Market-based Risk Allocation for Multi-agent Systems *

(Extended Abstract)

Masahiro Ono and Brian C. Williams
Computer Science and Artificial Intelligence Laboratory
Massachusetts Institute of Technology
{hiro_ono, williams}@mit.edu

ABSTRACT

This paper proposes *Market-based Iterative Risk Allocation* (*MIRA*), a new market-based decentralized optimization algorithm for multi-agent systems under stochastic uncertainty, with a focus on problems with continuous action and state space. In large coordination problems, from power grid management to multi-vehicle missions, multiple agents act collectively in order to maximize the performance of the system, while satisfying mission constraints. These optimal action plans are particularly susceptible to risk when uncertainty is introduced. We present a decentralized optimization algorithm that minimizes the system cost while ensuring that the probability of violating mission constraints is below a user-specifie upper bound.

We build upon the paradigm of *risk allocation* [3], in which the planner optimizes not only the sequence of actions, but also its allocation of risk among state constraints. We extend the concept of risk allocation to multi-agent systems by highlighting risk as a resource that is traded in a computational market. The equilibrium price of risk that balances the supply and demand is found by an iterative price adjustment process called *t tonnement* (also known as *Walrasian auction*). Our work is distinct from the classical tâtonnement approach in that we use Brent's method to provide fast guaranteed convergence to the equilibrium price. The simulation results demonstrate the efficien y and optimality of the proposed decentralized optimization algorithm.

Categories and Subject Descriptors

I.2 [Artificial Intelligence]: Distributed Artificia Intelligence

General Terms

Algorithms

Keywords

Chance constrained optimal planning, Decentralized optimization, Tâtonnement, Walrasian auction, Continuous resource allocation

1. RISK ALLOCATION

Risk is summable. Hence, it can be considered as a resource. This result is derived from Boole's inequality:

$$\Pr\left[\bigcup_i F_i\right] \leq \sum_i \Pr\left[F_i\right] \tag{1}$$

Assume that F_i in the above inequality represents the event that the ith agent in a multi-agent system fails. We defin *system failure* as an event in which at least one of the agents in the system fails. Then the left-hand side of (1) means the probability of system failure. It is upper-bounded by the right-hand side, which is the sum of the individual probabilities that each agent may fail.

The user of the system wants to limit the probability of system failure to a constant probability S (*risk bound*). This constraint is called *joint chance constraint*:

$$\Pr\left[\bigcup_i F_i\right] \leq S \tag{2}$$

Using Boole's inequality (1), it can be easily shown that the following condition is a sufficien condition of the original joint chance constraint (2):

$$\forall_i \quad \Pr\left[F_i\right] \leq \Delta_i \tag{3}$$

$$\wedge \quad \sum_i \Delta_i \leq S \tag{4}$$

Eq.(3) constrains the probability that each individual agent may fail (*individual chance constraints*). Eq.(4) states that the sum of the risk bounds of all individual chance constraints must not exceed the risk bound of the original joint chance constraint S. Here, the analogue to resource allocation is found: S is the total amount of resource (i.e., risk), which is distributed among the agents in the system; Δ_i is the amount of resource allocated to the ith agent.

Once the risk is allocated to each agent, a joint chance constraint over multiple agents (2) is decomposed into individual chance constraints over individual agents (3), which can be handled by each individual agent in a distributed manner. Then, a challenging problem is how to fin the globally optimal risk allocation without a centralized computation.

2. THE MIRA ALGORITHM

We have developed a *Market-based Iterative Risk Allocation (MIRA)* algorithm that find the globally optimal risk allocation through a tâtonnement-like process. In the MIRA algorithm, each agent demands risk in a computational market in order to improve its own performance (i.e., reduce its own cost function), while the user supplies a fi ed amount of risk by specifying the risk bound S. MIRA

*This research is funded by The Boeing Company grant MIT-BA-GTA-1

Cite as: Market-based Risk Allocation for Multi-agent Systems (Extended Abstract), Masahiro Ono and Brian Williams, *Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, van der Hoek, Kaminka, Lespérance, Luck and Sen (eds.), May, 10–14, 2010, Toronto, Canada, pp. 1423-1424

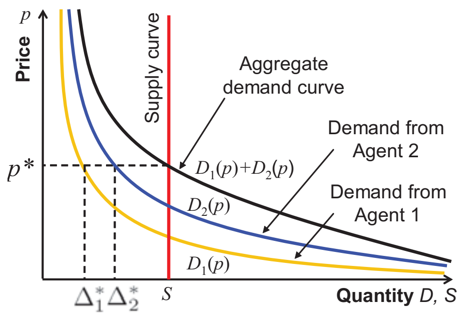

Figure 1: Market-based risk allocation in a system with two agents.

iteratively adjusts the price to fin the equilibrium price, where the demand and supply are balanced.

Agents are price takers. Given the price, each agent computes the optimal amount of risk to take (i.e., *demand for risk*) by solving an optimization problem that minimizes the sum of its own cost and the payment for the risk. The optimal action sequence and the internal risk allocation are also determined by solving the optimization problem, just as in the single-agent case [4]. The demand from each agent can be seen as a function of the price of risk (*demand curve*) as in Figure 1. Typically, the higher the price is, the less each agent demands. Each agent has a different demand curve according to its sensitivity to risk. The price must be adjusted so that the total demand (*aggregate demand*) becomes equal to the supply, which is a constant. The equilibrium price $p^\star$ lies at the intersection of the aggregate demand curve and the supply curve. The optimal risk allocation for the two agents corresponds to their demands at the equilibrium price ($\Delta_1^\star$ and $\Delta_2^\star$ in Figure 1).

Algorithm 1 presents MIRA algorithm. The algorithm terminates when the price of risk reaches the equilibrium where aggregate demand $\sum_i D_i(p)$ is balanced with the supply S. It follows from the KKT conditions for optimality that, if the cost functions of all agents are convex, the risk allocation is globally optimized at the equilibrium price, although each agent only minimizes its own cost [5]. See [5] for the proof.

3. RELATED WORK

Past work on risk allocation [3][4] focused on single-agent problems with multiple state constraints, where optimal risk allocation over the constraints is obtained by a centralized algorithm. Although tâtonnement has drawn less attention than auctions, it has been successfully applied to various problems such as the distribution of heating energy in an offic building [7] and the control of electrical power fl w [2]. Classical tâtonnement uses a simple linear price update rule, but the convergence of price can only be guaranteed under a quite restrictive condition [6]. Its slow convergence is also an issue. We employ Brent's method [1] to provide guaranteed convergence with a superlinear rate of convergence by exploiting the fact that a risk, which is treated as a resource in our problem formulation, is a scalar value.

Algorithm 1 Market-based Iterative Risk Allocation

```
Fix S;        //Total supply of risk
Initialize p;        //Price of risk
while |∑_i D_i(p) − S| > ε and p > 0 do
    Auctioneer announces p;
    Each agent computes its demand for risk D_i(p)
    Each agent submits its demand to the auctioneer;
    The auctioneer updates p using Brent's method;
end while
```

Table 1: Comparison of the computation time of three optimization algorithms. Values are the average of 10 runs with randomly generated problems. Computations of the distributed algorithms were conducted parallelly.

Number of agents	Computation time [sec]		
	Centralized	Decentralized (linear update)	MIRA
2	13.9	80.6	6.4
4	63.8	540.5	18.1
8	318.5	797.8	37.5

4. SIMULATION RESULT

The computation times of the following three algorithms are compared: 1) centralized optimization, 2) decentralized optimization with the linear price update rule (classical tâtonnement), and 3) MIRA, which updates the price using Brent's method.

Table 1 shows the result. The computation time of the centralized algorithm quickly grows as the problem size increases. Decentralized optimization with a linear price increment is even slower than the centralized algorithm, although the growth rate of computation time is slower. MIRA, the proposed algorithm, outperforms the other two for all problem sizes. The advantage of MIRA becomes clearer as the problem size increases. See [5] for the parameters and simulation settings used.

5. ACKNOWLEDGMENTS

Thanks to Michael Kerstetter, Scott Smith, and the Boeing Company for their support.

6. REFERENCES

[1] K. E. Atkinson. *An Introduction to Numerical Analysis, Second Edition*. John Wiley & Sons, 1989.

[2] J. C. Jacobo, D. De Roure, and E. H. Gerding. An agent-based electrical power market. In *Proceedings of AAMAS-08: Demo Papers*, 2008.

[3] M. Ono and B. C. Williams. An efficien motion planning algorithm for stochastic dynamic systems with constraints on probability of failure. In *Proceedings of AAAI-08*, 2008.

[4] M. Ono and B. C. Williams. Iterative risk allocation: A new approach to robust model predictive control with a joint chance constraint. In *Proceedings of 47th IEEE Conference on Decision and Control (CDC-08)*, 2008.

[5] M. Ono and B. C. Williams. Risk allocation for multi-agent systems using tâtonnement. Technical report, MIT Computer Science and Artificia Intelligence Laboratory, 2009.

[6] J. Tuinstra. *Price Dynamics in Equilibrium Models: The Search for Equilibrium and the Emergence of Endogenous Fluctuations*. Kluwer Academic Publishers, 2000.

[7] H. Voos. Agent-based distributed resource allocation in technical dynamic systems. In *Proceedings of IEEE Workshop on Distributed Intelligent Systems: Collective Intelligence and Its Applications (DIS 2006)*, 2006.