

Robust Shortest Path Planning and Semicontractive Dynamic Programming

Dimitri P. Bertsekas[†]

Abstract

In this paper we consider shortest path problems in a directed graph where the transitions between nodes are subject to uncertainty. We use a minimax formulation, where the objective is to guarantee that a special destination state is reached with a minimum cost path even under the worst possible instance of the uncertainty. Problems of this type arise, among others, in planning and pursuit-evasion contexts, and in model predictive control. Our analysis makes use of the recently developed theory of abstract semicontractive dynamic programming models. We investigate questions of existence and uniqueness of solution of the optimality equation, existence of optimal paths, and the validity of various algorithms patterned after the classical methods of value and policy iteration, as well as a new Dijkstra-like algorithm for problems with nonnegative arc lengths.

1. INTRODUCTION

In this paper, we discuss shortest path problems that embody a worst-case view of uncertainty. These problems relate to several other types of problems arising in stochastic and minimax control, model predictive control, Markovian decision processes, planning, sequential games, robust and combinatorial optimization, and solution of discretized large-scale differential equations. Consequently, our analysis and algorithms relate to a large body of existing theory. However, in this paper we rely on a recently developed abstract dynamic programming theory of semicontractive problems, and capitalize on general results developed in the context of this theory [Ber13]. We first discuss informally these connections and we survey the related literature.

Relations with Other Problems and Literature Review

The closest connection to our work is the classical shortest path problem where the objective is to reach a destination node with a minimum length path from every other node in a directed graph. This is a fundamental problem that has an enormous range of applications and has been studied extensively (see e.g., the surveys [Dre69], [GaP88], and many textbooks, including [Roc84], [AMO89], [Ber98], [Ber05]). The assumption is that at any node x , we may determine a successor node y from a given set of possible successors, defined by the arcs (x, y) of the graph that are outgoing from x .

In some problems, however, following the decision at a given node, there is inherent uncertainty about the successor node. In a stochastic formulation, the uncertainty is modeled by a probability distribution over the set of successors, and our decision is then to choose at each node one distribution out of a given set

[†] Dimitri Bertsekas is with the Dept. of Electr. Engineering and Comp. Science, and the Laboratory for Information and Decision Systems, M.I.T., Cambridge, Mass., 02139.

of distributions. The resulting problem, known as stochastic shortest path problem (also known as transient programming problem), is a total cost infinite horizon Markovian decision problem, with a substantial analytical and algorithmic methodology, which finds extensive applications in problems of motion planning, robotics, and other problems where the aim is to reach a goal state with probability 1 under stochastic uncertainty (see [Pal67], [Der70], [Pli78], [Whi82], [BeT89], [BeT91], [Put94], [HCP99], [HiW05], [JaC06], [LaV06], [Bon07], [Ber12], [BeY13], [CCV13], [YuB13a]). Another important area of application is large-scale computation for discretized versions of differential equations (such as the Hamilton-Jacobi-Bellman equation, and the eikonal equation); see [GoR85], [Fal87], [KuD92], [BGM95], [Tsi95], [PBT98], [Set99a], [Set99b], [Vla08], [AIM12], [ChV12], [AnV13], [Mir14].

In this paper, we introduce a sequential minimax formulation of the shortest path problem, whereby the uncertainty is modeled by set membership: at a given node, we may choose one subset out of a given collection of subsets of nodes, and the successor node on the path is chosen from this subset by an antagonistic opponent. Our principal method of analysis is dynamic programming (DP for short). Related problems have been studied for a long time, in the context of control of uncertain discrete-time dynamic systems with a set membership description of the uncertainty (starting with the theses [Wit66] and [Ber71], and followed up by many other works; see e.g., the monographs [BaB91], [KuV97], [BIM08], the survey [Bla99], and the references given there). These problems are relevant for example in the context of model predictive control under uncertainty, a subject of great importance in the current practice of control theory (see e.g., the surveys [MoL99], [MRR00], and the books [Mac02], [CaB04], [RaM09]; model predictive control with set membership disturbances is discussed in the thesis [Ker00] and the text [Ber05], Section 6.5.2).

Sequential minimax problems have also been studied in the context of sequential games (starting with the paper [Sha53], and followed up by many other works, including the books [BaB91], [FiV96], and the references given there). Sequential games that involve shortest paths are particularly relevant; see the works [PaB99], [GrJ08], [Yu11], [BaL13]. An important difference with some of the works on sequential games is that in our minimax formulation, we assume that the antagonistic opponent knows the decision and corresponding subset of successor nodes chosen at each node. Thus in our problem, it would make a difference if the decisions at each node were made with advance knowledge of the opponent's choice ("min-max" is typically not equal to "max-min" in our context). Generally shortest path games admit a simpler analysis when the arc lengths are assumed nonnegative or positive (as is done for example in the recent works [GrJ08], [BaL13]), when the problem inherits the structure of negative DP (see [Str66], or the texts [Put94], [Ber12]) or abstract monotone increasing abstract DP models (see [Ber77], [BeS78], [Ber13]). However, our formulation and line of analysis is based on the recently introduced abstract semicontractive DP model of [Ber13], and allows negative as well as nonnegative arc lengths. A further specialized analysis of our problem is possible when the arc lengths are restricted to be nonnegative, similar to the related paper [Ber14], but this lies beyond the scope of this paper.

An important application of our shortest path problems is in pursuit-evasion (or search and rescue) contexts, whereby a team of "pursuers" are aiming to reach one or more "evaders" that move unpredictably. Problems of this kind have been studied extensively from different points of view (see e.g., [Par76], [MHG88], [BaS91], [HKR93], [BFS94], [BBF95], [GLL99], [VKS02], [LaV06], [AHK07], [BBH08], [BaL13]). For our shortest path formulation to be applicable to such a problem, the pursuers and the evaders must have perfect information about each others' positions, and the Cartesian product of their positions (the state of the system) must be restricted to the finite set of nodes of a given graph, with known transition costs (i.e., a "terrain map" that is known a priori).

We may deal with pursuit-evasion problems with imperfect state information and set-membership

uncertainty by means of a reduction to perfect state information, which is based on set membership estimators and the notion of a sufficiently informative function, introduced in the thesis [Ber71] and in the subsequent paper [BeR73]. In this context, the original imperfect state information problem is reformulated as a problem of perfect state information, where the states correspond to subsets of nodes of the original graph (the set of states that are consistent with the observation history of the system, in the terminology of set membership estimation). Thus, since X has a finite number of nodes, the reformulated problem still involves a finite (but much larger) number of states, and may be dealt with using the methodology of this paper. Note that the problem reformulation just described is also applicable to general minimax control problems with imperfect state information, not just to pursuit-evasion problems.

Our work is also related to the subject of robust optimization (see e.g., the book [BGN09] and the recent survey [BBC11]), which includes minimax formulations of general optimization problems with set membership uncertainty. However, our emphasis here will revolve around the destination node and the requirement for termination, which is the salient feature and the essential structure of shortest path problems. Moreover, a difference with other works on robust shortest path selection (see e.g., [YuY98], [BeS03], [MoG04], [EGB12]) is that in our work the uncertainty about transitions or arc cost data at a given node is decoupled from the corresponding uncertainty at other nodes. This allows a DP formulation of our problem.

Because the problem of this paper has not been considered earlier, at least as formulated here, our results are new. The line of analysis is also new, and is based on the connection with the recent theory of abstract semicontractive DP mentioned earlier. In addition to simpler proofs, a major benefit of this abstract line of treatment is deeper insight into the structure of our problem, and the nature of our analytical and computational results. Several related problems, involving for example an additional stochastic type of uncertainty, admit a similar treatment. Some of these problems are described in the last section, and their analysis and associated algorithms are subjects for further research.

Robust Shortest Path Problem Formulation

To formally describe our problem, we consider a graph with a finite set of nodes $X \cup \{t\}$ and a finite set of directed arcs $\mathcal{A} \subset \{(x, y) \mid x, y \in X \cup \{t\}\}$, where t is a special node called the *destination*. At each node $x \in X$ we may choose a control or action u from a nonempty set $U(x)$, which is a subset of a finite set U . Then a successor node y is selected by an antagonistic opponent from a nonempty set $Y(x, u) \subset X \cup \{t\}$ and a cost $g(x, u, y)$ is incurred. The destination node t is absorbing and cost-free, in the sense that the only outgoing arc from t is (t, t) and we have $g(t, u, t) = 0$ for all $u \in U(t)$.

A policy is defined to be a function μ that assigns to each node $x \in X$ a control $\mu(x) \in U(x)$. We denote the finite set of all policies by \mathcal{M} . A *possible path under a policy μ starting at node $x_0 \in X$* is an arc sequence of the form

$$p = \{(x_0, x_1), (x_1, x_2), \dots\}, \quad (1.1)$$

such that $x_{k+1} \in Y(x_k, \mu(x_k))$ for all $k \geq 0$. The set of all possible paths under μ starting at x is denoted by $P(x, \mu)$. The length of a path $p \in P(x, \mu)$ is defined by

$$L_{p, \mu} = \sum_{k=0}^{\infty} g(x_k, \mu(x_k), x_{k+1}),$$

if the series above is convergent, and more generally by

$$L_{p, \mu} = \limsup_{m \rightarrow \infty} \sum_{k=0}^m g(x_k, \mu(x_k), x_{k+1}), \quad (1.2)$$

if it is not. Note that the length of p depends on the corresponding policy, and not just on the arcs of p .

For completeness, we also define the length of a portion

$$\{(x_i, x_{i+1}), (x_{i+1}, x_{i+2}), \dots, (x_m, x_{m+1})\}$$

of a path $p \in P(x_0, \mu)$, consisting of a finite number of consecutive arcs, by

$$\sum_{k=i}^m g(x_k, \mu(x_k), x_{k+1}).$$

When confusion cannot arise we will also refer to such a finite-arc portion as a path. However, in the absence of a clear implication to the contrary, a path in what follows will mean an infinite arc sequence of the form (1.1). Of special interest are *cycles*, that is, paths of the form $\{(x_i, x_{i+1}), (x_{i+1}, x_{i+2}), \dots, (x_{i+m}, x_i)\}$. Paths that do not contain any cycle other than the self-cycle (t, t) are called *simple*.

A path p is said to be *terminating* if it has the form

$$p = \{(x_0, x_1), (x_1, x_2), \dots, (x_m, t), (t, t), \dots\}, \quad (1.3)$$

where m is a positive integer, and x_0, \dots, x_m are nondestination nodes. Since $g(t, u, t) = 0$ for all $u \in U(t)$, for a policy μ , the length of a terminating path $p \in P(x_0, \mu)$ of the form (1.3), is given by

$$L_{p, \mu} = g(x_m, \mu(x_m), t) + \sum_{k=0}^{m-1} g(x_k, \mu(x_k), x_{k+1}),$$

and is equal to the length of its initial portion that consists of the first $m + 1$ arcs.

An important characterization of a policy μ is provided by the subset of arcs [other than (t, t)] that may be included in paths corresponding to μ :

$$\mathcal{A}_\mu = \cup_{x \in X} \{(x, y) \mid y \in Y(x, \mu(x))\}. \quad (1.4)$$

We say that \mathcal{A}_μ is *destination-connected* if for each $x \in X$ there exists a terminating path in $P(x, \mu)$. We say that μ is *proper* if the subgraph of arcs \mathcal{A}_μ is acyclic (i.e., contains no cycles). Thus μ is proper if and only if all the paths in $\cup_{x \in X} P(x, \mu)$ are simple and hence terminating (equivalently μ is proper if and only if \mathcal{A}_μ is destination-connected and has no cycles). The term “proper” is consistent with a similar term in stochastic shortest path problems, where it indicates a policy under which the destination is reached with probability 1, see e.g., [Pal67], [BeT89], [BeT91]. If μ is not proper, it is called *improper*, in which case the subgraph of arcs \mathcal{A}_μ must contain a cycle; see the examples of Fig. 1.1.

For a proper μ , we associate with every $x \in X$ the worst-case path length over the finite set of possible paths starting from x , which is denoted by

$$J_\mu(x) = \max_{p \in P(x, \mu)} L_{p, \mu}, \quad x \in X. \quad (1.5)$$

Thus $J_\mu(x)$ is the length of the *longest* path from x to t in the subgraph of arcs \mathcal{A}_μ . Since this subgraph is acyclic, $-J_\mu(x)$ can be obtained by solving the shortest path problem obtained when the signs of the arc lengths $g(x, \mu(x), y)$, $(x, y) \in \mathcal{A}_\mu$, are reversed. In particular, J_μ is real-valued, and it is the unique solution

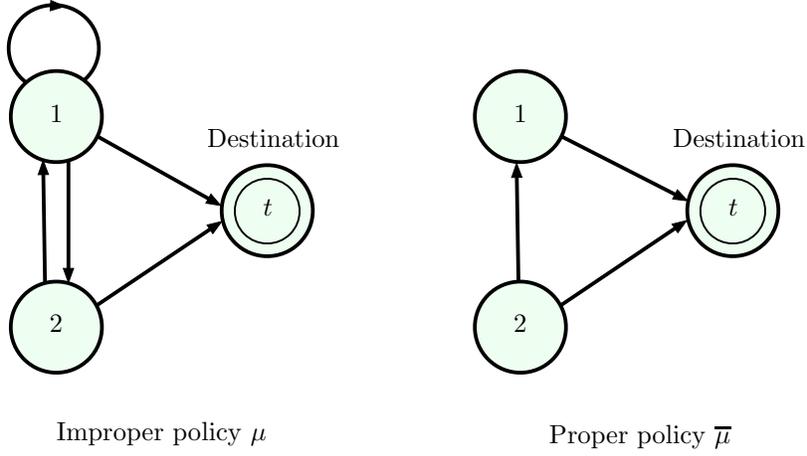


Figure 1.1. Graph representation of a robust shortest path problem with $X = \{1, 2\}$, two controls at node 1, and one control at node 2. There are two policies, μ and $\bar{\mu}$, corresponding to the two controls at node 1. The figure shows the subgraphs of arcs \mathcal{A}_μ and $\mathcal{A}_{\bar{\mu}.t}$

of the optimality equation (or Bellman equation) for the longest path problem associated with the proper policy μ and the acyclic subgraph of arcs \mathcal{A}_μ :

$$J_\mu(x) = \begin{cases} \max_{y \in Y(x, \mu(x))} [g(x, \mu(x), y) + J_\mu(y)] & \text{if } x \in X, \\ 0 & \text{if } x = t. \end{cases} \quad (1.6)$$

We extend the definition of J_μ to any policy μ , proper or improper, as follows:

$$J_\mu(x) = \limsup_{k \rightarrow \infty} \sup_{p \in P(x, \mu)} L_{p, \mu}^k, \quad x \in X, \quad (1.7)$$

where $L_{p, \mu}^k$ is the sum of the lengths of the first k arcs in the path p .

Our problem is to find an optimal proper policy, i.e., one that minimizes $J_\mu(x)$ over all proper μ , simultaneously for all $x \in X$, under assumptions that parallel those for the classical shortest path problem. We refer to this as the problem of *robust shortest path selection* (RSP for short). Note that in our problem, *reaching the destination starting from any node is a requirement*, regardless of the choices of the hypothetical antagonistic opponent. This is significant since there may exist improper policies μ and nodes x such that $J_\mu(x) < \inf_{\mu': \text{proper}} J_{\mu'}(x)$, because some paths in \mathcal{A}_μ contain cycles with negative length. While our analysis with address this possibility, our assumptions for the most part will exclude it (see Assumption 2.1 in the next section).

Of course for the RSP problem to have a feasible solution and thus be meaningful, there must exist at least one proper policy, and this may be restrictive for a given problem. We note that the existence of a proper policy may be verified by applying a reachability algorithm, which constructs the sequence $\{N_k\}$ of sets

$$N_{k+1} = N_k \cup \{x \in X \cup \{t\} \mid \text{there exists } u \in U(x) \text{ with } Y(x, u) \subset N_k\},$$

starting with $N_0 = \{t\}$ (see [Ber71], [BeR71]). A proper policy exists if and only if this algorithm stops with a final set $\cup_k N_k$ equal to $X \cup \{t\}$. If there is no proper policy, this algorithm will stop with $\cup_k N_k$ equal to the strict subset of $X \cup \{t\}$ of nodes starting from which there exists a terminating path under some policy. The problem may then be reformulated over the reduced graph consisting of the node set $\cup_k N_k$.

One may deal with cases where the existence of a proper policy is unknown by introducing for every x an artificial “termination action” \bar{u} into $U(x)$ [i.e., a \bar{u} with $Y(x, \bar{u}) = \{(x, t)\}$], associated with very large length [i.e., $g(x, \bar{u}, t) = \bar{g} \gg 1$]. Then the policy $\bar{\mu}$ that selects the termination action at each x is proper and has cost function $J_{\bar{\mu}}(x) \equiv \bar{g}$. In the problem thus reformulated the optimal cost over proper policies will be unaffected for all nodes x for which there exists a proper policy μ with $J_{\mu}(x) < \bar{g}$. Since for a proper μ , the cost $J_{\mu}(x)$ is bounded above by the number of nodes in X times the largest arc length, a suitable value of \bar{g} is readily available.

In Section 2 we will formulate RSP in a way that the semicontractive DP framework can be applied. In Section 3, we will describe briefly this framework and we will quote the results that will be useful to us. In Section 4, we will develop our main analytical results for RSP. In Section 5, we will discuss algorithms of the value and policy iteration type, by specializing corresponding algorithms of semicontractive DP, and by adapting available algorithms for stochastic shortest path problems. Among others, we will give a Dijkstra-like algorithm for problems with nonnegative arc lengths, which terminates in a number of iterations equal to the number of nodes in the graph, and has low order polynomial complexity.

2. MINIMAX FORMULATION

In this section we will reformulate RSP as a minimax problem, whereby given a policy μ , an antagonistic opponent selects a successor node $y \in Y(x, \mu(x))$ for each $x \in X$, with the aim of maximizing the lengths of the resulting paths. We will initially assume the following.

Assumption 2.1:

- (a) There exists at least one proper policy.
- (b) For every improper policy μ , all cycles in the subgraph of arcs \mathcal{A}_{μ} have positive length.

The preceding assumption parallels and generalizes the typical assumptions in the classical deterministic shortest path problem, i.e., the case where $Y(x, \mu)$ consists of a single arc. Then condition (a) is equivalent to assuming that each node is connected to the destination with a path, while condition (b) is equivalent to assuming that all directed cycles in the graph have positive length. Later in Section 4 we will replace “positive length” with “nonnegative length” in condition (b) above, in which case the condition is automatically satisfied when all arc lengths $g(x, u, y)$ are nonnegative.

We consider the problem of finding

$$J^*(x) = \min_{\mu \in \mathcal{M}} J_{\mu}(x), \tag{2.1}$$

and a policy attaining the minimum above, simultaneously for all $x \in X$. Note that *the minimization is over all policies, not just the proper policies*. However, under Assumption 2.1 it will turn out that improper policies cannot attain the minimum above, so minimizing $J_{\mu}(x)$ over all policies is equivalent to minimizing it over just the proper policies.

Embedding Within an Abstract DP Model

We will now reformulate the minimax problem of Eq. (2.1) more abstractly, by expressing it in terms of the mapping that appears in Bellman's equation (1.6), thereby bringing to bear the theory of abstract DP. We denote by $E(X)$ the set of functions $J : X \mapsto [-\infty, \infty]$, and by $R(X)$ the set of functions $J : X \mapsto (-\infty, \infty)$. Note that since X is finite, $R(X)$ can be viewed as a finite-dimensional Euclidean space. We introduce the mapping $H : X \times U \times E(X) \mapsto [-\infty, \infty]$ given by

$$H(x, u, J) = \max_{y \in Y(x, u)} [g(x, u, y) + \tilde{J}(y)], \quad (2.2)$$

where for any $J \in E(X)$ we denote by \tilde{J} the function of $E(X \cup \{t\})$ given by

$$\tilde{J}(y) = \begin{cases} J(y) & \text{if } y \in X, \\ 0 & \text{if } y = t. \end{cases} \quad (2.3)$$

We consider for each policy μ , the mapping $T_\mu : E(X) \mapsto E(X)$, defined by

$$(T_\mu J)(x) = H(x, \mu(x), J), \quad x \in X, \quad (2.4)$$

and we note that for a proper policy μ , the fixed point equation $J_\mu = T_\mu J_\mu$ is identical to the optimality equation (1.6). We also consider the mapping $T : E(X) \mapsto E(X)$ defined by

$$(TJ)(x) = \min_{u \in U(x)} H(x, u, J), \quad x \in X, \quad (2.5)$$

also equivalently written as

$$(TJ)(x) = \min_{\mu \in \mathcal{M}} (T_\mu J)(x), \quad x \in X. \quad (2.6)$$

We denote by T^k and T_μ^k the k -fold compositions of the mappings T and T_μ with themselves, respectively.

Let us introduce the zero function, which we denote by \bar{J} :

$$\bar{J}(x) \equiv 0, \quad x \in X.$$

Using Eqs. (2.2)-(2.4), we see that for any $\mu \in \mathcal{M}$ and $x \in X$, $(T_\mu^k \bar{J})(x)$ is the result of the k -stage DP algorithm that computes the length of the longest path under μ that starts at x and consists of k arcs, i.e.,

$$(T_\mu^k \bar{J})(x) = \sup_{p \in P(x, \mu)} L_{p, \mu}^k, \quad x \in X.$$

Thus the definition (1.7) of J_μ can be written in the alternative and equivalent form

$$J_\mu(x) = \limsup_{k \rightarrow \infty} (T_\mu^k \bar{J})(x), \quad x \in X. \quad (2.7)$$

The results that we will show under Assumption 2.1 generalize the main analytical results for the classical deterministic shortest path problem. These, stated in abstract form, are the following:

- (a) J^* is the unique fixed point of T within $R(X)$, and we have $T^k J \rightarrow J^*$ for all $J \in R(X)$.
- (b) Only proper policies can be optimal, and there exists an optimal proper policy, i.e., a proper μ^* with $J_{\mu^*} = J^*$.
- (c) A policy μ is optimal if and only if it attains the minimum in the equation $J^* = TJ^*$, i.e., $T_\mu J^* = TJ^*$.

Proofs of these results from first principles are quite complex. However, fairly easy proofs can be obtained by embedding RSP, and its more abstract version of minimizing the function J_μ of Eq. (2.7) over $\mu \in \mathcal{M}$, within the abstract semicontractive DP framework introduced in [Ber13]. In particular, we will use general results for this framework, which we will summarize in the next section.

3. SEMICONTRACTIVE DP MODELS

We will now view the problem of minimizing over $\mu \in \mathcal{M}$ the cost function J_μ , as given in the abstract form (2.7), as a special case of a semicontractive DP model. We first provide a brief review of this model, with a notation that corresponds to the one used in the preceding section.

The starting point is a set of states X , a set of controls U , and a control constraint set $U(x) \subset U$ for each $x \in X$. For the general framework of this section, X and U are arbitrary sets. A policy is a mapping $\mu : X \mapsto U$ with $\mu(x) \in U(x)$ for all $x \in X$, and the set of all policies is denoted by \mathcal{M} . For each policy μ , we are given a mapping $T_\mu : E(X) \mapsto E(X)$ that is monotone in the sense that for any two $J, J' \in E(X)$,

$$J \leq J' \quad \Rightarrow \quad T_\mu J \leq T_\mu J'.$$

The cost function of μ is defined as

$$J_\mu(x) = \limsup_{m \rightarrow \infty} (T_\mu^m \bar{J})(x), \quad x \in X,$$

where \bar{J} is some given function in $E(X)$. The objective is to find

$$J^*(x) = \inf_{\mu \in \mathcal{M}} J_\mu(x)$$

for each $x \in X$, and a policy μ such that $J_\mu = J^*$, if one exists. Based on the correspondences with Eqs. (2.2)-(2.7), it can be seen that the shortest path problem of the preceding section is the special case of the problem of this section, where X and U are finite sets, T_μ is defined by Eq. (2.4), and \bar{J} is the zero function.

In contractive models, the mappings T_μ are assumed to be contractions, with respect to a common weighted sup-norm and with a common contraction modulus, in the subspace of functions in $E(X)$ that are bounded with respect to the weighted sup-norm. These models have a strong analytical and algorithmic theory, which dates to [Den67]; see also [BeS78], Ch. 3, and recent extensive treatments given in Chapters 1-3 of [Ber12], and Ch. 2 of [Ber13]. In semicontractive models, only some policies have a contraction-like property. This property is captured by the notion of S -regularity of a policy introduced in [Ber13]. More specifically, given a set of functions $S \subset E(X)$, we say that a policy μ is S -regular if:

- (a) $J_\mu \in S$ and $J_\mu = T_\mu J_\mu$.
- (b) $\lim_{k \rightarrow \infty} T_\mu^k J = J_\mu$ for all $J \in S$.

A policy that is not S -regular is called S -irregular. Roughly, μ is S -regular if J_μ is an asymptotically stable equilibrium point of T_μ within S . An important case of an S -regular μ is when S is a complete subset of a metric space and T_μ is a contraction with respect to the metric of that space.

There are several different choices of S , which may be useful depending on the context, such as for example $R(X)$, $E(X)$, $\{J \in R(X) \mid J \geq \bar{J}\}$, $\{J \in E(X) \mid J \geq \bar{J}\}$, and others. There are also several sets of assumptions and corresponding results, which are given in [Ber13] and will be used to prove our analytical results for the RSP problem. In this paper, we will use $S = R(X)$, but for ease of reference, we will quote results from [Ber13] for a general set S .

We give below an assumption relating to semicontractive models, which is Assumption 3.2.1 of [Ber13].

Assumption 3.1: In the semicontractive model of this section with a set $S \subset R(X)$ the following hold:

(a) S contains \bar{J} , and has the property that if J_1, J_2 are two functions in S , then S contains all functions J with $J_1 \leq J \leq J_2$.

(b) The function \hat{J} given by

$$\hat{J}(x) = \inf_{\mu: S\text{-regular}} J_\mu(x), \quad x \in X,$$

belongs to S .

(c) For each S -irregular policy μ and each $J \in S$, there is at least one state $x \in X$ such that

$$\limsup_{k \rightarrow \infty} (T_\mu^k J)(x) = \infty.$$

(d) The control set U is a metric space, and the set

$$\{\mu(x) \mid (T_\mu J)(x) \leq \lambda\}$$

is compact for every $J \in S$, $x \in X$, and $\lambda \in \mathfrak{R}$.

(e) For each sequence $\{J_m\} \subset S$ with $J_m \uparrow J$ for some $J \in S$ we have

$$\lim_{m \rightarrow \infty} (T_\mu J_m)(x) = (T_\mu J)(x), \quad \forall x \in X, \mu \in \mathcal{M}.$$

(f) For each function $J \in S$, there exists a function $J' \in S$ such that $J' \leq J$ and $J' \leq TJ'$.

The following two propositions are given in [Ber13] as Prop. 3.2.1 and Lemma 3.2.4, respectively. Our analysis will be based on these two propositions.

Proposition 3.1: Let Assumption 3.1 hold. Then:

(a) The optimal cost function J^* is the unique fixed point of T within the set S .

(b) We have $T^k J \rightarrow J^*$ for all $J \in S$. Moreover, there exists an optimal S -regular policy.

(c) A policy μ^* is optimal if and only if $T_{\mu^*} J^* = TJ^*$.

(d) For any $J \in S$, if $J \leq TJ$ we have $J \leq J^*$, and if $J \geq TJ$ we have $J \geq J^*$.

Proposition 3.2: Let Assumption 3.1(b),(c),(d) hold. Then:

- (a) The function \hat{J} of Assumption 3.1(b) is the unique fixed point of T within S .
- (b) Every policy μ satisfying $T_\mu \hat{J} = T \hat{J}$ is optimal within the set of S -regular policies, i.e., μ is S -regular and $J_\mu = \hat{J}$. Moreover, there exists at least one such policy.

The second proposition is useful for situations where only some of the conditions of Assumption 3.1 are satisfied. This happens when \hat{J} , the infimum of J_μ over just the S -regular policies μ , is different from J^* , which is the infimum over all policies.

4. SEMICONTRACTIVE MODELS AND SHORTEST PATH PROBLEMS

We will now apply the preceding two propositions to the minimax formulation of the RSP problem: minimizing over all $\mu \in \mathcal{M}$ the cost $J_\mu(x)$ as given by Eq. (1.7) for both proper and improper policies. We will first derive some preliminary results. The following proposition clarifies the properties of J_μ when μ is improper.

Proposition 4.1: Let μ be an improper policy and let J_μ be its cost function as given by Eq. (1.7).

- (a) If all cycles in the subgraph of arcs \mathcal{A}_μ have nonpositive length, $J_\mu(x) < \infty$ for all $x \in X$.
- (b) If all cycles in the subgraph of arcs \mathcal{A}_μ have nonnegative length, $J_\mu(x) > -\infty$ for all $x \in X$.
- (c) If all cycles in the subgraph of arcs \mathcal{A}_μ have zero length, J_μ is real-valued.
- (d) If there is a positive length cycle in the subgraph of arcs \mathcal{A}_μ , we have $J_\mu(x) = \infty$ for at least one node $x \in X$. More generally, for each $J \in R(X)$, we have $\limsup_{k \rightarrow \infty} (T_\mu^k J)(x) = \infty$ for at least one $x \in X$.

Proof: Any path with a finite number of arcs, can be decomposed into a simple path, and a finite number of cycles (see e.g., the path decomposition theorem of [Ber98], Prop. 1.1, and Exercise 1.4). Since there are only a finite number of simple paths under μ , their length is bounded above and below. Thus in part (a) the length of all paths with a finite number of arcs is bounded above, and in part (b) it is bounded below, implying that $J_\mu(x) < \infty$ for all $x \in X$ or $J_\mu(x) > -\infty$ for all $x \in X$, respectively. Part (c) follows by combining parts (a) and (b).

To show part (d), consider a path p , which consists of an infinite repetition of the positive length cycle that is assumed to exist. Let $C_\mu^k(p)$ be the length of the path that consists of the first k cycles in p . Then $C_\mu^k(p) \rightarrow \infty$ and $C_\mu^k(p) \leq J_\mu(x)$ for all k , where x is the first node in the cycle, thus implying that $J_\mu(x) = \infty$. Moreover for every $J \in R(X)$ and all k , $(T_\mu^k J)(x)$ is the maximum over the lengths of the k -arc

paths that start at x , plus a terminal cost that is equal to either $J(y)$ (if the terminal node of the k -arc path is $y \in X$), or 0 (if the terminal node of the k -arc path is the destination). Thus we have,

$$(T_\mu^k \bar{J})(x) + \min \left\{ 0, \min_{x \in X} J(x) \right\} \leq (T_\mu^k J)(x).$$

Since $\limsup_{k \rightarrow \infty} (T_\mu^k \bar{J})(x) = J_\mu(x) = \infty$ as shown earlier, it follows that $\limsup_{k \rightarrow \infty} (T_\mu^k J)(x) = \infty$ for all $J \in R(X)$. **Q.E.D.**

Note that if there is a negative length cycle in the subgraph of arcs \mathcal{A}_μ , it is not necessarily true that for some $x \in X$ we have $J_\mu(x) = -\infty$. Even for x on the negative length cycle, the value of $J_\mu(x)$ is determined by the *longest* path in $P(x, \mu)$, which may be simple in which case $J_\mu(x)$ is a real number, or contain an infinite repetition of a positive length cycle in which case $J_\mu(x) = \infty$.

A key fact in our analysis is the following characterization of the notion of $R(X)$ -regularity and its connection to the notion of properness. It shows that proper policies are $R(X)$ -regular, but the set of $R(X)$ -regular policies may contain some improper policies, which are characterized in terms of the sign of the lengths of their associated cycles.

Proposition 4.2: (Relation of Regularity and Properness of a Policy) Consider the minimax formulation of the RSP problem, viewed as a special case of the abstract semicontractive DP model of Section 3.1 with T_μ given by Eqs. (2.2)-(2.4), and \bar{J} being the zero function. The following are equivalent for a policy μ :

- (i) μ is $R(X)$ -regular.
- (ii) The subgraph of arcs \mathcal{A}_μ is destination-connected and all its cycles have negative length.
- (iii) μ is either proper or else it is improper, all the cycles of the subgraph of arcs \mathcal{A}_μ have negative length, and $J_\mu \in R(X)$.

Proof: To show that (i) implies (ii), let μ be $R(X)$ -regular and to arrive at a contradiction, assume that \mathcal{A}_μ contains a nonnegative length cycle. Let x be a node on the cycle, consider the path p that starts at x and consists of an infinite repetition of this cycle, and let $L_{p,\mu}^k$ be the length of the first k arcs of that path. Let also J be a nonzero constant function, $J(x) \equiv r$, where r is a scalar. Then we have

$$L_{p,\mu}^k + r \leq (T_\mu^k J)(x),$$

since from the definition of T_μ , we have that $(T_\mu^k J)(x)$ is the maximum over the lengths of all k -arc paths under μ starting at x , plus r , if the last node in the path is not the destination. Since μ is $R(X)$ -regular, we have $\limsup_{k \rightarrow \infty} (T_\mu^k J)(x) = J_\mu(x) < \infty$, so it follows that for all scalars r , we have

$$\limsup_{k \rightarrow \infty} (L_{p,\mu}^k + r) \leq J_\mu(x) < \infty.$$

This relation must hold for all scalars r , so it follows that $\limsup_{k \rightarrow \infty} L_{p,\mu}^k = -\infty$, which contradicts the nonnegativity of the cycle of p . Thus all cycles of \mathcal{A}_μ have negative length. To show that \mathcal{A}_μ is destination-connected, assume the contrary. Then there exists some node $x \in X$ such that all paths in $P(x, \mu)$ contain an

infinite number of cycles. Since the length of all cycles is negative, as just shown, it follows that $J_\mu(x) = -\infty$, which contradicts the $R(X)$ -regularity of μ .

To show that (ii) implies (iii), we note that if μ is improper, the cycles of \mathcal{A}_μ have negative length by (ii), so there remains to show that $J_\mu \in R(X)$. Since by (ii) \mathcal{A}_μ is destination-connected, the set $P(x, \mu)$ contains a simple path for all $x \in X$. Moreover, each path in $P(x, \mu)$ that is not simple has smaller length than some simple path in $P(x, \mu)$, by the negativity of the cycle lengths. This implies that $J_\mu(x)$ is equal to the largest path length among simple paths in $P(x, \mu)$, so $J_\mu(x)$ is a real number for all $x \in X$.

To show that (iii) implies (i), we note that (iii) implies that $J_\mu \in R(X)$ if μ is improper, whereas we also have $J_\mu \in R(X)$ for all proper μ . Thus to complete the proof of $R(X)$ -regularity of μ , there remains to show that $(T_\mu^k J)(x) \rightarrow J_\mu(x)$ for all $x \in X$ and $J \in R(X)$. Indeed, from the definition of T_μ , we have

$$(T_\mu^k J)(x) = \sup_{p \in P(x, \mu)} [L_{p, \mu}^k + J(x_p^k)], \quad (4.1)$$

where x_p^k is the node reached after k arcs along the path p , and $J(t)$ is defined to be equal to 0. Thus as $k \rightarrow \infty$, for every path p that contains an infinite number of cycles (each necessarily having negative length), the sequence $L_{p, \mu}^k + J(x_p^k)$ approaches $-\infty$. It follows that for sufficiently large k , the supremum in Eq. (4.1) is attained by one of the simple paths in $P(x, \mu)$, so $x_p^k = t$ and $J(x_p^k) = 0$. Thus the limit of $(T_\mu^k J)(x)$ does not depend on J , and is equal to the limit of $(T_\mu^k \bar{J})(x)$, i.e., $J_\mu(x)$. **Q.E.D.**

We illustrate the preceding proposition with a two-node example involving an improper policy.

Example 4.1:

Let $X = \{1\}$, and consider the policy μ with $Y(1, \mu(1)) = \{1, t\}$, which is improper since its subgraph of arcs \mathcal{A}_μ contains the self-cycle $(1, 1)$; cf. Fig. 4.1. Let

$$g(1, \mu(1), 1) = a, \quad g(1, \mu(1), t) = 0.$$

Then,

$$J_\mu(1) = \begin{cases} \infty & \text{if } a > 0, \\ 0 & \text{if } a \leq 0. \end{cases}$$

Consistently with Prop. 4.2, the following hold:

- (a) For $a > 0$, the cycle $(1, 1)$ has positive length, and μ is $R(X)$ -irregular because $J_\mu(1) = \infty$.
- (b) For $a = 0$, the cycle $(1, 1)$ has zero length, and μ is $R(X)$ -irregular because for a function $J \in R(X)$ with $J(1) > 0$,

$$\limsup_{k \rightarrow \infty} (T_\mu^k J)(x) = J(1) > 0 = J_\mu(1).$$

- (c) For $a < 0$, the cycle $(1, 1)$ has negative length, and μ is $R(X)$ -regular because $J_\mu \in R(X)$ and for all $J \in R(X)$,

$$\limsup_{k \rightarrow \infty} (T_\mu^k J)(1) = 0 = J_\mu(1).$$

We now show one of our main results.

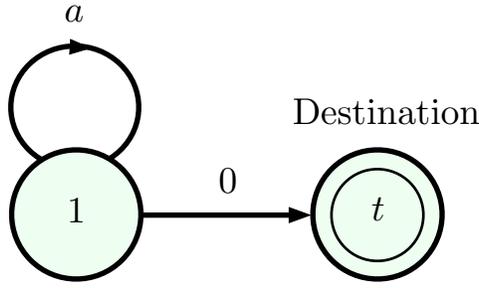


Figure 4.1. The subgraph of arcs \mathcal{A}_μ corresponding to an improper policy μ , for the case of a single node 1 and a destination node t . The arcs lengths are shown in the figure.

Proposition 4.3: Let Assumption 2.1 hold. Then:

- (a) The minimax-optimal cost function J^* of Eq. (2.1) is the unique fixed point of T within $R(X)$.
- (b) A policy μ^* is optimal (i.e., $J_{\mu^*} = J^*$) if and only if $T_{\mu^*}J^* = TJ^*$. Moreover, there exists an optimal policy and every optimal policy is proper.
- (c) We have $T^k J \rightarrow J^*$ for all $J \in R(X)$.
- (d) For any $J \in R(X)$, if $J \leq TJ$ we have $J \leq J^*$, and if $J \geq TJ$ we have $J \geq J^*$.

Proof: We verify the parts (a)-(f) of Assumption 3.1 with $S = R(X)$. The result then will follow from Prop. 3.1. To this end we note that:

- (1) Part (a) is satisfied since $S = R(X)$.
- (2) Part (b) is satisfied since by Assumption 2.1(a), there exists at least one proper policy, which by Prop. 4.2 is $R(X)$ -regular. Moreover, for each $R(X)$ -regular policy μ , we have $J_\mu \in R(X)$. Since the number of all policies is finite, it follows that $\hat{J} \in R(X)$.
- (3) To show that part (c) is satisfied, note that since by Prop. 4.2 every $R(X)$ -irregular policy μ must be improper, so by Assumption 2.1(b), the subgraph of arcs \mathcal{A}_μ contains a cycle of positive length; by Prop. 4.1(d), this implies that for each $J \in R(X)$, we have $\limsup_{k \rightarrow \infty} (T_\mu^k J)(x) = \infty$ for at least one $x \in X$.
- (4) Part (d) is satisfied since $U(x)$ is a finite set.
- (5) Part (e) is satisfied since X is finite and T_μ is a continuous function mapping the finite-dimensional space $R(X)$ into itself.
- (6) To show that part (f) is satisfied, we note that by applying Prop. 3.2 with $S = R(X)$, we have that \hat{J} is the unique fixed point of T within $R(X)$. It follows that for each $J \in R(X)$, there exists a sufficiently large scalar $r > 0$ such that the function J' given by

$$J' = \hat{J} - re, \quad \forall x \in X, \quad (4.2)$$

where e is the unit function, $e(x) \equiv 1$, satisfies $J' \leq J$ as well as

$$J' = \hat{J} - re = T\hat{J} - re \leq T(\hat{J} - re) = TJ', \quad (4.3)$$

where the inequality holds in view of Eqs. (2.2) and (2.5), and the fact $r > 0$.

Thus all parts of Assumption 3.1 with $S = R(X)$ are satisfied, and Prop. 3.1 applies, with $S = R(X)$. Since under Assumption 2.1, improper policies are $R(X)$ -irregular [cf. Prop. 4.1(d)] and so cannot be optimal, the minimax formulation of Section 2 is equivalent to RSP, and the conclusions of Prop. 3.1 are precisely the results we want to prove. **Q.E.D.**

The following variant of the two-node Example 4.1 illustrates what may happen in the absence of Assumption 2.1(b), when there may exist improper policies that involve a nonpositive length cycle.

Example 4.2:

Let $X = \{1\}$, and consider the improper policy μ with $Y(1, \mu(1)) = \{1, t\}$ and the proper policy $\bar{\mu}$ with $Y(1, \bar{\mu}(1)) = \{t\}$ (cf. Fig. 4.2). Let

$$g(1, \mu(1), 1) = a \leq 0, \quad g(1, \mu(1), t) = 0, \quad g(1, \bar{\mu}(1), t) = 1.$$

Then it can be seen that under both policies, the longest path from 1 to t consists of the arc $(1, t)$. Thus,

$$J_\mu(1) = 0, \quad J_{\bar{\mu}}(1) = 1,$$

so the improper policy μ is optimal for the minimax problem (2.1). To explain what is happening here, we consider two different cases:

- (1) $a = 0$: In this case, the optimal policy μ is both improper and $R(X)$ -irregular, but with $J_\mu(1) < \infty$. Thus the conditions of both Props. 3.1 and 4.3 do not hold because Assumptions 3.1(c) and Assumption 2.1(b) are violated.
- (2) $a < 0$: In this case, μ is improper but $R(X)$ -regular, so there are no $R(X)$ -irregular policies. Then all the conditions of Assumption 3.1 are satisfied, and Prop. 3.1 applies. Consistent with this proposition, there exists an optimal $R(X)$ -regular policy (i.e., optimal over both proper and improper policies), which however is improper and hence not an optimal solution for RSP.

We next discuss modifications of Prop. 4.3, which cover the two cases of the preceding example.

The Case of Improper Policies with Negative Length Cycles

We note that by Prop. 4.2, the set of $R(X)$ -regular policies includes not just proper policies, but also some improper ones (those μ for which \mathcal{A}_μ is destination-connected and all its cycles have negative length). As a result we can weaken Assumption 2.1 as long as it still implies Assumption 3.1 so we can use Prop. 3.1 to obtain corresponding versions of our main result of Prop. 4.3. Here are two such weaker versions of Assumption 2.1.

Assumption 4.1: Every policy μ is either proper or else it is improper and its subgraph of arcs \mathcal{A}_μ is destination-connected with all cycles having negative length.

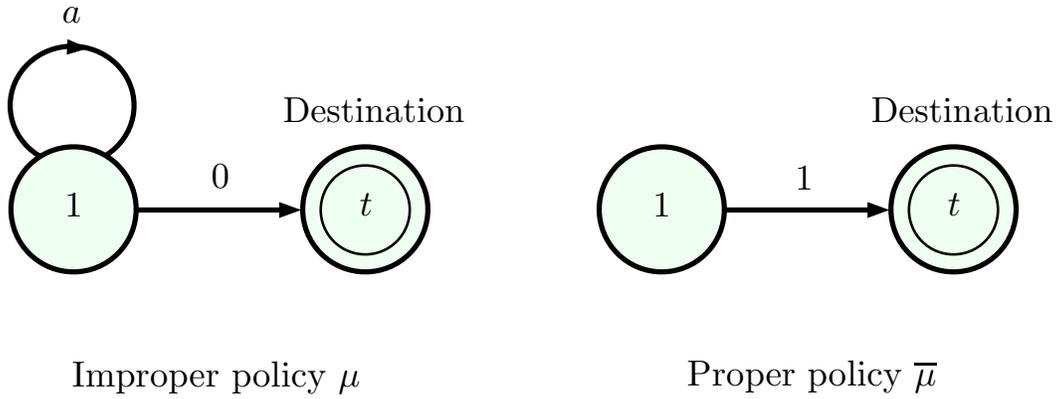


Figure 4.2. A counterexample involving a single node 1 in addition to the destination t . There are two policies, μ and $\bar{\mu}$, with corresponding subgraphs of arcs \mathcal{A}_μ and $\mathcal{A}_{\bar{\mu}}$, and arc lengths shown in the figure. The improper policy μ is optimal when $a \leq 0$. It is $R(X)$ -irregular if $a = 0$, and it is $R(X)$ -regular if $a < 0$.

From Prop. 4.2, it follows that the preceding assumption is equivalent to all policies being $R(X)$ -regular. The next assumption is weaker in that it allows policies μ that are $R(X)$ -irregular, as long as some cycle of \mathcal{A}_μ has positive length.

Assumption 4.2:

- (a) There exists at least one $R(X)$ -regular policy.
- (b) For every $R(X)$ -irregular policy μ , some cycle in the subgraph of arcs \mathcal{A}_μ has positive length.

By essentially repeating the proof of Prop. 4.3, we see that Assumption 4.2 implies Assumption 3.1, so that Prop. 3.1 applies. Then we obtain the following variant of Prop. 4.3.

Proposition 4.4: Let either Assumption 4.1 or (more generally) Assumption 4.2 hold. Then:

- (a) The minimax-optimal cost function J^* of Eq. (2.1) is the unique fixed point of T within $R(X)$.
- (b) A policy μ^* is optimal (i.e., $J_{\mu^*} = J^*$) if and only if $T_{\mu^*} J^* = T J^*$. Moreover, there exists an optimal policy and every optimal policy is $R(X)$ -regular.
- (c) We have $T^k J \rightarrow J^*$ for all $J \in R(X)$.
- (d) For any $J \in R(X)$, if $J \leq T J$ we have $J \leq J^*$, and if $J \geq T J$ we have $J \geq J^*$.

Note that the optimal $R(X)$ -regular policy μ^* of part (b) above, while minimax-optimal, it may not be proper, in which case it will not solve the RSP problem (cf. Example 4.2 with $a < 0$). Thus, in a given

problem, one would have to additionally prove that optimal $R(X)$ -regular policies are proper in order for the proposition to fully apply to RSP.

The Case of Improper Policies with Zero Length Cycles

In some problems, it may be easier to guarantee nonnegativity rather than positivity of the lengths of cycles corresponding to improper policies, which is required by Assumption 2.1(b). This is true for example in the important case where all arc lengths are nonnegative, i.e., $g(x, u, y) \geq 0$ for all $x \in X$, $u \in U(x)$, and $y \in Y(x, u)$, as in case (1) of Example 4.2. Let us consider the following relaxation of Assumption 2.1.

Assumption 4.3:

- (a) There exists at least one proper policy.
- (b) For every improper policy μ , all cycles in the subgraph of arcs \mathcal{A}_μ have nonnegative length.

For an analysis under the preceding assumption, we will use a perturbation approach that was introduced in Section 3.2.2 of [Ber13]. The idea is to consider a scalar $\delta > 0$ and a δ -perturbed problem, whereby each arc length $g(x, u, y)$ with $x \in X$ is replaced by $g(x, u, y) + \delta$. As a result, a nonnegative cycle length corresponding to an improper policy as per Assumption 4.3(b) becomes strictly positive, so Assumption 2.1 is satisfied for the δ -perturbed problem, and Prop. 4.3 applies. We thus see that J_δ^* , the optimal cost function of the δ -perturbed problem, is the unique fixed point of the mapping T_δ given by

$$(T_\delta J)(x) = \min_{u \in U(x)} H_\delta(x, u, J), \quad x \in X,$$

where $H_\delta(x, u, J)$ is given by the equation for H [cf. Eq. (2.2)], with δ added to $g(x, u, t)$ and $g(x, u, y)$. Moreover there exists an optimal proper policy μ_δ for the δ -perturbed problem, which by Prop. 4.3(b), satisfies the optimality equation

$$T_{\mu_\delta, \delta} J_\delta^* = T_\delta J_\delta^*,$$

where $T_{\mu, \delta}$ is the mapping that corresponds to a policy μ in the δ -perturbed problem:

$$(T_{\mu, \delta} J)(x) = H_\delta(x, \mu(x), J), \quad x \in X.$$

We have the following proposition.

Proposition 4.5: Let Assumption 4.3 hold. Then

$$\lim_{\delta \downarrow 0} J_\delta^* = \min_{\mu: \text{proper}} J_\mu,$$

where J_δ^* is the optimal cost function of the δ -perturbed problem. Moreover, there exists $\bar{\delta} > 0$ such that for all $\delta \in (0, \bar{\delta}]$, an optimal policy for the δ -perturbed problem is an optimal proper policy for RSP.

Proof: For all $\delta > 0$, consider an optimal proper policy μ_δ of the δ -perturbed problem. We have

$$\min_{\mu: \text{proper}} J_\mu \leq J_{\mu_\delta} \leq J_{\mu_\delta, \delta} = J_\delta^* \leq J_{\mu', \delta} \leq J_{\mu'} + N\delta, \quad \forall \mu' : \text{proper},$$

where N is the number of nodes of X (since an extra δ cost is incurred in the δ -perturbed problem every time a path goes through a node $x \neq t$, and any path under a proper μ' contains at most N nodes $x \neq t$). By taking the limit as $\delta \downarrow 0$ and then the minimum over all μ' that are proper, it follows that

$$\min_{\mu: \text{proper}} J_\mu \leq \lim_{\delta \downarrow 0} J_\delta^* \leq \min_{\mu: \text{proper}} J_\mu,$$

so $\lim_{\delta \downarrow 0} J_\delta^* = \min_{\mu: \text{proper}} J_\mu$.

To prove the last part of the proposition, we note that for every proper policy μ we have $\lim_{\delta \downarrow 0} J_{\mu, \delta} = J_\mu$. Hence if a proper μ is not optimal for RSP, it is also nonoptimal for the δ -perturbed problem for all $\delta \in [0, \delta_\mu]$, where δ_μ is some positive scalar. Let $\bar{\delta}$ be the minimum δ_μ over the nonoptimal proper policies μ . Then for $\delta \in (0, \bar{\delta}]$, an optimal policy for the δ -perturbed problem cannot be nonoptimal for RSP. **Q.E.D.**

Note that we may have $J^*(x) < \min_{\mu: \text{proper}} J_\mu(x)$ for some x , but in RSP, only proper policies are admissible, so by letting $\delta \downarrow 0$ we approach the optimal solution of interest. This happens for instance in Example 4.2 when $a = 0$. The following example describes an interesting problem, where Prop. 4.5 applies.

Example 4.3: (Minimax Search Problems)

Consider searching a graph with node set $X \cup \{t\}$, looking for an optimal node $x \in X$ at which to stop. At each $x \in X$ we have two options: (1) stopping at a cost $s(x) \geq 0$, which will stop the search by moving to t , or (2) continuing the search by choosing a control $u \in U(x)$, in which case we will move to a node y , chosen from within a given set of nodes $Y(x, u)$ by an antagonistic opponent, at a cost $g(x, u, y) \geq 0$. An interesting special case is when searching is cost-free [i.e., $g(x, u, y) \equiv 0$], but may lead in the future to nodes where a higher stopping cost will become inevitable. Then a policy that never stops is optimal but improper. If we then introduce a small perturbation $\delta > 0$ to the costs $g(x, u, y)$, we will make the lengths of all cycles positive, so we may use Prop. 4.5 to find an optimal proper policy.

5. COMPUTATIONAL METHODS

We will now discuss computational methods that are patterned after the classical DP algorithms of value iteration and policy iteration (VI and PI for short, respectively). In particular, the methods of this section are motivated by specialized stochastic shortest path algorithms.

5.1. Value Iteration Algorithms

We have already shown as part of Prop. 4.3 that under Assumption 2.1, the VI algorithm, which sequentially generates $T^k J$ for $k \geq 0$, converges to the optimal cost function J^* for any starting function $J \in R(X)$. We can extend this convergence property to asynchronous versions of VI based on the monotonicity and fixed point properties of the mapping T . This has been known since [Ber82] (see also [Ber83], [BeT89]), and we refer to the discussions in Sections 2.6.1, and 3.3.1 of [Ber13], which apply in their entirety when specialized to the RSP problem of this paper.

It turns out that for our problem, under Assumption 2.1, the VI algorithm terminates finitely when initialized with $J(x) = \infty$ for all $x \in X$ [it can be seen that in view of the form (2.5) of the mapping T , the VI algorithm is well-defined with this initialization]. In fact the number of iterations for termination is no more than N , where N is the number of nodes in X . This is consistent with a similar result for stochastic shortest path problems ([Ber12], Section 3.4.1), which relies on the assumption of acyclicity of the graph of possible transitions under an optimal policy. Because this assumption is restrictive, finite termination of the VI algorithm is an exceptional property in stochastic shortest path problems. However, in the minimax case of this paper, under Assumption 2.1 an optimal policy μ^* exists and is proper [cf. Prop. 4.3(b)], so the graph of possible transitions under μ^* is acyclic, and it turns out that finite termination of VI is guaranteed to occur.

To show the finite termination property just described, let μ^* be an optimal proper policy, consider the sets of nodes X_0, X_1, \dots , defined by

$$X_0 = \{t\},$$

$$X_{k+1} = \left\{ x \notin \cup_{m=0}^k X_m \mid y \in \cup_{m=0}^k X_m \text{ for all } y \in Y(x, \mu^*(x)) \right\}, \quad k = 0, 1, \dots, \quad (5.1)$$

and let $X_{\bar{k}}$ be the last of these sets that is nonempty. Then in view of the acyclicity of the subgraph of arcs \mathcal{A}_{μ^*} , we have

$$\cup_{m=0}^{\bar{k}} X_m = X \cup \{t\}.$$

We will now show by induction that starting from $J(x) \equiv \infty$ for all $x \in X$, the iterates $T^k J$ of VI satisfy

$$(T^k J)(x) = J^*(x), \quad \forall x \in \cup_{m=1}^k X_m, \quad k = 1, \dots, \bar{k}. \quad (5.2)$$

Indeed, it can be seen that this is so for $k = 1$. Assume that $(T^k J)(x) = J^*(x)$ if $x \in \cup_{m=1}^k X_m$. Then, since $TJ \leq J$ and T is monotone, $(T^k J)(x)$ is monotonically nonincreasing, so that

$$J^*(x) \leq (T^{k+1} J)(x), \quad \forall x \in X. \quad (5.3)$$

Moreover, by the induction hypothesis, the definition of the sets X_k , and the optimality of μ^* , we have

$$(T^{k+1} J)(x) \leq H(x, \mu^*(x), T^k J) = H(x, \mu^*(x), J^*) = J^*(x), \quad \forall x \in \cup_{m=1}^{k+1} X_m, \quad (5.4)$$

where the first equality follows from the form (2.2) of H and the fact that for all $x \in \cup_{m=1}^{k+1} X_m$, we have $y \in \cup_{m=1}^k X_m$ for all $y \in Y(x, \mu^*(x))$ by the definition (5.1) of X_{k+1} . The two relations (5.3) and (5.4) complete the induction proof.

Thus under Assumption 2.1, the VI method when started with $J(x) = \infty$ for all $x \in X$, will find the optimal costs of all the nodes in the set $\cup_{m=1}^k X_m$ after k iterations; cf. Eq. (5.2). In particular, all optimal costs will be found after $\bar{k} \leq N$ iterations, where N is the number of nodes in X . This indicates that the behavior of the VI algorithm, when initialized with $J(x) = \infty$ for all $x \in X$, is similar to the one of the Bellman-Ford algorithm for deterministic shortest path problems. Still each iteration of the VI algorithm requires as many as N applications of the mapping T at every node. Thus it is likely that the performance of the VI algorithm can be improved with a suitable choice of the initial function J , and with an asynchronous implementation that uses a favorable order of selecting nodes for iteration, “one-node-at-a-time” similar to the Gauss-Seidel method. This is consistent with the deterministic shortest path case, where there are VI-type algorithms, within the class of label-correcting methods, which are faster than the Bellman-Ford algorithm and even faster than efficient implementations of the Dijkstra algorithm for some

types of problems; see e.g., [Ber98]. For the RSP problem, it can be seen that the best node selection order is based on the sets X_k defined by Eq. (5.1), i.e., iterate on the nodes in the set X_1 , then on the nodes in X_2 , and so on. In this case, only one iteration per node will be needed. While the sets X_k are not known, an algorithm that tries to approximate the optimal order could be much more efficient than the standard “all-nodes-at-once” VI method that computes the sequence $T^k J$, for $k \geq 0$ (for an example of an algorithm of this type for stochastic shortest path problems, see [PBT98]). The development of such more efficient VI algorithms is an interesting subject for further research, which, however, is beyond the scope of the present paper.

Example 5.1:

Let us illustrate the VI method for the problem of Fig. 5.1. The optimal policy is shown in this figure, and it is proper, since Assumption 2.1 is satisfied. The table gives the iteration sequence of two VI methods, starting with $J_0 = (\infty, \infty, \infty, \infty)$. The first method is the all-nodes-at-once method $J_k = T^k J_0$, which finds J^* in four iterations. In this example, we have $X_0 = \{t\}$, $X_1 = \{1\}$, $X_2 = \{4\}$, $X_3 = \{3\}$, $X_4 = \{2\}$, and the assertion of Eq. (5.2) may be verified. The second method is the asynchronous VI method, which iterates one-node-at-a-time in the (most favorable) order 1, 4, 3, 2. The second method also finds J^* in four iterations and with four times less computation.

We finally note that in the absence of Assumption 2.1 or Assumption 4.1, it is possible that the VI sequence $\{T^k J\}$ will not converge to J^* starting from any J with $J \neq J^*$. This can be seen with a simple deterministic shortest path problem involving a zero length cycle, a simpler version of Example 4.2. Here there is a single node 1, aside from the destination t , and two choices at 1: stay at 1 at cost 0, and move to t at cost 1. Then we have $J^* = 0$, while T is given by

$$TJ = \min\{J, 1\}.$$

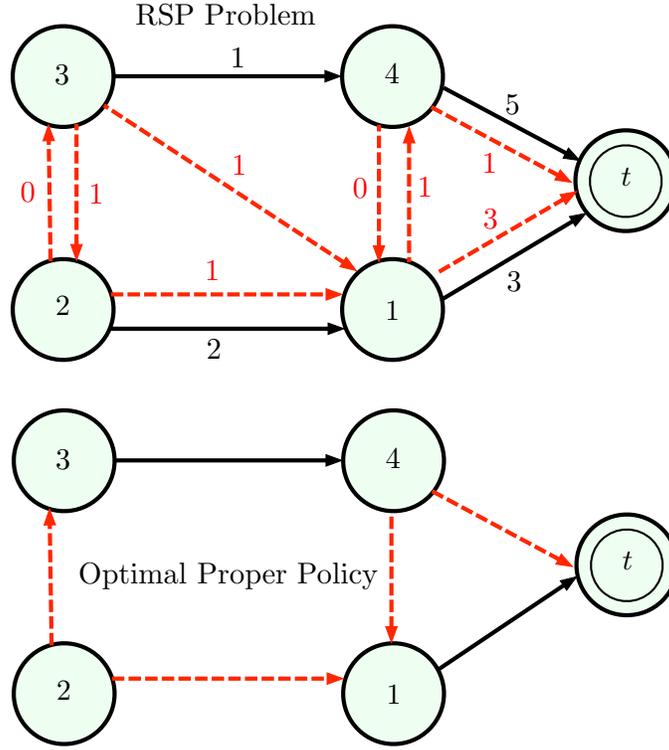
It can be seen that the set of fixed points of T is $(-\infty, 1]$, and contains J^* in its interior. Starting with $J \geq 1$, the VI sequence converges to 1 in a single step, while starting at $J \leq 1$ it stays at J . In the case of Example 4.2 with $a = 0$, the situation is somewhat different but qualitatively similar. There it can be verified that $J^* = 1$, the set of fixed points is $[0, 1]$, $\{T^k J\}$ will converge to 1 starting from $J \geq 1$, will converge to 0 starting from $J \leq 0$, and will stay at J starting from $J \in [0, 1]$.

5.2. Policy Iteration Algorithms

The development of PI algorithms for the RSP problem is straightforward given the connection with semi-contractive models. Briefly, under Assumption 2.1, based on the analysis of Section 3.3.2 of [Ber13], there are two types of PI algorithms. The first is a natural form of PI that generates proper policies exclusively. Let μ_0 be an initial proper policy (there exists one by assumption). At the typical iteration k , we have a proper policy μ_k , and first compute J_{μ_k} by solving a longest path problem over the corresponding acyclic subgraph of arcs \mathcal{A}_{μ_k} . We then compute a policy μ_{k+1} such that $T_{\mu_{k+1}} J_{\mu_k} = T J_{\mu_k}$, by minimizing over $u \in U(x)$ the expression $H(x, u, J_{\mu_k})$ of Eq. (2.2), for all $x \in X$. We have

$$J_{\mu_k} = T_{\mu_k} J_{\mu_k} \geq T J_{\mu_k} = T_{\mu_{k+1}} J_{\mu_k} \geq \lim_{m \rightarrow \infty} T_{\mu_{k+1}}^m J_{\mu_k} = J_{\mu_{k+1}}, \tag{5.5}$$

where the second inequality follows from the monotonicity of $T_{\mu_{k+1}}$. Given that μ_k is proper, it follows in view of Prop. 4.1(d), that the next policy μ_{k+1} is proper, and has improved cost over μ_k . Therefore the



Iteration #	All-Nodes-at-Once VI Method	One-Node-at-a-Time VI Method
0	$J_0 = (\infty, \infty, \infty, \infty)$	$J_0 = (\infty, \infty, \infty, \infty)$
1	$J_1 = TJ_0 = (3, \infty, \infty, 5)$	$J_1 = (3, \infty, \infty, \infty)$
2	$J_2 = T^2J_0 = (3, 5, 6, 3)$	$J_2 = (3, \infty, \infty, 3)$
3	$J_3 = T^3J_0 = (3, 5, 4, 3)$	$J_3 = (3, \infty, 4, 3)$
4	$J_4 = T^4J_0 = (3, 4, 4, 3) = J^*$	$J_4 = (3, 4, 4, 3) = J^*$

Figure 5.1: An example RSP problem and its optimal policy. At each $X = \{1, 2, 3, 4\}$ there are two controls: one (shown by a solid line) where $Y(x, u)$ consists of a single element, and another (shown by a broken line) where $Y(x, u)$ has two elements. Arc lengths are shown next to the arcs. Both the all-nodes-at-once and the one-node-at-a-time versions of VI terminate in four iterations, but the latter version requires four times less computation per iteration.

sequence of policies $\{\mu_k\}$ is well-defined and proper, and the corresponding sequence $\{J_{\mu_k}\}$ is nonincreasing. Under Assumption 2.1, it is then possible to show that J_{μ_k} converges to J^* in a finite number of iterations. The reason is that from Eq. (5.5), we have that at the k th iteration, either strict improvement $J_{\mu_k}(x) > (TJ_{\mu_k})(x) \geq J_{\mu_{k+1}}(x)$ is obtained for at least one node $x \in X$, or else $J_{\mu_k} = TJ_{\mu_k}$, which implies that $J_{\mu_k} = J^*$ [since J^* is the unique fixed point of T within $R(X)$, by Prop. 4.3(a)] and μ_k is an optimal proper policy.

Unfortunately, when there are improper policies, the preceding PI algorithm is somewhat limited,

because an initial proper policy may not be known, and also because when asynchronous versions of the algorithm are implemented, it is difficult to guarantee that all the generated policies are proper. There is another algorithm, combining value and policy iterations, which has been developed in [BeY10], [BeY12], [YuB13a], [YuB13b] for a variety of DP models, including discounted, stochastic shortest path, and abstract, and is described in Sections 2.6.3 and 3.3.2 of [Ber13]. This algorithm updates a cost function J and a policy μ , but it also maintains an additional function V , which acts as a threshold to keep J bounded and the algorithm convergent. The algorithm not only can tolerate the presence of improper policies, but can also be operated in asynchronous mode, whereby the value iterations, policy evaluation operations, and policy improvement iterations are performed one-node-at-a-time without any regularity. The algorithm is valid even in a distributed asynchronous environment, and in the presence of communication delays between processors. The specialization of this algorithm to RSP under Assumption 2.1 is straightforward, and will be presented briefly in its asynchronous form, but without communication delays between processors.

We consider a distributed computing system with m processors, denoted $1, \dots, m$, a partition of the node set X into sets X_1, \dots, X_m , and an assignment of each subset X_ℓ to a processor $\ell \in \{1, \dots, m\}$. The processors collectively maintain two functions $J_k(x)$ and $V_k(x)$, $x \in X$, and a policy $\mu_k \in \mathcal{M}$, $k = 0, 1, \dots$. We denote by $\min[V_k, J_k]$ the function in $E(X)$ that takes values $\min[V_k(x), J_k(x)]$ for $x \in X$. The initial conditions $J_0(x), V_0(x), \mu_0(x)$, $x \in X$, are arbitrary. For each processor ℓ , there are two infinite disjoint subsets of times $\mathcal{K}_\ell, \overline{\mathcal{K}}_\ell \subset \{0, 1, \dots\}$, corresponding to local (within the subset X_ℓ) policy improvement and policy evaluation iterations by that processor, respectively. More specifically, at each time k and for each processor ℓ , we have one of the following three possibilities:

- (a) *Local policy improvement*: If $k \in \mathcal{K}_\ell$, processor ℓ sets for all $x \in X_\ell$,

$$J_{k+1}(x) = V_{k+1}(x) = \min_{u \in U(x)} H(x, u, \min[V_k, J_k]), \quad (5.6)$$

and sets $\mu_{k+1}(x)$ to a u that attains the above minimum.

- (b) *Local policy evaluation - Value iteration*: If $k \in \overline{\mathcal{K}}_\ell$, processor ℓ sets for all $x \in X_\ell$,

$$J_{k+1}(x) = H(x, \mu_k(x), \min[V_k, J_k]), \quad (5.7)$$

and leaves V and μ unchanged, i.e., for all $x \in X_\ell$,

$$V_{k+1}(x) = V_k(x), \quad \mu_{k+1}(x) = \mu_k(x).$$

- (c) *No local change*: If $k \notin \mathcal{K}_\ell \cup \overline{\mathcal{K}}_\ell$, processor ℓ leaves J , V , and μ unchanged, i.e., for all $x \in X_\ell$,

$$J_{k+1}(x) = J_k(x), \quad V_{k+1}(x) = V_k(x), \quad \mu_{k+1}(x) = \mu_k(x).$$

In view of the form (2.2) of the mapping H , the local policy improvement iteration (5.6) involves the solution of a static minimax problem, where the minimizing player chooses $u \in U(x)$ and the maximizing player chooses $y \in Y(x, u)$. The local policy evaluation iteration (5.7) involves a maximization over $y \in Y(x, \mu_k(x))$.

The function V_k in Eqs. (5.6)-(5.7) is reminiscent of a stopping cost in optimal stopping problems. The use of V_k is essential for the asymptotic convergence of the algorithm to optimality, i.e., $J_k \rightarrow J^*$, $V_k \rightarrow J^*$, and for finite convergence of $\{\mu_k\}$ to an optimal proper policy. Without V_k the algorithm may potentially

oscillate (there is an important counterexample that documents this type of phenomenon, given in [WiB93]; see also the discussion in [BeY10], [BeY12], [Ber12]).

Note that the preceding algorithm includes as a special case a one-node-at-a-time asynchronous PI algorithm, whereby each node is viewed as a processor by itself, and at each iteration a single node is selected and a local policy improvement or local policy evaluation of the form (5.6) or (5.7), respectively, is performed just at that node (see the discussion of Section 2.6 of [Ber12] or Section 2.6 of [Ber13]). This Gauss-Seidel type of algorithm is often considerably faster than all-nodes-at-once versions. The comparative evaluation of PI algorithms that use different initial conditions $J_0(x), V_0(x), \mu_0(x)$, $x \in X$, and different orders of local policy improvement and policy evaluation iterations remains a subject for further research and experimentation.

5.3. A Dijkstra-Like Algorithm for Nonnegative Arc Lengths

One of the most important computational approaches for the classical deterministic shortest path problem with nonnegative arc lengths is based on Dijkstra’s algorithm, whereby the shortest distances of the nodes to the destination are determined one-at-a-time in nondecreasing order. When properly implemented, this approach yields shortest path methods with excellent computational complexity and practical performance (see e.g., [AMO89], [Ber98]).

Dijkstra’s algorithm has been extended to continuous-space shortest path problems in [Tsi95], and finds extensive application in a variety of large-scale computational problems involving the eikonal and other equations; see [Set99a], [Set99b]. For recent work in this area, see [ChV12], [AnV13], [CCV13], which quote a lot of other references. Dijkstra’s algorithm has also been extended to finite-state stochastic shortest path problems, through the notion of a “consistently improving optimal policy” (introduced in the 2001 2nd edition of the author’s DP book, and also described in its 4th edition, [Ber12], Section 3.4.1). Roughly, with such a policy, from any node we may only go to a node of no greater optimal cost. While existence of a consistently improving optimal policy is a restrictive condition, the associated Dijkstra-like algorithm has found application in some special contexts, including large-scale continuous-space shortest path problems, where it is naturally satisfied; see [Vla08].

We propose a Dijkstra-like algorithm that is patterned after the Dijkstra-like stochastic shortest path algorithm, but requires less restrictive conditions because an optimal proper policy has the essential character of a consistently improving policy when the arc lengths are nonnegative. Related Dijkstra-like algorithms were proposed recently in [GrJ08], [BaL13], under the assumption that all arc lengths are strictly positive. When the problem is deterministic, i.e., $Y(x, u)$ consists of a single node for all (x, u) , our algorithm coincides with Dijkstra’s algorithm for the classical deterministic shortest path problem. We will assume the following:

Assumption 5.1:

- (a) There exists at least one proper policy.
- (b) For every improper policy μ , all cycles in the subgraph of arcs \mathcal{A}_μ have positive length.
- (c) All arc lengths are nonnegative.

Parts (a) and (b) of the preceding assumption are just Assumption 2.1, under which the favorable results of Prop. 4.3 apply to RSP with both nonnegative and negative arc lengths. The arc length nonnegativity assumption of part (c) is characteristic of Dijkstra-like algorithms, and provides additional structure, which forms the basis for the following algorithm.

Our Dijkstra-like algorithm maintains and updates a subset of nodes denoted V , and a number $J(x)$ for each $x \in X \cup \{t\}$, called the *label of x* . Initially,

$$V = \{t\}, \quad J(x) = \begin{cases} 0 & \text{if } x = t, \\ \infty & \text{if } x \in X. \end{cases}$$

At any given point in the algorithm, let W be the set

$$W = \{x \mid J(x) < \infty, x \notin V\}. \quad (5.8)$$

The algorithm terminates when V is empty. The typical iteration, assuming V is nonempty, is as follows.

Typical Iteration of the Dijkstra-Like Algorithm:

We remove from V a node y^* such that

$$J(y^*) = \min_{y \in V} J(y),$$

and add it in W , i.e., replace W with $W \cup \{y^*\}$. For every $x \notin W$, we let

$$\hat{U}(x) = \{u \in U(x) \mid Y(x, u) \subset W \text{ and } y^* \in Y(x, u)\}, \quad (5.9)$$

and we update $J(x)$ and V according to the following two cases:

- (1) If $\hat{U}(x)$ is nonempty and $J(x) > \min_{u \in \hat{U}(x)} \max_{y \in Y(x, u)} [g(x, u, y) + J(y)]$, we set

$$J(x) = \min_{u \in \hat{U}(x)} \max_{y \in Y(x, u)} [g(x, u, y) + J(y)], \quad (5.10)$$

and we place x in V if it is not already there.

- (2) Otherwise, we leave $J(x)$ and V unchanged.

Note that at each iteration of the preceding algorithm, the single node y^* exits V , and enters the set W of Eq. (5.8). Thus W is the set of nodes that have entered V at some previous iteration but are not currently in V . Moreover, from the definition of the algorithm, once a node enters W it stays in W and never returns to V . Also, upon entrance of a node into V , its label changes from ∞ to some nonnegative number. In the terminology of Dijkstra-like algorithms, W is the set of nodes that are “permanently labeled,” and V is the set of nodes that are “candidates” for permanent labeling. We will show that all the nodes x will enter W in order of nondecreasing $J(x)$, and at the time of entrance, $J(x) = J^*(x)$.

Proposition 5.1: Let Assumption 5.1 hold. Then at the end of an iteration of the Dijkstra-like algorithm, we have $J(x') \geq J(x)$ for all $x' \notin W$ and $x \in W$.

Proof: We use induction on the iteration count. Clearly the assertion holds at the end of the initial iteration since then $W = \{t\}$, $J(t) = 0$, and according to the formula (5.10) for changing labels and the nonnegativity of the arc lengths, we have $J(x) \geq 0$ for all $x \in X$. Assume that the assertion holds for iteration $k-1$. Let $J(x)$ and $\tilde{J}(x)$ denote the node labels at the start and the end of iteration k , respectively. Then by the minimum label rule for selection of y^* , we have

$$J(x') \geq J(y^*) \geq J(x) = \tilde{J}(x), \quad \forall x' \notin W \cup \{y^*\}, x \in W \cup \{y^*\}, \quad (5.11)$$

where the equality holds because the labels of all $x \in W \cup \{y^*\}$ will not change in iteration k . During iteration k the labels of nodes $x' \notin W \cup \{y^*\}$ will change, if $\hat{U}(x') \neq \emptyset$, according to Eq. (5.10), so that

$$\begin{aligned} \tilde{J}(x') &= \min \left[J(x'), \min_{u \in \hat{U}(x')} \max_{y \in Y(x', u)} [g(x', u, y) + J(y)] \right] \\ &\geq \min [J(x'), J(y^*)] \\ &\geq J(x) \\ &= \tilde{J}(x), \quad \forall x' \notin W \cup \{y^*\}, x \in W \cup \{y^*\}, \end{aligned}$$

where the first inequality holds because $g(x', u, y^*) \geq 0$, and $y^* \in Y(x', u)$ for all $u \in \hat{U}(x')$, and the second inequality and second equality hold because of Eq. (5.11). The induction proof is complete. **Q.E.D.**

Since no node will enter V twice, while exactly one node exits V at each iteration, the algorithm will terminate after no more than $N+1$ iterations, where N is the number of nodes in X . The next proposition shows that V will become empty after exactly $N+1$ iterations, at which time W must necessarily be equal to $X \cup \{t\}$.

Proposition 5.2: Let Assumption 5.1 hold. The Dijkstra-like algorithm will terminate after exactly $N+1$ iterations with $V = \emptyset$ and $W = X \cup \{t\}$.

Proof: Assume the contrary, i.e., that the algorithm will terminate after a number of iterations $k < N+1$. Then upon termination, W will have k nodes, V will be empty, and the set

$$\bar{V} = \{x \in X \mid x \notin W\}$$

will have $N+1-k$ nodes and thus be nonempty. Let $\bar{\mu}$ be the proper policy, which is assumed to exist by Assumption 5.1(a). For each $x \in \bar{V}$ we cannot have $Y(x, \bar{\mu}(x)) \subset W$, since then x would have entered

V prior to termination, according to the rules of the algorithm. Thus for each $x \in \bar{V}$, there exists an arc $(x, y) \in Y(x, \bar{\mu}(x))$ with $y \in \bar{V}$. This implies that the subgraph of arcs $\mathcal{A}_{\bar{\mu}}$ contains a cycle of nodes in \bar{V} , thus contradicting the properness of $\bar{\mu}$. **Q.E.D.**

Proposition 5.3: Let Assumption 5.1 hold. Upon termination of the Dijkstra-like algorithm, we have $J(x) = J^*(x)$ for all $x \in X$.

Proof: The proof is long, but follows the general pattern of corresponding induction-based proofs of Dijkstra’s algorithm for the classical deterministic shortest path problem. For $k = 0, 1, \dots$, let $V_k, W_k, J_k(x)$ denote the sets V, W , and the labels $J(x)$ at the start of iteration k , respectively, and let y_k^* denote the minimum label node to enter W during iteration k . Thus, we have $W_{k+1} = W_k \cup \{y_k^*\}$,

$$J_{k+1}(x) = J_k(x), \quad \forall x \in W_{k+1},$$

and

$$J_{k+1}(x) = \begin{cases} \min \left[J_k(x), \min_{u \in \hat{U}_k(x)} \max_{y \in Y(x, u)} [g(x, u, y) + J_k(y)] \right] & \text{if } x \in V_{k+1}, \\ \infty & \text{if } x \notin W_{k+1} \cup V_{k+1}, \end{cases} \quad (5.12)$$

where

$$\hat{U}_k(x) = \{u \in U(x) \mid Y(x, u) \subset W_k \text{ and } y^* \in Y(x, u)\}.$$

For each k and $x \in X \cup \{t\}$ consider the sets of policies

$$M_k(x) = \{\mu : \text{proper} \mid \text{the nodes of all paths } p \in P(x, \mu) \text{ lie in } W_k \cup \{x\}\}.$$

Note that $M_{k+1}(x) \supset M_k(x)$ since $W_{k+1} \supset W_k$, and that from the rule for a node to enter V , we have

$$M_k(x) = \emptyset \quad \text{if and only if} \quad x \notin W_k \cup V_k, \quad (5.13)$$

[the reason why $M_k(x) \neq \emptyset$ for all $x \in W_k \cup V_k$ is that for entrance of x in V at some iteration there must exist $u \in U(x)$ such that $Y(x, u)$ is a subset of $W \cup \{y^*\}$ at that iteration].

We will prove by induction that for all $x \in X \cup \{t\}$, we have

$$J_k(x) = \begin{cases} \min_{\mu \in M_k(x)} \max_{p \in P(x, \mu)} L_{p, \mu} & \text{if } x \in W_k \cup V_k, \\ \infty & \text{if } x \notin W_k \cup V_k. \end{cases} \quad (5.14)$$

[In words, we will show that at the start of iteration k , $J_k(x)$ is the shortest “minimax” distance from x to t , using proper policies, which generate paths that start at x and go exclusively through W_k . The idea is that since nodes in W_k have smaller labels than nodes not in W_k and the arc lengths are nonnegative, it would not be optimal to use a path that moves in and out of W_k . This is similar to proofs for Dijkstra’s algorithm for the deterministic shortest path problem.] Equation (5.14) will imply that upon termination, when $M_{N+1}(x)$ is equal to the set of all proper policies, we have

$$J_{N+1}(x) = \min_{\mu: \text{proper}} \max_{p \in P(x, \mu)} L_{p, \mu}, \quad \forall x \in X,$$

which will prove the proposition. We will often use the fact that for all $x \in W_{k+1} \cup V_{k+1}$ and $\mu \in M_{k+1}(x)$ we have

$$\max_{p \in P(x, \mu)} L_{p, \mu} = \max_{y \in Y(x, \mu(x))} \left[g(x, \mu(x), y) + \max_{p' \in P(y, \mu)} L_{p', \mu} \right]. \quad (5.15)$$

This is just the optimality equation for the longest path problem associated with μ and the subgraph of arcs A_μ .

Initially, for $k = 0$, we have $W_0 = \emptyset$, $V_0 = \{t\}$, $J_0(t) = 0$, $J_0(x) = \infty$ for $x \neq t$, so Eq. (5.14) holds. Assume that Eq. (5.14) holds for some k . We will show that it holds for $k + 1$.

For $x \notin W_{k+1} \cup V_{k+1}$, we have $J_{k+1}(x) = \infty$, since such x have never entered V so far, and therefore their label was never reduced from the initial value ∞ . This proves Eq. (5.14) with k replaced by $k + 1$ and $x \notin W_{k+1} \cup V_{k+1}$.

For $x = y_k^*$, from the definition (5.13), the set of policies $M_{k+1}(y_k^*)$ is equal to $M_k(y_k^*)$, so by using also the induction hypothesis and the fact $J_{k+1}(y_k^*) = J_k(y_k^*)$, it follows that

$$J_{k+1}(y_k^*) = J_k(y_k^*) = \min_{\mu \in M_k(y_k^*)} \max_{p \in P(y_k^*, \mu)} L_{p, \mu} = \min_{\mu \in M_{k+1}(y_k^*)} \max_{p \in P(y_k^*, \mu)} L_{p, \mu}.$$

This proves Eq. (5.14) with k replaced by $k + 1$ and $x = y_k^*$.

For $x \in W_k \cup V_{k+1}$, we write

$$\min_{\mu \in M_{k+1}(x)} \max_{p \in P(x, \mu)} L_{p, \mu} = \min[E_1, E_2],$$

where

$$E_1 = \min_{\mu \in M_k(x)} \max_{p \in P(x, \mu)} L_{p, \mu},$$

which is equal to $J_k(x)$ by the induction hypothesis, and

$$E_2 = \min_{\mu \in M_{k+1}(x)/M_k(x)} \max_{p \in P(x, \mu)} L_{p, \mu}.$$

[The set $M_{k+1}(x)/M_k(x)$ may be empty, so here and later we use the convention that the minimum over the empty set is equal to ∞ .] Thus for $x \in W_k \cup V_{k+1}$, we have

$$\min_{\mu \in M_{k+1}(x)} \max_{p \in P(x, \mu)} L_{p, \mu} = \min \left[J_k(x), \min_{\mu \in M_{k+1}(x)/M_k(x)} \max_{p \in P(x, \mu)} L_{p, \mu} \right], \quad (5.16)$$

and we need to show that the right-hand side above is equal to $J_{k+1}(x)$. To estimate the second term of the right-hand side, we consider the two separate cases where $x \in W_k$ and $x \in V_{k+1}$.

Assume first that $x \in W_k$. Then for each $\mu \in M_{k+1}(x)/M_k(x)$, the subgraph of arcs A_μ must contain y_k^* , so

$$\max_{p \in P(x, \mu)} L_{p, \mu} \geq \max_{p' \in P(y_k^*, \mu)} L_{p', \mu} \geq J_k(y_k^*) \geq J_k(x),$$

where the first inequality follows in view of the nonnegativity of the arc lengths, the second inequality follows because paths in $P(y_k^*, \mu)$ go exclusively through W_k and thus can also be generated by some policy in $M_k(y_k^*)$, and the last inequality follows since nodes enter W in order of nondecreasing label (cf. Prop. 5.1). Thus the expression (5.16) is equal to $J_k(x)$, and hence is also equal to $J_{k+1}(x)$ (since $x \in W_k$). This proves Eq. (5.14) with k replaced by $k + 1$ and $x \in W_k$.

Assume next that $x \in V_{k+1}$. Then to estimate the term $\min_{\mu \in M_{k+1}(x)/M_k(x)} \max_{p \in P(x,\mu)} L_{p,\mu}$ in Eq. (5.16), we write

$$M_{k+1}(x)/M_k(x) = \tilde{M}_{k+1}(x) \cup \hat{M}_{k+1}(x),$$

where

$$\tilde{M}_{k+1}(x) = \{\mu \in M_{k+1}(x)/M_k(x) \mid y_k^* \notin Y(x, \mu(x))\},$$

is the set of policies for which there exists a path $p \in P(x, \mu)$ that passes through y_k^* after more than one transition, and

$$\hat{M}_{k+1}(x) = \{\mu \in M_{k+1}(x)/M_k(x) \mid y_k^* \in Y(x, \mu(x))\},$$

is the set of policies for which there exists a path $p \in P(x, \mu)$ that moves to y_k^* at the first transition.

For all $\mu \in \tilde{M}_{k+1}(x)$, we have using Eq. (5.15),

$$\begin{aligned} \max_{p \in P(x,\mu)} L_{p,\mu} &= \max_{y \in Y(x,\mu(x))} \left[g(x, \mu(x), y) + \max_{p' \in P(y,\mu)} L_{p',\mu} \right] \\ &\geq \max_{y \in Y(x,\mu(x))} \left[g(x, \mu(x), y) + \min_{\mu' \in M_{k+1}(y)} \max_{p' \in P(y,\mu')} L_{p',\mu'} \right] \\ &= \max_{y \in Y(x,\mu(x))} \left[g(x, \mu(x), y) + J_{k+1}(y) \right] \\ &= \max_{y \in Y(x,\mu(x))} \left[g(x, \mu(x), y) + J_k(y) \right], \end{aligned}$$

where the last equality holds because we have shown that $J_{k+1}(y) = J_k(y)$ for all $y \in W_k$. Therefore, since for all $\mu \in \tilde{M}_{k+1}(x)$ there exists a $\mu' \in M_k(x)$ such that $\mu(x) = \mu'(x)$, we have by taking the minimum over $\mu \in \tilde{M}_{k+1}(x)$ in the preceding relation,

$$\min_{\mu \in \tilde{M}_{k+1}(x)} \max_{p \in P(x,\mu)} L_{p,\mu} \geq \min_{\mu' \in M_k(x)} \max_{y \in Y(x,\mu'(x))} \left[g(x, \mu'(x), y) + J_k(y) \right] = \min_{\mu' \in M_k(x)} \max_{p \in P(x,\mu')} L_{p',\mu'} = J_k(x). \quad (5.17)$$

For all $\mu \in \hat{M}_{k+1}(x)$, we have

$$\max_{p \in P(x,\mu)} L_{p,\mu} = \max_{y \in Y(x,\mu(x))} \left[g(x, \mu(x), y) + \max_{p' \in P(y,\mu)} L_{p',\mu} \right];$$

cf. Eq. (5.15). Moreover, for all $\mu \in \hat{M}_{k+1}(x)$, we have $\mu(x) \in \hat{U}_k(x)$ by the definitions of $\hat{M}_{k+1}(x)$ and $\hat{U}_k(x)$. It follows that

$$\begin{aligned} \min_{\mu \in \hat{M}_{k+1}(x)} \max_{p \in P(x,\mu)} L_{p,\mu} &= \min_{\mu \in \hat{M}_{k+1}(x)} \max_{y \in Y(x,\mu(x))} \left[g(x, \mu(x), y) + \max_{p' \in P(y,\mu)} L_{p',\mu} \right] \\ &= \min_{u \in \hat{U}_k(x)} \max_{y \in Y(x,u)} \left[g(x, u, y) + \min_{\mu \in \hat{M}_{k+1}(x)} \max_{p' \in P(y,\mu)} L_{p',\mu} \right] \\ &= \min_{u \in \hat{U}_k(x)} \max_{y \in Y(x,u)} \left[g(x, u, y) + \min_{\mu \in M_{k+1}(y)} \max_{p' \in P(y,\mu)} L_{p',\mu} \right] \\ &= \min_{u \in \hat{U}_k(x)} \max_{y \in Y(x,u)} \left[g(x, u, y) + J_{k+1}(y) \right] \\ &= \min_{u \in \hat{U}_k(x)} \max_{y \in Y(x,u)} \left[g(x, u, y) + J_k(y) \right], \end{aligned} \quad (5.18)$$

where the third equality holds because for $y \in Y(x, \mu(x))$, the collections of paths $P(y, \mu)$ under policies μ in $\hat{M}_{k+1}(x)$ and $M_{k+1}(y)$, are identical, and the last equality holds because we have already shown that $J_{k+1}(y) = J_k(y)$ for all $y \in W_k \cup \{y_k^*\}$. Thus from Eqs. (5.16)-(5.18) we obtain

$$\min_{\mu \in M_{k+1}(x)} \max_{p \in P(x, \mu)} L_{p, \mu} = \min \left[J_k(x), \min_{u \in \hat{U}_k(x)} \max_{y \in Y(x, u)} [g(x, u, y) + J_k(y)] \right].$$

Combining this equation with the update formula (5.10) for the node labels, we have

$$\min_{\mu \in M_{k+1}(x)} \max_{p \in P(x, \mu)} L_{p, \mu} = J_{k+1}(x),$$

thus proving Eq. (5.14) with k replaced by $k + 1$ and $x \in V_{k+1}$. This completes the induction proof of Eq. (5.14) and concludes the proof. **Q.E.D.**

Since the algorithm terminates in $N+1$ iterations, and each iteration requires at most $O(AL)$ operations, where A is the number of arcs and L is the number of elements of U , the complexity of the algorithm is bounded by $O(NAL)$. This complexity estimate may potentially be improved with the use of efficient data structures of the type used in efficient implementations of Dijkstra's algorithm in deterministic shortest path problems to expedite the selection of the minimum label node [i.e., $y^* \in \arg \min_{y \in V} J(y)$]. However, we have not investigated this issue. It is also unclear how the Dijkstra-like algorithm compares with the finitely terminating VI algorithm and other asynchronous VI and PI algorithms discussed in Sections 5.1, 5.2.

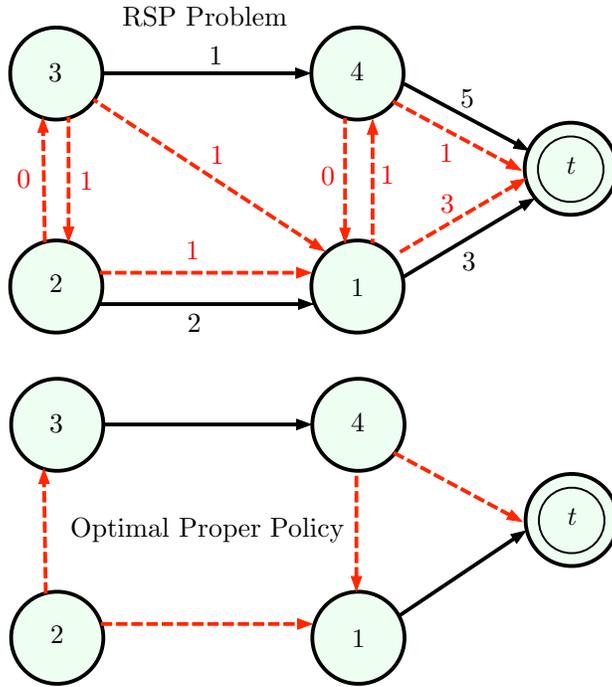
Example 5.2:

We illustrate our Dijkstra-like algorithm with the RSP problem of Example 5.1, shown again for convenience in Fig. 5.2. The table gives the iterations of the algorithm, and the results are consistent with Props. 5.1-5.3, as the reader may verify.

6. CONCLUDING REMARKS AND FURTHER RESEARCH

We have considered shortest path problems with set membership uncertainty, and we have shown that they can be fruitfully analyzed in the context of abstract semicontractive models. We have thus proved the existence and uniqueness of the solution of Bellman's equation, and obtained conditions for optimality of a proper policy. Moreover, we have discussed the properties of algorithms of the value and policy iteration type, and we have proposed a finitely terminating Dijkstra-like algorithm for problems with nonnegative arc lengths. The comparative evaluation and the efficient implementation of these algorithms for specific types of applications, such as for example pursuit-evasion, is an interesting subject for further investigation.

In this paper we have covered the important case of nonnegative arc lengths via the perturbation analysis of Section 4. However, there is an alternative line of analysis, which is based on the fact that when the arc lengths are nonnegative we have $T\bar{J} \geq \bar{J}$, bringing to bear the theory of semicontractive-monotone increasing DP models given in Section 4.4 of [Ber13], which embody the essential structure of negative DP (see [Str66], or the texts Put94, [Ber12]). This theory relies on the assumption of existence of an optimal S -regular policy, and is somewhat different in character from the analysis of this paper.



Iteration #	V at Start of Iteration	Labels at Start of Iteration	Min. Label Node out of V
0	$\{t\}$	$(\infty, \infty, \infty, \infty, 0)$	t
1	$\{1, 4\}$	$(3, \infty, \infty, 5, 0)$	1
2	$\{2, 4\}$	$(3, 5, \infty, 3, 0)$	4
3	$\{2, 3\}$	$(3, 5, 4, 3, 0)$	3
4	$\{2\}$	$(3, 4, 4, 3, 0)$	2

Figure 5.2: The iterations of the Dijkstra-like algorithm for the RSP problem of Example 5.1. The nodes exit V and enter W in the order 1, 4 3, 2.

Let us mention some interesting stochastic extensions of our RSP problem that involve an additional random variable at each stage.[†] In one extension of this type, when at node x , we choose control $u \in U(x)$, then a value of a random variable z is selected from the finite set $\{1, \dots, n\}$ with probabilities $p_{xz}(u)$, and then the successor node is chosen by an antagonistic opponent from a set $Y_z(x, u) \subset X \cup \{t\}$. To analyze this problem using a semicontractive model, the mapping H of Eq. (2.2) should be replaced by

$$H(x, u, J) = \sum_{z=1}^n p_{xz}(u) \max_{y \in Y_z(x, u)} [g(x, u, y) + \tilde{J}(y)], \quad (6.1)$$

[†] This type of stochastic problem arises among others in a context of discretization of the state space of a continuous-space minimax control problem, where randomization in the discretized problem's dynamics is introduced in order to reduce the error between the optimal cost function of the original continuous-space problem and the optimal cost function of its discretized version (see, [Tsi95], [PBT98], [Set99a], [Set99b]).

where

$$\tilde{J}(y) = \begin{cases} J(y) & \text{if } y \in X, \\ 0 & \text{if } y = t. \end{cases}$$

A formulation as an abstract DP problem is then possible with an appropriately modified mapping T_μ , similar to Section 2, and the semicontractive theory may be applied similar to Section 4. This analysis and the associated algorithms are, however, beyond our scope. Note that when $n = 1$, from Eq. (6.1) we obtain the mapping (2.2) of the RSP problem of Sections 1 and 2, while when $Y_z(x, u)$ consists of a single node, we obtain the mapping associated with a standard finite-state stochastic shortest path problem (see e.g., [Pal67], [BeT91], [Ber12]). Thus the semicontractive model that is based on the mapping (6.1) generalizes both of these shortest path problems. In the case where $g(x, u, y) \geq 0$ for all $x \in X$, $u \in U(x)$, $z \in \{1, \dots, n\}$, and $y \in Y_z(x, u)$, we have $T\bar{J} \geq \bar{J}$, and the theory of monotone increasing models of Sections 4.3, 4.4 of [Ber13] can be used to provide a first layer of analysis without any further assumptions.

In another extension of RSP, which involves randomization, when at node x , we choose control $u \in U(x)$, then a variable w is chosen by an antagonistic opponent from a set $W(x, u)$, and then a successor node $y \in X \cup \{t\}$ is chosen according to probabilities p_{wy} . Here, to apply the line of analysis of the present paper, the mapping H of Eq. (2.2) should be replaced by

$$H(x, u, J) = \max_{w \in W(x, u)} \left[g(x, u, w) + \sum_{y \in X \cup \{t\}} p_{wy} \tilde{J}(y) \right],$$

where

$$\tilde{J}(y) = \begin{cases} J(y) & \text{if } y \in X, \\ 0 & \text{if } y = t. \end{cases} \quad (6.2)$$

A somewhat more complex variation is given by

$$H(x, u, J) = \sum_{z=1}^n p_{xz} \max_{w \in W(z, u)} \left[g(z, u, w) + \sum_{y \in X \cup \{t\}} p_{wy} \tilde{J}(y) \right], \quad (6.3)$$

where, for each $x \in X$, z is a random variable taking values from the finite set $\{1, \dots, n\}$ with probabilities p_{xz} [the two models based on Eqs. (6.2) and (6.3) coincide if $p_{xx} = 1$ for all $x \in X$]. The resulting models again combine elements of RSP and a standard finite-state stochastic shortest path problem. They may also be viewed as instances of a generalized aggregation model of the type introduced in Section 7.3.7 of [Ber12].[†] Again the semicontractive theory may be applied, but the corresponding analysis is a subject for further research.

Finally let us consider algorithms with approximations. While we have assumed a finite number of nodes, there are many problems of practical interest where the number of nodes is extremely large, and the algorithms of Section 5 are very time-consuming. This is particularly so in minimax control problems with imperfect state information, which are reformulated as problems of perfect state information using the sufficiently informative function approach of [Ber71], [BeR73]. In such cases one may consider minimax analogs of approximate DP or reinforcement learning approaches (see e.g., the books [BeT96], [SuB98],

[†] In the context of generalized aggregation, for the mapping of Eq. (6.2), we have a high-dimensional RSP problem, whose states are represented by w , and a lower dimensional “aggregate” RSP problem, whose states are represented by the nodes x . Then the p_{wy} play the role of aggregation probabilities in the terminology of [Ber12]. A similar interpretation may be given for the mapping of Eq. (6.3).

[Pow11], [Ber12]). The development of such methods is an interesting subject for further research. In what follows, we discuss the possibility for approximate solution using a rollout algorithm, which parallels related algorithms for finite horizon and stochastic shortest path problems (see e.g., [Ber05], [Ber12]).

To define a rollout algorithm, we assume that we have a proper policy μ , called the *base policy*. We then define the *rollout policy*, denoted $\bar{\mu}$, as follows: for each $u \in U(x)$ and each $y \in Y(x, u)$, we compute $J_\mu(y)$, and we set

$$\bar{\mu}(x) \in \arg \min_{u \in U(x)} \max_{y \in Y(x, u)} [g(x, u, y) + J_\mu(y)], \quad x \in X. \quad (6.4)$$

The computation of $\bar{\mu}(x)$ can be done on-line, only for the nodes x that are encountered in the process of control. Moreover, assuming that $J_\mu(y)$ is precomputed for all y , and that the sets $U(x)$ and $Y(x, u)$ have a reasonable number of elements, the computation in Eq. (6.4) can be performed very fast. The same is true if for any $u \in U(x)$, $J_\mu(y)$ can be efficiently computed on-line for all $y \in Y(x, u)$. It can be seen that the rollout policy $\bar{\mu}$ is just the policy obtained from the base policy using a single policy iteration. In particular, under Assumption 2.1, the rollout policy improves on the base policy in the sense that

$$J_{\bar{\mu}}(x) \leq J_\mu(x), \quad \forall x \in X.$$

This is a well-known property of rollout algorithms for finite horizon and stochastic shortest path problems, and follows by writing for all $x \in X$

$$\begin{aligned} (T_{\bar{\mu}}J_\mu)(x) &= (TJ)(x) \\ &= \min_{u \in U(x)} \max_{y \in Y(x, u)} [g(x, u, y) + J_\mu(y)] \\ &\leq \max_{y \in Y(x, \mu(x))} [g(x, \mu(x), y) + J_\mu(y)] \\ &= (T_\mu J_\mu)(x) \\ &= J_\mu(x). \end{aligned}$$

Applying $T_{\bar{\mu}}$ repeatedly to both sides of the inequality $T_{\bar{\mu}}J_\mu \leq J_\mu$, we obtain [cf. Eq. (5.5)] that $J_{\bar{\mu}} \leq J_\mu$ and that $\bar{\mu}$ is proper.

As an example of such a rollout algorithm, consider a pursuit-evasion problem with state $x = (z_1, z_2)$, where z_1 is the location of the minimizer/pursuer and z_2 is the location of the maximizer/evader, in a two- or three-dimensional space. Then a suitable base policy μ is for the pursuer is to follow a shortest path from z_1 to z_2 under the assumption that the evader will stay at his current location z_2 at all future times. To do this for all (z_1, z_2) requires the solution of an all-pairs shortest path problem, which is possible in $O(N^3)$ time using the Floyd-Warshall algorithm [AMO89], [Ber98], where N is the number of possible values of z_1 and z_2 . Having precomputed $\mu(x)$ for all $x = (z_1, z_2)$ with this shortest path computation, the maximization

$$\max_{y \in Y(x, u)} [g(x, u, y) + J_\mu(y)]$$

that is needed for the on-line computation of the rollout control $\bar{\mu}(x)$ in Eq. (6.4) is a tractable longest path computation in a graph of N nodes. Note that the preceding algorithm can be adapted for the imperfect information case where the pursuer knows z_2 imperfectly. This is possible by using a form of assumed certainty equivalence: the pursuer's base policy and the evader's maximization can be computed by using an estimate of the current location z_2 instead of the unknown true location.

In the preceding pursuit-evasion example, the choice of the base policy was facilitated by the special structure of the problem. Generally, however, finding a suitable base policy whose cost function J_μ can be

conveniently computed is an important problem-dependent issue. We leave this issue as a subject for further research in the context of more specialized problems. Finally, let us note that a rollout algorithm may be well-suited for on-line suboptimal solution in cases where data may be changing or be revealed during the process of control.

7. REFERENCES

- [AHK07] L. Alfaro, L. T. Henzinger, T., and O. Kupferman, O., 2007. “Concurrent Reachability Games,” *Theoretical Computer Science*, Vol. 386, pp. 188-217.
- [AMO89] Ahuja, R. K., Magnanti, T. L., and Orlin, J. B., 1989. “Network Flows,” in *Handbooks in Operations Research and Management Science*, Vol. 1, Optimization, Nemhauser, G. L., Rinnooy-Kan, A. H. G., and Todd M. J. (eds.), North-Holland, Amsterdam, pp. 211-369.
- [AIM12] Alton, K., and Mitchell, I. M., 2012. “An Ordered Upwind Method with Precomputed Stencil and Monotone Node Acceptance for Solving Static Hamilton-Jacobi Equations,” *J. of Scientific Computing*, Vol. 51, pp. 313-348.
- [AnV13] Andrews, J., and Vladimirsky, A., 2013. “Deterministic Control of Randomly-Terminated Processes,” arXiv:1310.7161v1 [math.OC].
- [BBC11] Bertsimas, D., Brown, D. B., and Caramanis, C., 2011. “Theory and Applications of Robust Optimization,” *SIAM Rev.*, Vol. 53, pp. 464-501.
- [BBF95] Bardi, M., Bottacin, S., and Falcone, M., 1995. “Convergence of Discrete Schemes for Discontinuous Value Functions of Pursuit-Evasion Games,” *New trends in dynamic games and applications*, Vol. 3, pp. 273-304, *Annals Intern. Soc. Dynamic Games*, Birkhauser, Boston.
- [BBH08] Bopardikar, S. D., Bullo, F., and Hespanha, J. P., 2008. “On Discrete-Time Pursuit-Evasion Games With Sensing Limitations,” *IEEE Trans. on Robotics*, Vol. 24, pp. 1429-1439.
- [BFS94] Bardi, M., Falcone, M., and Soravia, P., 1994. “Fully Discrete Schemes for the Value Function of Pursuit-Evasion Games,” *New trends in dynamic games and applications*, Vol. 1, pp. 89-105, *Annals Intern. Soc. Dynamic Games*, Birkhauser, Boston.
- [BGM95] Bertsekas, D. P., Guerriero, F., and Musmanno, R., 1995. “Parallel Shortest Path Methods for Globally Optimal Trajectories,” *High Performance Computing: Technology, Methods, and Applications*, (J. Dongarra et al., Eds.), Elsevier.
- [BGN09] Ben-Tal, A., El Ghaoui, L., and Nemirovski, A., 2009. *Robust Optimization*, Princeton University Press, Princeton, N. J.
- [BaB91] Basar, T., and Bernhard, P., 1991. *H-Infinity Optimal Control and Related Minimax Design Problems: A Dynamic Game Approach*, Birkhauser, Boston.
- [BaL13] Bardi, M., and Lopez, J. P. M., 2013. “A Dijkstra-Type Algorithm for Dynamic Games,” Sorbonne Universites, Archive Ouverte HAL-UPMC: hal-00881218.
- [BaS91] Bardi, M., and Soravia, P., 1991. “Approximation of Differential Games of Pursuit-Evasion by Discrete-Time Games,” *Differential games-developments in modeling and computation (Espoo, 1990)*, pp. 131-143, *Lecture Notes in Control and Information Sci.*, 156, Springer, Berlin.
- [BeR71] Bertsekas, D. P., and Rhodes, I. B., 1971. “Recursive State Estimation for a Set-Membership

- Description of the Uncertainty,” *IEEE Trans. Automatic Control*, Vol. AC-16, pp. 117-128.
- [BeR73] Bertsekas, D. P., and Rhodes, I. B., 1973. “Sufficiently Informative Functions and the Minimax Feedback Control of Uncertain Dynamic Systems,” *IEEE Trans. on Automatic Control*, Vol. AC-18, pp. 117-124.
- [BeS78] Bertsekas, D. P., and Shreve, S. E., 1978. *Stochastic Optimal Control: The Discrete Time Case*, Academic Press, N. Y.; may be downloaded from <http://web.mit.edu/dimitrib/www/home.html>
- [BeS03] Bertsimas, D., and Sim, M., 2003. “Robust Discrete Optimization and Network Flows,” *Math. Programming*, Vol. 98, pp. 49-71.
- [BeT89] Bertsekas, D. P., and Tsitsiklis, J. N., 1989. *Parallel and Distributed Computation: Numerical Methods*, Prentice-Hall, Englewood Cliffs, N. J; republished by Athena Scientific, Belmont, MA, 1997.
- [BeT96] Bertsekas, D. P., and Tsitsiklis, J. N., 1996. *Neuro-Dynamic Programming*, Athena Scientific, Belmont, MA.
- [BeT91] Bertsekas, D. P., and Tsitsiklis, J. N., 1991. “An Analysis of Stochastic Shortest Path Problems,” *Math. of OR*, Vol. 16, pp. 580-595.
- [BeY10] Bertsekas, D. P., and Yu, H., 2010. “Asynchronous Distributed Policy Iteration in Dynamic Programming,” *Proc. of Allerton Conf. on Communication, Control and Computing*, Allerton Park, Ill, pp. 1368-1374.
- [BeY12] Bertsekas, D. P., and Yu, H., 2012. “Q-Learning and Enhanced Policy Iteration in Discounted Dynamic Programming,” *Math. of OR*, Vol. 37, pp. 66-94.
- [BeY13] Bertsekas, D. P., and Yu, H., 2013. “Stochastic Shortest Path Problems Under Weak Conditions,” *Lab. for Information and Decision Systems Report LIDS-P-2909*, MIT.
- [Ber71] Bertsekas, D. P., 1971. “Control of Uncertain Systems With a Set-Membership Description of the Uncertainty,” Ph.D. Dissertation, Massachusetts Institute of Technology, Cambridge, MA (available in scanned form from the author’s www site).
- [Ber77] Bertsekas, D. P., 1977. “Monotone Mappings with Application in Dynamic Programming,” *SIAM J. on Control and Optimization*, Vol. 15, pp. 438-464.
- [Ber82] Bertsekas, D. P., 1982. “Distributed Dynamic Programming,” *IEEE Trans. Aut. Control*, Vol. AC-27, pp. 610-616.
- [Ber83] Bertsekas, D. P., 1983. “Asynchronous Distributed Computation of Fixed Points,” *Math. Programming*, Vol. 27, pp. 107-120.
- [Ber98] Bertsekas, D. P., 1998. *Network Optimization: Continuous and Discrete Models*, Athena Scientific, Belmont, MA.
- [Ber05] Bertsekas, D. P., 2005. *Dynamic Programming and Optimal Control*, Vol. I, 3rd Edition, Athena Scientific, Belmont, MA.
- [Ber12] Bertsekas, D. P., 2012. *Dynamic Programming and Optimal Control*, Vol. II, 4th Edition: Approximate Dynamic Programming, Athena Scientific, Belmont, MA.
- [Ber13] Bertsekas, D. P., 2013. *Abstract Dynamic Programming*, Athena Scientific, Belmont, MA.
- [Ber14] Bertsekas, D. P., 2014. “Infinite-Space Shortest Path Problems and Semicontractive Dynamic Programming,” *Lab. for Information and Decision Systems Report*, MIT.

- [BIM08] Blanchini, F., and Miani, S., 2008. *Set-Theoretic Methods in Control*, Birkhauser, Boston, MA.
- [Bla99] Blanchini, F., 1999. "Set Invariance in Control – A Survey," *Automatica*, Vol. 35, pp. 1747-1768.
- [Bon07] Bonet, B., 2007. "On the Speed of Convergence of Value Iteration on Stochastic Shortest-Path Problems," *Math. of Operations Research*, Vol. 32, pp. 365-373.
- [CCV13] Clawson, Z., Chacon, A., and Vladimirov, A., 2013. "Causal Domain Restriction for Eikonal Equations," arXiv preprint arXiv:1309.2884.
- [CaB04] Camacho, E. F., and Bordons, C., 2004. *Model Predictive Control*, 2nd Edition, Springer-Verlag, New York, N. Y.
- [ChV12] Chacon, A., and Vladimirov, A., 2012. "Fast Two-scale Methods for Eikonal Equations," *SIAM Journal on Scientific Computing*, Vol. 34, pp. A547-A577.
- [Den67] Denardo, E. V., 1967. "Contraction Mappings in the Theory Underlying Dynamic Programming," *SIAM Review*, Vol. 9, pp. 165-177.
- [Der70] Derman, C., 1970. *Finite State Markovian Decision Processes*, Academic Press, N. Y.
- [Dre69] Dreyfus, S. E., 1969. "An Appraisal of Some Shortest-Path Algorithms," Vol. 17, pp. 395-412.
- [EGB12] Ermon, S., Gomes, C., Selman, B., and Vladimirov, A., 2012. "Probabilistic Planning with Non-Linear Utility Functions and Worst-Case Guarantees," *Proc. of the 11th Intern. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, Valencia, Spain.
- [FiV96] Filar, J., and Vrieze, K., 1996. *Competitive Markov Decision Processes*, Springer, N. Y.
- [Fal87] Falcone, M., 1987. "A Numerical Approach to the Infinite Horizon Problem of Deterministic Control Theory," *Appl. Math. Opt.*, Vol. 15, pp. 1-13.
- [GLL99] Guibas, L. J., Latombe, J.-C., LaValle, S. M., D. Lin, and Motwani, R. 1999. "A Visibility-Based Pursuit-Evasion Problem," *Intern. J. of Computational Geometry and Applications*, Vol. 9, pp. 471-493.
- [GaP88] Gallo, G., and Pallottino, S., 1988. "Shortest Path Algorithms," *Annals of Operations Research*, Vol. 7, pp. 3-79.
- [GoR85] Gonzalez, R., and Rofman, E., 1985. "On Deterministic Control Problems: An Approximation Procedure for the Optimal Cost, Parts I, II," *SIAM J. Control Optimization*, Vol. 23, pp. 242-285.
- [GrJ08] Grune, L., and Junge, O., 2008. "Global Optimal Control of Perturbed Systems," *J. of Optimization Theory and Applications*, Vol. 136, pp. 411-429.
- [HCP99] Hernandez-Lerma, O., Carrasco, O., and Perez-Hernandez. 1999. "Markov Control Processes with the Expected Total Cost Criterion: Optimality, Stability, and Transient Models," *Acta Appl. Math.*, Vol. 59, pp. 229-269.
- [HKR93] Hu, T. C., Kahng, A. B., and Robins, G., 1993. "Optimal Robust Path Planning in General Environments," *IEEE Trans. on Robotics and Automation*, Vol. 9.
- [HiW05] Hinderer, K., and Waldmann, K.-H., 2005. "Algorithms for Countable State Markov Decision Models with an Absorbing Set," *SIAM J. of Control and Optimization*, Vol. 43, pp. 2109-2131.
- [JaC06] James, H. W., and Collins, E. J., 2006. "An Analysis of Transient Markov Decision Processes," *J. Appl. Prob.*, Vol. 43, pp. 603-621.
- [Ker00] Kerrigan, E. C., 2000. *Robust Constraint Satisfaction: Invariant Sets and Predictive Control*, Ph.D. Dissertation, Univ. of Cambridge, England.

- [KuD92] Kushner, H. J., and Dupuis, P. G., 1992. *Numerical Methods for Stochastic Control Problems in Continuous Time*, Springer-Verlag, N. Y.
- [KuV97] Kurzhanski, A., and Valyi, I., 1997. *Ellipsoidal Calculus for Estimation and Control*, Birkhauser, Boston, MA.
- [LaV06] LaValle, S. M., 2006. *Planning Algorithms*, Cambridge University Press, N. Y.
- [MHG88] Meggido, N., Hakimi, S. L., Garey, M. R., and Johnson, D. S., 1988. "The Complexity of Searching a Graph," *J. of ACM*, Vol. 35, pp. 18-44.
- [MRR00] Mayne, D. Q., Rawlings, J. B., Rao, C. V., and Scolaert, P. O. M., 2000. "Constrained Model Predictive Control: Stability and Optimality," *Automatica*, Vol. 36, pp. 789-814.
- [Mac02] Maciejowski, J. M., 2002. *Predictive Control with Constraints*, Addison-Wesley, Reading, MA.
- [Mir14] Mirebeau, J-M., 2014. "Efficient Fast Marching with Finsler Metrics," Vol. 126, pp. 515-557.
- [MoG04] Montemanni, R., and Gambardella, L. M., 2004. "An Exact Algorithm for the Robust Shortest Path Problem with Interval Data," *Computers and Operations Research*, Vol. 31, pp. 1667-1680.
- [MoL99] Morari, M., and Lee, J. H., 1999. "Model Predictive Control: Past, Present, and Future," *Computers and Chemical Engineering*, Vol. 23, pp. 667-682.
- [PBT98] Polymenakos, L. C., Bertsekas, D. P., and Tsitsiklis, J. N., 1998. "Efficient Algorithms for Continuous-Space Shortest Path Problems," *IEEE Trans. on Aut. Control*, Vol. 43, pp. 278-283.
- [PaB99] Patek, S. D., and Bertsekas, D. P., 1999. "Stochastic Shortest Path Games," *SIAM J. on Control and Optimization*, Vol. 36, pp. 804-824.
- [Pal67] Pallu de la Barriere, R., 1967. *Optimal Control Theory*, Saunders, Phila; republished by Dover, N. Y., 1980.
- [Par76] Parsons, T. D., 1976. "Pursuit-Evasion in a Graph," in *Theory and Applications of Graphs*, Y. Alavi and D. Lick, Eds., Springer-Verlag, Berlin, pp. 426-441.
- [Pli78] Pliska, S. R., 1978. "On the Transient Case for Markov Decision Chains with General State Spaces," in *Dynamic Programming and Its Applications*, M. L. Puterman (ed.), Academic Press, N. Y.
- [Pow11] Powell, W. B., 2011. *Approximate Dynamic Programming: Solving the Curses of Dimensionality*, 2nd edition, J. Wiley and Sons, Hoboken, N. J.
- [Put94] Puterman, M. L., 1994. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*, J. Wiley, N. Y.
- [RaM09] Rawlings, J. B., and Mayne, D. Q., 2009. *Model Predictive Control Theory and Design*, Nob Hill Publishing, Madison, WI.
- [Roc84] Rockafellar, R. T., 1984. *Network Flows and Monotropic Programming*, Wiley-Interscience, N. Y.
- [Set99a] Sethian, J. A., 1999. *Level Set Methods and Fast Marching Methods Evolving Interfaces in Computational Geometry, Fluid Mechanics, Computer Vision, and Materials Science*, Cambridge University Press, N. Y.
- [Set99b] Sethian, J. A., 1999. "Fast Marching Methods," *SIAM Review*, Vol. 41, pp. 199-235.
- [Str66] Strauch, R., 1966. "Negative Dynamic Programming," *Ann. Math. Statist.*, Vol. 37, pp. 871-890.
- [SuB98] Sutton, R. S., and Barto, A. G., 1998. *Reinforcement Learning*, MIT Press, Cambridge, MA.

- [Tsi95] Tsitsiklis, J. N., 1995. "Efficient Algorithms for Globally Optimal Trajectories," *IEEE Trans. Automatic Control*, Vol. AC-40, pp. 1528-1538.
- [VKS02] Vidal, R., Kim, H. J., Shim, D. H., and Sastry, S., 2002. "Probabilistic Pursuit-Evasion Games: Theory, Implementation, and Experimental Evaluation," *IEEE Trans. on Robotics and Automation*, Vol. 18, pp. 662-669.
- [Vla08] Vladimirov, A., 2008. "Label-Setting Methods for Multimode Stochastic Shortest Path Problems on Graphs," *Math. of Operations Research*, Vol. 33, pp. 821-838.
- [WiB93] Williams, R. J., and Baird, L. C., 1993. "Analysis of Some Incremental Variants of Policy Iteration: First Steps Toward Understanding Actor-Critic Learning Systems," Report NU-CCS-93-11, College of Computer Science, Northeastern University, Boston, MA.
- [Wit66] Witsenhausen, H. S., 1966. "Minimax Control of Uncertain Systems," Ph.D. Dissertation, Massachusetts Institute of Technology, Cambridge, MA.
- [Whi82] Whittle, P., 1982. *Optimization Over Time*, Wiley, N. Y., Vol. 1, 1982, Vol. 2, 1983.
- [YuB13a] Yu, H., and Bertsekas, D. P., 2013. "Q-Learning and Policy Iteration Algorithms for Stochastic Shortest Path Problems," *Annals of Operations Research*, Vol. 208, pp. 95-132.
- [YuB13b] Yu, H., and Bertsekas, D. P., 2013. "A Mixed Value and Policy Iteration Method for Stochastic Control with Universally Measurable Policies," Lab. for Information and Decision Systems Report LIDS-P-2905, MIT, July 2013.
- [YuY98] Yu, G., and Yang, J., 1998. "On the Robust Shortest Path Problem," *Computers and Operations Research*, Vol. 25, pp. 457-468.
- [Yu11] Yu, H., 2011. "Stochastic Shortest Path Games and Q-Learning," Lab. for Information and Decision Systems Report LIDS-P-2875, MIT.