

Viewing Design as a Cooperative Task

Katia P. Sycara

The Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15213
katia@cs.cmu.edu

ABSTRACT

Design can be modeled as a multi-agent planning task where several agents that possess different expertise and evaluation criteria cooperate to produce a design. The differences may result in conflicts that have to be resolved during design. The process by which conflict resolution is achieved is negotiation. In this paper, we propose a model of group problem solving among cooperating experts that supports negotiation. The model incorporates accessing information from a case memory of existing designs, communication of design rationale, evaluation and critiquing of design decisions. Incremental design modifications are performed based on constraint relaxation and comparison of utilities.

INTRODUCTION

The design task can be described as taking a set of functional specifications and constraints, and producing an artifact representation whose behavior, when manufactured, conforms to the given specification description. Design constraints involve not only physical laws and domain principles that determine the behavior of the device but also restrictions and interactions arising from concerns such as cost, ease of manufacturing, ease of assembly and ease of maintaining the artifact. Taking these concerns into consideration in the initial design stages is known as concurrent engineering [DICE 89, Sriram 88]². Concurrent engineering is a team effort that requires the collaboration of teams of specialists representing relevant perspectives. Typically, each specialist has expertise in one area pertaining to the design, limited knowledge of the constraints and intentions of the other specialists and in general different evaluation criteria of the design. Hence, inconsistencies and conflicting views may occur. The specialists interact in a cooperative manner in order to effect tradeoffs and resolve inconsistencies. The final design is a compromise that has been agreed upon by all concerned agents³. Existing approaches to modeling concurrent design have primarily focused on investigating architectures for communication between various experts [DICE 89, Lander 88, Talukdar 88], or on conflict detection [Robinson 87, Sriram 88]. In this paper, we present a model of the group problem solving process of a team of concurrent engineering cooperating experts. The proposed team design model is inspired by our work in an adversarial domain, namely conflict resolution in labor management disputes [Sycara 87]. The model is currently being implemented in the CADET system that integrates Case-Based Reasoning, Qualitative Reasoning, and Constraint Propagation in the domain of mechanical design [Sycara 89a, Sycara 89b].

¹This research has been supported by DARPA and AFOSR under contract number F49620-90-C-0003.

²It has been recently recognized that the practice of concurrent engineering is very advantageous in terms of reducing the time of a product from "concept to market". Industry is rapidly moving to adopt the practice.

³We use the words "agent", "expert" and "specialist" interchangeably for the purposes of this paper.

The team design problem has the following characteristics:

- The global goal is to produce a design that is synthesized from contributions of different expertise, concerns and constraints
- During the design process, conflicts in the form of constraint violations could arise. If these conflicts are not resolved in a satisfactory manner, infeasible designs will occur.
- Disparate evaluations of (partial or complete) designs could surface as a result of different criteria used to evaluate designs from different perspectives. Typically, these criteria cannot be simultaneously and optimally satisfied. The design decisions that optimize one set of criteria could conflict with those that optimize another set. If these conflicts do not get resolved in a satisfactory fashion, design suboptimalities occur.
- The global goal is achieved by making the best tradeoffs on conflicting design goals and constraints.
- Because of the presence of conflicting constraints, goals and possibly evaluation criteria, it is impossible for each expert to optimize the overall design using only local information.
- Backtracking, resulting from infeasible designs, can be a major problem since it may result in invalidating design decisions that other agents have made.

As a result of the above characteristics, the final successful design can be viewed as a compromise⁴. that incorporates tradeoffs such as cost, ease of manufacturing and assembly, reliability and maintainability. The process through which design decisions are made by a team of experts is *negotiation*. Typically in manufacturing enterprises (e.g., [Bond 89]), after initial study of the design specifications, an initial design (called an initial cartoon) is created by the main designer. Subsequently, the specialists meet to negotiate proposed changes and tradeoffs with respect to the initial design. Compromises are suggested and discussed. The suggested compromises initiate further studies that necessitate additional meetings to reconcile continuing problem areas. This iterative process continues until all parties reach agreement⁵. Depending on particular decisions concerning tradeoffs, different designs will be produced. For example, the valve for a water tap could be a metallic threaded part or a plastic plug valve with a hole. There is a tradeoff between the low cost of the plastic valve and the high durability of the metal valve. Despite the difficulty of applying negotiation techniques, recorded conflicts and their resolution, namely previous similar design cases, provide a foundation for rationalizing designs.

Negotiation enters the design process at the following points:

- When the result of design decisions is an infeasible design (i.e. when constraint violations have been identified).
- When a design is feasible but suboptimal.
- When alternate approaches can achieve similar functional results.

⁴Pruitt [Pruitt 81] has identified two types of negotiation which are used by expert human negotiators to seek acceptable solutions and which may be applicable to machine agents: (a) compromise negotiation where each party makes concessions on its demands to facilitate agreement, and (b) integrative negotiation where the most important goals of each party are used to form innovative solutions, relinquishing, if necessary, secondary goals. In our view, both these negotiation types result in *compromise* solutions, in other words in partial goal satisfaction. Moreover, in typical negotiations a goal of secondary importance to one agent could be of primary importance to another because of different local evaluation functions. My use of the word "compromise" encompasses both these negotiation types.

⁵Throughout this process, time and cost determine the number of iterations allowed.

Negotiation is a process in which the parties iteratively exchange proposals and proposal justifications until an agreement is reached. During the negotiation process, the feasibility and desirability of proposed tradeoffs is evaluated and may result in incremental design adaptations. Thus, a negotiation model must be (a) iterative, (b) include mechanisms to incorporate feedback of the parties concerning evaluation of a proposal (partial design) from their point of view, (c) include criteria for judging whether progress in the negotiation is being made, and (d) incorporate negotiation protocols for the exchange of proposals, arguments and justifications of the proposed design decisions. The proposed model incorporates the above characteristics.

AN EXAMPLE OF EXPERT INTERACTION

Consider the process of designing a turbine blade. Some of the dominant specialties are aerodynamics, structural engineering, manufacturing and marketing. The blade design team operates within constraint ranges specified by the design team for the aircraft engine. The blade design team incorporates various concerns. The concern of aerodynamics is aerodynamic efficiency; for structural engineering it is reliability and safety; for manufacturing, it is ease and cost of manufacturing and testing; for marketing it is overall cost and customer satisfaction. The two variables of concern in a turbine blade that we consider are: (a) root radius, and (b) blade length. From the perspective of structural design, the bigger the root radius, the better since it decreases stress concentration. From the perspective of aerodynamics, the smaller the root radius, the better, since it increases aerodynamic efficiency. Concerning the length of the blade, from the point of view of structural design, the shorter the blade, the lower the tensile stresses; from the point of view of aerodynamics the longer the blade, the better the aerodynamics. On the other hand, if the blade is shorter, it makes for a lighter engine which is a desirable characteristic for aerodynamic efficiency. Thus, we see that the aerodynamics expert needs to make tradeoffs internal to its perspective. From the point of view of marketing, aerodynamic efficiency lowers the cost of operation of the aircraft, thus making it more attractive to customers. From the point of view of manufacturing, it is easier to manufacture shorter blades with bigger root radii.

The following is a simplified example dialogue of the various concerned perspectives in an attempt to arrive at a mutually satisfactory turbine blade design:

Aerodynamics (using case based reasoning) suggests particular values x and y for length and root radius of the blade. The suggested x and y values are within acceptable constraint ranges.

Structural engineering evaluates these values from its point of view and suggests values x' and y' where $x' < x$ and $y' > y$ (i.e., shortening the length and increasing the root radius) to increase safety.

Aerodynamics counters by saying that the values structural engineering suggested would considerably decrease aerodynamic efficiency.

Structural engineering counters that shorter blade makes engine lighter, thus also increasing efficiency.

Marketing says aerodynamic efficiency sells the product since it is less costly to operate.

Manufacturing supports structural by saying it is easier to manufacture short blades with big root radius.

Aerodynamics suggests that the materials engineering expert could try to investigate new materials that make the blade lighter, thus alleviating weight considerations.

Manufacturing says that new materials take lots of time to test and debug.

Structural engineering adds that new materials may introduce safety hazards that could go undetected.

This example illustrates the exchange of proposals for values of length and radius of the blade as well as the exchange of arguments and justifications in critiquing the proposed values. The nature of the resulting design will depend on (a) the ranges of various artifact constraints, (b) which constraints can be relaxed and in what ways, (c) the relative importance of various artifact-dependent goals (e.g., aerodynamic efficiency), (d) relative importance of various artifact-independent goals (e.g., safety), and (e) the way in which particular variable values contribute to the achievement of the goals.

THE MODEL

At the start of the team design process, an initial design is generated and presented to each expert, who evaluates it from its own point of view and registers its reactions (evaluations, objections and suggestions). At each negotiation iteration, the input is the set of conflicting concerns, violated constraints of the various design agents and the context of the design (e.g., constraints that have been handed to the design team from others). The final output is either a single agreed upon design or an indication of failure if the negotiating agents did not reach agreement within a particular number of iterations. The final output is reached through iterations of the following tasks: (a) proposal of an initial design, (b) arguments to support justification and critiquing of the design and (c) incremental modification and improvement. These tasks are performed using knowledge of existing designs and their characteristics, knowledge of physical laws and constraints, traces of design decisions made so far, and models of the expertise and concerns of the participating agents. The agents have access to a common case base of previous designs but because of their specialized knowledge each one accesses and evaluates a design from a particular *view*.

Negotiation, as seen in Figure 1, is performed through integration of Case-Based Reasoning (CBR) [Kolodner et al. 85, Sycara 87], use of multi-attribute utilities, called Preference Analysis, and constraint relaxation. These methods are employed in all negotiation tasks, namely in generation of an initial proposal, repair of a rejected proposal to formulate a counterproposal, and communication of justifications and objections. The process interleaves local computation and communication of computation results to other agents. One of the important concerns in distributed problem solving is minimization of communication overhead.

Use of previous design cases offers a reasoner (a) suggestions of how tradeoffs and resolutions have been made in the past, (b) failure avoidance advice (since failures are recorded in the design case), and (c) possible modifications and repairs to design suboptimalities. In addition, for group problem solving, having a memory of past problem solving experiences (successes and failures) has the following advantages [Sycara 89c]:

- Case-based inference minimizes the need for information exchange, thus minimizing communication overhead.
- Anticipating and avoiding problems through reasoning from past failures helps the agents minimize the exchange of proposals that will be rejected, thus minimizing backtracking.
- Reasoning from previous cases helps in recognizing problems or opportunities in a proposed design.
- If the repair of a past failure is also stored in memory, computation by each agent is




Figure 1: The Negotiation Process

minimized.

An integral part of the negotiation process is the ability of the agents to communicate their views and preferences as well as influence the decisions of other agents. Thus, a group problem solving model must have the ability to support (a) representing and maintaining belief models, (b) reasoning about agents' beliefs, and (c) influencing other agents' beliefs. The knowledge needed to perform these tasks is an agent's *belief* and *preference* structure. The belief structure of an agent, represented in a goal graph, consists of a collection of goals, goal importance and relationships among goals. The preference structure of an agent records its utilities associated with various potential decisions. The values in the goal graph are based on the constraints on the design and company policy. For example, after an airplane crash whose cause is traced to a defective engine, the importance of the safety goal increases for the marketing agent⁶.

We represent an agent's belief structure as a directed acyclic graph where each node represents an agent's goal. Edges of the graph linking two goals represent the relationship between goals in terms of how one affects (positively or negatively) the achievement of the other. For example, aerodynamic efficiency positively affects lower operation costs. Associated with each node is:

⁶This has been observed in real situations, such as one of last year's airplane accidents involving an engine made by GE. The company gave increased importance to testing procedures and design practices to increase engine reliability.

- a *sign* (+ or -) that denotes the desirability of an increase or decrease in that goal
- the *values* by which the attribute/goal should be increased or decreased
- the *importance* that the agent attaches to the goal
- the *feasibility* as perceived by the agent of achieving the goal

Directed edges connect subgoals to the higher level goals to which they contribute. A *contribution value* is associated with each directed edge denoting the contribution of the subgoal to the higher level goal. Contribution values range from -100% to +100%. A positive value means that the subgoal supports the achievement of the higher level goal by the denoted percentage. A negative contribution value has the interpretation that the subgoal is detrimental to the higher level goal. Sink nodes⁷ are the highest level goals of an agent.

A path from node X to node Y in a goal graph constitutes a causal/justification chain that provides an explanation of the change in Y in terms of the change in X, assuming no other change has occurred in the rest of the graph. For example, lengthening a turbine blade results in increasing aerodynamic efficiency, which in turn results in lower operating costs, thus resulting in increased marketability of the blade. By traversing goal graphs a reasoner can answer the following queries:

- Which goals are supported by a set of design decisions?
- Which design decisions are justified by a set of goals?

In addition to an agent's beliefs, the representation includes an estimate of its utilities for each attribute in the goal graph. Utilities express the *preference structure* of an agent [Sycara 88]. Moreover, utilities express the tradeoff structure among various attribute values associated with alternative designs. The (possibly nonlinear) utilities of individual attributes are combined to give an overall utility, the payoff, of an alternative. Being able to compare different alternatives enables a reasoner to choose the alternative that affords the maximum payoff. An integration algorithm traverses the belief structure to determine which way goal values should be moved to increase payoff and thus the acceptability of a resolution. Moreover, goal graph traversal allows an agent to discover alternative design decisions that support important goals thus leading to innovative designs.

THE NEGOTIATION PROTOCOL

The agents interact through message passing. The messages that the negotiating agents exchange contain the following information:

- The proposed design
- Justifications of design decisions
- Agreement or disagreement with the proposal
- Requests for additional information, such as with which issue in the proposed design the agent disagrees.
- Reasons for disagreement.
- Utilities/preferences of the agents associated with disagreed upon issues.

⁷Sink nodes have no out-going edges

We present in detail the communication protocol used in our system⁸.

1. Agent1 communicates to agent2 a design proposal, as well as arguments and justifications in support of the proposal.
2. Agent2 uses the arguments and justifications communicated by agent1 to possibly modify its goal graph (e.g., change importance of goals, including possibly abandoning goals).
3. Agent2 evaluates the proposal from its point of view (using its constraints and utilities).
4. If the proposal satisfies agent2's local constraints and gives it payoff above a threshold, it communicates ACCEPT to agent1.
5. If not, agent2 generates a counterproposal by whatever problem solving means it has at its disposal (e.g., CBR, constraint relaxation).
6. Agent2 evaluates the counterproposal. If the counterproposal gives agent2 payoff above the threshold, agent2 communicates to agent1:
 - The PORTION/ISSUES of the proposal that have been modified
 - The REASON for modifying the previous proposal (e.g., value1 violates some of the agent2's hard constraints, a set of proposed values does not contribute enough to higher level goals of agent2).
 - The COUNTERPROPOSAL and its PAYOFF.
 - ARGUMENTS and JUSTIFICATIONS in favor of the counterproposal.
7. If the counterproposal does not give agent2 payoff above the threshold, agent2 goes to step 5.
8. If agent2 has exhausted all counterproposals it can generate through the methods of step 5, it traverses its goal graph to see whether there is another way to satisfy its higher level goals.
 - If there is, it generates a counterproposal and goes to step 6.
 - If there is not, it communicates FAILURE to agent1 (who now has to generate a modification and/or look for alternative ways in *its* goal graph).

CONCLUDING REMARKS

Design can be viewed as a multi-agent planning process involving multiple conjunctive and potentially conflicting goals. The agents have different expertise (e.g., mechanisms, hydraulics, assembly, testing) which results in viewing and evaluating designs using different, possibly conflicting criteria. The process of negotiation is used to propose and examine design decisions involving various tradeoffs. Negotiation is performed recursively at all stages of design and involves different design teams. We have proposed a model of group problem solving by cooperating experts that supports negotiation. The model is based on (a) knowledge of previous designs, (b) communication of design rationale, justifications and objections to proposed design decisions, (c) constraint propagation and relaxation, and (d) traversal of goal graphs.

⁸For simplicity, the protocol is presented for two agents, agent1, who initiates an initial design and agent2, who evaluates the design and possibly generates a counterproposal. The protocol generalizes to more than one agent that evaluates and suggests modifications.

REFERENCES

- [Bond 89] Bond, A.H., and Ricci, R.J., "Cooperation in Aircraft Design," *Proceedings of the MIT/JSME Workshop on Cooperative Product Development*, MIT, Cambridge, Mass., November 1989.
- [DICE 89] DICE: Initiative in Concurrent Engineering, "Red Book of Functional Specifications for the DICE Architecture," Tech. report, Concurrent Engineering Research Center, West Virginia University, February 1989.
- [Kolodner et al. 85] Kolodner, J.L., Simpson, R.L., and Sycara-Cyranski, K., "A Process Model of Case-Based Reasoning in Problem Solving," *Proceedings of IJCAI-85*, Los Angeles, CA, 1985, pp. 284-290.
- [Lander 88] Lander, S. and Lesser, V., "Negotiation to Resolve Conflicts Among Design Experts," *Proceedings of the AAAI-88 Workshop on AI in Design*, AAAI, St. Paul, MN., 1988.
- [Pruitt 81] Pruitt, D. G., *Negotiation Behavior*, Academic Press, New York, N.Y., 1981.
- [Robinson 87] Robinson, W., "Towards the formalization of specification design, Master's Thesis," Tech. report, University of Oregon, 1987.
- [Sriram 88] Sriram, D., Logcher, R., and Groleau, N., "Cooperative Engineering Design," *Proceedings of the AAAI-88 Workshop on AI in Design*, AAAI, St. Paul, MN., 1988.
- [Sycara 87] Sycara, K., *Resolving Adversarial Conflicts: An Approach Integrating Case-Based and Analytic Methods*, PhD dissertation, School of Information and Computer Science Georgia Institute of Technology, 1987.
- [Sycara 88] Sycara, K., "Utility Theory in Conflict Resolution," *Annals of Operations Research*, Vol. 12, 1988, pp. 65-84.
- [Sycara 89a] Sycara, K. and Navinchandra, D., "Integrating Case-Based Reasoning and Qualitative Reasoning in Engineering Design," in *Artificial Intelligence in Engineering Design*, J. Gero, ed., Computational Mechanics Publications, 1989.
- [Sycara 89b] Sycara, K. and Navinchandra D., "A Process Model of Experience-Based Design," *Proceedings of the Eleventh Annual Conference of the Cognitive Science Society*, Ann Arbor, MI., 1989.
- [Sycara 89c] Sycara, K., "Multi-Agent Compromise via Negotiation," in *Distributed Artificial Intelligence, Volume II*, M. Huhns and L. Gasser, ed., Pitman Publishing Ltd and Morgan Kaufmann, 1989.
- [Talukdar 88] Talukdar, S., Elfes, A., and Papanikolopoulos, N., "Concurrent design, simultaneous engineering and distributed problem solving," *Proceedings of the AAAI-88 Workshop on AI in Design*, AAAI, St. Paul, MN., 1988.