

# NEGOTIATIONS

Leen-Kiat Soh

Fall 2016

CSCE990 Advanced Multiagent Systems Seminar  
Department of Computer Science and Engineering  
University of Nebraska

# Introduction

- This presentation is based on a set of papers published in late 1990s and is meant to cover a range of negotiation research directions that took place at that time, serving as an informative, foundational overview for students who are interested in negotiations to pursue more recent papers in the 2000s/2010s

# Introduction

- Agents **negotiate** to exchange information
  - Each agent maintains its own information base and experience
- When an agent needs to collaborate with another agent to solve a problem, the **(initiating)** agent negotiates with that other **(responding)** agent
  - Negotiation involves exchanges of information until a deal is agreed or rejected
- Multiagent negotiations can be used to perform task allocations (collaborative agents), resource allocations (cooperative agents), knowledge distribution (learning agents), etc.

# Bazaar and Bayesian Negotiation 1

- Work by Katia Sycara around 1996-1998
- Bazaar is powered by a Bayesian belief network
- Given each response from the counterpart, an agent updates its beliefs regarding the counterpart's **reservation price**
  - A supplier's reservation price is the price below which the supplier agent will not accept an offer
  - A customer's reservation price is the price above which the customer agent will not accept an offer
- The negotiation is based on making new offers based on previous offers -- **imitative**

# Bazaar and Bayesian Negotiation 2

- Unique characteristic: Bayesian belief network
- Reliance on solely the offer or counter-offer to drive the updates in the belief network constrains the approach
- Has to come up with the initial conditional probabilities to build the belief network

# Negotiation Decision Models 1

- Work by Peyman Faratin around 1998
- A formal model of negotiation between autonomous agents
  - Multi-party, multi-issue, and single-encounter negotiations among competitive and cooperative agents

# Negotiation Decision Models 2

- Negotiation **issues** are such as price, volumes, duration, quality, etc.
  - If A1 is trying to sell an item to A2 and both are negotiating about the price, then the price is the negotiation issue
- Negotiation **tactics** are the set of functions that determine how to compute the value of an issue by considering a single criteria
  - Time, resources, previous offers, counter offers, etc.

# Negotiation Decision Models 3

- A negotiation strategy:
  - A linear combination of various tactics generates the value for an issue at the current time step of a negotiation
  - Different types of negotiation behavior can be obtained by weighting the tactics differently
  - A negotiation strategy is any function that is based on the matrix of the weights and the mental state of the agent



# Negotiation Decision Models 4

- Time-dependent tactics
  - Polynomial, exponential, Boulware tactics (stand firm until the time is almost exhausted), conceder
- Resource-dependent tactics
  - Dynamic-deadline tactics
    - The greater the number of agents who are negotiating with agent A for a particular service, the lower the pressure on agent A to reach an agreement with any specific individual
    - The longer the negotiation, the greater the pressure on A to come to an agreement
  - Resource estimation tactics in which diminishing resource (based on time) is factored into Boulware tactics

# Negotiation Decision Models 5

- Behavior-dependent tactics
  - **Relative tit-for-tat:** an agent reproduces, in *percentage* terms, the behavior that its counterpart performed several steps ago
  - **Random absolute tit-for-tat:** similar, but uses *absolute* increase or decrease
  - **Averaged tit-for-tat:** an agent computes the *average* of percentages of changes in a window of a certain size of its counterpart's history when determining its new offer

# Argumentative Negotiation Models 1

- Work by Parsons and Jennings, around 1996-1998
- Uses bridge rules with their logic and theories
  - A bridge rule is a rule of inference with premises and conclusions in different units
- Incorporates multi-context Belief-Desire-Intention (BDI) agents
- Incorporates formalism to construct arguments to evaluate offers and counter-offers

# Argumentative Negotiation Models 2

- Why arguing? About what?
  - About what to do, what to give up, **AND**
  - About *why* one agent has to do it, *why* one agent has to give up a certain resource
- Parsons and Jennings make a strong case for **sending over the reasoning rules** of an agent to another agent
  - So that the other agent knows exactly what is behind it's counterpart's claims
  - Instead of "tell me the importance", it is "give me your reasoning process and I will be the judge of that".
- Makes sense in an agency of diverse and competitive agents

# Argumentative Negotiation Models 3

- Allows diversification and efficient *knowledge* (not just information) distribution within an agency
- In an agency of **homogeneous** agents, this model may not be necessary
- Truly argumentative: **agents not only argue about what evidence there is to support their claims, but also how one should interpret the evidence**

# Argumentative Negotiation Models 4

- There are several scenarios where an agent B does not agree with the arguments of an agent A:
  - Arguments **directly** conflicts with the objective of B
  - Arguments **partially** conflicts with the objective of B
  - Arguments are acceptable but B does **not** have the resources to accept the deal

# Game Theory 1

- Work by Sarit Kraus, around 1995-
- Game theoretic techniques suitable for:
  - Agents are **self-motivated** and try to **maximize their own benefits**
  - Designers of the agents may agree in advance on regulations and protocols for the agents' interaction
  - Number of agents small
  - Agents can communicate and have computational capabilities
- Game theory models are highly abstract representations of classes of real life situations that involve individuals who have **different goals or preferences**
  - Each active entity is a **player**

# Game Theory 2

- Diplomat (Kraus and Lehmann 1995)
  - Uses a 1-to-N negotiation approach
  - Uses a **notice board** that keeps track of the common expected profits of each negotiation so that the agent can make a selection among the N **concurrent** negotiated deals
  - Uses the notice board to record contradictions among the deals
  - Based on the profits and contradictions, computes the **intention of honoring a deal**
  - Each agent keeps track of a **loyalty** measure of its neighbors



# Model of Alternative Offers 1

- Work by Sarit Kraus, around 1993-1997
- A negotiation model on hostage crisis
  - Modification of Rubinstein's model of alternative offers which focuses on
    - the passage of time
    - the preferences of the players for different agreements
    - for opting out of the negotiations
  - Both players can opt out
  - Takes into account the effect of time: one player gains over time, the other loses over time

# Model of Alternative Offers 2

- Assumptions
  - Agents care only about the nature of the outcome (agreement of opting out) and the time at which the outcome is reached
  - Agents do not care about the sequence of offers and counter-offers that lead to the outcome
    - Cost of negotiation not important!
  - No agent regrets either making an offer that was rejected or rejecting an offer
    - Cost of making or rejecting an offer not important!
  - Disagreement is the worst outcome!
  - The initiator gains over time; the responder loses over time (hostage scenario)

# Model of Alternative Offers 3

- Shows that if there is an **agreement zone**, an agreement will be reached in the **first** or the **second** step of the negotiation!
- Why? Because the players either **have the complete knowledge of each other to begin with** or **have the complete knowledge of each other after the first offer**
- **My opinions:** **CAUTION**
  - Does not apply to general multiagent negotiations
  - Agents are not competitive (not wanting to gain)
  - The initiating agent is driven to keep the negotiation alive as long as possible – not quite practical in other domains

# Model of Alternative Offers 3

- Investigation under time constraints (started by Kraus and Wilkenfeld 1991)

## Situation 1

Two agents that *lose over time* need to share a common resource and each agent knows *all* relevant information about the other agent. They have no alternative but to negotiate until an agreement is reached.

## Strategy

The agents will reach the perfect equilibrium (a deal) in *one* handshake (or after the first offer)

# Model of Alternative Offers 4

## Situation 2

Two agents that *lose over time* need to cooperate to satisfy a common goal and each agent knows *all* relevant information about the other agent. Each agent can *unilaterally leave* the negotiations.

## Strategy

The agents will reach the perfect equilibrium (a deal) in *one* handshake (or after the first offer)

# Model of Alternative Offers 5

## Situation 3

Two agents need to share a resource. One of the agents already has access to the resource and is using it during the negotiation process. It is thus *gaining over time*. The other agent is waiting to use the resource and *loses over time*. Both agents have *full* information and can *unilaterally leave* the negotiations.

## Strategy

### CAUTION

The agreement is guaranteed at the latest in the second step since even though the agent that has the upperhand prefers to continue the negotiation indefinitely, it is afraid that the counterpart might threaten to opt out at any given time!

# Model of Alternative Offers 5

## Situation 4

Two agents need to share a resource. One of the agents already has access to the resource and is using it during the negotiation process. It is thus *gaining over time*. The other agent is waiting to use the resource and *loses over time*. agents *do not have complete info.* about each other and can *unilaterally leave* the negotiations.

## Strategy

### CAUTION

The agents conduct a *sequential equilibrium* negotiation. The agents will reach the perfect equilibrium (a deal) in at most two steps.

# Model of Alternative Offers 6

## Situation 5

Several agents need to cooperate to satisfy a common goal. *All of them are losing over time*, have *full* information about each other, and can *unilaterally leave* the negotiations.

## Strategy

At each time period, one agent makes an offer to all the other agents. Each of the agents accepts the offer, or rejects it, or opts out. If the offer is accepted by all, then the negotiation ends and the agreement is implemented. Opting out by any agent ends the negotiation. After a rejection by at least one agent, another agent must make a counter offer. Each has *veto* power!!!



# Model of Alternative Offers 7

- The **sequential equilibrium** (Kreps and Wilson 1982) involves a sequence of strategies and a system of belief with the following properties:
  - Each agent has a belief about its counterpart's type
  - At each step, the strategy for an agent A is optimal given its current belief and its counterpart's possible strategies in the sequential equilibrium
  - The history is a sequence of previous proposals and responses

# Model of Alternative Offers 8

- My opinions: **CAUTION**
  - Agents are *implicitly* cooperative even though they are non-cooperative
  - In a full information situation, an agent knows exactly what the counterpart wants and thus makes an offer accordingly
  - In a partial information situation, an agent learns exactly what the counterpart wants after the first offer
  - The approach is particularly useful when inter-agent communication is **minimal**, individual agent knowledge is small, and less **intelligent** as a result

# United Negotiation Protocol 1

- Work by Zlotkin and Rosenschein, around 1989-1995
- A **theoretical** negotiation model for rational agents in general non-cooperative domains
- A **united negotiation protocol** (UNP) for conflict resolutions and agents taking partial cooperative steps
  - Agents coexist in a single system and are **predisposed towards cooperative** activity; some **notion of global utility** to maximize
  - Agents are **self-serving**, have their **own utility functions** and **no global notion of utility**; **disparate goals**

# United Negotiation Protocol 2

- Assumptions (Zlotkin and Rosenschein 1989)
  - Utility maximizer: Each agent wants to maximize its expected utility
  - Complete knowledge
  - No history: Each negotiation stands alone
  - Fixed goals: Agents do not change their pre-defined goals
  - Bilateral negotiation: No 1-to-N negotiations, always a pair
  - Symmetric abilities: key to their exchange of plans, everybody can perform the same set of tasks
  - Deterministic world

# United Negotiation Protocol 3

- The *utility* of an agent from a deal is simply the difference between the cost of achieving its goal alone and its expected part of the deal
  - A deal is *individual rational* if the *utility* of that deal to *each partner of the deal* is not negative
  - A deal is *pareto-optimal* if there does *not* exist another deal that dominates it—there does *not* exist another deal that is *better for one of the agents and not worse for the other* (no incentive to back out of a deal!)
  - The *negotiation set* is the set of all the deals that are *both individual rational and pareto-optimal*

# United Negotiation Protocol 4

- A **necessary** condition for the negotiation set to be non-empty is that there is **no contradiction** between the two agents' goals
- But, even so, a **conflict** is still possible! A **conflict** is where any **joint plan that satisfies the union of goals** will cost one agent (or both) more than it would have spent achieving his own goal in isolation—i.e. **no deal is individual rational**

# United Negotiation Protocol 4

- Conflicts
  - A joint plan satisfies the sum condition if the sum of all costs for all agents to perform the task individually is **greater than or equal to** the sum of all costs for all agents to perform the parts of the joint plan
  - A joint plan satisfies the min condition if the minimum cost for an agent to perform the task individually is **greater than or equal to** the minimum cost for an agent to perform the joint plan
  - When these two conditions are true, the agents are in cooperative situations!

# United Negotiation Protocol 5

- In non-conflict situations, if neither the min condition nor the sum condition are true, then in order for the agents to cooperate, then at least one of the agents will have to do more than if it were alone in the world and achieved only its own goals
  - An agent has to redefine its utility of the “worth of a goal”, as the maximum expected cost that an agent is willing to pay in order to achieve its goal
  - So, now, the utility of an agent from a deal is the difference between the worth of a goal and the agent’s cost of its expected part of the deal



# United Negotiation Protocol 6

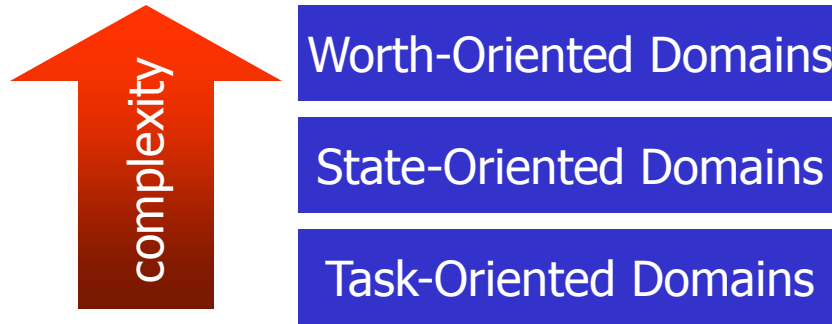
- Given this redefinition of utility, the agents may now be involved in three possible interactions:
  - A **cooperative** situation in which there **exists a deal in the negotiation set that is preferred by an agent over achieving his goal alone**; so that agent welcomes the existence of the other agents
  - A **compromise** situation in which **there are individual rational deals for an agent**. However, an agent would prefer to be **alone** in the world, and to accomplish its goal **alone**
  - A **conflict** situation in which the **negotiation set is empty**

# United Negotiation Protocol 7

- When a **conflict** arises, a conflict resolution via coin toss based on weights is conducted
  - So, one agent sacrifices to help another agent achieve that agent's goals
  - When this occurs, we have a **semicooperative** deal (reluctant!)
- In **cooperative** and **compromise** situations, the agents negotiate on deals that are **mixed joint plans** (or **cooperative** deals)

# United Negotiation Protocol 8

- Extensions



- In TOD, agents have nonconflicting jobs to do and these jobs/tasks can be **re-distributed** among the agents; **objective is to re-distribute tasks to everyone's mutual benefit if possible**
- In SOD, actions in these domains can have **side effects**, where an agent performing one task might hinder or help another agent; **objective is to develop joint plans or schedules for the agents**
- In WOD, agents strive for **better states** through a decision-theoretic formulation; **objective is to have joint plans or schedules and goal relaxation**

# United Negotiation Protocol 9

- Three types of Task-Oriented Domains (Rosenschein and Zlotkin 1994)
  - **Subadditive** where tasks may be inter-related and by combining them may cost less
  - **Concave** where an additional task adds less cost to a task X than to a task Y
  - **Modular** where tasks are not related and cannot be less costly done if combined

# United Negotiation Protocol 10

- Behavior of a deceitful agent (Rosenschein and Zlotkin 1994)
  - Letter hiding to withhold information
  - Phantom letters to cheat
  - Decoy task to mislead

# United Negotiation Protocol 11

- Attributes of standards for agents (Zlotkin and Rosenschein 1996) for designers
  - **Efficiency**; agents should not squander resources when they come to an agreement
  - **Stability**; **no** designer should have an incentive to **deviate from agreed-upon strategies** (heterogeneous agents)
  - **Simplicity**; should have **low computational demands** on the agents and require **little communication overhead**
  - **Distribution**; preferably the interaction rules will **not require a central decision maker**
  - **Symmetry**; **no** designer wants the negotiation process to be arbitrarily **biased** against its agent

# Partial Global Planning 1 **CAUTION**

- Work by Victor Lesser and Ed Durfee, around 1991-1994
- A partial global planning that gives an agent the ability to
  - **Represent** its own expected interpretation activities (of sensory data)
  - **Communicate** about these expectations with others
  - **Model the collective activities** of multiple systems
  - **Propose changes** to one or more systems' interpretation activities to improve group performance
  - **Modify its planned local activities** in accordance with the more coordinated proposal

# Partial Global Planning 2

- Object of negotiation: to combine information from several agents in order to reach an outcome for the entire group of agents
- Agents are organized hierarchically:
  - **Subordinate agents**: sends the plan information to supervisors, receives the modified plans and can adopt them
  - **Supervisor agents**: reviews the plan information from subordinates, makes modifications, and sends the modified plan information back
  - **Peer agents**: sends information to another agent and receives modified plan information back and can carry out its tasks locally



# Partial Global Planning 3

- When messages are passed, the **goals**, the **long-term strategies**, and the **rating of a plan** are exchanged
- A framework called **TEAM** that implements the **cooperative search** and **conflict resolution** among heterogeneous, reusable, expert agents, based on the partial global planning approach
- Two types of heterogeneity
  - **Logical** where agents may have different long-term knowledge (expertise), goals, views or perspectives, constraints or preferences, or criteria for evaluating solutions
  - **Implementational** where agents may have different knowledge representations, languages, architectures, inference engines, etc.

# Partial Global Planning 4

- A set of requirements for the agents
  - Agents sets are **dynamically** formed by grouping agents with the specific expertise required for the problem
  - Agents do **not have prior knowledge** of what other agents will be included in the set and what their capabilities will be **WOW!**
  - The agents sets are **cooperative**—agents are not hostile and will not intentionally mislead or otherwise try to sabotage another agent's reasoning
  - Agents are **willing** to contribute both knowledge and solutions to other agents as appropriate and to accept solutions that are not locally optimal in order to find a mutually-acceptable solution

# Partial Global Planning 5

- Two basic approaches to negotiated search
  - When an agent recognizes a conflict with another agent in an existing solution, it extends its local search until a solution is found that does not conflict
    - Heuristic search, case-based search, and searching for alternate goal expansions
  - An agent relaxes some requirement on a solution, thereby expanding its local search space
    - Relaxing or relinquishing constraints, relaxing or relinquishing goals, manipulating constraints, or manipulating evaluation criteria

# Partial Global Planning 6

- A **general negotiated search** is an **opportunistic** search augmented by the **communication and assimilation of conflict information**
- The search behaves in the following manner:
  - One or more agents **produce base proposals**
  - Other agents **critique the partial solutions** created from those proposals
  - If a **conflict** is detected, any **constraining information** is **communicated** to other agents
  - Agents that receive conflict information **attempt to assimilate** that information

# Partial Global Planning 7

- The search behaves in the following manner, cont'd:
  - If an agent has successfully assimilated conflict information from another agent and later attempts to generate a proposal, the **new proposal will avoid that conflict**
  - If the required number of solutions is not found within a specified number (the **relaxation threshold**) of agent cycles, each agent will apply the negotiated-search operator *relax-solution-requirement* to expand the solution space
  - Agents **continue in a cycle of search and relaxation** until an acceptable solution is found or until further relaxation is impossible (at which point a failure is declared)

# Partial Global Planning 8

- The TEAM Framework (Lander and Lesser 1992)
  - Agents **communicate** through a **shared memory** (a **blackboard**)
  - A **framework controller** controls the framework
  - During run time, there are two phases: an **agent cycle** and a **framework cycle** **ART?**
  - During the agent cycle, each agent is invoked **sequentially**. Each uses information in shared memory to **choose** applicable negotiated-search operators and **add** them to its agenda. It then **invokes its highest-priority operators** and **returns** the result
  - During the framework cycle, **the controller is invoked to update** the shared memory based on messages from agents and to **propagate** the effect of changes to shared memory objects

# Partial Global Planning 9

- The TEAM Framework, cont'd
  - Each agent maintains a **local** knowledge base; **no consistency enforcement across agents**
  - When there are **explicit inconsistencies** in the local and the received knowledge, an agent
    - **overrides** local knowledge and, or
    - **relaxes** local requirements
    - **ignores** conflicting external knowledge, or
  - A **flexibility value** is attached to each piece of information to which degree an agent is willing to relax that information

**GOOD!!**

# Partial Global Planning 10

- The TAEMS Framework (Decker and Lesser 1994)
  - An extendable family of coordination mechanisms, called **generalized partial global planning**
  - TAEMS stands for **T**ask **A**nalysis, **E**nvironment **M**odeling, and **S**imulation
  - Represents coordination problems in a **formal, domain-independent way**
  - Attaches a **negotiability index** to every commitment that indicates the **difficulty in rescheduling** a taskgroup if the commitment is broken

**GOOD!!**



- There is **no explicit, direct negotiation** involved in the above approach
  - Agents exchange information through a blackboard, and re-plan based on the critiques and conflicts
  - Then they repost their solutions to the blackboard
  - The solutions are refined until acceptable by all agents or otherwise
- However, the negotiated search addresses **relaxation** and **commitment sharing**, very important topics in agent-based negotiations

# A Logical Model: ANA 1

- Work by Kraus and Sycara (1998)
- A **logical model** of the mental states of the agents based on a representation of their **beliefs, desires, intentions,** and **goals**
- **Argumentation** is an **iterative process** emerging from exchanges among agents to **persuade** each other and **bring about a change in intentions**
- Automated **Negotiation Agent** or ANA

# A Logical Model: ANA 2

- In order to negotiate effectively, an agent needs the ability to
  - Represent and maintain a model of its own beliefs, desires, goals, and intentions
  - Reason with other agents' beliefs, desires, goals, and intentions (communication required)
  - Influence other agents' beliefs, intentions, and behavior (persuasion required)
- The mental state of an agent is thus characterized by the notions of beliefs, goals, desires, intentions, and local preferences (that determines the degrees of desires, intentions, and goals)

# A Logical Model: ANA 3

- Agent characteristics
  - **Bounded**: a bounded agent does **not believe in falsehood**
  - **Omniscient**: an omniscient agent has beliefs that are closed under inferences, i.e., **its beliefs are consistent**
  - **Knowledgeable**: an agent is knowledgeable if **its beliefs are correct**
  - An agent will **keep believing in something** until it observes something otherwise, or it is memoryless, or it does not forget anything
  - **Confident**: an agent is confident if it **believes that it will succeed in carrying out its intended actions**

# A Logical Model: ANA 4

- Agent structure
  - A **mental state** (beliefs, desires, goals, and intentions)
  - **Characteristics** (agent type, capabilities, belief verification capabilities)
  - **Inference rules** (mental state update, argument generation, argument selection, request evaluation)

# A Logical Model: ANA 5

- Six argument types
  - **Threats** to produce goal adoption or abandonment on the part of the responder
  - **Entice** the responder with a **promise of a future reward**
  - **Appeal to past promise**
  - **Appeal to precedents as counterexamples** to convey to the responder a **contradiction** btw. what it says and the past actions
  - **Appeal to prevailing practice** to convey to the responder that the proposed action will further its goals since **it has furthered others' goals in the past**
  - **Appeal to self-interest** to convince a responder that taking this action will achieve a **high-importance goal**

# A Logical Model: ANA 6

- Use meta-rules to generate arguments, to select arguments, and to evaluate requests
- Here are some meta-rules for **argument generation**:
  - If the opponent is a memoryless agent, then do not choose “Appeal to Past Promise”
  - If the agent received a request from the opponent in the past, which included a future reward argument, and if that reward was the intended action right now, then choose “Appeal to Past Promise”
  - Use “Appeal to Self Interest” when the agent believes the opponent is not aware of the implications—therefore this should not be used on a knowledgeable or reasonable agent

# A Logical Model: ANA 7

- Here are some meta-rules for argument generation, cont'd:
  - Use “Appeal to Prevailing Practice” if you can find a third agent that has done this before
  - Use “Counterexample” if the agent knows about the past history of the responder



# A Logical Model: ANA 8

- Here are some meta-rules for argument generation, cont'd:
  - Use “Threat” if the agent can find something that the opponent hates
    - Obtain the list of desires of the opponent
    - Consider first the desires with the highest preference value for the opponent which are not included in the agent’s own desires set
    - Try to find a desire which involves action that the agent can perform while its opponent cannot
    - Then find a contradicting action to that desire
    - If the agent cannot, then choose “Appeal to Self Interest”
    - Otherwise, threaten!

# A Logical Model: ANA 9

- Here are some meta-rules for argument generation, cont'd:
  - Use “Promise of a Future Reward” if the agent can find something the opponent likes
    - A intends that B do X
    - B says no
    - A finds something that B might interested and then makes another offer: “Would you please do X for me if I promise to do Y for you?”

## INTERESTING!

- The responsibility lies with A, not B!
- Usually: B comes up with the things that it wants A to do.
- What's the difference??

# A Logical Model: ANA 10

- Argument selection rules
  - To choose from a set of potential arguments the **best** argument to communicate to its opponent
  - Try to use the **weakest argument** first, and if it does not succeed, follow with **stronger arguments**
  - **Severity order** is: **IMPORTANT!**
    - An Appeal to Prevailing Practice
    - A Counter-Example
    - An Appeal to Past Promise
    - An Appeal to Self Interest
    - A Promise of a Future Reward
    - A Threat

# A Logical Model: ANA 11

- Request evaluation rules
  - To evaluate a request (or a counter-request) to decide whether to accept or reject it
  - Computes three parameters: **IMPORTANT!**
    - **Collision\_flag**: indicate whether the results of the requested action conflict with the agent's current goals
    - **Convincing\_factor**: indicate how convincing the argument is given the requested action
    - **Acceptance\_value**: indicate the overall preference of the results of the requested action as opposed to all the other desires of the agent; also includes the *opponent's reliability for keeping promises* and *the opponent's percentage of threat executing*

# Others 1

- Contract Net (Smith and Davis, circa 1980) **IMPORTANT!**
- Negoplan (Matwin 1989)
- Multistage Negotiation (Conry 1988, 1991)
- Recursive Negotiation Model (Laasri et al. 1992)
- TRACONET (Sandholm and Lesser, circa 1993)
- Kasbah (Chavez and Maes 1996)
- Negotiation Dialogs (Walton and Krabbe 1995)
- Case-based Argumentative Negotiation (Soh and Tsatsoulis, circa 2000)
- Dynamic Coalition Formation (Soh and Tsatsoulis, circa 2000)

# References 1

- Chavez, A. and P. Maes (1996). Kasbah: An Agent Marketplace for Buying and Selling Goods, in *Proceedings of the 1<sup>st</sup> International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology*, 75-90.
- Chavez, A., A. Moukas and P. Maes (1997). Challenger: A Multi-agent System for Distributed Resource Allocation, in *Proceedings of the International Conference on Autonomous Agents*, Marina Del Ray, CA.
- Conry, S. E., R. A. Meyer, and V. R. Lesser (1988). Multistage Negotiation in Distributed Planning, in A. H. Bond and L. Gaser (eds.), *Readings in Distributed Artificial Intelligence*, 367-384, San Mateo, CA: Morgan Kaufmann.
- Conry, S. E., K. Kuwabara, V. R. Lesser, and R. A. Meyer (1991). Multistage Negotiation for Distributed Satisfaction, *IEEE Transactions on Systems, Man, and Cybernetics, Special Issue on Distributed Artificial Intelligence*, **21**(6):1463-1477.
- Decker, K. and V. Lesser (1993). A One-Shot Dynamic Coordination Algorithm for Distributed Sensor Networks, in *Proceedings of AAI-93*, pp. 210-216.
- Decker, K. and V. Lesser (1994). Designing a Family of Coordination Algorithms, in *Proceedings of the 13<sup>th</sup> International Workshop on Distributed Artificial Intelligence*, July.
- Doran, J. E., S. Franklin, N. R. Jennings, and T. J. Norman (1997). On Cooperation in Multi-Agent Systems, *The Knowledge Engineering Review*, **12**(3):309-314.
- Durfee, E. H. and V. R. Lesser (1991). Partial Global Planning: A Coordination Framework for Distributed Hypothesis Formation, *IEEE Transactions on Systems, Man, and Cybernetics*, **21**(5):1167-1183.
- Dworman, G., S. O. Kimbrough, and J. D. Laing (1996). Bargaining by Artificial Agents in Two Coalition Games: A Study in Genetic Programming for Electronic Commerce, in *Proceedings of AAI Genetic Programming Conference*.
- Elvang-Goransson, M., P. Krause, and J. Fox (1993). Dialectic Reasoning with Inconsistent Information, in *Proceedings of the 9<sup>th</sup> Conference on Uncertainty in Artificial Intelligence*.
- Faratin, P., C. Sierra, and N. R. Jennings (1998). Negotiation Decision Functions for Autonomous Agents, *International Journal of Robotics and Autonomous Systems*, **24**(3-4):159-182.
- Fox, J., P. Krause, and S. Ambler (1992). Arguments, Contradictions and Practical Reasoning, in *Proceedings of the 10<sup>th</sup> European Conference on Artificial Intelligence*.

# References 2

- Harsanyi, J. C. (1977). *Rational Behavior and Bargaining Equilibrium in Games and Social Situations*, Cambridge: Cambridge University Press.
- Jennings, N. R. (1995). Controlling Cooperative Problem Solving in Industrial Multi-Agent Systems Using Joint Intentions, *Artificial Intelligence*, **75**:195-240.
- Jennings, N. R., P. Faratin, M. J. Johnson, T. J. Norman, P. O'Brien, and M. E. Wiegand (1996). Agent-Based Business Process Management, *International Journal of Cooperative Information Systems*, **5**(2-3):105-130.
- Jennings, N. R., S. Parsons, P. Noriega, and C. Sierra (1998). On Argumentation-Based Negotiation, in Proceedings of International Workshop on Multi-Agent Systems, Boston, MA.
- Kraus, S. (1997a). Negotiation and Cooperation in Multi-Agent Environments, *Artificial Intelligence*, **94**(1-2):79-97. Special Issue on economic Principles of Multi-Agent Systems.
- Kraus, S. (1997b). Beliefs, Time, and Incomplete Information in Multiple Encounter Negotiations among Autonomous Agents, *Annals of Mathematics and Artificial Intelligence*, **20**(1-4):111-159.
- Kraus, S., E. Ephrati, and D. Lehmann (1991). Negotiation in a Non-Cooperative Environment, *Journal of Experimental and Theoretical Artificial Intelligence*, **3**(4):255-282.
- Kraus, S. and D. Lehmann (1995). Designing and Building a Negotiating Automated Agent, *Computational Intelligence*, **11**(1):132-171.
- Kraus, S., M. Nirkhe, and K. Sycara (1993). Reaching Agreements through Argumentation: A Logic Model (Preliminary Report), in *Proceedings of the Workshop on Distributed Artificial Intelligence*.
- Kraus, S., K. Sycara, and A. Evenchik (1998). Reaching Agreements through Argumentation: A Logical Model and Implementation, *Artificial Intelligence Journal*, **104**(1-2):1-69.
- Kraus, S. and V. S. Subrahmanian (1995). Multiagent Reasoning with Probability, Time, and Beliefs, *International Journal of Intelligent Systems*, **10**(5):459-499.
- Kraus, S. and J. Wilkenfeld (1993). A Strategic Negotiations Model with Applications to an International Crisis, *IEEE Transactions on Systems, Man, and Cybernetics*, **23**(1):313-323.
- Kraus, S. and J. Wilkenfeld (1991a). Negotiations over Time in a Multiagent Environment: Preliminary Report, in *Proceedings of IJCAI-91*, Australia, 56-61.

# References 3

- Kraus, S. and J. Wilkenfeld (1991b). The Function of Time in Cooperative Negotiations, in *Proceedings of AAAI-91*, California, 179-184.
- Kraus, S., J. Wilkenfeld, and G. Zlotkin (1995). Multiagent Negotiation under Time Constraints, *Artificial Intelligence*, **75**:297-345.
- Krause, P., S. Ambler, M. Elvang-Gøransson, and J. Fox (1995). A Logic of Argumentation for Reasoning under Uncertainty, *Computational Intelligence*, **11**:113-131.
- Kreps, D. and R. Wilson (1982). Sequential Equilibria, *Econometrica*, **50**:863-894.
- Lâasri, B., H. Lâasri, S. Lander, and v. Lesser (1992). A Generic Model for Intelligent Negotiating Agents, *International Journal of Intelligent and Cooperative Information Systems*, **1**(2):291-317.
- Lander, S. and V. Lesser (1992). Customizing Distributed Search Among Agents with Heterogeneous Knowledge, in *Proceedings of the 1<sup>st</sup> International Conference on Information and Knowledge Management*, November, Baltimore, MD, 335-344.
- Lander, S. and V. Lesser (1993). Understanding the Role of Negotiation in Distributed Search among Heterogeneous Agents, in *Proceedings of the 12<sup>th</sup> International Workshop on DAI*, Hidden Valley, PA.
- Liu, J. S. and K. Sycara (1997). Coordination of Multiple Agents for Production Management, *Annals of Operations Research*, **75**:235-289.
- Luce, R. D. and H. Raiffa (1957). *Games and Decisions*, New York: Wiley.
- Loui, R. (1987). Defeat among Arguments: A System of Defeasible Inference, *Computational Intelligence*, **3**:100-106.
- Maes, P. (1996). Agents that Reduce Work and Information Overload, *Communications of the ACM*, **37**(7):31-40.
- Matos, N., C. Sierra, and N. R. Jennings (1998). Determining Successful Negotiation Strategies: An Evolutionary Approach, in *Proceedings of the 3<sup>rd</sup> International Conference on Multi-Agent Systems (ICMAS-98)*, Paris, France, 182-189.
- Matwin, S., S. Szpakowicz, Z. Koperczak, G. E. Kersten, and W. Michalowski (1989). Negoplan: An Expert System Shell for Negotiation Support, *IEEE Expert*, **4**(4):50-62.
- Nash, J. F. (1950). The Bargaining Problem, *Econometrica*, **28**:155-162.



# References 4

- Noriega, P. and C. Sierra (1996). Towards Layered Dialogical Agents, in *Proceedings of the 3<sup>rd</sup> International Workshop on Agents Theories, Architectures and Languages*, 157-171.
- Osborne, M. J. and A. Rubinstein (1994). *A Course in Game Theory*, Cambridge, MA: MIT Press.
- Parsons, S. and N. R. Jennings (1996). Negotiation through argumentation—A Preliminary Report, in *Proceedings of the International Conference on Multi-Agent Systems*, Kyoto, Japan.
- Parsons, S., C. Sierra, and N. R. Jennings (1998). Agents that Reason and Negotiate by Arguing, *Journal of Logic and Computation*, **8**(3):261-292.
- Prasad, M. V. N., V. R. Lesser, and S. E. Lander (1997). Learning Organizational Roles for Negotiated Search in a Multi-Agent System, *International Journal of Human Computer Studies, Special Issue on Evolution and Learning in Multiagent Systems*.
- Pruitt, D. G. (1981). *Negotiation Behavior*, Academic Press.
- Rao, A. and M. Georgeff (1991). Modeling Rational Agents within a BDI-Architecture, in *Proceedings of the 2<sup>nd</sup> International Conference on Principles of Knowledge Representation and Reasoning*, 473-484.
- Rao, A. and M. Georgeff (1995). BDI Agents: From Theory to Practice, in *Proceedings of the 1<sup>st</sup> International Conference on Multi-Agent Systems*, 312-319.
- Rosenschein, J. S. and G. Zlotkin (1994a). *Rules of Encounter*, Cambridge, MA: The MIT Press.
- Rosenschein, J. S. and G. Zlotkin (1994b). Designing Conventions for Automated Negotiation, *AI Magazine*, **15**(3):29-46.
- Rubinstein, A. (1982). Perfect Equilibrium in a Bargaining Model, *Econometrica*, **50**(1):97-109.
- Sandholm, T. (1993). An Implementation of the Contract Net Protocol Based on Marginal Cost Calculations, in *Proceedings of the 11<sup>th</sup> National Conference on Artificial Intelligence*, Washington, D.C., 256-262.
- Sandholm, T. and V. Lesser (1995a). Issues in Automated Negotiation and Electronic Commerce: Extending the Contract Net Framework, in *Proceedings of the 1<sup>st</sup> International Conference on Multiagent Systems (ICMAS-95)*, San Francisco, CA.
- Sandholm, T. and V. Lesser (1995b). Coalition Formation among Bounded Rational Agents, in *Proceedings of the 14<sup>th</sup> International Joint Conference on Artificial Intelligence (IJCAI-95)*, Montral, Canada.

# References 5

- Shehory, O. and S. Kraus (1995a). Coalition Formation Among Autonomous Agents: Strategies and Complexity, in *From Reactions to Cognition*, Berlin:Springer, 57-72.
- Shehory, O. and S. Kraus (1995b). Task Allocation via Coalition Formation Among Autonomous Agents, in *Proceedings of IJCAI-95*, Montreal, Quebec, Canada, 655-661.
- Shehory, O. and S. Kraus (1996a). A Kernel-Oriented Model for Coalition-Formation in General Environments: Implementation and Results, in *Proceedings of AAAI-96*, Portland, OR, 134-140.
- Shehory, O. and S. Kraus (1996b). Emergent Cooperative Goal-Satisfaction in Large Scale Automated-Agent Systems, in *Proceedings of ECAI-96*, Budapest, Hungary, 544-548.
- Sierra, C., P. Faratin, and N. R. Jennings (1997a). A Service-Oriented Negotiation Model Between Autonomous Agents, in *Proceedings of the 8<sup>th</sup> European Workshop on Modelling Autonomous Agents in a Multi-Agent World*, 17-35.
- Sierra, C., N. R. Jennings, P. Noriega, and S. Parsons (1997b). A Framework for Argumentation-Based Negotiation, in *Proceedings of the 4<sup>th</sup> International Workshop on Agent Theories, Architectures and Languages*, 167-182.
- Smith, R. G. (1979). A Framework for Distributed Problem Solving, in *Proceedings of the Sixth International Joint Conference on Artificial Intelligence*, August, Cambridge, MA, 863-841.
- Smith, R. G. (1980). The Contract Net Protocol: High-Level Communication and Control in a Distributed Problem Solver, *IEEE Transactions on Computer*, **C-29**(12):1104-1113.
- Smith, R. G. and R. Davis (1981). Frameworks for Cooperation in Distributed problem Solving, *IEEE Transactions on Systems, Man, and Cybernetics*, **SMC-11**(1):61-70. (Also published in *Readings in Distributed Artificial Intelligence*, A. h. Bond and L. Gasser (eds.), Morgan Kaufmann, 333-356.)
- Smith, R. G. and R. Davis (1983). Negotiation as a Metaphor for Distributed Problem Solving, *Artificial Intelligence*, **20**:63-109.
- Sycara, K. (1989a). Argumentation: Planning Other Agents' Plans, in Proceedings of the 11<sup>th</sup> International Joint Conference on Artificial Intelligence, 517-523, Detroit, MI.
- Sycara, K. (1989b). Multi-Agent Compromise via Negotiation, in *Distributed Artificial Intelligence*, Gasser, L. and M. Huhns (eds.), Los Altos, CA: Morgan Kaufman.

# References 6

- Sycara, K. (1990a). Persuasive Argumentation in Negotiation, *Theory and Decision*, **28**(3):203-242.
- Sycara, K. (1990b). Negotiation Planning: An AI Approach, *European Journal of Operational Research*, **46**(2):216-234.
- Sycara, K. (1992). The PERSUADER, in *The Encyclopedia of Artificial Intelligence*, Shapiro, D. (ed.), New York, NY: John Wiley and Sons.
- Zeng, D. and K. Sycara (1996). Bayesian Learning in Negotiation, in *Working Notes of the AAAI Spring Symposium on Adaptation, Co-Evolution and Learning in Multiagent Systems*, Stanford, CA.
- Zeng, D. and K. Sycara (1997). Benefits of Learning in Negotiation, in *Proceedings of AAAI-97*.
- Zeng, D. and K. Sycara (1998). Bayesian Learning in Negotiation, *International Journal of Human-Computer Studies*, **48**:125-141.
- Walton, D. N. and E. C. W. Krabbe (1995). *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning*, Albany, NY: State University of New York Press.
- Zlotkin, G. and J. S. Rosenschein (1989). Negotiation and Task Sharing Among Autonomous Agents in Cooperative Domains, in *Proceedings of the Eleventh International Joint Conference of Artificial Intelligence (IJCAI)*, Detroit, MI, August, 912-917.
- Zlotkin, G. and J. S. Rosenschein (1991). Cooperation and Conflict Resolution via negotiation among Autonomous Agents in Noncooperative Domains, *IEEE Transactions on Systems, Man, and Cybernetics, Special Issue on Distributed Artificial Intelligence*, **21**(6):1317-1324.
- Zlotkin, G. and J. S. Rosenschein (1996a). Mechanism Design for Automated Negotiation, and Its Application to Task Oriented Domains, *Artificial Intelligence*, **86**(2):195-244.
- Zlotkin, G. and J. S. Rosenschein (1996b). Mechanisms for Automated Negotiation in State Oriented Domains, *Journal of Artificial Intelligence Research*, **5**:163-238.
- Zlotkin, G. and J. S. Rosenschein (1996c). Compromise in Negotiation: Exploiting Worth Functions Over States, *Artificial Intelligence*, **84**(1-2):151-176.