Towards Safe Coordination in Multi-agent Systems

Anita Raja[1], Michael Barley[2], and Xiaoqin Shelley Zhang[3]

[1] Department of Software and Information Systems, The University of North Carolina at Charlotte, Charlotte, NC 28223, anraja@uncc.edu
[2] Department of Computer Science, University of Auckland. Auckland, NZ mbar098@cs.auckland.ac.nz
[3] Department of Computer Science, The University of Massachusetts Darmouth, North Dartmouth, MA 02747, x2zhang@umassd.edu

Abstract. Conservative design is the ability of an individual agent to ensure predictability of its overall performance even if some of its actions and interactions may be inherently less predictable or even completely unpredictable. In this paper, we describe the importance of conservative design in cooperative multi-agent systems and briefly characterize the challenges that need to be addressed to achieve this goal.

1 Introduction

Uncertainty is ubiquitous in complex MAS operating in open environments. A safe multi-agent system is composed of agents that are equipped with "conservative design" [2] capabilities. We define conservative design in cooperative multi-agent systems as the ability of an individual agent to ensure predictability of its overall performance even if some of its actions and interactions may be inherently less predictable or even completely unpredictable. An essential feature of conservative design is the ability of agents to efficiently handle risk. Risk is the potential for realizing unwanted negative consequences of an event [5].

Establishing environmental predictability and enforcing risk management measures are first-class goals of real-world organizations. Corporations are motivated by the need to maintain their reputations while maximizing profit. Reputation can be a precious commodity and corporations can charge more for their products because they have a reputation, say, for reliability. When corporations depend on other organizations that are less reliable than themselves, they must come up with plans that enable them to guarantee that they will still live up to their reputations. In practice, the customers of a reputable corporation assume that the corporation has accounted for the risk of its suppliers failing to deliver within the given constraints, has made contingency plans, and will deliver the product as agreed upon.

In this paper, we describe this concept of conservative design in multi-agent systems specifically as it relates to coordination with other agents. We identify a number of challenge areas along with examples that have to be addressed to achieve this goal.

2 Conservative Design

Predictability of both the environment and of the outcomes of agent performance are essential for building safe systems. One way of ensuring conservative design in agents is to identify levels of risk in an agent's goals and to incorporate reasoning about risk into agent planning. More specifically, an important challenge problem for safe coordination in multi-agent systems is the ability to transform an unpredictable and highly complex world into a simpler world where even if the actions are not entirely deterministic, at least there are predictable bounds on the probabilities of their outcomes. The following is an example of a multi-agent application that motivates the need for conservative design. We also identify the questions that need to be addressed by the multi-agent system.

Consider the supply chain scenario described in Figure 1. Each agent represents a company/organization involved in the supply chain. The *PCManufacturer*, the two Chip Producers (*ChipProducer1* and *ChipProducer2*) and the two Transporters (*TransporterA* and *TransporterB*) are represented by individual agents. The *PCManufacturer* makes a commitment to the customer to deliver the product within a deadline DL. This means the *PCManufacturer* has to negotiate commitments with one of the two Chip Producers and one of the two Transporters such that there is enough time to assemble and deliver the final product to the customer before the deadline DL. It may be the case that *PCManufacturer* may have to take some risky actions to ensure completing the goal within the deadline. Agent *PCManufacturer* has the achievement goal of satisfying the customer's request within the given constraints. Its maintenance goal could be to keep the costs low and allow for risk within a certain threshold while handling the customer's request.

Suppose the *ChipProducer1* agent is unable to keep its initial commitment to the *PCManufacturer* agent. *PCManufacturer* may choose to re-negotiate a new commitment instead. The re-negotiation is caused by the local dynamics of *ChipProducer1*, and it may result in a change of the content of the existing commitment, such as the quality or the quantity of the parts, the cost, or the delivery time of the parts [8]. Resources invested in searching for alternate commitments in real-time could result in high costs. The *PCManufacturer* agent could avoid these costs by estimating the probability that *ChipProducer1* will need to re-negotiate its commitment and include this information in the decision process while choosing the Chip Producer agent. Furthermore, if the probability of re-negotiation by *ChipProducer1* is within an acceptable threshold, the costs of online re-negotiation can be kept low by identifying contingent commitments. The following are some questions that will help determine the risk of choosing *ChipProducer1* as the supplier:

1. How trustworthy is the agent based on its reputation? How likely is it to drop its commitment?
2. Is it useful to enforce a decommitment penalty for this particular subcontract?
3. How likely is it for the agent to complete the task within the expected performance profile?

Fig. 1. A Supply Chain Scenario

4. Should payment for product be a function of the performance profile of end product delivered?
5. Are there enough resources to establish commitments with both *ChipProducer1* and *ChipProducer2*? This will allow *PCManufacturer* to store one of the chips while using the other in case both agents deliver on schedule. This reduces the risk of not having a chip available by the delivery deadline.

Determination of risk enables a multi-agent system to identify predictable and less-predictable parts of its solution space. The multi-agent system, depending on its goals and its environment, may have to spend resources in ensuring that the less-predictable parts of the solution space are made more predictable in order to provide performance guarantees. We now describe some factors that will facilitate risk handling and leverage the uncertainty in multi-agent environments.

2.1 Identifying Risk Thresholds

One way that corporations enforce risk management is by establishing legal thresholds on the amount of risk that they can assume or are willing to tolerate. For instance, the boards of directors of a company cannot act ”irresponsibly” or ”recklessly”. They also may have thresholds on the risk associated with their products. The Federal Aviation Agency (FAA) may set risk thresholds on the products they allow for airplane parts manufacturers. An example of such a risk threshold could be the average number of air miles flown between maintenance checks. The plane would not be allowed on the market until it has been shown to have a risk lower than the legal threshold. The thresholds can be fuzzy and may depend on a number of variables. Yet they are useful measures that help achieve a entity's overall goal which is to make its environment deterministic, so that planning and scheduling can be made simpler.

Our goal in this work is to build agents that want to maintain their reputation within a multi-agent system while minimizing the cost of ensuring predictability of the environment. The cost of dealing with the unwanted consequences of an event is called *risk impact*. For instance, if the power goes down in a manufacturing company for a number of hours, there is an impact on the productivity for that day. Predictability of the environment involves pre-computing and controlling risk impact for different situations. Suppose risk is defined by m attributes and x_i is one such attribute contributing to risk; p_i is the probability of this attribute occurring; imp_i is the impact (cost) of the attribute x_i. Then,

$$Risk = \sum_{i=1}^{m} p_i * imp_i$$

It is possible to extend this function to contexts where there is uncertainty in the impact. In other words, if a particular risk attribute occurs, various impact outcomes are possible depending on the current environmental state. For example, if the power goes down at 10 a.m. as opposed to 10 p.m., the resulting impact would be very different. Suppose there are are n possible impacts for each attribute x_i and each impact p_{i_j} occurs with probability imp_{i_j}, then

$$Risk = \sum_{i=1}^{m} \sum_{j=1}^{n} p_{i_j} * imp_{i_j}$$

A risk threshold τ may have to be determined in a situation specific fashion and conservative design means that the agents have the following maintenance goal:

$$Risk \ < \ \tau$$

The following are examples of risk attributes in the supply chain scenario described earlier:

1. Chip Producer agent delivers product x units after deadline $d1$ established in a commitment.
2. Transport agent delivers product y units after deadline $d2$ established in a commitment.
3. Storage costs of redundant orders result in higher total costs, thereby lowering total profit.

Consider a situation where *TransporterA* promises to deliver the product by time 10 with cost 20, however there is a probability 0.4 it may be delayed to time 20. *TransporterB* promises to deliver by time 10 with cost 30 and with probability 0.05 it may be delayed to time 12. The following are some issues that affect *PCManufacturer*'s decision process:

1. The risk threshold for *PCManufacturer* could be a function of the tightness of the deadline, the cost constraints and the importance of maintaining its reputation. Suppose the price negotiated with the customer for the final product is 60. If the deadline is sufficiently beyond time 12, say time 25, *PCManufacturer* would choose *TransporterA* and pay 20 cost units. If the deadline is 18, then *PCManufacturer* would be willing to pay the higher cost of 30 cost units to have *TransporterB* deliver the Chip with 100% guarantee by time 12.
2. It is also possible for *PCManufacturer* to look at the possible product delivery times and try to move the schedule a little ahead. For instance, the agent could anticipate a potential product request, initiate negotiation to complete the product early and pay for storage. This will reduce the risk threshold of failure to have the product part on time but does increase the cost by paying for storage.

Identifying risk attributes and risk thresholds are important problems that need to be addressed to accurately handle uncertainty.

2.2 Using Risk to Prune Solution Space

There are multiple choices in an agent's coordination process including with which agent to cooperate, when to cooperate and how to cooperate. When an agent has to coordinate with several other agents, the order of the coordination is also an important factor. An agent uses its performance criteria to make its coordination choices. The performance criteria is defined by multiple factors including risk threshold, utility and cost. In order to maintain its reputation and reliability, the agent could use risk as a factor to prune the search space. In other words, if a coordination solution carries a risk measure higher than the risk threshold acceptable by that the agent, this solution will not be considered as a candidate. After deleting all solutions with risk beyond the threshold, the agent can evaluate the rest of the solutions using a utility/cost measure, meaning, the agent can select the solution with the highest utility/lowest cost. Another approach to evaluate those solutions is to use a multi-dimensional weighted function that combines both the utility and the risk measurement. For example, given the following performance measures of three solution options:

$$Solution1 : utility = 100, risk = 10$$

$$Solution2 : utility = 120, risk = 15$$

$$Solution3 : utility = 145, risk = 25$$

Suppose the acceptable risk threshold is 20, Solutions 1 and 2 are valid candidates, while Solution 3 is eliminated. The agent can then use a multi-dimensional function to evaluate the remaining two solutions. The following is an example of a possible multi-dimensional evaluation function:

$$F = w1 * utility - w2 * risk$$

where $w1$ and $w2$ represent the importance of utility and risk are to the agent. This simple weighted function would allow the agent to balance utility and risk in its decision-making process. If $w2$ is higher than $w1$, the agent tries to find the solution with the highest utility while minimizing risk.

Thus we conjecture that an agent can use risk to reduce its problem solving costs. It has to first enumerate all possible solutions, use its risk threshold to prune the solution space, and then evaluate the remaining solutions using the utility measure or a multi-dimensional function that combines utility and risk.

2.3 Contingencies in Coordination

In mission-critical environments, failure can lead to catastrophic consequences. This means that it is not sufficient for an agent to minimize risk in coordination. For instance, the uncertainty in agent outcomes and environment characteristics could necessitate dropping previous commitments and establishing new commitments [7]. However, online re-coordination can be prohibitively expensive. While negotiating, it will be useful for agents both to determine the probability of a commitment completing successfully and to identify contingency commitment contracts in case of commitment failure. There has been significant previous work in studying contingencies for planning and scheduling [1, 3, 4]. We are interested in extending some of these ideas to coordination.

Consider an example where agent *A* should complete task *T1* by deadline 80. Suppose Task *T1* can be achieved by two alternate ways *P1* and *P2* where *P1* is a high risk solution with high utility while *P2* is a low risk solution with low utility. Specifically, *P1* requires that agent *B* completes an enabling action *M1*. However agent *B* is usually heavily loaded and has the reputation of dropping its commitments regularly.

One way of handling this uncertainty would be for agent *A* to establish a contingent commitment that consists of two parts: the details of a commitment along with time of completing the contracted task, and a time at which the contractor agent will confirm the commitment. This confirmation time will lie between the time of establishing the commitment and the time for completing the commitment.

In the above example, suppose the current time is 12 and a commitment between agent *A* and agent *B* is feasible with a probability of 0.4. The contingent commitment would be as follows: Agent *B* commits to complete action *M1* by time 43; and will confirm the commitment at time 25. Agent *A* will choose plan P2 as the contingent plan. In the case that agent *B* confirms the commitment positively at time 25, agent *A* will continue with Plan *P1.* If on the other hand, agent *B* states that it has to drop its commitment, then agent *A* will resort to plan *P2.* Additionally, there will a high penalty for agent *B* if it confirms commitment at time 25 but fails to complete action *M1* by time 43.

An alternative way of handling uncertainty would be leveled commitment contracts [6]. These are contracts where each party can decommit by paying a predetermined penalty. This would allow agents to accommodate events which unfolded since the contract was entered into while allowing the other agents to use the penalty to make alternate plans.

A critical question is to determine how much contingency planning for coordination is enough? The contingent plans themselves may require contingencies in case of failure. The challenge is then to design agents that are equipped with the ability to do cost-benefit tradeoffs for the depth of contingency in coordination.

3 Conclusions

In this paper, we defined the concept of conservative design in multi-agent systems. This is based on the premise that safety, like security, should not be an after-thought but an integral part of the design of the agent's decision making capabilities. We also identified factors that contribute to analyzing and responding to risk in agent environments. As future work, we plan to formalize the representation of risk using the supply chain scenario as an example. Then we plan to implement some of the reasoning processes described in this paper. We see this as first step towards conservative design in multi-agent systems.

References

1. Draper, D., Hanks, S., Weld, D.: Probabilistic planning with information gathering and contingent execution. In: Proceedings of the Second International Conference on Artificial Intelligence Planning Systems (AIPS-94). (1994) 31–36.
2. Mead, C., Conway, L.: Introduction to VLSI Systems. Addison-Wesley, Reading, MA (1980).
3. Onder, N., Pollack, M.: Contingency selection in plan generation. In: Proceedings of the Fourth European Conference on Planning. (1997).
4. Raja, A., Wagner, T., Lesser, V.: Reasoning about Uncertainty in Design-to-Criteria Scheduling. In: Working Notes of the AAAI 2000 Spring Symposium on Real-Time Systems, Stanford. (2000).
5. Rowe, W.D.: An Anatomy of Risk. Robert E. Krieger Publishing Co., Malabar, FL (1988).
6. Sandholm, T., Lesser, V.: Leveled commitment contracting: A backtracking instrument for multiagent systems. AI Magazine (2000) 89–100.
7. Xuan, P., Lesser, V.R.: Handling uncertainty in multi-agent commitments. Technical Report UM-CS-1999-005 (1999).
8. Zhang, X., Lesser, V., Wagner, T.: A layered approach to complex negotiations. Web Intelligence and Agent Systems: An International Journal **2** (2004) 91–104.