Game Theory and Planning (PhD seminar)

Georg Moser & Radu Prodan

Institute of Computer Science @ UIBK

Winter 2010

Topics

- Learning, regret minimisation and equilibria
- Computation of market equilibria by convex programming
- Graphical games
- Mechanism design
- Combinatorial auctions
- Routing games
- Load balancing or job allocation schemes
- Price of anarchy and the design of scalable resource allocation mechanisms
- Cascading behaviour in networks: algorithmic and economic issues
- Sponsored search auctions

discussed on October 13

Schedule

lectures	seminar talks
October 6 (GM)	January 14
October 13 (RP)	January 15
October 20 (GM)	
November 3 (RP)	

Seminar Talk

the seminar talks should present the main results obtained with respect to the language studied

Seminar Report

a short, but detailed overview of the material covered in the talk has to be handed in (maximum 5 pages); deadline February 20, 2011

Overview

(Algorithmic) game theory is conceivable as the study of mathematical models of conflict and cooperation between intelligent and rational decision-makers.
Today many applications of computer science involve autonomous decision-makers with conflicting objectives.
One application domain were computer science profi ts from game theoretic knowledge is networks in general and scheduling in particular.

Games in Extensive Form

Example

- player 1 and 2 put 1€ in a pot
- player 1 draws a card, which is either red or black
- player 1 looks at this card in private and can either raise or pass
- if player 1 passes, then she shows the card
  - if the card is red, then player 1 wins the pot
  - if the card is black, then player 2 wins the pot
- if player 1 raises, she adds another euro
- player 2 can meet or fold
  - if player 2 folds the game ends and player 1 wins the pot
  - if player 2 meets she has to add 1€
- the games continues as above

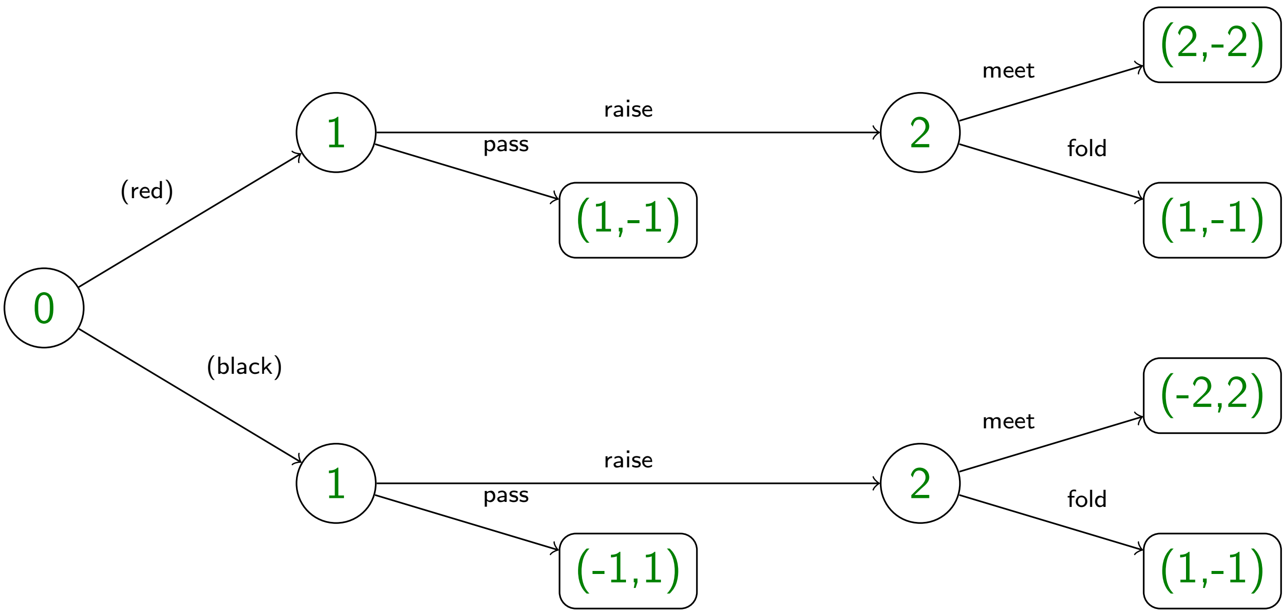

Definition

- node 0 is a chance node
- nodes 1,2 are decision nodes
- the path representing the actual events is called path of play

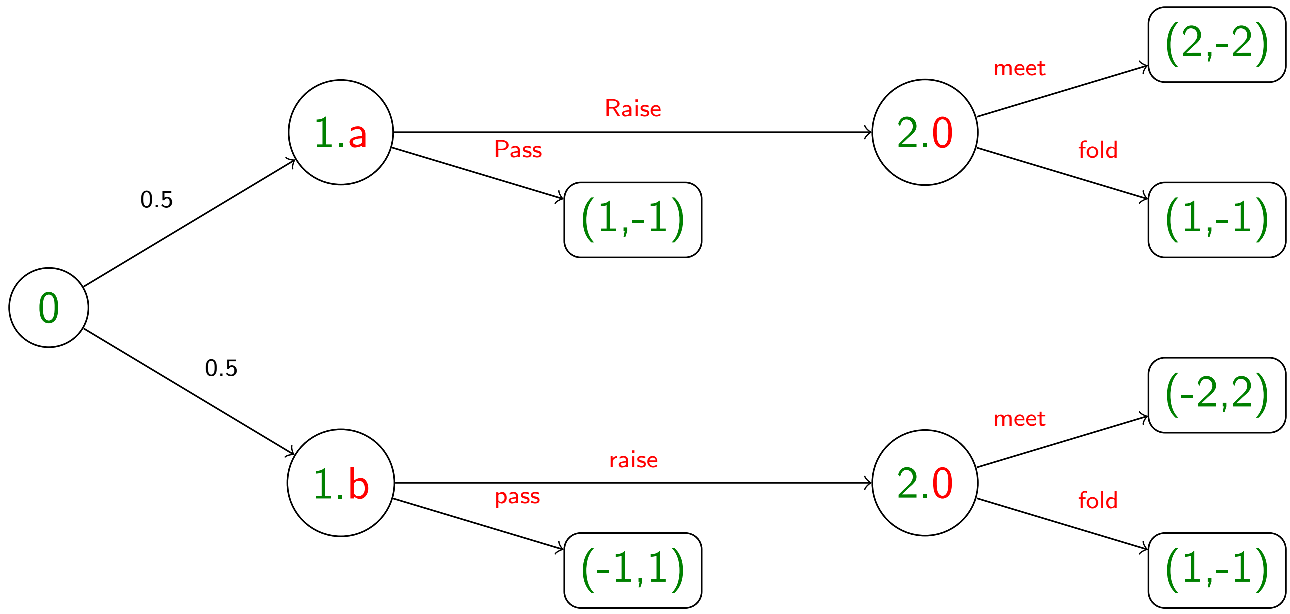

Definition

each decision nodes has two labels

1. the player label
2. the information label

Requirement

the set of move-labels following two nodes must be the same if the two nodes are controlled by the same player in the same information state

n-Person Extensive-Form Game

Definition

an n-person extensive-form game Γ^e is a labelled tree, where also edges are labelled such that

1. each nonterminal node has player label in $\{0, 1, \dots, n\}$
   nodes labelled with 0 are called chance nodes
   nodes labelled within $\{1, \dots, n\}$ are called decision nodes
2. edges leaving chance nodes (also called alternatives) are labelled with probabilities that sum up to 1
3. player nodes have a second label, the information label reflecting the information state
4. each alternative at a player node has a move label
5. each terminal node is labelled with $(u_1, \dots, u_n)$, the payoff

Definition (cont'd)

6. $\forall$ player i,
   $\forall$ nodes x y z controlled by i,
   $\forall$ alternatives b at x
   - suppose y and z have the same information state
     y follows x and b
   - $\exists$ node w, $\exists$ alternative c at w
     such that z follows w and c
   - and w is controlled by player i
     w has the same information label as x
     c has the same move label as b

Question

what does the last condition mean?

Answer

it asserts perfect recall: whenever a player moves, she remembers all the information she knew earlier

No Perfect Recall

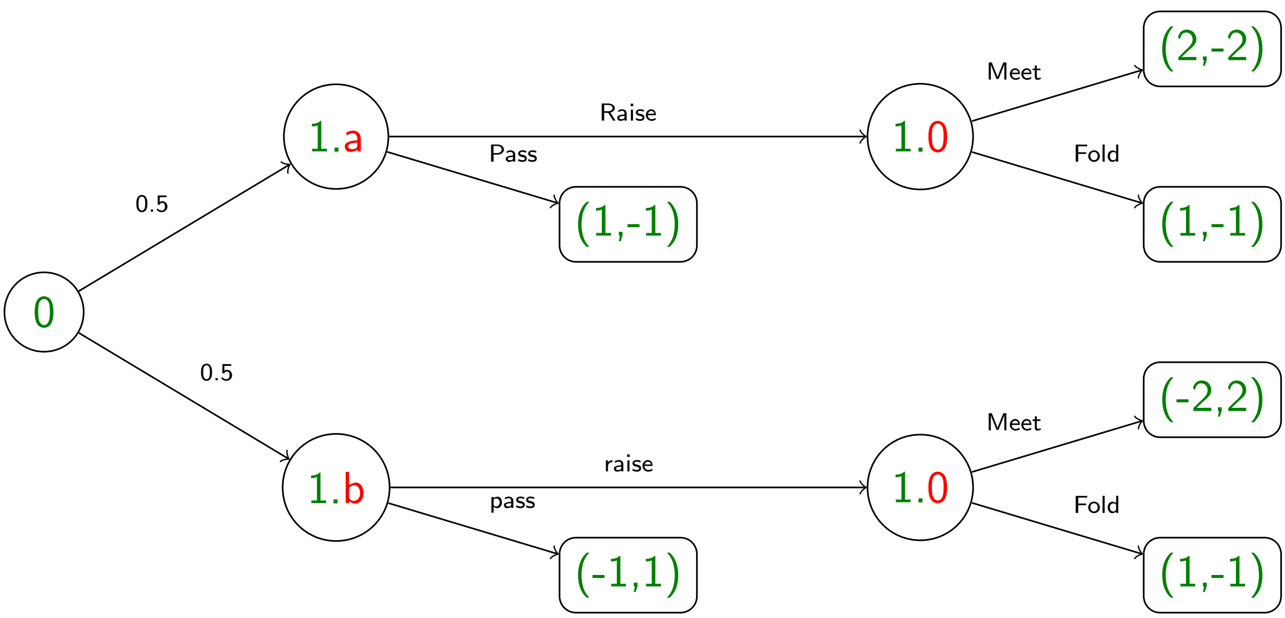

Perfect Information Games

Definition

if no two nodes have the same information state, we say the game has perfect information

Definition

- S_i is the set of information states per player i
- D_s is the set of possible moves at $s \in S_i$
- the set of strategies for player i is

$$\prod_{s \in S_i} D_s = \underbrace{D_s \times D_s \times \cdot \times D_s}_{S_i\text{-times}}$$

Example

the set of strategies for player 1: $\{Rr, Rp, Pr, Pp\}$

Influencing Your Opponent

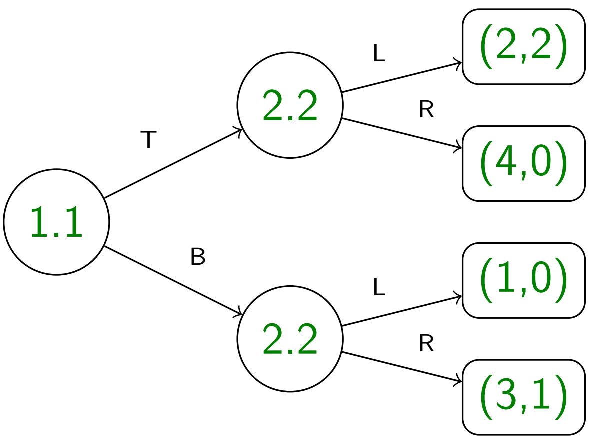

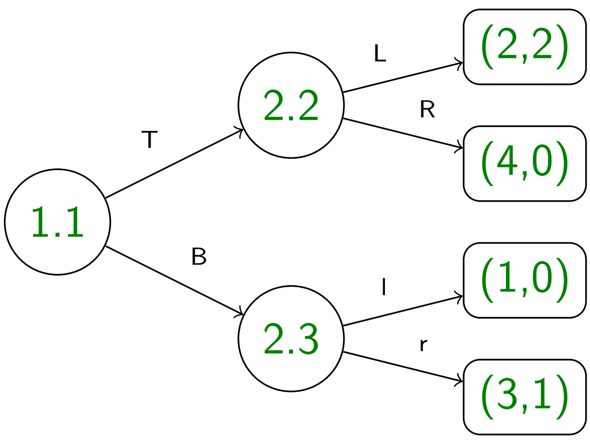

Observation
player 1 profits more, if she chooses T

player 2 doesn't know player 1's choice

Observation
player 1 profits more, if she chooses B

player 2 does know player 1's choice

Example

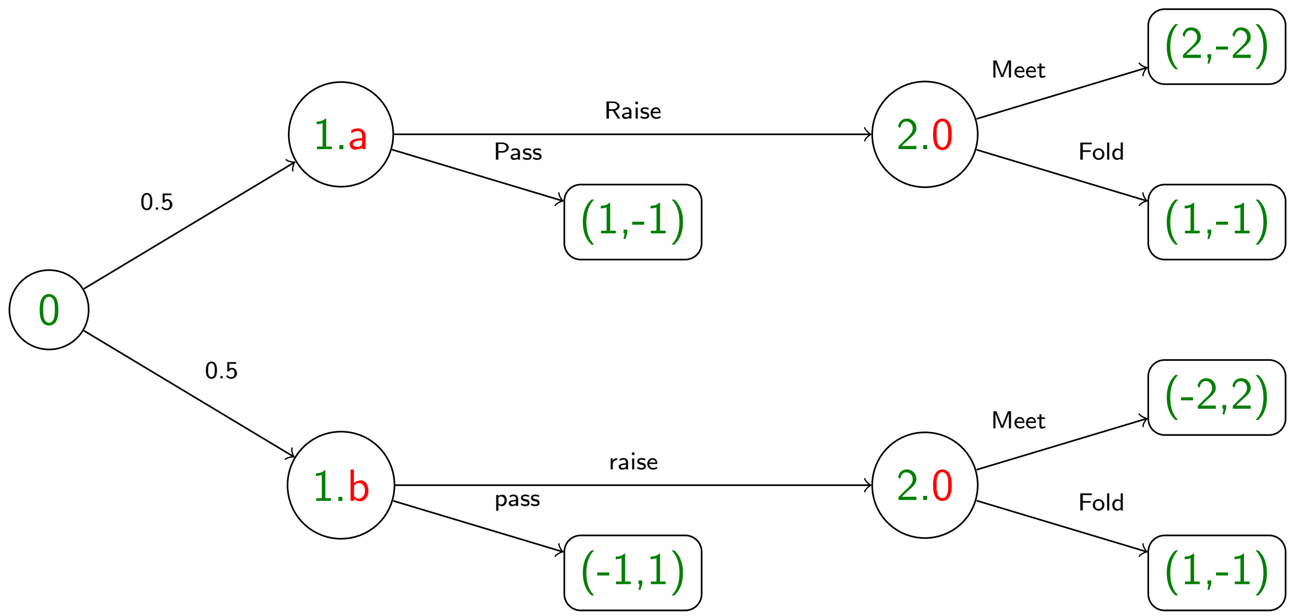

Strategies

$$\underbrace{\{Rr, Rp, Pr, Pp\}}_{\text{for player 1}} \qquad \underbrace{\{M, F\}}_{\text{for player 2}}$$

Strategic-Form Games

Definition

a strategic-form game is a tuple $\Gamma = (N, (C_i)_{i \in N}, (u_i)_{i \in N})$ such that

1. N is the set of players
2. for each i: C_i is the set of strategies of player i
3. for each i: $u_i \colon \prod_{i \in N} C_i \to \mathbb{R}$ is the expected utility payoff

a strategic-form game is finite if N and each C_i is finite

Example

consider the card game, suppose player 1 plans to use strategy Rp and player 2 plans to use M

$$u_1(Rp, M) = 2 \cdot \frac{1}{2} + -1 \cdot \frac{1}{2} = 0.5$$

$$u_2(Rp, M) = -2 \cdot \frac{1}{2} + 1 \cdot \frac{1}{2} = -0.5$$

Definition

given a game Γ^e in extensive form, we define the normal representation as strategic-form game $\Gamma = (N, (C_i)_{i \in N}, (u_i)_{i \in N})$:

1. $N = \{1, \dots, n\}$, if Γ^e is an n-person game
2. for each i: C_i denotes the strategies of each player as defined above
3. we define the expected utility payoff u_i
   - set $C = \prod_{i \in N} C_i$
   - let x be a node in Γ^e
   - let $c \in C$ denote a given strategy profile
   - let $P(x|c)$ denotes the probability that the path of play goes through x, if c is chosen
   - let Ω^* denote the set of all terminal nodes
   - for $x \in \Omega^*$, $w_i(x)$ denotes the payoff for player i
   - set
   $$u_i(c) = \sum_{x \in \Omega^*} P(x|c) w_i(x)$$

Example

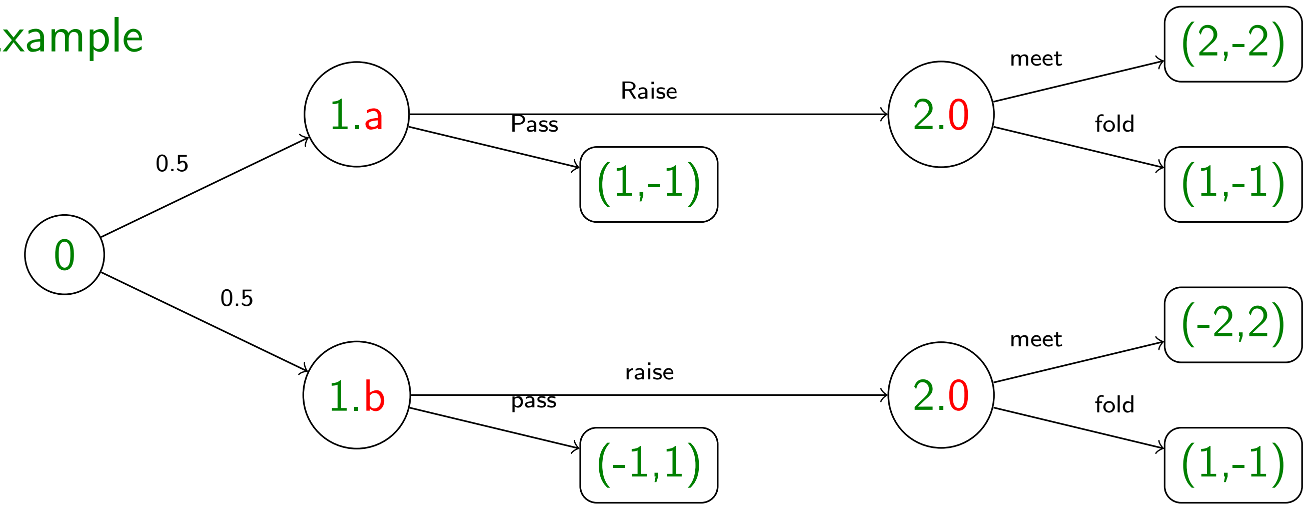

Definition

the normal representation

C_1	C_2: M	C_2: F
Rr	$0, 0$	$1, -1$
Rp	$0.5, -0.5$	$0, 0$
Pr	$-0.5, 0.5$	$1, -1$
Pp	$0, 0$	$0, 0$

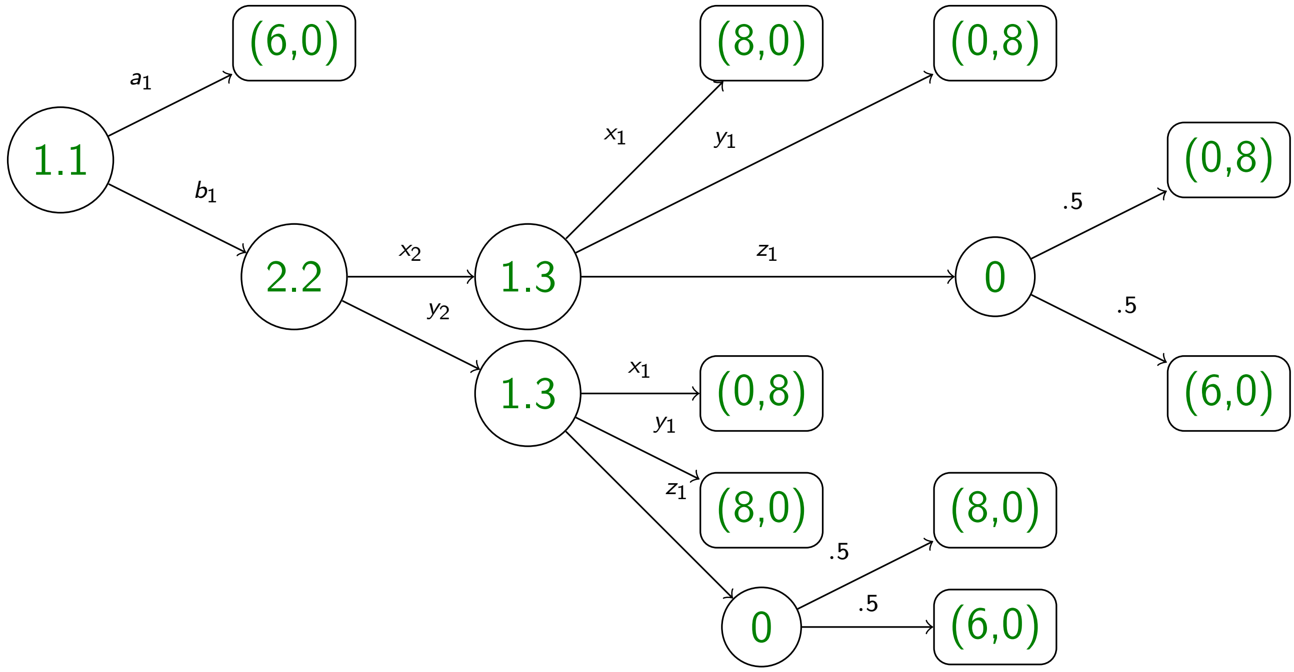

C_1	C_2: x_2	y_2
a_1x_1	$6, 0$	$6, 0$
a_1y_1	$6, 0$	$6, 0$
a_1z_1	$6, 0$	$6, 0$

C_1	C_2: x_2	y_2
b_1x_1	$8, 0$	$0, 8$
b_1y_1	$0, 8$	$8, 0$
b_1z_1	$3, 4$	$7, 0$

(Fully) Reduced Normal Representation

Definition

let $\Gamma = (N, (C_i)_{i\in N}, (u_i)_{i\in N})$, we say d_i and e_i in C_i, are payoff equivalent if

$$u_j(c_{-i}, d_i) = u_j(c_{-i}, e_i) \quad \text{for all } c_{-i} \in C_{-i},\ j \in N$$

Example

strategies a_1x_1, a_1y_1, a_1z_1 are payoff equivalent

Definition

identifying payoff equivalent strategies yields the purely reduced normal representation

Example

C_1	C_2: x_2	y_2
$a_1\cdot$	$6, 0$	$6, 0$
$b_1 x_1$	$8, 0$	$0, 8$
$b_1 y_1$	$0, 8$	$8, 0$
$b_1 z_1$	$3, 4$	$7, 0$

Example

consider the following game

C_1	C_2: x_2	y_2	z_2
x_1	3,0	0,2	0,3
y_1	2,0	1,1	2,0
z_1	0,3	0,2	3,0

the unique Nash equilibrium is (y_1, y_2) as it is the best-response to all other strategies