensive Form Games are represented by *Kuhn trees*. Just like for Perfect Information Games, each node in a Kuhn tree belongs to a single player or to nature and payoffs are annotated at the leafs. Furthermore the set of nodes for each player is divided into equivalence classes called information sets. Intuitively, the player can not distinguish between positions in the game belonging to the same information set. The set of possible actions must therefor be identical for all positions in a particular information set. We usually require *perfect recall* in Extensive Form Games: players do not forget information the once knew. To solve an Extensive Form Game we may convert it into a Matrix Game and solve this using linear programming, but in general this will give an exponential blowup in the size of the problem. Even more disturbing, the size of the representation of a mixed strategy found this way will may be exponential in the size of the Kuhn tree. By this reason strategies of Extensive Form Games are commonly described by *behavior strategies*. These are probability distributions on the actions of each information set in the tree, and the size of such a strategy is therefor at most linear in the number of edges in the tree. Kuhn proved behavior strategies to be *behaviorally equivalent* to mixed strategies for Extensive Form Games of perfect recall[3]. Given a behavior strategy for a player, the *realisation weight* of a sequence of moves in the game is the product of the probabilities assigned to the moves by the strategy. A *realisation plan* is then a map from the set of all such sequences to the realization weights of the sequences. Koller, Megiddo and von Stengel showed how to deduce the corresponding behavior strategy given such a realization plan, and they showed how to find a maximin realisation plan given a two-player, zero-sum game in extensive form, using linear programming[1].

3 Equilibrium refinements

To study weaknesses of the Nash equilibrium concept for two-person zero-sum games(equivalent to the maximin/minimax concept), we define the game *Guess-the-Ace*:

- A deck of card is shuffled.
- Either ace of spades is the top card or not.
- Player one does not know if ace of spades is the top card or not.
- Player one may chose to end the came, and no money will be exchanged, or he may chose to keep playing.
- If player one chose to keep playing, player 2 tries to guess if ace of spades is the top card. If he guesses correctly, player one pays him $1000, otherwise no money is exchanged.

The Kuhn tree for this game together with a behavior strategy for both players is depicted in figure 1. One would think that player two should play the game randomly with probability $\frac{1}{52}$ for making the guess that ace of spades is the top card and hence $\frac{51}{52}$ for making the guess that this is not the case. But in fact the illustrated strategy is a minimax strategy for player 2, because player one's single maximin strategy is to always stop the game. These strategies are also easilly seen to be Nash equilibria (as they should be): player one would definitely not risk to deviate from his strategy, which would make player two win if ace of spades is the top card. And player two has no incentive to change her strategy, since she never really gets part of the game.

But if player two plays by this strategy, he will only rarely take advantage of mistakes made by player one. We would like to avoid counterintuitive behavior like this. This motivates the equilibrium concepts covered in the following subsections.

3.1 Subgame perfection

One intuitive requirement for equilibria capturing rational behavior is that of *sequential rationality*. A strategy is sequentially rational if starting in any information set, the strategy from that point on is a best response to the strategies of the other players.

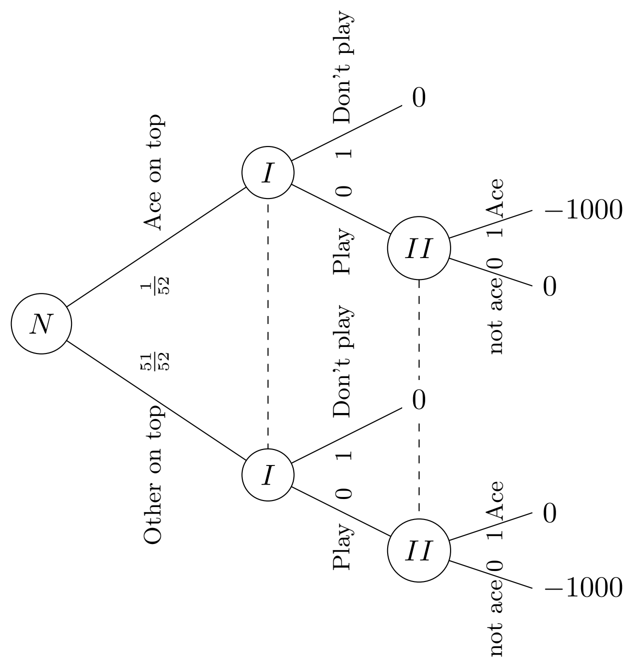

Figure 1: Nash equilibrium found by Gambit by KMvS algorith

Subgame perfection[8] was the first attempt to capture this notion of rationality. A subgame of a game in extensive form is a game starting in an information set consisting of a single node and containing all successors of this node. Furthermore, if a node in a particular information set is part of the subgame then all other nodes of the information set must be part of the subgame. A Nash equilibrium of a game is said to be *subgame perfect*, if it induces Nash equilibria in all subgames of the game. Subgame perfect Nash equilibria can be computed efficiently by a simple algorithm:

- Solve each subgame seperately and remember the strategies.
- Replace each subgame by its computed value.
- Solve the new game and combine the computed strategy with the strategies for the relevant subgames.

But having a second look at the behavior stategies in figure 1 we realise that the strategy for player two *is* subgame perfect, since the only subgames of the game are the game itself and the trivial subgames consisting of a single leaf.

3.2 Trembling hand perfection

If a player makes a mistake deviating from his strategy due to trembling, we would like the strategies of the other players to take advantage of this mistake. The concept of *trembling hand perfect equilibria* tries to capture this notion of rationality by enforcing positive probability on all actions in all information sets[9].

3.2.1 Extensive form

A *perturbed game* of a game in extensive form is a game in which we associate a tremble parameter $\epsilon > 0$ to each information set and disallow behavior probabilities smaller than this parameter, modelling a notion of tremble in each information set. A limit point of such perturbed games as the tremble parameter goes to zero is called an *extensive form trembling hand perfect equilibrium*[9].

Reviewing the computeded strategies for Guess-the-Ace, we see that player two's strategy is not trembling hand perfect as she could improve her strategy to take advantage of mistakes made by player one.

It is apparently still an open problem if extensive form trembling hand perfect equilibria for a given two-player, zero-sum extensive form game with perfect recall can be computed in polynomial time.

3.2.2 Normal form

The *normal form trembling hand perfect equilibria*[9] of an extensive form game G is defined to be the extensive form trembling hand perfect equilibria of the game G' derived from G by transforming G into strategic form (normal form) and then transforming it back into extensive form with just one information set for each player. For a two-player game, a Nash equilibrium is normal form perfect if and only if it consists of two undominated strategies.

3.2.3 Normal form vs. extensive form

In some games it is impossible to achieve an equilibrium that is both normal form trembling hand perfect and extensive form trembling hand perfect[4]. To show this, we introduce Mertens' voting game:

- Two players must elect one of them to perform an effortless task. The task may be performed either correctly or incorrectly.
- If it is performed correctly, both players receive a payoff of 1 (thus it is a non-zero-sum game), otherwise both players receive a payoff of 0.
- The election is by a secret vote:
  - If both players vote for the same player, that player gets to perform the task.
  - If each player votes for himself, the player to perform the task is chosen at random but is not told that he was elected this way.
  - If each player votes for the other, the task is performed by somebody else, with no possibility of it being performed incorrectly.

The Kuhn tree for this game is depicted in figure 2. In any extensive-form trembling hand perfect equilibrium of this game, at least one of the players believes that he is at least as likely as the other player to perform the task incorrectly and hence votes for the other player, playing by a dominated strategy. Thus any extensive form trembling hand perfect equilibrium in Mertens' voting game is not a normal form trembling hand perfect equilibrium.

This also illustrates that in the perturbed games leading to an extensive form trembling hand perfect equlibrium, the players implicitly agree on the relative magnitudes of the future trembles.

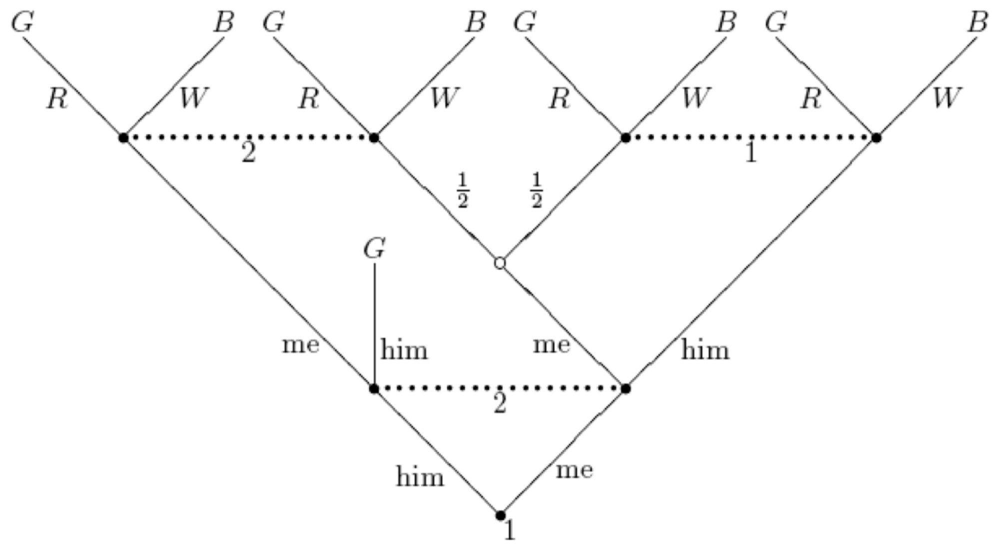

Figure 2: Kuhn tree for Mertens' voting game

Clearly such agreements are unwarranted. By this reason the concept of extensive form trembling hand perfectness was abandomed in favour of normal form trembling hand perfection. It is an open problem, if there exists a zero-sum game for which the sets of extensive form and normal form trembling hand perfect equilibria are disjoint.

To compute a normal form perfect equilibrium of a game in extensive form, we compute the value of the game using the algorithm by Koller, Megiddo and von Stengel, and then select a behavior plan achieving this value and also maximizing the payoff against some fixed fully mixed strategy of the opponent.

But studying this game:

- Player two may chose to either stop the game of give player one a dollar;
- if player one gets the dollar, he may either stop the game or give player two the dollar back;
- if player two gets the dollar back, he may either stop the game or give player one two dollars,

we realise that normal form perfect equilibria are not guaranteed to be sequentially rational, since it is part of a normal form perfect equilibrium for player one to give the dollar back if he gets it.

3.3 Sequential equilibria

Sequential equilibrium[2] is a further refinement of subgame perfect equilibrium, and it is itself refined by extensive form trembling hand perfect equilibrium. In addition to assigning a strategy to each player, a *sequential equilibrium* also assigns a *belief* to each player. A belief is a probability distribution on the nodes in each of the player's information sets. A profile of strategies and beliefs is called an *assessment* for the game. Intuitively an assessment is a sequential equilibrium if its strategies are sensible given its beliefs and its beliefs are sensible given its strategies.

We study this game:

- Player one either stops the game or asks player two for a dollar.
- Player two can either refuse or give player two a dollar.

It is part of a sequential equilibrium for player one to stop the game and not ask player two for a dollar. This strategy is dominated by the strategy in which he asks player two for the dollar, thus a sequential equilibrium may use dominated strategies. Intuitively, a sequential equilibrium reacts correctly to mistakes made in the past, but does not anticipate mistakes that may be made in the future.

3.4 Quasi-perfect equilibria

Quasi-perfect equilibrium[10] is a further refinement of both sequential equilibrium and normal form trembling hand perfect equilibrium. Intuitively a player playing by a strategy from a quasi-perfect equilibrium takes observed as well as potential future mistakes of his opponents into account but assumes that he himself will not make a mistake in the future, even if he observes that he has done so in the past.

To define quasi perfect equilibrium formally we need to define the notion of *local best responses* and *ϵ-quasi-perfect strategy profiles.*

Definition 1 *An action a in information set h is a* local best response*, if there is a plan ρ for completing play after taking a, so that best possible payoff is achieved among all strategies agreeing with ρ except possibly at h and afterwards.*

Using this definition of local best response, we define ϵ-quasi perfect strategy profiles:

Definition 2 *A strategy profile is an* ϵ-quasi-perfect strategy profile *if it satisfies that if some action is not a local best response, it is taken with probability at most ϵ.*

Finally we are ready to define the notion of quasi-perfect equilibrium:

Definition 3 *A quasi-perfect equilibrium is a limit point of ϵ-quasi- perfect strategy profiles as ϵ goes to zero.*

Miltersen and Sørensen showed how to find quasi-perfect equilibria of zero-sum games by modifying the linear programs of Koller, Megiddo and von Stengel using symbolic perturbations, ensuring that a quasi-perfect equilibrium is computed[5]. They define a *perturbed game* in the following way:

Definition 4 *Let G be a two-player game of perfect recall and let $\epsilon > 0$ be a parameter. We define the perturbed game $G(\epsilon)$ to be a game of exactly the same structure as G (i.e., same information sets, actions and payoffs) but with a restriction on the realization plans allowed: In a valid realization plan for either player in $G(\epsilon)$, the realization weight of any sequence σ of actions that can be played must be at least $\epsilon^{|\sigma|}$ where $|\sigma|$ is the number of actions in the sequence σ.*

And they show this lemma:

Lemma 5 *Let G be a two-player game of perfect recall. For any $\epsilon > 0$, any Nash equilibrium in behavior strategies of $G(\epsilon)$ is an ϵ-quasi-perfect strategy profile for G.*

Using this lemma they show how to compute quasi-perfect equilibria by computing Nash equilibria of $G(\epsilon)$ for ϵ going to zero symbolically.

This approach genaralizes to non-zero-sum games by using complementarity programs.

3.5 Normal form proper equilibria

Normal form proper equilibrium[7] further refines quasi-perfect equilibrium. Intuitively a player playing by a strategy from a normal form proper equilibrium assumes that other players may make mistakes, but he also assumes that they make the mistakes in a rational manner in which more costly mistakes are made with significantly smaller probability than less costly ones.

We define normal form proper equilibrium formally:

Definition 6 *An* ϵ-proper strategy profile *is two fully mixed strategies, so that for any two pure strategies* i, j *belonging to the same player, if* j *is a worse response than* i *to the mixed strategy of the other player, then* $Pr(j) \leq Pr(i)$.

Definition 7 *A* normal form proper equilibrium *is a limit point of* ϵ*-proper strategy profiles as* ϵ *goes to zero.*

Miltersen and Sørensen designed an algorithm based on the algorithm by Koller, Megiddo and von Stengel finding a normal form proper equilibrium for games with imperfect information[6]:

- Solve the original game using the algorithm by Koller, Megiddo and van Stengel.
- Identify the inequalities that may be satisfied with slack in some optimal solution (these are the inequalities indexed by action sequences containing mistakes).
- Find the maximin possible slack in those inequalities.
- Freeze this slack in those inequalities.

References

[1] D. Koller, N. Megiddo, and B. Von Stengel. Fast algorithms for finding randomized strategies in game trees. *Proceedings of the twenty-sixth annual ACM symposium on Theory of computing*, pages 750—-759, 1994.

[2] D.M. Kreps and R. Wilson. Sequential equilibria. *Econometrica: Journal of the Econometric Society*, 50:863—-894, 1982.

[3] HW Kuhn. Extensive games and the problem of information. *Classics in Game Theory*, pages 46–68, 1953.

[4] J.F. Mertens. Two examples of strategic equilibrium. *Games and Economic Behavior*, 8:378—-388, 1995.

[5] P. B. Miltersen and T. B. Sørensen. Computing a quasi-perfect equilibrium of a two-player game. *Economic Theory*, 42:175—-192, 2010.

[6] P.B. Miltersen and T.B. Sørensen. Fast algorithms for finding proper strategies in game trees. In *Proceedings of the nineteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 874—-883. Society for Industrial and Applied Mathematics, 2008.

[7] RB Myerson. Refinements of the Nash equilibrium concept. *International Journal of Game Theory*, 7:73—-80, 1978.

[8] R. Selten. Spieltheoretische behandlung eines oligopolmodells mit nachfrageträgheit. *Zeitschrift für die gesamte Staatswissenschaft*, 121:301—-324, 1965.

[9] R. Selten. Reexamination of the perfectness concept for equilibrium points in extensive games. In *International journal of game theory*, volume 4, pages 25—-55. Springer, 1975.

[10] E. Van Damme. *Stability and perfection of Nash equilibria*. Springer, 1991.