0368.4170: Cryptography and Game Theory Ran Canetti and Alon Rosen

Lecture 8

December 9, 2009 *Scribe: Naama Ben-Aroya*

Last Week

- 2 player zero-sum games (min-max)
- Mixed NE (existence, complexity)
- ϵ-NE
- Correlated equilibrium

Today

- Using crypto to implement correlated equilibrium
  - Extensive form game
  - Computational game
  - Computational NE
  - Actual implementation

1 Background

We want to use cryptography to solve a game-theoretic problem which arises naturally in the area of two party strategic games. The standard game-theoretic solution concept for such games is that of an equilibrium, which is a pair of self-enforcing strategies making each players strategy an optimal response to the other players strategy. It is known that for many games the expected equilibrium payoffs can be much higher when a trusted third party (a mediator) assists the players in choosing their moves (correlated equilibria), than when each player has to choose its move on its own (Nash equilibria). It is natural to ask whether there exists a mechanism that eliminates the need for the mediator yet allows the players to maintain the high payoffs offered by mediator-assisted strategies.

2 Extensive Form Game

An extensive game is an explicit description of the sequential structure of the decision problems encountered by the players in a strategic situation. The model allows us to study solutions in which each player can consider his plan of action not only at the beginning of the game but also at any point of time at which he has to make a decision. Today we will discuss extensive games with perfect information, in which each player is perfectly informed about the players' previous actions at each point in the game.

Definition 1 (Extensive Form Game) *An* extensive form game *is a tuple* $G = (N, H, P, (A_i), (U_i))$ *where:*

- N *is the set of players.*
- H *is a (possibly infinite) set of (finite) history sequence s.t.* $\epsilon \in H$. *A* $h \in H$ *is* terminal *if* $\{a|(h,a) \in H\} = \emptyset$. *(The set of terminal histories is denoted by* Z*).*
- $P : (H \setminus Z) \to N$ *is a function that assigns a "next" player to every* $h \in H \setminus Z$.

- $\forall i \in N$, A_i *is a function that assigns* $\forall h \in H \setminus Z$ *a finite set* $A_i(h)$ *of actions available to player* $i = P(h)$.
- $U_i : Z \to \mathbb{R}$ *utility for player* $i \in N$.

We interpret such a game as follows. After any nonterminal history h player $P(h) = i$ chooses an action from the set $A_i(h) = \{a|(h, a) \in H\}$. The empty history is the starting point of the game; we sometimes refer to it as the initial history. At this point player $P(\epsilon)$ chooses a member of $A(\epsilon)$. For each possible choice a^0 from this set, player $P(a^0)$ subsequently chooses a member of the set $A(a^0)$; this choice determines the next player to move, and so on. A history after which no more choices have to be made is terminal. Note that a history may be an infinite sequence of actions. As in the case of a strategic game we often specify the players' preferences over terminal histories by giving payoff functions that represent the preferences.

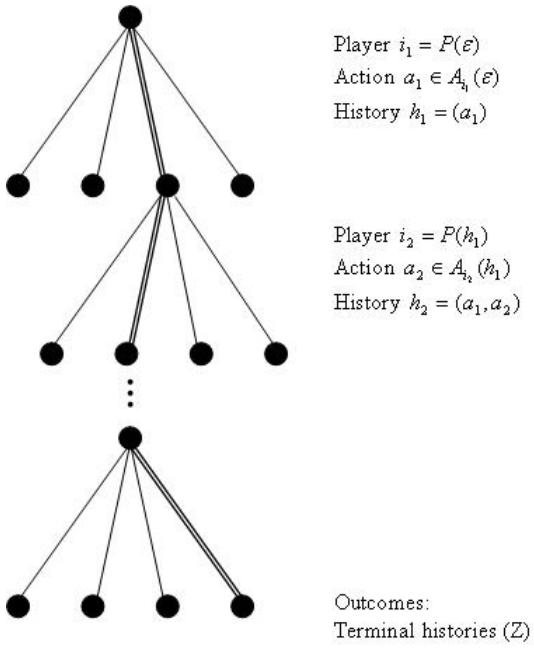

Figure 1: A schematic representation of an extensive form game

Figure 1 suggests an alternative definition of an extensive game in which the basic component is a tree. In this formulation each node corresponds to a history and any pair of nodes that are connected corresponds to an action. The leaves refer to histories after which no more choices have to be made; Indicating terminal histories.

This is a convenient representation of an extensive form game. The circle at the top of the diagram represents the initial history ϵ (the starting point of the game). $P(\epsilon) = i_1$ (player i_1 makes the first move). The line segments that emanate from the circle correspond to the members of $A(\epsilon)$ (the possible actions of player i_1 at the initial history); Each line segment leads to a circle indicating the sequel history. The highlighted segment indicates the selected action of i_1: a_1. $P(a_1) = i_2$ (Player i_2 makes the second move). The line segments that emanate from that circle correspond to the members of $A(a_1)$ (the possible actions of player i_2 at the current history). And so on, until reaching a terminal history - which is the outcome of that game.

Simultaneous Moves To model situations in which players move simultaneously after certain histories, we can modify the definition of an extensive game as follows. An extensive game with simultaneous moves is a tuple $G = (N, H, P, (U_i))$ where N, H, and U_i for each $i \in N$ are the same as in Definition 1, P is a function that assigns to each nonterminal history a set of players ($P : (H \setminus Z) \to I \subseteq N$), and H and P jointly satisfy the condition that for every nonterminal history h there is a collection $\{A_i(h)\}_{i \in P(h)}$ of sets for which $A(h) = \{a | (h, a) \in H\} = \times_{i \in P(h)} A_i(h)$.

A history in such a game is a sequence of vectors; the components of each vector a^k are the actions taken by the players whose turn it is to move after the history $(a^l)_{l=1}^{k-1}$. The set of actions among which each player $i \in P(h)$ can choose after the history h is $A_i(h)$; the interpretation is that the choices of the players in $P(h)$ are made simultaneously.

Strategies A strategy of a player in an extensive game is a plan that specifies the action chosen by the player for every history after which it is his turn to move.

Definition 2 (Strategy) *A* strategy *for player $i \in N$ is a function that takes $h \in H \setminus Z$ (s.t. $P(h) = i$) and outputs $a_i \in A_i(h)$.*

Notes:

1. A strategy completely defines the actions taken by a player $\forall h \in H \setminus Z$, even for histories that, if the strategy is followed, are never reached.
2. As in a strategic game we can define a mixed strategy to be a probability distribution over the set of (pure) strategies.
3. In the computational setting s is usually PPT.

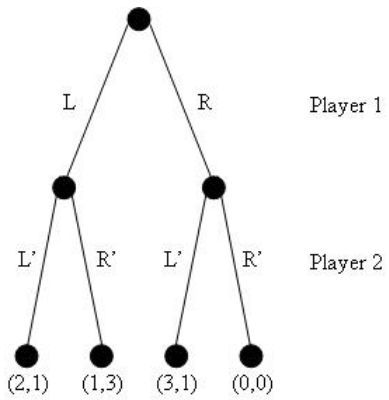

Figure 2: Example 1

Examples

1. Two players game. The formal definition of this extensive form game:
   - $N = \{1, 2\}$
   - $H = \{(\epsilon), (L), (R), (L, L'), (L, R'), (R, L'), (R, R')\}$
   - $Z = \{(L, L'), (L, R'), (R, L'), (R, R')\}$
   - $P(\epsilon) = 1$; $P(L) = P(R) = 2$

- $A_1(\epsilon) = \{L, R\}$; $A_2(R) = A_2(L) = \{L', R'\}$
- $U_1(L, L') = 2$; $U_2(L, L') = 1$; $U_1(L, R') = 1$; $U_2(L, R') = 3$; ...

A strategy s_1 for the first player will define which action to chose in the initial history (ϵ), i.e. $s_1(\epsilon) \in A_1(\epsilon) = \{L, R\}$. A strategy s_2 for the second player will define which action to chose in any optional history such that player 2 is the next player, i.e. $s_2(L) \in A_2(L)$, and $s_2(R) \in A_2(R)$. In this game there are 2 NE: (L, R'), (R, L').

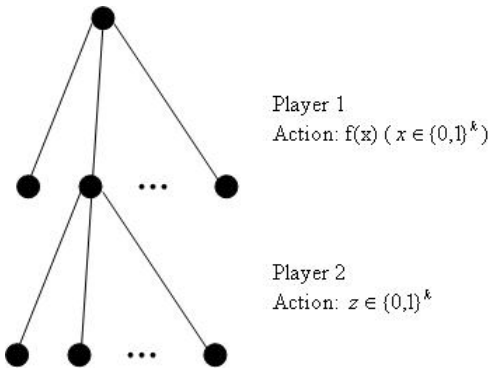

Figure 3: Example 2: OWP game (interactive version)

2. The one-way permutation game: the first player chooses an arbitrarily $x \in \{0,1\}^k$ and sends $f(x)$, where $f : \{0,1\}^k \to \{0,1\}^k$ is a permutation. The second player is challenged to come up with a z such that $z = x$. The formal definition of this extensive form game:

   - $N = \{1, 2\}$
   - $H = \{\epsilon\} \cup \{0,1\}^k \cup \{0,1\}^k \times \{0,1\}^k$
   - $Z = \{0,1\}^k \times \{0,1\}^k$
   - $P(\epsilon) = 1$; $\forall h \in \{0,1\}^k$, $P(h) = 2$
   - $A_1(\epsilon) = \{0,1\}^k$; $\forall h \in \{0,1\}^k$, $A_2(h) = \{0,1\}^k$
   - $U_2(x, z) = \begin{cases} 1 & , z = x \\ -1 & , \text{otherwise} \end{cases}$

3 Implementing Mediator (2-player game)

In the language of cryptography, we ask if we can design a two party game to eliminate the trusted third party from the original game. We consider an extended game, in which the players first exchange messages, and then choose their moves and execute them simultaneously as in the original game. The payoffs are still computed as a function of the moves, according to the same payoff function as in the original game.

It is very easy to see that every Nash equilibrium payoff of the extended game is also a correlated equilibrium payoff of the original game. We would like to show that for players which are computationally bounded to probabilistic polynomial time, any correlated equilibrium payoff of the original game can always be achieved by some Nash equilibrium of the extended game.

In the cryptographic setting, parties are assumed to be indifferent to negligible change in their utilities. We will define a computational Nash equilibrium as a pair of efficient strategies where no polynomially bounded player can gain a non-negligible advantage by not following its strategy:

Definition 3 (Computational Game) *A* computational game *is a sequence* $G^{(k)} = (N, (A_i^{(k)}), (U_i^{(k)}))$ *where:*

- N *is the set of players.*
- $\forall k \in \mathbb{N}$, $\forall i \in N$, $A_i^{(k)}$ *is the (finite) set of actions available to player* i.
- $A_i^{(k)} = A_i \cup \{\perp\}$ *where* $\perp$ *means abort.*
- $U_i^{(k)} : A_1^{(k)} \times ... \times A_n^{(k)} \to \mathbb{R}$ *utility of player* i.

Definition 4 (ϵ-Computational NE) *A strategy profile* $s = (s_1, ..., s_n)$ *is said to be in* ϵ-computational NE (ϵ-CNE) *in a computational game* $G^{(k)}$ *if* $\forall i \in N$, $\forall s_i' \in PPT$, $\forall k \in \mathbb{N}$*:*

$$U_i^{(k)}(s_i, s_{-i}) \geq U_i^{(k)}(s_i', s_{-i}) - \epsilon(k) \tag{1}$$

where: $U_i^{(k)}(s_i, s_{-i})$ *denotes the expected utility of player* i *that follows strategy* s_i.

The strategies of both players are restricted to probabilistic polynomial time. Also, since we are talking about a computational model, the definition must account for the fact that the players may break the underlying cryptographic scheme with negligible probability, thus gaining some advantage in the game. In the definition, we denote by $\epsilon : \mathbb{N} \to \mathbb{R}$ some function that is negligible in k. We notice that the new "philosophy" for both players is still to maximize their expected payoff, except that the players will not change their strategy if their gain is negligible.

We will now define the extended game: Let G be any two player game and M be a correlated equilibrium for G. Let π be a 2-party protocol for "securely" sampling M.

- Stage 1:
  - The players are given the security parameter and they freely exchanging messages (i.e. executing protocol π).
  - Histories are possible transcripts of π.
  - Actions are possible messages (all string) of π.
  - Question: What if a player aborts? (Add $\perp$ to the available actions).
- Stage 2: Play G.
  - Let (a_1, a_2) be (local) recommendation of π (implementing M).
  - Players decide whether to play according to (a_1, a_2).

The final payoff u_i' of the extended game are just the corresponding payoff of the original game applied to the players' simultaneous moves at the second stage.

The idea of getting rid of the mediator is now very simple. Consider a correlated equilibrium s of the original game. We can view the mediator as a trusted party who securely computes a probabilistic function s. Thus, to remove it we can have the two players execute a cryptographic protocol π that "securely" computes the function s. The strategy of each player would be to follow the protocol π, and then play the action a that it got from π.

Definition 5 (Secure Protocol) *Protocol* π *is said to* secure *if for every probabilistic polynomial time algorithm* π' *(representing a real model adversary strategy) there exists a probabilistic polynomial time algorithm* π'' *(representing an ideal model adversary strategy) such that*

$$\{\text{real}^{\pi_i', \pi_{-i}}(1^k)\}_{k \in \mathbb{N}} \underset{comp.}{\approx} \{\text{ideal}^{\pi_i'', \pi_{-i}}(1^k)\}_{k \in \mathbb{N}} \tag{2}$$

where:

- $\{\text{real}^{\pi_i', \pi_{-i}}(1^k)\}_{k \in \mathbb{N}}$ - *the joint output of players in real world execution.*
- $\{\text{ideal}^{\pi_i'', \pi_{-i}}(1^k)\}_{k \in \mathbb{N}}$ - *the joint output of players in ideal world execution.*

According to standard definitions of secure protocols, π is secure if the above output pair can be simulated in an ideal world. This ideal world is almost exactly the model of the trusted mediator. The security of π implies that the output distribution in the execution of the protocol in the real world is indistinguishable from that of the ideal world.

Theorem 6 *Let G be a 2-player game, M be a correlated equilibrium for G, and suppose $\pi = (\pi_1, \pi_2)$ is a "secure" protocol for computing M. Then, the following strategy is a CNE in the game G' (the extended game):*

1. *Play according to (π_1, π_2).*
2. *Let (a_1, a_2) be π's output - play (a_1, a_2).*

Moreover, the payoff of both players are exactly the same as the ones they would receive by playing according to M in G.

That is, any correlated equilibrium payoff of G can be achieved using a computational Nash equilibrium of G'. Thus, the mediator can be eliminated if the players are computationally bounded and can communicate prior to the game.

Proof:

Observation 7 *By the definition of π_i, $\forall k \in \mathbb{N}$:*

$$U_i^{(k)}(s_i, s_{-i}) = U_i^{(k)}(a_i, a_{-i} | a_i) \tag{3}$$

Where:

- $U_i^{(k)}(a_i, a_{-i} | a_i)$ - *expected utility of $i \in N$ given that it plays a_i after having received a_i, and others play a_{-i}.*
- $a_i = \text{real}^{\pi_1, \pi_2}(1^k)$
- $s =$
  1. *Play protocol π.*
  2. *Act according to output of π.*

Now let $s_i' = \pi_i'$ (s_i decides what action to take based on the transcript of π_i') be some PPT strategy, and consider π_i'' that is guaranteed by definition of security of π. (π_i'' works in ideal world)

Observation 8 *Working in ideal world is equivalent to playing G with M. By definition of correlated equilibrium:*

$$U_i^{(k)}(a_i'', a_{-i} | a_i) \leq U_i^{(k)}(a_i, a_{-i} | a_i) \tag{4}$$

Where:

- $a_i'' = \text{ideal}^{\pi_i'', \pi_{-i}}(1^k)$
- a_i = *output of M.*

Claim 9 *Suppose that $U_i^{(k)}$ can be computed in poly$^{(k)}$-time, then there exist a negligible function $\epsilon(k)$ s.t. $\forall k \in \mathbb{N}$:*

$$U_i^{(k)}(a_i', s_{-i}) - \epsilon(k) < U_i^{(k)}(a_i'', a_{-i}|a_i) \tag{5}$$

Where:

- a_i = *output of correlated device.*
- $a_i' = \text{real}^{\pi_i', \pi_{-i}}(1^k)$
- $a_i'' = \text{ideal}^{\pi_i'', \pi_{-i}}(1^k)$

Putting it all together:

$$\begin{aligned} U_i^{(k)}(s_i', s_{-i}) - \epsilon(k) &< U_i^{(k)}(a_i'', a_{-i}|a_i) && (6)\\ &\leq U_i^{(k)}(a_i, a_{-i}|a_i) && (7)\\ &= U_i(s_i, s_{-i}) && (8) \end{aligned}$$

Where equations 6, 7 and 8 follow from equations 5, 4 and 3 respectively. ■

The above description does not completely specify the strategy of the players. The models of Game Theory and Cryptography are significantly different. Two-party cryptographic protocols always assume that at least one player is honest, while the other player could be malicious. In the game-theory settings, both players are selfish and rational: they deviate from the protocol if they benefit from it. A full specification of a strategy must also indicate what a player should do if the other player deviates from its strategy (does not follow the protocol π). While cryptography does not address this question, it is crucial to resolve it in our setting, since in game theory - you always have to play: No matter what happens inside π, both players eventually have to take simultaneous actions, and receive the corresponding payoff.

Question: What if a player "aborts"? What if player 1 gets output and "aborts"? (Protocol is not "fair").

Possible solution: Implement a "punishment for deviation": punish the cheating player to his min-max level (the smallest payoff that one player can "force" the other player to have). Let s_2^* be a min-max (possibly mixed) strategy against player 1, that is, s_2^* minimizes $\underset{s_1}{max} U_1(s_1, s_2^*)$. s_2^* "punishes" player 1 by giving it the lowest possible utility, assuming player 1 plays best response to s_2^*.

We will claim that it is never in any player advantage to deviate from the protocol / abort:

Claim 10 *Let $s = (s_1, s_2)$ be a correlated equilibrium. Let v_i be the min-max payoff of player i, then for every a_1, a_2 it holds that $U_i(a_i, s_{-i}|a_i) \geq v_i$.*
(Where $U_i(a_i, s_{-i}|a_i)$ is the expected utility of player i when its recommended action is a_i)

Proof: Let $s_2 = s_2(a_1)$ be player 2's marginal distribution conditions on player 1's recommendation being a_1. Since M is a correlated equilibrium, a_1 is best response to s_2. If player 1 aborts then player 2 will play s_2^*. By definition of s_2^*, player 1's response to s_2^* cannot give better utility then it best response to s_2. So, player 1 always does worse by aborting. ■

Example:

	L	R	P
U	6,6	2,7	$-\infty, -\infty$
D	7,2	0,0	$-\infty, -\infty$
P	$-\infty, -\infty$	$-\infty, -\infty$	$-\infty, -\infty$

In this game, the best strategy (that will achieve the correlated equilibrium) is to play (U,L), (U,R) or (D,L) (each with probability $\frac{1}{3}$) and gain (expected) payoff of 5. If player 1 deviates from that strategy, the "punishment" strategy of player 2 will make sure that player 1 will get the smallest possible utility ($-\infty$), even if it will cost him his worst payoff ($-\infty$).