Multilateral Mediated Negotiation Protocols with Feedback

Reyhan Aydoğan, Koen V. Hindriks and Catholijn M. Jonker

Multilateral Negotiation

- Negotiation among more than two participants
  - Four friends negotiating on their holiday
  - Three political parties negotiating on a new regulation
- All parties mutually agree on the final decision/outcome

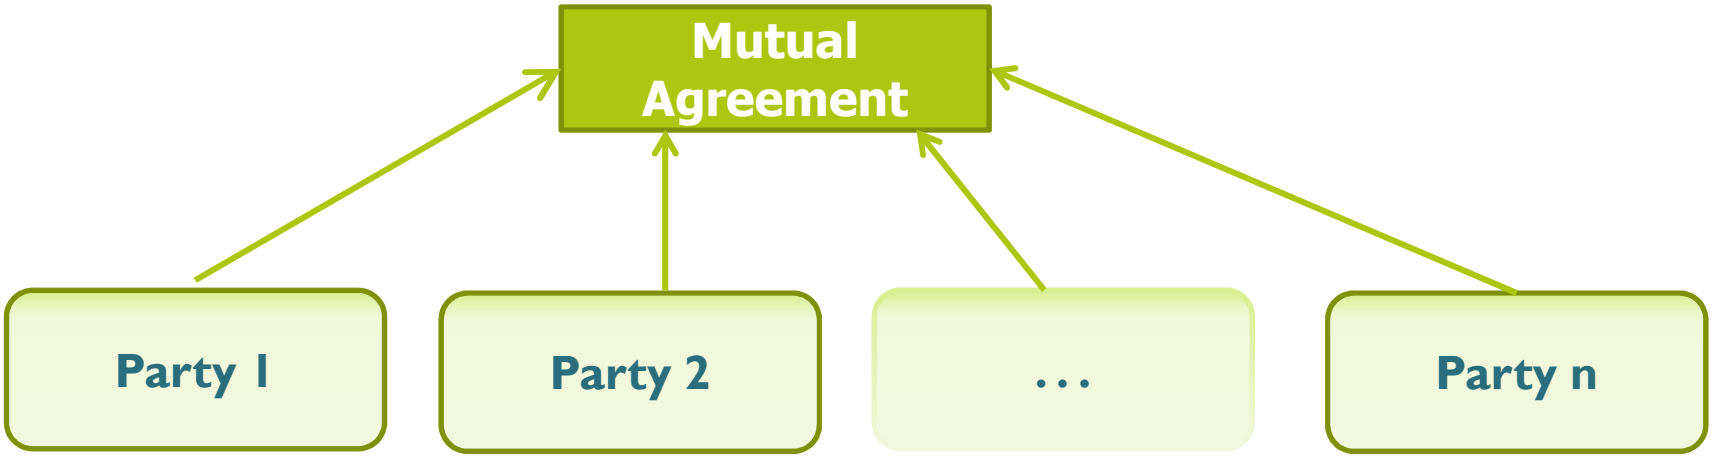

Multilateral Negotiation Protocol

- Protocol: governs the interaction between parties
  - How do the participants interact?
  - What are the valid actions for each party?
  - When does the negotiation end?
  - How is the final decision made?
- As a starting point, taking the mediated single text negotiation protocol [Klein *et al.*, 2003]
- Proposing two variants of that protocol
  - Based on feedbacks and preference modelling

Mediated Single Text Negotiation Protocol

Mediator generates an offer and asks negotiation agents for their votes either to accept or to reject this offer.

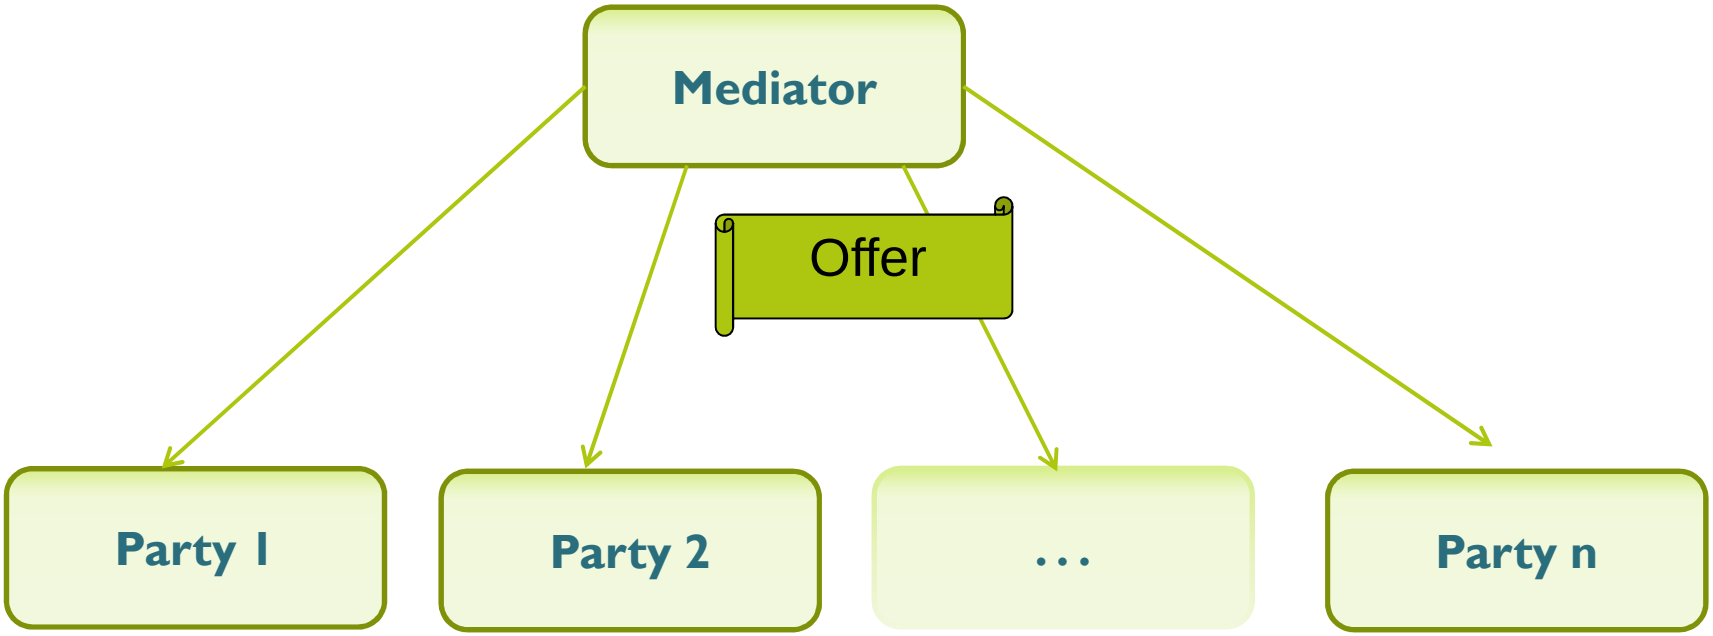

Mediated Single Text Negotiation Protocol

Negotiating agents send their votes for the current bid according to their acceptance strategy.

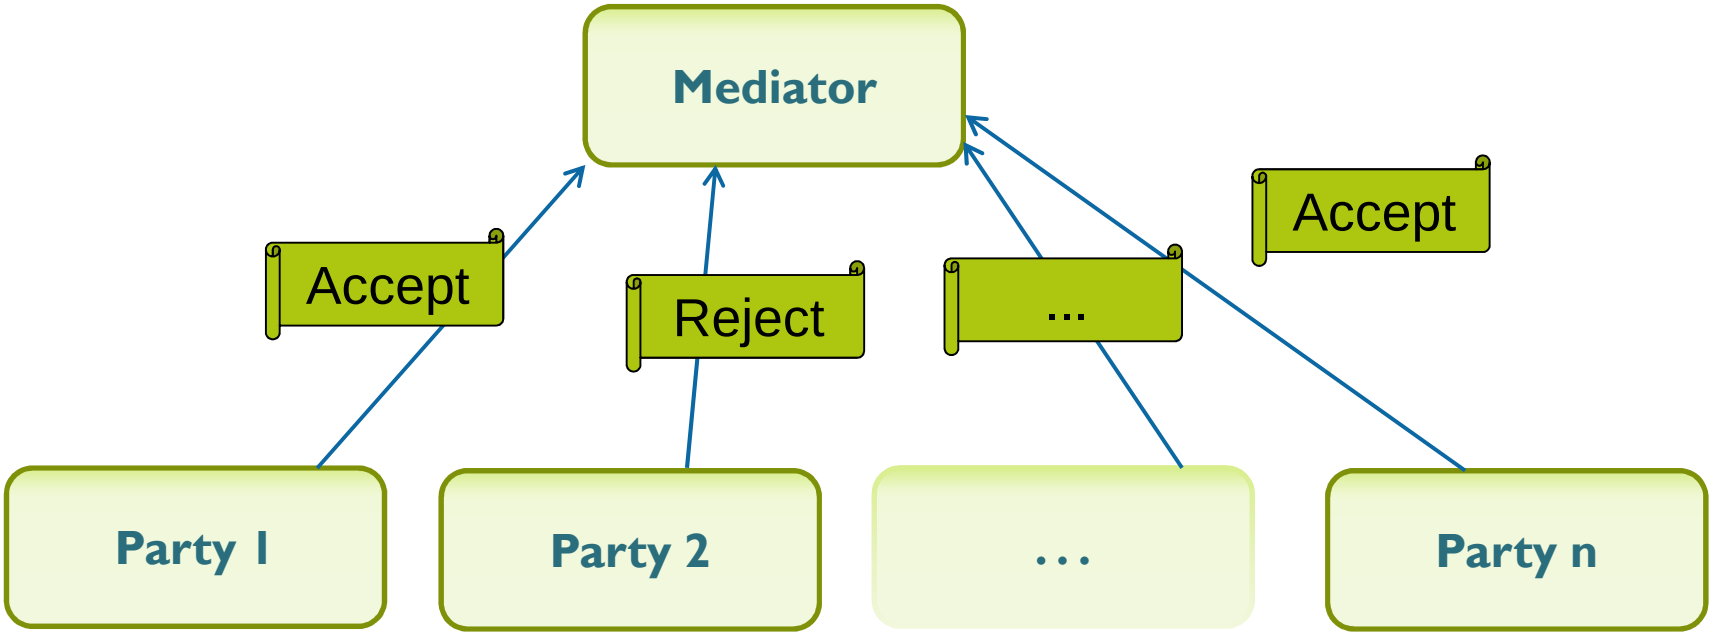

Mediated Single Text Negotiation Protocol

Mediator modifies the most recently accepted bid by exchanging one value arbitrary and asks negotiating agents' votes again

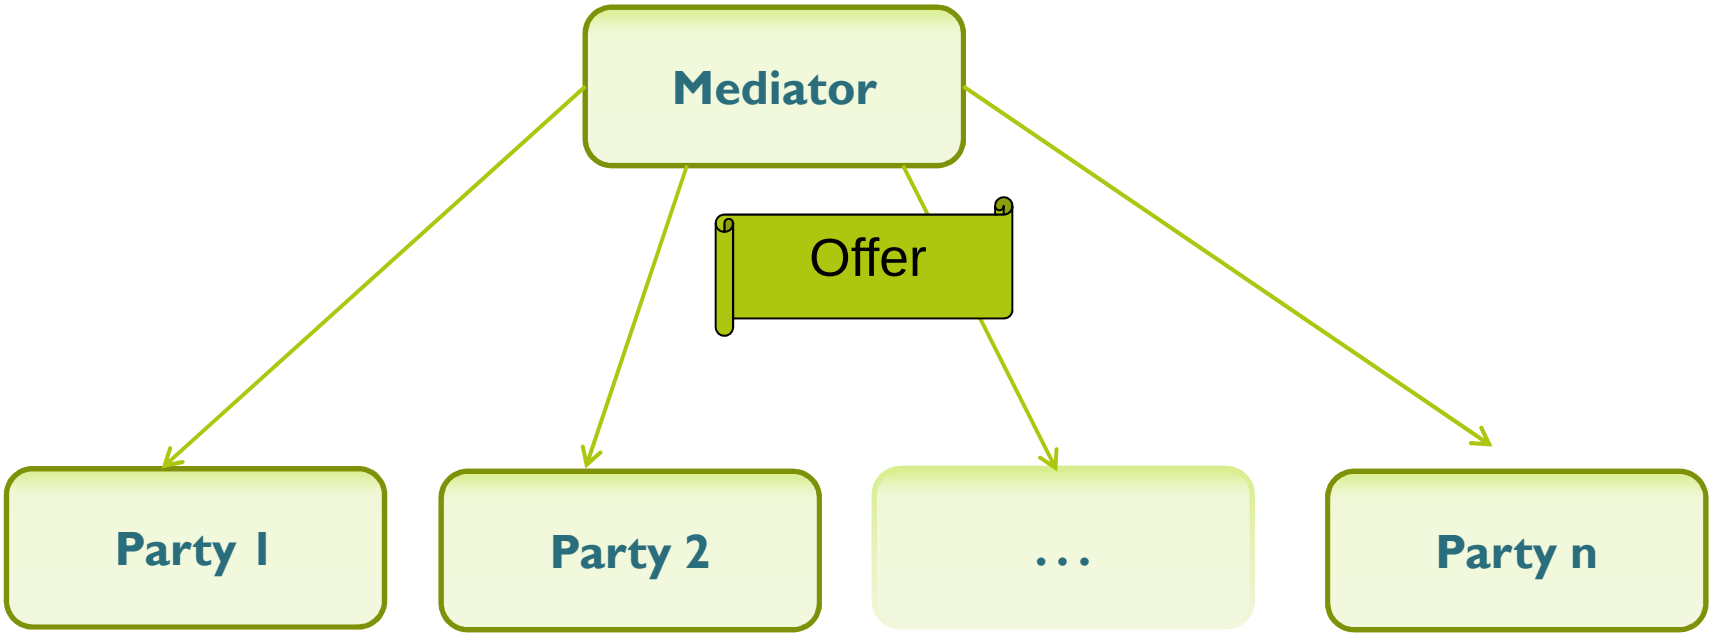

This process continues iteratively until reaching a predefined number of bids.

Mediated Single Text Negotiation: Mediator

- In the first round, the mediator
  - generates its first bid randomly
    - **E.g. Bid: (Paris, 1-week holiday, 3 star hotel)**
  - asks the negotiating agents to vote for this bid (accept/reject)
  - labels the bid as the most recently accepted bid if all negotiating agents vote as "accept"
    - **E.g. MRA Bid: (Paris, 1-week holiday, 3 star hotel)**

Mediated Single Text Negotiation Mediator

- In further rounds, the mediator
  - Modifies the most recently accepted bid by exchanging one value with another randomly in the bid
    - **MRA Bid: (Paris, 1-week holiday, 3 star hotel)**
    - **New Bid: (Rome, 1-week holiday, 3-star hotel)**
  - Asks the negotiating agents to vote for this bid (accept/reject)
  - Updates the most recently accepted bid if all negotiating agents vote as "accept"
    - **MRA Bid: (Rome, 1-week holiday, 3 star hotel)**
  - Continue generating offers and asking other agents' votes until reaching a predefined number of rounds.

Mediated Single Text Negotiation Hill-Climber Agent

- Accept a bid if its utility is higher than the utility of the most recently accepted bid
  - **MRA Bid= (Antalya, 1-week, 3 star-hotel),**
  - **Bid_6= (Antalya, 1-week, 5 star-hotel),**
  - **$U(Bid_6)=0.95 > U(\text{MRA Bid})=0.87 \rightarrow$ ACCEPT**
- Problem:
  - If the utility of initial bid is quite high for one of the agents, that agent may not accept other bids even though those bids might be better for the majority.

Mediated Single Text Negotiation: Annealer Agent

- Calculates the probability of acceptance for the current bid:

$$P(\text{accept}) = \min(1, e^{-\Delta U/T})$$

T: Virtual temperature gradually declines over time

- Higher probability for acceptance
  - The utility difference is small
  - Virtual temperature is high

- Tendency to accept individually worse bids earlier so the agents find win-win bids later

Proposed Mediated Negotiation

- The agents give a feedback such as “better”, “worse” and “same”
  - Comparing the current bid with the previous one
- Based on those feedbacks, the mediator can generate better bids for all of the agents
  - Modelling the preferences of each agents by building up preference graphs
  - Applying a heuristic to estimate the utility of a bid for each agent
  - Generating the bids according to the estimated utilities

Mediator:
Feedback Based Preference Modelling

- During the negotiation, mediator
  - Mutates its previous bid by flipping one of the issues
    - **Previous bid: (Paris, One-week, 3-star hotel)**
    - **Current bid: (Barcelona, One-week, 3-star hotel)**
  - Gets feedback from the negotiating agents
    - E.g. "**Better": Barcelona> Paris**

- Modelling each agent's preferences
  - Assumption: No preferential interdependency & total preorder
  - Constructing a preference graph for each issue
    - $M_i = \{PG_1, PG_2, \ldots. PG_n\}$ if we have n issues
    - Nodes: denote the values of the given issue
    - Edges: show the *improving flips*; from less preferred to more preferred

Extracting more preferential information from the graph

- By using three feedbacks,
  - Feedback 1: Barcelona is better than Paris.
  - Feedback 2: Paris is same with Rome.
  - Feedback 3: Budapest is worse than Rome

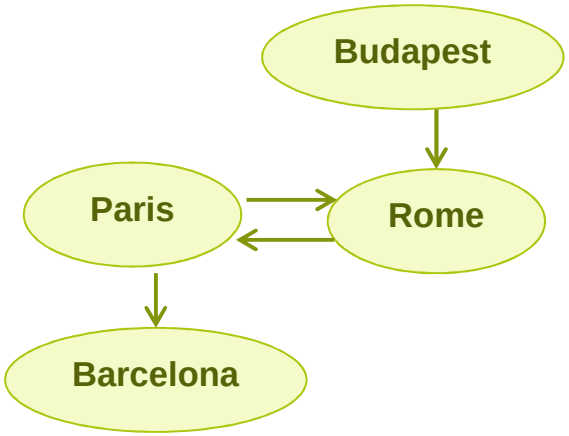

- By applying "transitivity", we are also able to compare the following value pairs:
  - Barcelona is better than Rome.
  - Barcelona is better than Budapest.
  - Paris is better than Budapest.

- Partial Graph: cannot compare each value pair.
  - Applying a heuristic similar to depth in our work with CP-net

Scoring Each Value in the Preference Graph

- Assigning a score to each node and updating that score during the negotiation,
  - If x is better than y, the score of x will be higher than that of y.
  - If x is the same with y, the score of those nodes will be the same.

- Assume that x is the previous value and y is the current value
  - If y does not exist in the graph,
    - Score (y) ← Score (x) + 1 when feedback is better
    - Score (y) ← Score (x) – 1 when feedback is worse
    - Score (y) ← Score (x) when feedback is same
  - Otherwise,
    - If there is any inconsistency in scoring according to the given feedback, update the scores to resolve the inconsistency

Resolving inconsistency in scoring according to the feedback

- Y = {Y1, Y2, Y3, Y4, Y5, Y6}
- Current feedback says **Y1>Y6**

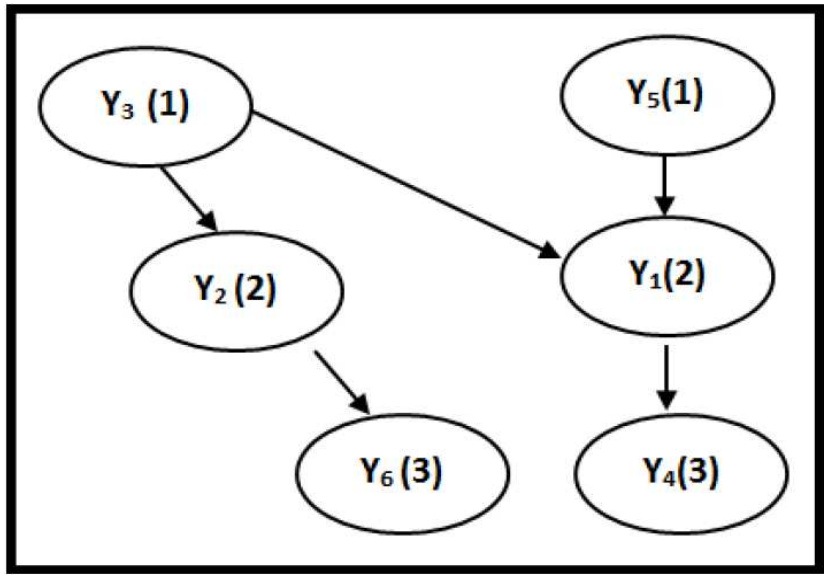

The graph before the feedback: The score of Y6 (3) is greater than the score of Y1 (2).

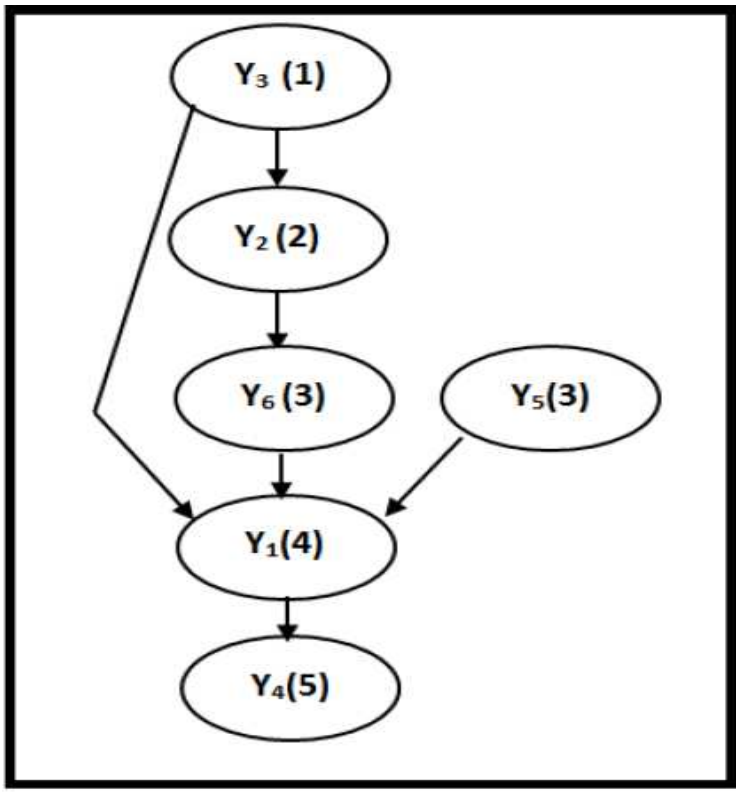

After updating the graph wrt the given feedback: The score of Y6 (3) is lower than the score of Y1(4).

How does the mediator use the estimated scores?

- Aim is to increase the social welfare and find the values that are better for all agents
- Scores are used to find the values giving the Nash product (maximizing the product)
  - The estimated scores are scaled between zero and one (0,1].
- Assume that we have three agents and their estimated score for $D(X) = \{x_1, x_2, x_3\}$
  - M1 (first agent): $EU(x_1)=1:0$; $EU(x_2)=0:66$; $EU(x_3)= 0:33$.
  - M2 (second agent): $EU(x_1)=0:5$; $EU(x_2)= 1$; $EU(x_3)=1$.
  - M3 (third agent): $EU(x_1)=0:33$; $EU(x_2)=0:66$; $EU(x_3)=1$.
- Products:
  - $P(x_1)=0.17$; $P(x_2)=\mathbf{0.44}$; $P(x_3)=0.33$;

Protocol -1: Feedback Based Protocol

- Phase-1: *Searching* – change only one issue value at a time according to the following heuristics:
  - Unused Values: randomly choose the values that have not been used before.
  - Incomparable Values: randomly choose the values that could not be compared with the previous issue value.
  - Random Values: randomly choose any issue value that may improve the bid for all agents
- Phase-2: *Exploitation*
  - Nash Values: randomly choose an issue and select the value for that issue whose product of the estimated utility is the maximum (Nash)

Protocol -1: Feedback Based Protocol

- During the negotiation, the mediator keeps
  - *"Last recent better bid"*
- If none of the agents' feedbacks is "worse", update the current bid as "last recent better bid".
- When reaching the deadline, the last recent better bid is taken as a negotiation outcome.

Protocol-2: Feedback & Voting Based Protocol

- Phase 1: *Searching and Learning*
  - Same with the feedback based protocol
  - Unused values, incomparable values, random values are used to make a new bid
  - If there is no such values, pass the second phase

- Phase 2: Voting with estimated Nash bids
  - Generates Nash bids maximizing the product of the estimated utilities for all agents
  - Asks agents' vote to either accept or reject
  - Updates most recently accepted bid
  - After generating all Nash bids, the mediator finalizes the negotiation with most recently accepted bid

Experiments

- Using party domain consisting of six issues
  - # of possible outcome 3072
- Creating five different group and each group negotiates 100 times in each protocol setting

Group	Agents	Maximum Product of Utilities (Nash Product)
Group-1	(A1-A2-A3)	0.76
Group-2	(A1-A4-A5)	0.61
Group-3	(A2-A4-A6)	0.50
Group-4	(A3-A5-A6)	0.64
Group-5	(A1-A7-A6)	0.78

- Metric: Average product of utilities of the agents
- Different deadline durations:
  - 50 rounds, 250 rounds and 500 rounds

Results: when the deadline is 50

Average product of utilities of the agents

Group	Hill-Climber	Annealer	Feedback	Feedback & Voting*
Group-1	0.42	0.42	0.65	**0.71**
Group-2	0.37	0.40	**0.48**	**0.47**
Group-3	0.25	0.23	**0.30**	**0.30**
Group-4	0.53	0.46	0.62	**0.64**
Group-5	0.47	0.48	**0.56**	**0.57**
Overall:	0.41	0.40	**0.52**	**0.54**

* It completes the negotiation in 30 rounds on average.

- *Feedback* and *Feedback & Voting* protocols outperforms others on average.

Results: when the deadline is 250

Average product of utilities of the agents

Group	Hill Climber	Annealer	Feedback	Feedback & Voting*
Group-1	0.42	0.61	0.65	**0.71**
Group-2	0.37	**0.52**	**0.51**	0.47
Group-3	0.25	**0.35**	0.31	0.31
Group-4	0.53	0.54	**0.64**	**0.64**
Group-5	0.47	**0.66**	0.57	0.57
Overall:	0.41	**0.54**	**0.53**	**0.54**

It completes the negotiation in 30 rounds on average.

- The performance of Annealer increases drastically when the number of rounds increases.
- Note that Feedback & Voting ends negotiation in 30 rounds.

Results: when the deadline is 500

Average product of utilities of the agents

- The performance of Annealer is better than ours.
- Feedback & Voting completes negotiation in only 30 rounds.

Group	Hill Climber	Annealer	Feedback	Feedback & Voting*
Group-1	0.42	0.66	0.66	**0.71**
Group-2	0.37	**0.55**	0.51	0.47
Group-3	0.25	**0.40**	0.31	0.31
Group-4	0.53	0.56	**0.64**	**0.64**
Group-5	0.47	**0.69**	0.57	0.57
Overall:	0.41	**0.57**	0.54	0.54

*It completes the negotiation in 30 rounds on average.

- When both time and performance are concerned, feedback & voting protocol is a promising protocol that results in reasonably good agreements in a short time.