Market-based Risk Allocation for Multi-agent Systems [*]

Masahiro Ono and **Brian C. Williams**

Massachusetts Institute of Technology
77 Massachusetts Avenue
Cambridge, MA 02139
hiro_ono@mit.edu, williams@mit.edu

Abstract

This paper proposes *Market-based Iterative Risk Allocation* (*MIRA*), a new market-based distributed planning algorithm for multi-agent systems under uncertainty. In large coordination problems, from power grid management to multi-vehicle missions, multiple agents act collectively in order to optimize the performance of the system, while satisfying mission constraints. These optimal plans are particularly susceptible to risk when uncertainty is introduced. We present a distributed planning algorithm that minimizes the system cost while ensuring that the probability of violating mission constraints is below a user-specified level.

We build upon the paradigm of *risk allocation* (Ono & Williams 2008), in which the planner optimizes not only the sequence of actions, but also its allocation of risk among each constraint at each time step. We extend the concept of risk allocation to multi-agent systems by highlighting risk as a commodity 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 *tâtonnement* (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 efficiency of the proposed distributed planner.

Introduction

Motivation

There is an increasing need for multi-agent systems that perform optimal planning under uncertainty. An example is planning and control of power grid systems. A power grid consists of a numbers of generators and electric transformers whose control should be carefully planned in order to maximize efficiency. A significant issue in power grid planning is the uncertainty in demand for energy by consumers. As the use of renewable energy, such as solar and wind power, become more popular, uncertainty in supply increases due to weather conditions. Another example is the Autonomous

---

Ocean Sampling Network (AOSN), which consists of multiple automated underwater vehicles (AUVs), robotic buoys, and aerial vehicles. AOSN should maximize science gain while being exposed to external disturbances, such as tides and currents.

In order to deploy AI planning algorithms on such systems, we need a robust plan execution capability. Robust execution often involves 1) handling continuous system dynamics, 2) handling uncertainty in the environment, 3) operating the system at a risk level that the user find acceptable, and 4) scaling to multi-agent system.

To address the four problems, we developed *Market-based Iterative Risk Allocation* (MIRA), a multi-agent optimal planning algorithm that operates within user-specified risk bounds. MIRA optimally allocates risk among agents, and computes optimal control sequence for each agent in a distributed manner.

Approach

Planning under uncertainty, and risk allocation   When planning actions under uncertainty, there is always a risk of failure that should be avoided. However, in many cases, performance can be improved only by taking extra risk. For example, we can reach a destination faster by driving a car at a faster speed and accepting a higher risk of an accident.

Without taking any risk, nothing can be done; however, no one dares to take unlimited risk. In many cases, people want to maximize performance, but with an upper-bound on the risk they take (*chance constraint*). For example, a race car driver would like to drive as fast as possible while limiting the probability of a crash to 0.1%. Therefore, we formulate the stochastic planning problem as an optimization problem with a chance constraint. With this formulation, (Ono & Williams 2008) showed that the planner should plan not only the sequence of actions but also the *risk allocation* in order to maximize the performance under a risk bound.

The example shown in Figure 1 illustrates the concept of risk allocation. A race car driver wants to plan a path to get to the goal as fast as possible. However, crashing into the wall leads to a fatal accident, so he wants to limit the probability of a crash to 0.1%. An intelligent driver would plan a path as shown in Figure 1, which runs mostly in the middle of the straightaway, but gets close to the wall at the corner. This is because taking a risk (i.e. approaching the
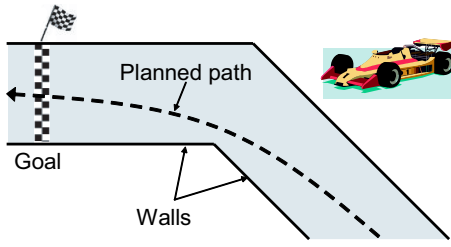
Figure 1: Risk allocation in a race car path planning scenario.

Figure 2: Multi-UAV fire-fighting scenario.

Figure 3: Demand and supply curves of the risk market with two agents. Note that we followed the economics convention of placing the price (independent variable) on the vertical axis.
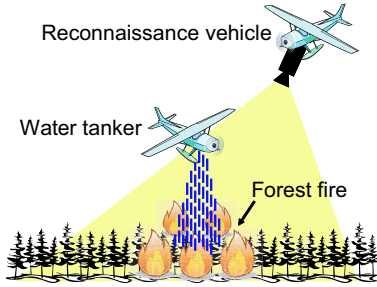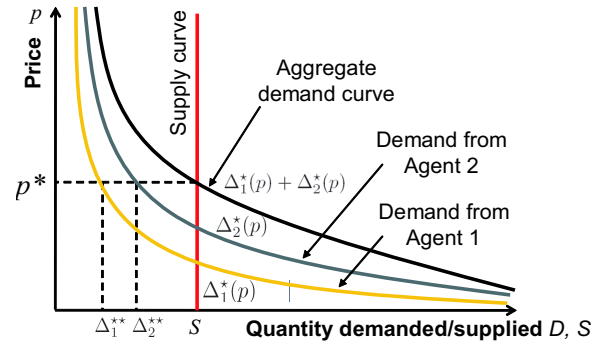
wall) at the corner results in a greater time saving than taking the same risk along the straightaway; in other words, the utility of taking risk is greater at the corner than the straightaway. Therefore the optimal plan allocates a large portion of risk to the corner, while allocating little to the straightaway. As illustrated by this example, *risk allocation* needs to be optimized across the constraints, in order to maximize the performance.

Distributed risk allocation for multi-agent system The concept of risk allocation can be naturally extended to multi-agent systems. Figure 2 shows an example of a multi-agent system with two unmanned air vehicles (UAVs), whose mission is to extinguish a forest fire. A water tanker drops water while a reconnaissance vehicle monitors the fire with its sensors. The loss of either vehicle results in a failure of the mission. Two vehicles are required to extinguish the fire as efficiently as possible, while limiting the probability of mission failure to a given risk bound, say, 0.1%. The water tanker can improve efficiency by flying at a lower altitude, but it involves risk. The reconnaissance vehicle can also improve the data resolution by flying low, but the improvement of efficiency is not as great as the water tanker. In such a case the optimal risk allocation is to allow the water tanker to take a large portion of risk by flying low, while keeping the reconnaissance vehicle at a high altitude to avoid risk. This is because the utility of taking risk (i.e. flying low) is greater for the water vehicle than for the reconnaissance vehicle.

Then, the question is how to find the optimal risk allocation between multiple vehicles in a distributed manner.

Tâtonnement: market-based risk allocation Our approach is to use the market-based mechanism to optimize the risk allocation between agents. In a computational competitive market, each agent demands risk in order to improve its own performance. However, it cannot take risk for free; it has to purchase it from the market at a given price.

Agents are price takers. Therefore the demand is a function of the price of risk (*demand curve*). The higher the price is, the less each agent demands. Each agent has a different demand curve according to its sensitivity to risk. On the other hand, the supply of risk is constant, since the upper-bound of total risk is given.

Figure 3 gives the graphical interpretation of the market-based risk allocation in a system with two agents, who have different demand curves. The aggregate demand curve is obtained by adding the two demand curves horizontally. The supply curve is a vertical line since it is constant. The equilibrium price $p^\star$ lies at the intersection of the aggregate demand curve and the supply curve. It can be proven that the optimal risk allocation for the two agents are their demands for the risk at the equilibrium price ($\Delta_1^{\star\star}$ and $\Delta_2^{\star\star}$ in Figure 3).

However, in practice, the demand curve is not available as a function. Therefore it is not easy to find the intersection of the demand curve and supply curve. Instead, the equilibrium price is found by the following intuitive iterative process, called *tâtonnement* or *Walrasian auction* (Tuinstra 2000):
- Increase the price if aggregate demand exceeds supply,
- Decrease the price if supply exceeds aggregate demand,
- Repeat until supply and demand are balanced.

Related Work

Market-based approach has recently been recognized as an effective tool for decentralized multi-agent systems in AI community (Wellman 1993)(MacKie-Mason *et al.* 2004). 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 office building (Voos 2006), and resource allocation in communication networks (Kelly, Maulloo, & Tan 1998). The convergence of

tâtonnement has been an issue in economics for a long time; with a simple linear price update rule, it can only be guaranteed under a quite restrictive condition (Tuinstra 2000). We solved this problem by applying a root-finding method called Brent's method(Atkinson 1989).

Derivation

Risk Allocation   The concept of risk allocation is derived from Boole's inequality:

$$\Pr\left[\bigcup_i F_i\right] \le \sum_i \Pr\left[F_i\right] \qquad (1)$$

Assume that F_i in the above inequality represents the event that the ith agent fails. Then the left hand side means the probability that at least one agent in the system fails (i.e. system failure). It is upper-bounded by the right hand side, which is the sum of the individual probabilities that each agent fails.

The user of the system limits the probability of system failure to S. This constraint is called *joint chance constraint*.

$$\Pr\left[\bigcup_i F_i\right] \le S \qquad (2)$$

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

$$\forall_i \quad \Pr\left[F_i\right] \ \le \ \Delta_i \qquad (3)$$
$$\wedge \quad \sum_i \Delta_i \ \le \ S \qquad (4)$$

Eq.(3) constrains the probability that each individual agent fails (*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 the resource allocation is found; S is the total amount of resource (i.e. risk), which is distributed to 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 agent Eq.(2) is decomposed into individual chance constraints over individual agents Eq.(3).

Distributed Optimization of Risk Allocation   The objective of our problem is to minimize the system cost, which is the total of the cost of all agents in the system, while limiting the probability of system failure (joint chance constraint). As explained above, the joint chance constraint Eq.(2) is implied by the individual chance constraints Eq.(3) and the total risk inequality Eq.(4). Therefore, our optimization is formulated as follows:

$$\min_{\Delta_{1:N}} \quad \sum_{i=1}^{N} J_i(\Delta_i) \qquad (5)$$
$$\text{s.t.} \quad (3)(4)$$

where J_i is the cost of ith agent, and N is the number of agents in the system. We assume that there is no coupling between agents through constraints. This formulation describes a centralized algorithm since the risk allocations of all agents are planned in one optimization problem. We omit the plant model (as linear constraints) and control limit constraints to keep the equations simple. See (Ono & Williams 2009) for the formulation with all constraints.

Solving the centralized optimization problem Eq.(5)(3)(4) is equivalent to solving the following N unconstrained optimization problems, since two formulations have the same Karush-Kuhn-Tucker (KKT) conditions for optimality.

$$\min_{\Delta_i} J_i(\Delta_i) + p\Delta_i \qquad (\text{for } i = 1 \cdots N) \qquad (6)$$

where $p \ge 0$ is the Lagrange multiplier. In order to be optimal, p and Δ_i must satisfy the following condition:

$$p\left(\sum_i \Delta_i - S\right) = 0 \qquad (7)$$

Since the optimization problems Eq.(6) contains only the variables related to the ith agent, it can be solved by each agent in a distributed manner.

Economic Interpretation   The interpretation of these mathematical manipulations becomes clear by regarding the Lagrange multiplier p as the *price of risk*. Each agent can improve the performance by taking risk Δ_i, but not for free. Note that a new term $p\Delta_i$ is added to the cost function Eq.(6). This is what the agent has to pay to take the amount of risk Δ_i. Given the price p, each agent computes the optimal *demand for risk* $\Delta_i^\star(p)$ by solving the optimization problem Eq.(6). The total amount of risk S can be interpreted as the *supply of risk*, which is given by the user.

In order to minimize the system cost, the price p must satisfy the condition Eq.(7). Such a price $p^\star$ is called the *equilibrium price*. The demand for risk of each agent at the equilibrium price $\Delta_i^\star(p^\star)$ is the optimal risk allocation for the agent.

Eq.(7) illustrates the relation between the equilibrium price $p^\star$, optimal demand $\Delta_i^\star(p^\star)$, and supply S; in the usual case where the equilibrium price is positive $p^\star > 0$, the aggregate demand $\sum_i \Delta_i^\star(p^\star)$ must be equal to the supply S, as illustrated in Figure 3; in a special case where the supply always exceeds the demand for all $p \ge 0$, the optimal price is zero $p^\star = 0$. If the aggregate demand always exceeds the supply for all $p \ge 0$, there is no solution that satisfies the constraint Eq.(4), and hence the problem is infeasible.

Finding Equilibrium Price   According to Eq.(7), the equilibrium price is the root of the following equation:

$$\sum_i \Delta_i^\star(p) - S = 0 \qquad (8)$$

The classical approach in economics is to iteratively adjust the price with the increment that is proportional to the excess demand $\sum_i \Delta_i^\star(p) - S$, until the price converges.

However, the convergence is guaranteed only under a strong condition called gross substitutability(Tuinstra 2000). The slow convergence is also an issue.

Our breakthrough is to use a root-finding algorithm called Brent's method. It is guaranteed to convergence at a super-liner convergence rate, by combining three methods: the bisection method, the secant method, and the inverse quadratic interpolation (Atkinson 1989). The only conditions for convergence is the continuity of the aggregate demand curve, which typically holds. As far as we know, the use of Brent's method for tâtonnement has not been discussed before. This is probably because adjusting price with such a complex method is not a natural model of the real-world economy. Nonetheless, this limitation is not relevant to our computational economy, since our objective is to obtain the optimal plan, not to model the real-world economy.

The Algorithm

The entire algorithm is summarized below. We call this algorithm as Market-based Iterative Risk Allocation (MIRA).

1. Sets the initial price p and announce it to all agents.
2. Each agent computes its optimal demand at the price $\Delta_i^\star(p)$ by solving Eq.(6), and bids it.
3. Terminate the algorithm if the aggregate demand is equal to the supply; otherwise, adjust the price p by computing one step of Brent's method, and announce it.
4. Go to Step 2.

The price converges to the equilibrium price $p^\star$ in this iterative process. The optimal risk allocation for each agent is its demand at the equilibrium price $\Delta_i^\star(p^\star)$.

Simulation

Simulations were conducted on a machine with Intel(R) Core(TM) i7 CPU clocked at 2.67 GHz and 8GB RAM. See (Ono & Williams 2009) for the used parameters.

To evaluate the efficiency of MIRA algorithm, the computation time of the following three methods were compared:

1. Centralized optimization,
2. Distributed optimization (tâtonnement) with a linear price increment, and
3. MIRA: distributed optimization (tâtonnement) with Brent's method.

Table 1 shows the results. The three algorithms were tested with different problem sizes - two, four, and eight agents. Each algorithm was run 10 times for each problem size with randomly generated constraints. The average running time is shown in the table. The computation of the distributed algorithms was conducted parallelly. Communication delay is not included in the result.

The computation time of the centralized optimization algorithm quickly grows as the problem size increases. Distributed optimization with a linear price increment is even slower than the centralized algorithm.

MIRA, the proposed algorithm, outperforms the other two for all problem sizes. The advantage of MIRA becomes clearer as the problem size increases. More simulation result is presented in (Ono & Williams 2009).

Table 1: Comparison of the computation time of three optimization algorithms.

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

Conclusion

We have developed Market-based Iterative Risk Allocation (MIRA), a multi-agent optimal planning algorithm that operates within user-specified risk bounds. The three key innovations that enabled MIRA were:

1. Extension of the concept of risk allocation to multi-agent system.
2. Derivation of distributed optimization method for multi-agent risk allocation using KKT conditions.
3. Introduction of Brent's method to tâtonnement as a price update rule.

The simulation result showed that MIRA achieved substantial speed-up compared to centralized optimization approach, particularly in a large problem.

Acknowledgments

References

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

Kelly, F.; Maulloo, A.; and Tan, D. 1998. Rate control in communication networks: Shadow prices, proportional fairness and stability. *J. of the Operational Research Society* 49:237–252.

MacKie-Mason, J. K.; Osepayshvili, A.; Reeves, D. M.; and Wellman, M. P. 2004. Price prediction strategies for market-based scheduling. In *Proc. of ICAPS-04*.

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

Ono, M., and Williams, B. C. 2009. Risk allocation for multi-agent systems using tâtonnement. Technical report, MIT CSAIL.

Tuinstra, J., ed. 2000. *Price Dynamics in Equilibrium Models: The Search for Equilibrium and the Emergence of Endogenous Fluctuations*.

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

Wellman, M. P. 1993. A market-oriented programming environment and its application to distributed multicommodity flow problems. *J. of Artificial Intelligence Research*.