



State Space Search for Risk-averse Agents

Patrice Perny, Olivier Spanjaard, Louis-Xavier Storme

► To cite this version:

Patrice Perny, Olivier Spanjaard, Louis-Xavier Storme. State Space Search for Risk-averse Agents. Twentieth International Joint Conference on Artificial Intelligence, Jan 2007, Hyderabad, India. pp.2353-2358. <hal-01311730>

HAL Id: hal-01311730

<https://hal.science/hal-01311730v1>

Submitted on 30 Jun 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

State Space Search for Risk-averse Agents

Patrice Perny

LIP6 - University of Paris 6
4 Place Jussieu
75252 Paris Cedex 05, France
patrice.perny@lip6.fr

Olivier Spanjaard

LIP6 - University of Paris 6
4 Place Jussieu
75252 Paris Cedex 05, France
olivier.spanjaard@lip6.fr

Louis-Xavier Storme

LIP6 - University of Paris 6
4 Place Jussieu
75252 Paris Cedex 05, France
louis-xavier.storme@lip6.fr

Abstract

We investigate search problems under risk in state-space graphs, with the aim of finding optimal paths for risk-averse agents. We consider problems where uncertainty is due to the existence of different scenarios of known probabilities, with different impacts on costs of solution-paths. We consider various non-linear decision criteria (EU, RDU, Yaari) to express risk averse preferences; then we provide a general optimization procedure for such criteria, based on a path-ranking algorithm applied on a scalarized valuation of the graph. We also consider partial preference models like second order stochastic dominance (SSD) and propose a multiobjective search algorithm to determine SSD-optimal paths. Finally, the numerical performance of our algorithms are presented and discussed.

1 Introduction

Various problems investigated in Artificial Intelligence can be formalized as shortest path problems in an implicit state space graph (e.g. path-planning for mobile robots, VLSI layout, internet searching). Starting from a given state, we want to determine an optimal sequence of admissible actions allowing transitions from state to state until a goal state is reached. Here, optimality refers to the minimization of one or several cost functions attached to transitions, representing distances, times, energy consumptions... For such problems, constructive search algorithms like A^* and A_e^* [Hart *et al.*, 1968; Pearl, 1984] for single objective problems or MOA* for multiobjective problems [Stewart and White III, 1991] have been proposed, performing the implicit enumeration of feasible solutions.

An important source of complexity in path-planning problems is the uncertainty attached to some elements of the problem. In some situations, the consequences of actions are not certain and the transitions are only known in probabilities. In some other, the knowledge of the current state is imperfect (partial observability). Finally, the costs of transitions might itself be uncertain. Although many studies concentrate on the two first sources of uncertainty (see the important literature on MDPs and POMDPs, e.g. Puterman, 1994; Kaebbling *et al.*, 1999), some others focus on the uncertainty attached to

transition-costs. For example, when costs are time dependent and representable by random variables, the SDA* algorithm has been introduced to determine the preferred paths according to the stochastic dominance partial order [Wellman *et al.*, 1995]. An extension of this algorithm specifically designed to cope with both uncertainty and multiple criteria has been proposed by Wurman and Wellman [1996].

We consider here another variation of the search problem under uncertainty, that concerns the search of “robust” solution-paths, as introduced by Kouvelis and Yu [1997]. Under total uncertainty, it corresponds to situations where costs of paths might depend on different possible scenarios (states of the world), or different viewpoints (discordant sources of information). Roughly speaking, the aim is to determine paths with “reasonable” cost in all scenarios. Under risk (i.e. when probabilities are known) this problem generalizes to the search of “low-cost/low-risk” paths. Let us consider a simple example:

Example 1 Consider the network pictured on Figure 1 where the initial state is 1 and the goal node is 6. Assume that only two scenarios with known probabilities p_1 and p_2 are relevant concerning the traffic, yielding two different sets of costs on the network. Hence, to each path P^i is associated a vector (x_1^i, x_2^i) , one cost per scenario: $P^1 = \langle 1, 3, 5, 6 \rangle$ with $x^1 = (5, 18)$, $P^2 = \langle 1, 3, 6 \rangle$ with $x^2 = (8, 15)$, $P^3 = \langle 1, 3, 4, 6 \rangle$ with $x^3 = (16, 15)$, $P^4 = \langle 1, 2, 5, 6 \rangle$ with $x^4 = (13, 10)$, $P^5 = \langle 1, 2, 6 \rangle$ with $x^5 = (16, 7)$, $P^6 = \langle 1, 2, 4, 6 \rangle$ with $x^6 = (20, 2)$. Using cost-distributions $X^i = (x_1^i, x_2^i; p_1, p_2)$, $i = 1, \dots, 6$, we want to determine solutions paths associated with low-risk cost-distributions.




Figure 1: A 2-scenarios problem and its representation

This simple problem might prove very hard to solve on larger instances due to the coexistence of two difficulties: the combinatorial nature of the solution space and the existence of several conflicting scenarios on costs. It is important to note that the vector-valued path problem introduced above cannot be reduced to a standard shortest path problem by linear scalarization of cost-vectors without loosing significant information. Assume for example that the arcs of the graph plotted on the left part of Figure 1 are valued according to their expected cost, so that each path P^i receives a weight $w(x^i, p) = p_1 x_1^i + p_2 x_2^i$. Then algorithm A^* used with such scalars weights might output P^1 or P^6 , depending on the relative value of p_1 and p_2 , but neither path P^4 nor P^2, P^5 . This can easily be shown using the right part of Figure 1 where the images of solution-paths are plotted in the valuation space; we can indeed see that P^2, P^4 and P^5 do not belong to the boundary of the convex hull (grey triangle) of the images of paths, thus being excluded from the set of potential winners, as long as a linear criterion is used. This is not satisfactory because P^4 presents a well-balanced profile and might be preferred to P^1 or P^6 by a risk-averse agent. Similarly he might prefer P^2 to P^1 or P^5 to P^6 , depending on probabilities.

Example 1 shows the limitations of linear aggregation functions in decision-making under risk on non-convex domains. To overcome the difficulty, we need to resort to more sophisticated decision criteria to compare cost distributions in term of risk, as those introduced in decision theory. These decision criteria escape linearity either by introducing a transformation of costs as in the Expected Utility Model (EU [von Neumann and Morgenstern, 1947]) or by introducing a probability-transformation as in Yaari's model [Yaari, 1987], or even both as in the Rank-Dependent Utility model (RDU [Quiggin, 1993]). Alternatively, partial comparison models including an idea of risk might be used when the agent's utility function is not known (e.g. Second-order Stochastic Dominance, SSD). The aim of this paper is to incorporate such models in search algorithms to determine low-risk solution paths in implicit graphs.

The paper is organized as follows: in Section 2, we introduce preliminary formal material as well as decision criteria modelling risk-sensitive decision behaviours. In Section 3, we propose a general optimization procedure to find the best paths with respect to such criteria. In Section 4, we propose a multiobjective search algorithm for the determination of SSD-optimal paths. Finally, numerical experiments of algorithms are given in Section 5.

2 Problem Formulation

2.1 Notations and Definitions

We consider a state space graph $G = (N, A)$ where N is a finite set of nodes (possible states), and A is a set of arcs representing feasible transitions between nodes. Formally, we have $A = \{(n, n'), n \in N, n' \in S(n)\}$ where $S(n) \subseteq N$ is the set of all successors of node n (nodes that can be reached from n by a feasible elementary transition). Then $s \in N$ denotes the source of the graph (the initial state), $\Gamma \subseteq N$ the subset of goal nodes, $\mathcal{P}(s, \Gamma)$ the set of all paths from s to a goal node $\gamma \in \Gamma$, and $\mathcal{P}(n, n')$ the set of all paths linking n to

n' . We call *solution-path* a path from s to a goal node $\gamma \in \Gamma$. Throughout the paper, we assume that there exists at least one solution-path.

Following a classical scheme in robust optimization [Kouvelis and Yu, 1997], we consider a finite set $S = \{s_1, \dots, s_m\}$ of possible scenarios, each having possibly a different impact on the transition-costs, and a scenario-dependent valuation $v : A \times S \rightarrow \mathbb{N}$ giving, for any arc $a \in A$ and any scenario $s \in S$ the cost $v(a, s)$ of the transition represented by a . Costs over a path are supposed to be additive, which allows valuation v to be extended from arcs to paths by setting, for any path P and any scenario s , $v(P, s) = \sum_{a \in P} v(a, s)$. In the sequel, we assume that the cost of every solution path is (upper) bounded by a positive constant M .

A cost-vector $x = (x_1, \dots, x_m) \in \mathbb{R}_+^m$ is associated to each path P in the graph in such a way that component $x_i = v(P, s_i)$. Let p_i denote the probability of scenario s_i , with $p_i \geq 0$ for $i = 1, \dots, m$ and $\sum_{i=1}^m p_i = 1$, then a path P with cost-vector x is represented by the distribution $(x_1, \dots, x_m; p_1, \dots, p_m)$. Let \mathcal{L} be the set of probabilistic distributions having a finite support in $[0, M]$. The cost of each path is a random variable X characterized by law $P_X \in \mathcal{L}$, defined for any $B \subseteq [0, M]$, by $P_X(B) = P(\{s \in S : X(s) \in B\})$. For any random variable X , the expected value of X is given by $E(X) = \sum_{i=1}^m p_i x_i$, the cumulative function F_X is given by $F_X(z) = P(\{s \in S : X(s) \leq z\})$ for all $z \in [0, M]$ and the associated decumulative function is denoted $G_X(z) = 1 - F_X(z)$.

2.2 Decision Criteria for Risk-Averse Agents

In the field of decision making under risk, the concept of risk-aversion has been widely investigated, first in the framework of EU theory and then in more general frameworks. Roughly speaking, risk-aversion amounts to preferring a solution with a guaranteed cost to any other risky solution with the same expected cost. This was formalized by Pratt and Arrow [Pratt, 1964; Arrow, 1965] that define *weak risk-aversion* for a weak preference relation \succsim on \mathcal{L} as follows:

Definition 1 *An agent is said to be weakly risk-averse if, for any distribution X in \mathcal{L} , he considers that $E(X)$ is as least as good as X , i.e. $E(X) \succsim X$.*

In EU theory, risk-aversion means that the agent's utility function u on payoffs is increasing and *concave*, the coefficient of risk-aversion of any agent being measured by $-u''(x)/u'(x)$ [Arrow, 1965]. In our context, the counterpart of EU is given by the expected weight function:

$$EW(X) = \sum_{i=1}^m p_i w(x_i) \quad (1)$$

where $w : [0, M] \rightarrow \mathbb{R}$ is a strictly increasing function such that $w(x_i)$ represents the subjective weight (disutility) attached to cost x_i by the agent. Criterion $EW(X)$ is to be minimized since it represents the disutility of any cost distribution X . In the EW model, risk aversion means choosing an increasing and *convex* w in Equation (1), so as to get $EW(E(X)) \leq EW(X)$ for all $X \in \mathcal{L}$.

Despite its intuitive appeal, EU theory does not explain all rational decision making behaviors (e.g. the violation of Savage's sure thing principle [Ellsberg, 1961]). This has led researchers to sophisticate the definition of expected utility. Among the most popular generalizations of EU, let us mention the rank dependent utility introduced by Quiggin [1993], which can be reformulated in our context as follows:

$$RDW(X) = w(x_{(1)}) + \sum_{i=1}^{m-1} \varphi(G_X(x_{(i)})) [w(x_{(i+1)}) - w(x_{(i)})] \quad (2)$$

where (\cdot) represents a permutation on $\{1, \dots, m\}$ such that $x_{(1)} \leq \dots \leq x_{(m)}$, φ is a non-decreasing probability transformation function, proper to any agent, such that $\varphi(0) = 0$ and $\varphi(1) = 1$, and w is a weight function assigning subjective disutility to real costs. This criterion can be interpreted as follows: the weight of a path with cost-distribution X is at least $w(x_{(1)})$ with probability 1. Then the weight might increase from $w(x_{(1)})$ to $w(x_{(2)})$ with probability mass $\varphi(G_X(x_{(1)}))$; the same applies from $w(x_{(2)})$ to $w(x_{(3)})$ with probability-mass $\varphi(G_X(x_{(2)}))$, and so on... When $w(z) = z$ for all z , RDW is known as Yaari's model [Yaari, 1987].

Weak risk-aversion can be obtained in Yaari's model by choosing a probability transformation such that $\varphi(p) \geq p$ for all $p \in [0, 1]$. This holds also for RDW provided function w is convex [Quiggin, 1993]. On the other hand, when φ is the identity function, then RDW boils down to EW . Indeed, considering probabilities $q_{(1)} = 1$ and $q_{(i+1)} = G_X(x_{(i)}) = \sum_{k=i+1}^m p_{(k)}$ for all $i = 1, \dots, m-1$, RDW criterion can be rewritten as follows:

$$RDW(X) = \sum_{i=1}^{m-1} [\varphi(q_{(i)}) - \varphi(q_{(i+1)})] w(x_{(i)}) + \varphi(q_{(m)}) w(x_{(m)}) \quad (3)$$

From this last equation, observing that $q_{(i)} - q_{(i+1)} = p_{(i)}$, we can see that RDW reduces to EW when $\varphi(z) = z$ for all z . Hence, RDW generalizes both EW and Yaari's model. For the sake of generality, we consider RDW in the sequel and investigate the following problem:

RDW Search Problem. *We want to determine a RDW -optimal distribution in the set of all cost distributions of paths in $\mathcal{P}(s, \Gamma)$.*

This problem is NP-hard. Indeed, choosing $w(x) = x$, $\varphi(0) = 0$ and $\varphi(x) = 1$ for all $x \in (0, M]$, we get $RDW(X) = x_{(m)} = \max_i x_i$. Hence RDW minimization in a vector valued graph reduces to the min-max shortest path problem, proved NP-hard by Murthy and Her [1992].

3 Search with RDW

As many other non-linear criteria, RDW breaks the Bellman principle and one cannot directly resort to dynamic programming to compute optimal paths. To overcome this difficulty, we propose an exact algorithm which proceeds in three steps: 1) *linear scalarization*: the cost of every arc is defined as the expected value of its cost distribution; 2) *ranking*: enumeration of paths by increasing order of expected costs; 3) *stopping condition*: stops enumeration when we can prove that a RDW -optimal distribution has been found. Step 2 can be performed by kA^* , an extension of A^* proposed by Galand and

Perny [2006] to enumerate the solution-paths of an implicit graph by increasing order of costs. Before expliciting step 3, we need to establish the following result:

Proposition 1 *For all non-decreasing probability transformations φ on $[0, 1]$ such that $\varphi(q) \geq q$ for all $q \in [0, 1]$, for all non-decreasing and convex weight functions w on $[0, M]$, for all $X \in \mathcal{L}$ we have: $RDW(X) \geq w(E(X))$*

Proof. Since $x_{(i+1)} \geq x_{(i)}$ for $i = 1, \dots, m-1$ and w is non-decreasing, we have: $w(x_{(i+1)}) - w(x_{(i)}) \geq 0$ for all $i = 1, \dots, m-1$. Hence, from Equation (2), $\varphi(q) \geq q$ for all $q \in [0, 1]$ implies that: $RDW(X)$

$$\begin{aligned} &\geq w(x_{(1)}) + \sum_{i=1}^{m-1} G_X(x_{(i)}) [w(x_{(i+1)}) - w(x_{(i)})] \\ &= [1 - G_X(x_{(1)})] w(x_{(1)}) + G_X(x_{(m-1)}) w(x_{(m)}) \\ &\quad + \sum_{i=2}^{m-1} [G_X(x_{(i-1)}) - G_X(x_{(i)})] w(x_{(i)}) \\ &= p_{(1)} w(x_{(1)}) + p_{(m)} w(x_{(m)}) + \sum_{i=2}^{m-1} p_{(i)} w(x_{(i)}) \\ &= EW(X) \geq w(E(X)) \text{ by convexity of } w. \quad \square \end{aligned}$$

Now, let $\{P^1, \dots, P^r\}$ denotes the set of elementary solution-paths in $\mathcal{P}(s, \Gamma)$, with cost distributions X^1, \dots, X^r , indexed in such a way that $E(X^1) \leq E(X^2) \leq \dots \leq E(X^r)$. Each distribution X^j yields cost $x_i^j = v(P^j, s_i)$ with probability p_i for $i = 1, \dots, m$. The sequence of paths $(P^j)_{j=1, \dots, r}$ can be generated by implementing the ranking algorithm of step 2 on the initial graph $G = (N, A)$, using a scalar valuation $v' : A \rightarrow \mathbb{R}_+$ defined by $v'(a) = \sum_{i=1}^m p_i v(a, s_i)$. Indeed, the value of any path P^j in this graph is $v'(P^j) = E(X^j)$ by linearity of expectation.

Now, assume that, during the enumeration, we reach (at step k) a path P^k such that: $w(E(X^k)) \geq RDW(X^{\beta(k)})$ where $\beta(k)$ is the index of a RDW -optimal path in $\{P^1, \dots, P^k\}$, then enumeration can be stopped thanks to:

Proposition 2 *If $w(E(X^k)) \geq RDW(X^{\beta(k)})$ for some $k \in \{1, \dots, r\}$, where $\beta(k)$ is the index of a RDW -minimal path in $\{P^1, \dots, P^k\}$, then $P^{\beta(k)}$ is a RDW -minimal solution-path in $\mathcal{P}(s, \Gamma)$.*

Proof. We know that $P^{\beta(k)}$ is RDW -minimal among the k -first detected paths. We only have to show that no other solution-path can have a lower weight according to RDW . For all $j \in \{k+1, \dots, r\}$ we have: $RDW(X^j) \geq w(E(X^j))$ thanks to Proposition 1. Moreover $E(X^j) \geq E(X^k)$ which implies $w(E(X^j)) \geq w(E(X^k)) \geq RDW(X^{\beta(k)})$. Hence $RDW(X^j) \geq RDW(X^{\beta(k)})$ which shows that $P^{\beta(k)}$ is RDW -minimal over $\mathcal{P}(s, \Gamma)$. \square

Propositions 1 and 2 show that the ranked enumeration of solution-paths performed at step 2 can be interrupted without losing the RDW -optimal solution. This establishes the admissibility of our 3-steps algorithm. Numerical tests performed on different instances and presented in Section 5 indicate that the stopping condition is activated early in the enumeration, which shows the practical efficiency of the proposed algorithm.

4 Dominance-based Search

Functions RDW and EW provide sharp evaluation criteria but require a precise knowledge of the agent's attitude

towards risk (at least to assess the disutility function). In this section we consider less demanding models yet allowing well-founded discrimination between some distributions.

4.1 Dominance Relations

A primary dominance concept to compare cost distributions in \mathcal{L} is the following:

Definition 2 For all $X, Y \in \mathcal{L}$, *Functional Dominance* is defined by: $X \text{ FD } Y \Leftrightarrow [\forall s \in S, X(s) \leq Y(s)]$

For relation FD and any other dominance relation \succeq defined in the sequel, the set of \succeq -optimal distributions in $L \subseteq \mathcal{L}$ is defined by: $\{X \in L : \forall Y \in L, Y \succeq X \Rightarrow X \succeq Y\}$.

When the probabilities of scenarios are known, functional dominance can be refined by *first order stochastic dominance* defined as follows:

Definition 3 For all $X, Y \in \mathcal{L}$, the *First order Stochastic Dominance relation* is defined by:

$$X \text{ FSD } Y \Leftrightarrow [\forall z \in [0, M], G_X(z) \leq G_Y(z)]$$

Actually, the usual definition of FSD involves cumulative distributions F_X applied to payoffs instead of decumulative functions G_X applied to costs. In Definition 3, $X \text{ FSD } Y$ means that X assigns no more probability than Y to events of type: “the cost of the path will go beyond z ”. Hence it is natural to consider that X is at least as good as Y when $X \text{ FSD } Y$.

Relation FSD is clearly related to the *EW* model since $X \text{ FSD } Y$ if and only if $EW(X) \leq EW(Y)$ for all increasing weight function w [Quiggin, 1993]. This gives a nice interpretation to Definition 3 within EU theory, with a useful consequence: if the agent is a EW-minimizer (with any increasing weight function w), then his preferred solutions necessarily belong to the set of FSD-optimal solutions. Now, an even richer dominance relation can be considered:

Definition 4 For all $X, Y \in \mathcal{L}$, the *Second order Stochastic Dominance relation* is defined as follows:

$$X \text{ SSD } Y \Leftrightarrow [\forall z \in [0, M], G_X^2(z) \leq G_Y^2(z)]$$

where $G_X^2(z) = \int_z^M G_X(y) dy$, for all $z \in [0, M]$.

Stochastic Dominance is acknowledged as a standard way of characterizing risk-averse behaviors independently of any utility model. For example, Rothschild and Stiglitz [1970] and Machina and Pratt [1997] provide axiomatic characterizations of SSD in terms of risk using “mean preserving spreads”. As a consequence, an agent is said to be *strongly risk-averse* if he prefers X to Y whenever $X \text{ SSD } Y$. Moreover, SSD has a natural interpretation within EU theory: $X \text{ SSD } Y$ if and only if $EW(X) \leq EW(Y)$ for all increasing and convex weight function w [Quiggin, 1993]. As a nice consequence, we know that whenever an agent is a risk-averse EW-minimizer, then his preferred solutions necessarily belong to the set of SSD-optimal solutions. The same applies to RDW provided $\phi(q) \geq q$ for all $q \in [0, 1]$ and function w is convex (this directly follows from a result of Quiggin [1993]). This shows that, even outside EU theory, SSD appears as a natural preference relation for risk-averse agents. It can be used as a first efficient filtering of risky paths.

Interestingly enough, relations FSD and SSD dominance relations can equivalently be defined by:

$$X \text{ FSD } Y \Leftrightarrow [\forall p \in [0, 1], \check{G}_X(p) \leq \check{G}_Y(p)] \quad (4)$$

$$X \text{ SSD } Y \Leftrightarrow [\forall p \in [0, 1], \check{G}_X^2(p) \leq \check{G}_Y^2(p)] \quad (5)$$

where \check{G}_X and \check{G}_X^2 are inverse functions defined by:

$$\begin{aligned} \check{G}_X(p) &= \inf\{z \in [0, M] : G_X(z) \leq p\} \text{ for } p \in [0, 1], \\ \check{G}_X^2(p) &= \int_0^p \check{G}_X(q) dq, \text{ for } p \in [0, 1] \end{aligned}$$

Since S is finite in our context, X is a discrete distribution; therefore G_X and \check{G}_X are step functions. Moreover G_X^2 and \check{G}_X^2 are piecewise linear functions. Function $\check{G}_X^2(z)$ is known as the Lorenz function. It is commonly used for inequality ordering of positive random variables [Muliere and Scarsini, 1989]. As an illustration, consider Example 1 with $p_1 = 0.4$ and $p_2 = 0.6$. We have: $\check{G}_{X^1}^2(p) = 18p$ for all $p \in [0, 0.6]$, $\check{G}_{X^1}^2(p) = 7.8 + 5p$ for all $p \in [0.6, 1]$, whereas $\check{G}_{X^5}^2(p) = 16p$ for all $p \in [0, 0.4]$, $\check{G}_{X^5}^2(p) = 3.6 + 7p$ for all $p \in [0.4, 1]$; hence $\check{G}_{X^5}^2(p) \leq \check{G}_{X^1}^2(p)$ for all p and therefore $X^5 \text{ SSD } X^1$. This confirms the intuition that path P^5 with cost (16, 7) is less risky than path P^1 with cost (5, 18).

The dominance relations introduced in this subsection being transitive, the sets of FD-optimal elements, FSD-optimal elements and SSD-optimal elements are not empty. Moreover, these sets are nested thanks to the following implications: $X \text{ FD } Y \Rightarrow X \text{ FSD } Y$ and $X \text{ FSD } Y \Rightarrow X \text{ SSD } Y$, for all distributions $X, Y \in \mathcal{L}$. In Example 1 we have $L = \{X^1, X^2, X^3, X^4, X^5, X^6\}$ and $p_1 = 0.4$ and $p_2 = 0.6$. Hence the set of FD-optimal elements is $\{X^1, X^2, X^4, X^5, X^6\}$, the set of FSD-optimal elements is the same and the set of SSD-optimal elements is $\{X^4, X^5, X^6\}$. The next section is devoted to the following:

SSD Search Problem. We want to determine all SSD-optimal distributions in the set of cost distributions of paths in $\mathcal{P}(s, \Gamma)$ and for each of them, at least one solution-path.

4.2 Problem Complexity

To assess complexity of the search, we first make explicit a link between SSD and Generalized Lorenz Dominance, as defined by Marshall and Olkin [1979]. Generalized Lorenz dominance, denoted GLD in the sequel, is based on the definition of Lorenz vector $L(x) = (L_1(x), \dots, L_m(x))$ for any vector $x = (x_1, \dots, x_m)$ where $L_k(x)$ is the sum of the k greatest components of x . Then, relation GLD is defined as follows: $x \text{ GLD } y$ if $L(x)$ Pareto dominates $L(y)$, i.e. $L_k(x) \leq L_k(y)$ for $k = 1, \dots, m$. Now, if $p_i = 1/m$ for $i = 1, \dots, m$ then SSD defined by Equation (5) on distributions reduces to Lorenz dominance on the corresponding cost vectors since $L_k(x) = m\check{G}_X^2(k/m)$. Hence, in the particular case of equally probable scenarios, the SSD search problem reduces to the search of Lorenz non-dominated paths, a NP-hard problem as shown by Perny and Spanjaard [2003]. This shows that the SSD search problem is also NP-hard.

4.3 The SSDA* Algorithm

Consider Example 1 and assume that the two scenarios have equal probabilities, we can see that the preferred subpath

from node 1 to node 5 is $P = \langle 1, 3, 5 \rangle$ with cost $x_P = (3, 10)$ which is preferred to path $P' = \langle 1, 2, 5 \rangle$ with cost $x_{P'} = (11, 2)$ since $(3, 10; 0.5, 0.5)$ SSD $(11, 2; 0.5, 0.5)$. Indeed, we are in the case of Subsection 4.2 (equally probable scenarios) with $L(x_P) = (10, 13)$ and $L(x_{P'}) = (11, 13)$ and obviously $(10, 13)$ Pareto dominates $(11, 13)$. Now, appending path $P'' = \langle 5, 6 \rangle$ with $x_{P''} = (2, 8)$ to P and P' respectively yields path $P^1 = P \cup P''$ with cost $x^1 = (5, 18)$ and path $P^4 = P' \cup P''$ with $x^4 = (13, 10)$. Hence $L(x^4)$ Pareto dominates $L(x^1)$, therefore $(13, 10; 0.5, 0.5)$ SSD $(5, 18; 0.5, 0.5)$ which constitutes a preference reversal and illustrates a violation of Bellman principle, thus invalidating a direct dynamic programming approach (optimal path P^4 would be lost during the search if P' is pruned at node 5 due to P).

However, the problem can be overcome knowing that: *i*) SSD-optimal paths are also FD-optimal; *ii*) FD-optimality satisfies the Bellman principle; *iii*) the set of scenarios being finite, FD-optimality on cost distributions is nothing else but Pareto-optimality on cost-vectors. SSD-optimal distributions might indeed be obtained in two stages: 1) generate FD-optimal solution-paths using Multiobjective A* (MOA*, Stewart and White III, 1991; Mandow and de la Cruz, 2005); 2) eliminate SSD-dominated solutions within the output set. However, FD-optimal solutions being often numerous, it is more efficient to focus directly on SSD-optimal solutions during the search. For this reason we introduce now a refinement of MOA* called SSDA* for the direct determination of SSD-optimal solutions.

As in MOA*, SSDA* expands vector-valued labels (attached to subpaths) rather than nodes. Note that, unlike the scalar case, there possibly exists several Pareto non-dominated paths with distinct cost-vectors to reach a given node; hence several labels can be associate to a same node n . At each step of the search, the set of generated labels is divided into two disjoint sets: a set OPEN of not yet expanded labels and a set CLOSED of already expanded labels. Whenever the label selected for expansion is attached to a solution path, it is stored in a set SOL. Initially, OPEN contains only the label attached to the empty subpath on node s , while CLOSED and SOL are empty. We describe below the essential features of the SSDA* algorithm.

Output: it determines the set of SSD-optimal solution-paths, i.e. solution-paths the distribution of which is SSD-optimal. If several paths have the same G^2 distribution, only one path among them is stored using standard bookkeeping techniques.

Heuristics: like in MOA*, a set $H(n)$ of heuristic cost-vectors is used at any node n since n may be on the path of more than one non-dominated solution. This set estimates the set $H^*(n)$ of non-dominated costs of paths from n to Γ .

Priority: to direct the search we use a *set-valued* label-evaluation function F defined in such a way that, $F(\ell)$, at any label ℓ , estimates the set $F^*(\ell)$ of non-dominated costs of solution paths extending the subpath associated with ℓ . This set $F(\ell)$ is computed from all possible combinations $\{g(\ell) + h : h \in H(n)\}$, where $g(\ell)$ denotes the value of the subpath associated with ℓ and n the node to which ℓ is attached. At each step of the search, SSDA* expands a label ℓ in OPEN such that $F(\ell)$ contains at least one SSD-optimal

cost-vector in $\bigcup_{\ell \in \text{OPEN}} F(\ell)$. Such a label can be chosen, for instance, so as to minimize EW with a convex w function. At goal nodes, this priority rule guarantees to expand only labels attached to SSD-optimal paths.

Pruning: the pruning of labels cannot be done directly with the SSD relation, as shown in the beginning of this subsection. The following pruning rules are used:

RULE 1: at node n , a label $\ell \in \text{OPEN}$ is pruned if there exists another label ℓ' at the same node n such that $g(\ell')$ FD $g(\ell)$. This rule is essentially the same as in MOA* and is justified by the fact that FD-optimality does satisfy the Bellman principle and FD dominance implies SSD dominance. Indeed, labels pruned like this necessarily lead to a FD-dominated paths and therefore cannot lead to SSD-optimal solution paths.

RULE 2: a label $\ell \in \text{OPEN}$ is pruned if for all $f \in F(\ell)$ there exists $\ell' \in \text{SOL}$ such that $g(\ell')$ SSD f . This rule allows an early elimination of uninteresting labels while keeping admissibility of the algorithm provided heuristic H is admissible, i.e. $\forall n \in N, \forall h^* \in H^*(n), \exists h \in H(n)$ s.t. h FD h^* . Indeed, if H is admissible, then for all $f^* \in F^*(\ell)$ there exists $f \in F(\ell)$ such that $f = g(\ell) + h$ FD $g(\ell) + h^* = f^*$, which implies that f SSD f^* and therefore $g(\ell')$ SSD f^* by transitivity of SSD.

Note that deciding whether X SSD Y can be performed in constant time. Indeed, since functions $\tilde{G}_X^2(p)$ and $\tilde{G}_Y^2(p)$ are piecewise linear as indicated in Section 2, their comparison amounts to test for Pareto dominance on the union set of break points of both functions, the cardinality of which is upper bounded by $2m$.

Termination: the process is kept running until the set OPEN becomes empty, i.e. there is no remaining subpath able to reach a new SSD-optimal solution path. By construction, SSDA* develops a subgraph of the one developed by MOA* and the termination derives from the termination of MOA*.

5 Numerical Tests

Various tests have been performed to evaluate the performance of algorithms on randomly generated graphs of different sizes. The number of nodes in these graphs varies from 1000 to 6000 and the number of arcs from 10^5 (for 1000 nodes) to $5 \cdot 10^6$ (for 6000 nodes). Cost vectors are integers randomly drawn within interval $[0, 100]$. Algorithms were implemented in C++. The computational experiments were carried out with a Pentium IV CPU 3.2GHz PC.

Table 1 presents the average performance of algorithms for different classes of instances, characterized by $\#_{nodes}$ (the number of nodes in the graph), and m (the number of scenarios). In each class, we give the average performance computed over 20 different instances. For every class, we give $\#_{SSD}$ the average number of SSD-optimal distributions and t_{SSD} the average time (in seconds) to solve the SSD search problem with SSDA*. Results given in Table 1 show that the average number of SSD-optimal distributions increases slowly with the size of the graph; moreover SSDA* computation times show a good efficiency (less than 15 seconds in worst cases). The two rightmost columns of Table 1 concern the performance in determining RDW-optimal paths with $w(z) = z^2$ and $\varphi(p) = p^{\frac{1}{2}}$. We give $\#_{Gen}$, the aver-

m	#nodes	#NSSD	t_{NSSD}	#Gen	t_{RDW}
2	1000	2.20	0.12	2.70	0.038
	3500	2.25	1.75	3.35	0.561
	6000	2.45	5.75	3.10	1.750
5	1000	5.10	0.25	14.90	0.05
	3500	5.70	4.14	33.05	0.75
	6000	6.60	13.69	30.95	2.36
10	1000	10.75	0.55	83.35	0.08
	3500	14.15	9.47	261.1	1.68
	6000	13.5	30.97	314.5	6.80

Table 1: Performance of the algorithms

age number of paths generated before reaching the stopping condition of Proposition 2, and t_{RDW} the average time of the search in seconds. Values obtained for #Gen show that path enumeration is stopped after a very reasonable number of iterations and computation times are about one second in worst cases. The gain in efficiency when compared to SSDA* is due to the preliminary scalarization of the graph valuation which avoids numerous Pareto-dominance tests during the exploration, but also to the fact that we only seek one RDW-optimal path among NSSD paths. We have performed other experiments which are not reported here to save space: when $\varphi(p) = p$ (EW model) or $w(z) = z$ (Yaari's model), the performance is even slightly better. Moreover, when convexity of w and concavity of φ are increased to enhance risk-aversion, e.g. with $w(z) = z^{10}$ and $\varphi(p) = p^{\frac{1}{10}}$, the performance is not significantly degraded.

6 Conclusion

We have provided efficient exact algorithms to determine low-risk/low-cost solution paths. Algorithm SSDA* proposed in Section 4 provides a subset of paths convenient for a risk-averse agent, without requiring the definition of a disutility function. Moreover, when a disutility criterion is known, more or less risky paths can be efficiently determined with the algorithm proposed in Section 3. In the future, it should be worth investigating optimization based on risk-sensitive models in other dynamic decision making problems, e.g. decision trees or Markov Decision Processes. In that direction, the main problem to deal with is the existence of dynamic inconsistencies induced by such nonlinear models. To face this difficulty, adapting the approaches proposed here to bypass the violation of Bellman principle might be of interest.

References

[Arrow, 1965] K. J. Arrow. *The Theory of Risk Aversion*. Ynjo Jahnsnin Saatio, Helsinki, 1965.

[Ellsberg, 1961] D. Ellsberg. Risk, ambiguity and the Savage axioms. *Quarterly J. of Economics*, 75(4):643–669, 1961.

[Galand and Perny, 2006] L. Galand and P. Perny. Compromise search in state space graphs. *Proc. of ECAI*, 2006.

[Hart et al., 1968] P. E. Hart, N. J. Nilsson, and B. Raphael. A formal basis for the heuristic determination of minimum cost paths. *IEEE Trans. Syst. and Cyb.*, SSC-4 (2):100–107, 1968.

[Kaebling et al., 1999] L.P. Kaebling, M. Littman, and A. Cassandra. Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101:99–134, 1999.

[Kouvelis and Yu, 1997] P. Kouvelis and G. Yu. *Robust discrete optimization and its applications*. Kluwer Academic Publisher, 1997.

[Machina and Pratt, 1997] M.J. Machina and J. W. Pratt. Increasing risk: Some direct constructions. *Journal of Risk and Uncertainty*, 14:103–127, 1997.

[Mandow and de la Cruz, 2005] L. Mandow and J.L. Prez de la Cruz. A new approach to multiobjective A* search. In *Proceedings of IJCAI-05*, pages 218–223, 2005.

[Marshall and Olkin, 1979] W. Marshall and I. Olkin. *Inequalities: Theory of Majorization and its Applications*. Academic Press, London, 1979.

[Muliere and Scarsini, 1989] P. Muliere and M. Scarsini. A note on stochastic dominance and inequality measures. *Journal of Economic Theory*, 49:314–323, 1989.

[Murthy and Her, 1992] I. Murthy and S. Her. Solving min-max shortest-path problems on a network. *Naval Research Logistics*, 39:669–689, 1992.

[Pearl, 1984] J. Pearl. *Heuristics. Intelligent Search Strategies for Computer Problem Solving*. Addison Wesley, 1984.

[Perny and Spanjaard, 2003] P. Perny and O. Spanjaard. An axiomatic approach to robustness in search problems with multiple scenarios. In *Proc. of UAI'03*, pages 469–476, 2003.

[Pratt, 1964] J. Pratt. Risk aversion in the small and in the large. *Econometrica*, 32(1):122–136, 1964.

[Puterman, 1994] M.L. Puterman. *Markov decision processes, discrete stochastic dynamic programming*. Wiley & Sons, 1994.

[Quiggin, 1993] J. Quiggin. *Generalized Expected Utility Theory: The Rank-Dependent Model*. Kluwer, 1993.

[Rotschild and Stiglitz, 1970] M. Rotschild and J. Stiglitz. Increasing risk: I. a definition. *Journal of Economic Theory*, 2:225–243, 1970.

[Stewart and White III, 1991] B. S. Stewart and C. C. White III. Multiobjective A*. *Journal of ACM*, 38(4):775–814, 1991.

[von Neumann and Morgenstern, 1947] J. von Neumann and O. Morgenstern. *Theory of games and economic behavior*. 2nd Ed. Princeton University Press, 1947.

[Wellman et al., 1995] M.P. Wellman, K. Larson, M. Ford, and P.R. Wurman. Path planning under time-dependent uncertainty. In *Proceedings of UAI*, pages 532–539, 1995.

[Wurman and Wellman, 1996] P.R. Wurman and M.P. Wellman. Optimal factory scheduling using stochastic dominance A*. In *Proceedings of UAI*, pages 554–563, 1996.

[Yaari, 1987] M.E. Yaari. The dual theory of choice under risk. *Econometrica*, 55:95–115, 1987.