A Flexible Reasoning Mechanism for the Trade-off of System Versus Local goals in Sensible Agents

K. S. Barber **T. H. Liu** **A. Goel** **S. Ramaswamy**
The Laboratory for Intelligent Processes and Systems
Electrical and Computer Engineering
The University of Texas at Austin
phone: (512) 471-6152
fax: (512) 471-3652

Abstract

Planning and control systems for highly dynamic and uncertain manufacturing environments require adaptive flexibility and decision-making capabilities. Modern distributed manufacturing systems assess the utility of planning and executing solutions for both system goals (e.g. minimize manufacturing production time for all parts or minimize WIP), and local goals (expedite part *A* production schedule or maximize machine *X* utilization). In this paper, an approach based on Sensible Agents, which dynamically alter their autonomy levels to use a variety of decision models (e.g. Expected Utility, Prospect, or Social Judgement), is presented. Sensible Agents provide for trade-off reasoning among system and local utilities that is flexible and responsive to an agent's abilities, situational context, and position in the organizational structure of the system.

Content Areas: multi-agent system, distributed manufacturing systems, distributed scheduling
Word Count: 1646
Acknowledgment: This research was supported in part by the Texas Higher Education Coordinating Board (#003658415).

Accepted by: CAICIM, Rensselaer's International Conference on Agile, Intelligent, and Computer-Integrated Manufacturing

October 7-9, Rensselaer Polytechnic Institute, New York

A Flexible Reasoning Mechanism for the Trade-off of System Versus Local goals in Sensible Agents

K. S. Barber T. H. Liu A. Goel S. Ramaswamy
The Laboratory for Intelligent Processes and Systems
Electrical and Computer Engineering
The University of Texas at Austin
Austin, TX 78712
barber@mail.utexas.edu

Abstract

Planning and control systems for highly dynamic and uncertain manufacturing environments require adaptive flexibility and decision-making capabilities. Modern distributed manufacturing systems assess the utility of planning and executing solutions for both system goals (e.g. minimize manufacturing production time for all parts or minimize WIP), and local goals (expedite part *A* production schedule or maximize machine *X* utilization). In this paper, an approach based on Sensible Agents, which dynamically alter their autonomy levels to use a variety of decision models (e.g. Expected Utility, Prospect, or Social Judgement), is presented. Sensible Agents provide for trade-off reasoning mechanisms among system and local utilities that is flexible and responsive to an agent's abilities, situational context, and position in the organizational structure of the system.

1. Introduction

Since manufacturing environments are inherently complex and dynamic, it is challenging to automate manufacturing tasks such as planning or scheduling. The use of agent-based systems offers significant advantages for automated manufacturing, including distribution of control and processing as well as adaptable automated or semi-automated problem solving. Since manufacturing domains require high flexibility and adaptability, we propose the Sensible Agent model (Barber, 1996) to promote dynamic forms of agent organization in response to system and local goal trade-offs, changing situational context, and changing agent resources.

Sensible Agents are goal-driven autonomous agents that operate on sets of goals with different autonomy levels associated with each goal. Three autonomy levels have been identified for Sensible Agents: ***(i)*** **command driven**- the agent does not plan and must obey orders given by another (master) agent, ***(ii)*** **consensus**- the agent works as a team member, sharing planning tasks with other agents, and ***(iii)*** **locally autonomous/ master**- the agent plans alone and may or

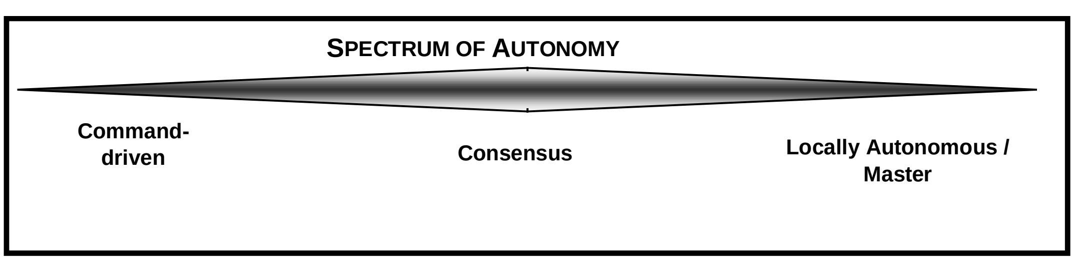

Figure 1: Three levels of autonomy levels

may not give orders to other agents. Figure 1 illustrates these autonomy levels in a spectrum (Martin and Barber, 1996)

One of the main issues in multi-agent system development is the need to balance system goals against local goals. The functional objectives of a system are the system goals for each agent. Local goals are perceived through an agent's reasoning activities based on local perspectives. As pointed out in (Gasser, 1991;Jennings, 1996), it is not practical (due to high communication load) to maintain a consistent global view for every agent. Agents should behave based on their local views and proper coordination mechanisms are required. Therefore, agents' local goals perceived through local views may conflict with each other and even with system goals due to the localized interpretations used by each agent.

Agents with local perspectives may need to pursue actions to fulfill local goals even if these actions may conflict with system goals, degrade system performance, or interfere with other agents' activities. Each agent in a system of Sensible Agents may or may not possess the same system goals. Each agent also owns its respective local goals. Local goals of one agent may be shared by others depending on the autonomy level of the agent(s) for their respective goals. A group of agents working to reach consensus on a plan for goal *A* must all share goal *A*. For example, in job shop scheduling, a system goal may be to perform line balancing and a local goal may be to increase utilization of a local resource to reduce the average cost of operation. Conflicts (i.e. number of jobs per line) may occur between these two goals even though they exist within a single agent. Similar problems could occur in many modern manufacturing lines; an automated manufacturing system must understand how to balance local needs against system needs.

This paper will first introduce some background research on trade-off reasoning in section II. Our approach based on the Sensible Agent Architecture will be discussed in section III with an example in Section IV. A short discussion and our conclusions are presented in Section V.

2. Trade-off Reasoning

A significant amount of research has investigated managing or reasoning about the trade-offs between multiple objectives [Vincke, 1989 #582](Vincke et al., 1989;Yu, 1985). Various theories are proposed to model or aid in the decision making process. For manufacturing environments in general, alternative solutions may not be enumerated and high uncertainty may be involved. The set of objects, decisions, or candidates to be explored during the decision procedure can be called action set *A*. This action set may be stable (defined as a priori and not open to change in the course of the procedure) or evolving (modified during the procedure due to either intermediary results which appear during the process or the changes of the natural environments). As pointed out by Vincke, the difficulty of such a multi-criteria problem lies "…in the fact that it is an ill-defined mathematical problem, i.e., it has no objective solution. There is generally no action which is better than all the others for all criteria considered simultaneously." For example, a ranking problem will result in an objective solution only if all the separate criteria considered each yield the same ranking. Therefore, it is different from classic optimization problems (which search for some kind of hidden truth or objective best solution). Such trade-off reasoning usually results in compromised solutions, which is highly dependent on the circumstances, methods, and preferences of decision-makers (Vincke et al., 1989).

Trade-off reasoning between system goals and local goals is a special case of multi-criteria decision making. The executors (agents) of system goals hold their own local goals which are generated during the planning process. Ideally, if such two kind of goals are consistent, the executors can function well and reach good satisfaction. When system goals and local goals result in different and conflicting plans, trade-off reasoning mechanisms are required to manage these inconsistencies. Various decision making approaches are available, some may favor system goals and the other may favor local goals. Roughly, trade-off reasoning between system goal and local goals can be classified into the following major approaches: ***(i)*** utility-based dynamic decision making, and ***(ii)*** ranking relations (heuristic priority rules).

Multiple Attribute Utility Theory

Multiple attribute utility theory (MAUT) (Hill et al., 1982;Keeney and Raiffa, 1976;Sycara, 1988) is based on the fundamental axiom that a decision-maker attempts to maximize some utility function $U = U(g_1, g_2, ...)$ which aggregates all the different viewpoints currently taken into account. Each parameter, g_i, represents some estimated value for a specific attribute. Such aggregation into a single numerical measure allows classic optimization algorithms to be applied to multi-criterion problems. This requires identification, evaluation and comparison of alternative solutions before the best solution is selected. Various decision models which implement utilities in different fashions have been proposed, including Expected Utility, Prospect, and Social Judgement Theory (Byrnes, 1998).

Expected utility (EU) relies on two parameters to support decision-making: 1) expectations, a belief about the likeliehood of the outcome, 2) values, a judgement about the desirability of that outcome. Expectations are implemented as probabilities. These expectations help the decision makers to select the best choice, which is the optimal combination of likelihood and desirability (Keeney and Raiffa, 1976).

Prospect theory is designed to modify EU to account for the following human decision behavior that EU cannot model: 1) the certainty effect, the tendency to prefer a sure thing over a risky outcome of equal expected value, 2) the reflection effect, the tendency for people to reverse their preferences when gains are replaced by loses, i.e. people prefers the riskier option (low probability of loses but may lose more) among equal expected value, 3) the isolation effect, the tendency to disregard common elements of option pairs and focus on elements that differentiate options. The overall value of an option is given by a function that combines a factor that represents the subjective worth of the outcome with a factor that weights outcomes in terms of importance. One major weakness of Prospect theory has been identified, however. Prospect theory may overweigh low probability actions, e.g. the tendency of some people to buy lottery tickets against overwhelming odds.

Social judgement theory proposes that there is a social cue guiding decision makers to make decision. There is causal structure to events in the world that adaptive individuals need to comprehend. Successful adoption requires that a person's cognitive representation of the world matches the causal structure of the world as close as possible. Such social cues serve as a form of signals about underlying causal structure, theoretically, effective use of cue can help someone attain desirable outcomes. Social judgement theorists try to construct regression equations that describe the cues and the associate weights. For example, people try to share their "cues" during voting process (Byrnes, 1998).

Ranking relations

Ranking relations differs from utility calculation approach by the amount of information provided by the analysis result (Roy, 1976). The large amount of information contained in the result of utility calculation is due to the theory's strong assumption (e.g. there exist a utility function, every attribute can be estimated numerically, and utility value of each attribute can be compared with each other) and to all the extra information demanded from the decision-maker (e.g. preference intensities/weight factor for each attribute). Using ranking relations can reduce the amount of information needed to make a decision. For a binary relation between two alternatives *a* and *b*, if there are enough arguments to conclude that *a* is better than *b*, and there are no reasons to refute that statement, there is no need to analyze *a* to any greater detail. Such relations represent the decision maker's established preferences and are not necessarily complete or transitive. Heuristic rules are easy to implement using this approach and virtual attributes like priorities and certainty factors can be used to build ranking relations and can be assigned either statically or dynamically. For example, in (Liu and Sycara, 1996) robots/agents follow a set of heuristic rules to make decisions selecting appropriate conflict resolution strategies to eliminate conflicts.

Preference modeling provides a formal representation for an agents current world model. In general, preference is modeled as a set of binary relations which has three categories: preference, indifference, and incomparability. We write:

$a\mathrm{P}b$ if *a* is preferred than *b*

$a\mathrm{I}b$ if *a* and *b* are indifference

$a\mathrm{J}b$ if *a* and *b* are incomparability

Heuristic relation rules and utility-based dynamic decision making can be used together as well. For example, researchers have proposed a three-stage heuristic reduction process to maximize system resource for multi-agent planning (Ephrati et al., 1995). The three stages include: ***(i)*** transformation from local to global utility measures, ***(ii)*** a global system assessment of the local agent's evaluations, and ***(iii)*** approximation algorithms. When global utility differs radically from individual agent utilities, agents can still maintain their commitments for planning.

3. Sensible Agent Architecture

The Sensible Agent Architecture (Figure 2), consists of four modules: ***(i)*** **Action Planner-** solves domain problems, and executes problem solution, ***(ii)*** **Conflict Resolution Advisor-** identifies, classifies and recommends proper strategies to resolve conflict between agents, ***(iii)*** **Autonomy Reasoner-** evaluates and assigns autonomy levels, and ***(iv)*** **Perspective Modeler-** maintains models (behavioral, goals and resources) of the agent's view of the environment, other agents and itself.

Autonomy levels are specified by the following four constructs: ***(i)*** **responsibility-** a measure of how much the agent must plan for a goal, ***(ii)*** **authority-** a measure of the agent's ability to access system resources in pursuit of a goal, ***(iii)*** **commitment-** a measure of the extent to which a goal must be achieved, and ***(iv)*** **independence-** a measure of how freely the agent can plan for a goal. Since each goal is assigned its own autonomy level, the autonomy constructs govern the interaction style that exists between agents with respect to their achievement of a specific goal.

Sensible Agents are goal-driven entities. As a result, it is necessary to formalize goals and goal hierarchies and their respective goal utilities. These issues are dealt with in more detail in (Goel et al., 1996;Martin and Barber, 1996). The following is a short description for the calculation of utilities for both system and local goals as well as the trade-offs between those goals.

Let A^k denote the set of agents in the system. Let g^k_j denote the j^{th} goal of A^k. The autonomy level of g^k_j is represented by a 4-tuple (R, C, a, i) where R is the responsibility distribution, C is the commitment index, a is the authority construct, and i is the independence index. For each autonomy construct, we define a function that measures the autonomy level impact on system and local utility for the goal to which it is assigned. For example, f_C is a function that returns a tuple $(f_{C,agent}, f_{C,system})_j$ using C for g^k_j, where $f_{C,agent}$ and $f_{C,system}$ denote the impact of the commitment index on the agent and the system, respectively (Goel et al., 1996). As shown in the utility evaluation and planning capabilities are encapsulated within the Action Planner.

The inherent system and local utilities of each goal an agent intends to achieve are altered by the set of functions described above. For example, an agent that is highly committed to manufacturing a part within a specific time period may increase system utility since other agents are able to plan with a higher degree of certainty concerning future events. Conversely, this may decrease local utility since the commitment similarly constrains the options an agent may have if this goal is in conflict with other goals.

The independence construct specifies the relative importance of system vs. local utility for a specific goal. An agent with a highly independent goal may be more prone to ignore system utility in favor of local utility. Thus, the domain-specific behavior of an agent on a manufacturing floor is abstracted away from the actual goals that an agent may be trying to accomplish. Goals can be added, deleted, or altered without changing the basic architecture of either the agent or the system.

The ability to define autonomy levels provides: ***(i)*** a mechanism to alter the values of system and local utility to reflect the disposition of the agent, ***(ii)*** a mechanism to control an agent's actions relative to those utility values, and ***(iii)*** a framework for supporting several decision making models based on situational context and organization.

4. Decision models associated with different autonomy levels

Under different autonomy levels, Sensible Agents behave differently due to the different interaction frameworks associated with each autonomy level. These different frameworks can change the appropriate decision model required for planning. Expected Utility models have already been implemented in multi-agent systems to assist in decision-making (Rosenschein and Zlotkin, 1994;Sycara, 1988).

Since Sensible Agents can dynamically alter their autonomy levels, this implicitly assumes that they have the ability to choose the proper decision model for a given situational context. To reason about the trade-offs among system and local goals, Sensible Agents should consider their roles in the organizational structure (represented as autonomy level and assigned by the Autonomy Reasoner). The autonomy level assigned to the goals an agent holds is the key to implementing the different decision models. These different models can be deployed for: ***(i)*** calculating system and local utility, and ***(ii)*** planning actions based on the calculated utilities.

At the command-driven level, Sensible Agents are not required to analyze all possible alternatives for decision making. Hence, the heuristic relations approach (priority rules) provided by their masters can be used to handle the trade-off reasoning among system goals and local goals. This reduces computing cost and the information required to make decisions – it is likely that command-driven agents do not have the information required and the ability to do a detailed utility analysis. For example, if a system goal has the highest priority, an agent could pursue the system goal without considering the impact of its actions on local goals. It can assume that the trade-off reasoning has already been done by its master agent.

For master or locally autonomous levels, Sensible Agents have to implement utility-based calculations. Either the Expected Utility model or Prospect Theory can be implemented based on domain characteristics. The large amount of information contained in the utility calculations

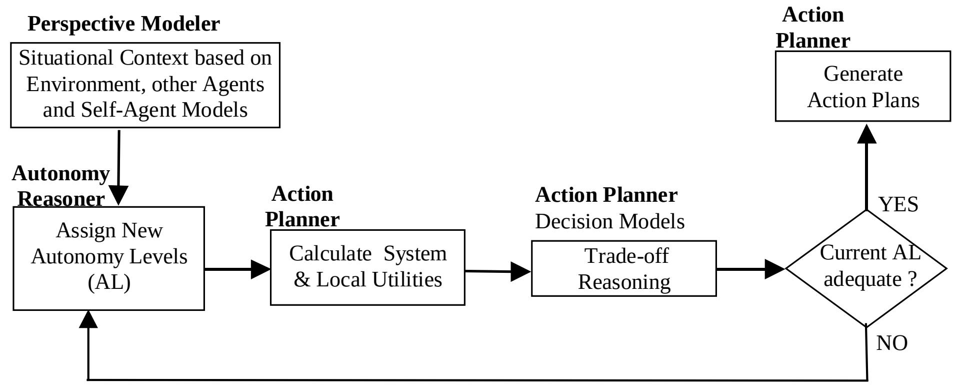

Figure 3: Decision Making Process for System versus Local Goals in Sensible Agents

may be useful to other agents within the system. For example, a master agent may use utility calculations to provide a set of priority rules governing the actions and choices of one of its command-driven agents. For example, to provide command-driven agents priority rules among different alternatives.

At consensus level, Sensible Agents can apply Social Judgement theory as their decision model. Agents within a consensus group may have the same social cues governing their interactions. These cues help agents in their decision making and implicitly coordinate their behaviors at the same time. For example, traffic laws in a multi-robot environments help robots to decide their actions when their local goals (e.g. take a shortest path from point *A* to point *B*) have certain conflict with system goals (e.g. maintain traffic flow).

Figure 3 illustrates the decision making process in Sensible Agents. As mentioned in the previous section, calculation of system and local utilities is a part of the reasoning process. After an autonomy level is assigned, the impact of the autonomy constructs on system and local utilities must be calculated. The Action Planner is then responsible for performing the trade-off reasoning using the appropriate decision model for the current situational context. If the current autonomy level and decision model do not provide adequate planning solutions, then the autonomy level must be recalculated by the Autonomy Reasoner.

The required data for the different decision models can be represented within the autonomy constructs of each goal. Moreover, the information gained during execution regarding the perceived performance of the agent can be used to dynamically alter the utility equations so that a Sensible Agent's capabilities may *improve* over time. Thus, Sensible Agents can use the autonomy constructs to make better decisions than those resulting from the use of a single decision making model.

5. Example Application for Job Shop Scheduling

In this section, we will use job shop scheduling to demonstrate the flexible decision making capabilities of Sensible Agents to reason between system and local goals. We will only show a discussion between a pair of system goal and local goal. A fully specified example is beyond the scope or space of this paper.

The general problem is to schedule production times for N jobs on M machines. Each job requires a sequence of processes (sequence of machines) and associated time with each process is known. The objectives for this problem can include: to finish these jobs before their due date, to minimize the number of tardy jobs, to minimize the average flow time, and to minimize the cost for completion of all jobs (Askin and Standridge, 1993).

Assume that the manufacturing floor has four machines (*w,x,y,z*) and four jobs that need to be scheduled (*A,B,C,D*). Assuming that the current time is 0, Table 1 shows the waiting time, due time and process path for each job with respect to *y*. Each of the four jobs listed are currently waiting for machine *y* in their process sequence. For example, *A* is shown as having already waited for 5 time units and is due at time 10. From the last column, we can see that *A* has already gone through machine *x* and has to go through machines *y* (which requires 3 time units) and *z* (also 3 time units). This table assumes that *y* has just finished its previous job and is ready to select the next one. *y* is required to make a decision about which job to take on next.

We define *y* to be controlled by a Sensible Agent which has the local goal: pick the next job to process. The system goals include: finishing all process within due time constraints, minimizing Work-in-process (WIP), and maximizing the production rate. Table 2 shows a list of common dispatching rules for job selection. We will now outline the decision-making process for the Sensible Agent controlling *y* at each of the three autonomy levels.

Table 1: Available Jobs for Machine *y*

Job	Waiting time	Due time	Machine sequence / time		
A	5	10	*x*, 5	*y*, 3	*z*, 3
B	2	15	*w*, 6	*x*, 4	*y*, 4
C	1	4	*z*, 3	*y*, 2	*w*, 2
D	0	20	*y*, 5	*x*, 1	*w*, 6

Table 2: Dispatching Rules

Name	Description
SPT	Shortest process time
EDD	Earliest due date
FCFS	First come, first served
LTWK	Least total work
LWKR	Least work remaining
MOPNR	Most operation/process remaining
MWKR	Most work/time remaining
WINQ	Subsequent machine has shortest queue
RANDOM	Random

Command-Driven

When Sensible Agents work at a command-driven level, heuristic relations approach (priority rules) are provided by their masters for decision making. This assumes that *y* is controlled by another agent that is acting as its master for the goal: pick the next job to process. For ranking relation approach, there are more than ten methods proposed for building relations. In this example, we just use the simplest one (priority rules) to illustrate our point. Detailed representations and discussion can be found in (Vincke et al., 1989). The command driven agent is given rule FCFS as a priority rule for determining how to fulfill its local goal. FCFS corresponds to a request to minimize the waiting time of queued jobs. Using FCFS, *A* has the longest waiting time and would be chosen as the next job.

However, the Sensible Agent must also make sure that its choice is consistent with its system goals (defined above). For this example, we will only use the system goal: finish all processes within due time constraints. Using dispatching rule EDD (Earliest Due Date), the agent would conclude that job C should be selected next. Since the local and system goals produce a conflicting choice, another priority rule must be provided by the master agent for the conflicting local goal. This priority rule must provide a relationship between the two dispatching rules used (FCFS and EDD) in order to allow the agent to make a decision.

Master/Locally Autonomous

As proposed in Section III, master or locally autonomous agents can use utility-based calculations for decision-making. We will use the simplest additive form of the Expected Utility model,

$$U(a) = \sum_{j=1}^{n} U_j(g_j(a)) \qquad \text{Eq (1)}$$

where a is the alternative under comparison, and g_j is the value function for attribute j. Preference among alternatives is represented as $a\mathrm{P}b \Leftrightarrow U(a) - U(b) > 0$.
If given the utility function (based on Covert rule (Askin and Standridge, 1993)),

$$U = \frac{\text{delay cost}}{\text{operation process time}} \qquad \text{Eq (2)}$$

$$\text{delay cost} = \begin{Bmatrix} 1 & \text{if slack} < 0 \\ 0 & \text{if slack} > \text{wait time} \\ 1 - \dfrac{\text{slack}}{\text{wait time}} & \text{otherwise} \end{Bmatrix} \qquad \text{Eq (3)}$$

Table 3 shows the calculations based on this utility function with the current situation. Job C has the highest utility value and is the preferred next job. However, it is obvious that a more

$$\text{slack} = \text{due time} - \text{current time} - \text{remaining process time} \qquad \text{Eq (4)}$$

$$\text{wait time} = 0.5 \times \text{number of remaining process} \qquad \text{Eq (5)}$$

complicated utility function may produce a different ordering for the various jobs. What is important is that the choice of autonomy level has an impact on the decision model used for planning.

Table 3. Utility Calculation

Job	Slack	Wait time	Delay cost	U
A	10-0-3-3 = 4	0.5 x 2 = 1	0	0
B	15-0-4 = 11	0.5 x 1 = 0.5	0	0
C	4-0-2-2 = 0	0.5 x 2 = 1	1	.25
D	20-5-1-6 = 8	0.5 x 3 = 1.5	0	0

Consensus

At consensus level, Sensible Agents can apply Social Judgement theory as their decision model. For example, agents in consensus may follow social decision rules like “select the job whose subsequent machine has the shortest queue” (corresponding to the WINQ dispatching rule (Askin and Standridge, 1993)). If machine y knows that machine w is now processing a job which requires another 5 time units to finish and machine z is currently idle, the selection would be job

A. The system goal of machine *y* is identical to the one specified in the command-driven example above. This would mean that job C would be chosen if only system goals were considered. Consensus level agents would not be able to choose an alternative solely due to system considerations because of the nature of their planning framework. In order for a consensus agent to choose another option, it would have to change its autonomy level and then replan.

This example has shown that under different autonomy levels and their associated decision models, Sensible Agents will behave differently. This difference in behavior can be modeled, controlled and analyzed using autonomy levels. This increases the flexibility of Sensible Agents by allowing them to understand which models are appropriate for which situational contexts. It is important to note that a real-world example would be much more complicated and finding a unique optimized solution is difficult, at best.

6. Conclusions

In this paper, we present a new approach for increasing the flexibility of agents' decision-making capabilities of reasoning about the trade-offs between system goals and local goals. Such reasoning is based on the calculation of system and local utilities of goals given the agent's current role in the organization (organization role is specified by the autonomy level). Autonomy level reasoning is based on an agent's abilities, situational context, and position in the organizational structure of the system. Autonomy levels can have an impact on the system and local utilities that are derived from goals and can help guide the decision making process when calculating trade-offs. Additionally, the autonomy constructs can support a variety of decision making models while extending their capabilities through experience-based learning. This increases the flexibility and utility of Sensible Agents-based systems for highly dynamic and uncertain environments.

7. References

Askin, R. G. and Standridge. 1993.*Modeling and Analysis of Manufacturing Systems*. New York: John Wiley & Sons, Inc.

Barber, K. S. 1996. The Architecture for Sensible Agents. In Proceedings of the *International Multidisciplinary Conference, Intelligent Systems: A Semiotic Perspective*, 49-54. Gaithersburg, MD: National Institute of Standards and Technology.

Byrnes, J. P. 1998.*The NATURE and DEVELOPMENT of DECISION MAKING: A Self-Regulation Model*. Mahwah, New Jersey: Lawrence Erlbaum Associates, Publishers.

Ephrati, E., Pollack, M. E., and Rosenschein, J. S. 1995. A Tractable Heuristic that Maximizes Global Utility Through Local Goal Combination. In Proceedings of the *First International Conference on Multi-Agents Systems*, 94-101. San Francisco, California: AAAI Press/ The MIT Press.

Gasser, L. 1991. Boundaries, Identity, and Aggregation: Plurity Issues in Multiagent Systems. In Proceedings of the *Proceedings of the 3rd European Workshop on Modelling Autonomous Agents in a Multi-Agents World*, 199-213. Kaiserslautern, Germany: Elsevier Science Publishers.

Goel, A., Liu, T. H., and Barber, K. S. 1996. Conflict Resolution in Sensible Agents. In Proceedings of the *International Multidisciplinary Conference on Intelligent Systems: A Semiotic Perspective*, 80-85. Gaithersburg, MD:

Hill, P. H., Bedau, H. A., Chechile, R. A., Crochetiere, W. J., Kellerman, B. L., Ounjian, D., Pauker, S. G., Pauker, S. P., and Rubin, J. Z. 1982.*Making Decisions: A Multidisciplinary Introduction*. Reading, Massachusetts: Addison-Wesley Publishing Company (Advanced Book Program).

Jennings, N. R. 1996. Coordination Techniques for Distributed Artificial Intelligence. in *Foundations of Distributed Artificial Intelligence*, *Sixth-Generation Computer Technology Series*, O'Hare, G. M. P. and Jennings, N. R., Eds. New York: John Wiley & Sons, Inc., 187-210.

Keeney, R. and Raiffa, H. 1976.*Decision with Multiple Objectives: Preferences and Value Trade-offs*: J. Wiley and Sons.

Liu, J. S. and Sycara, K. P. 1996. Multiagent Coordination in Tightly Coupled Task Scheduling. In Proceedings of the *Second International Conference on Multi-Agents Systems*, 181-188. Kyoto, Japan: AAAI Press/ The MIT Press.

Martin, C. E. and Barber, K. S. 1996. Multiple, Simultaneous Autonomy Levels for Agents-based Systems. In Proceedings of the *Fourth International Conference on Control, Automation, Robotics, and Vision*, 1318-1322. Westin Stamford, Singapore:

Rosenschein, J. S. and Zlotkin, G. 1994.*Rules of Encounter: Designing Conventions for Automated Negotiation Among Computers*. Cambridge, Massachusetts: The MIT Press.

Roy, B. 1976. From Optimization to Multicriteria Decision Aid: Three Main Operational Attitudes. in *Multiple Criteria Decision Making*, Thiriez, H. and Zionts, S., Eds.: Springer-Verlag, 1-32.

Sycara, K. P. 1988. Utility Theory In Conflict Resolution. *Annals of Operation Research* 12: 65-84.

Vincke, P., Gassner, M., and Roy, B. 1989.*Multicriteria Decision-aid*. Chichester: John Wiley & Sons.

Yu, P. L. 1985.*Multiple-Criteria Decison Making: Concepts, Techniques, and Extensions*. New York, NY: Plenum Press.