Operations Research Letters 43 (2015) 498–503

Contents lists available at ScienceDirect

Operations Research Letters

journal homepage: www.elsevier.com/locate/orl

Solvability in infinite horizon optimization

Timothy D. Lortz [a], Irina S. Dolinskaya [b,*], Archis Ghate [c], Robert L. Smith [d]

[a] Booz Allen Hamilton, Inc., Annapolis Junction, MD 20701, United States

[b] Department of Industrial Engineering and Management Sciences, Northwestern University, Evanston, IL 60208, United States

[c] Department of Industrial and Systems Engineering, University of Washington, Seattle, WA 98195, United States

[d] Department of Industrial and Operations Engineering, University of Michigan, Ann Arbor, MI 48109, United States

ARTICLE INFO

Article history:
Received 21 January 2015
Received in revised form
11 July 2015
Accepted 12 July 2015
Available online 17 July 2015

Keywords:
Dynamic programming
Forecast horizons
Coalescence

ABSTRACT

We give necessary and sufficient conditions for finite detection of an optimal initial decision for infinite horizon optimization under a broad set of assumptions and provide an algorithm that is guaranteed to solve every solvable problem under these assumptions. We illustrate the theory and algorithms developed with applications in production planning.

1. Introduction

Planners confronted by the task of making an optimal decision today cannot avoid the challenge that the quality of today's decision depends on the world one will face and the associated decisions one will make in the future. Key to this decision process is the length of the horizon over which decisions will be made. For example, the best current decision will preclude expensive actions whose investment must be recovered by savings over the long run if the study horizon is short. On the other hand, for firms that anticipate being in business for the long run, that horizon is typically indefinite with no predetermined end. In these situations, an infinite horizon model seems appropriate. However, unless one adopts the heroic assumption that the future will bring a world like the present, one is confronted by the task of forecasting an unlimited quantity of data. To attempt to overcome this, planners have traditionally adopted one of two alternatives. One can model the problem with stationary but stochastic data at least providing the illusion of a dynamic and changing world while the probability distributions remain stationary. Alternatively, the planner can adopt a limited lookahead and simply incorporate a nonstationary finite horizon model in place of the infinite horizon problem one actually confronts. The problem that arises here is the tendency of today's decision to be distorted by end of study effects. However, if those end of study distortions eventually stop, one can argue that one has found an optimal decision for today to the true underlying infinite horizon problem. In order to be certain a finite horizon optimal decision has in fact stabilized and will not change were we to lengthen it, we must somehow know that unrevealed data beyond the current horizon will not affect the optimality of the currently optimal decision in hand. Such a horizon is called a *forecast horizon*.

Considerable effort has been expended in the literature to establish conditions under which such a forecast horizon exists. Examples for which no forecast horizon exists can be found in the context of general infinite horizon optimization [1], production planning [4], and asset selling [6]. Classes of nonstationary infinite horizon optimization problems for which forecast horizons do exist have generally relied upon either uniqueness of the optimal immediate decision or monotonicity of that decision as a function of horizon length. For example, uniqueness of the first optimal decision has been established as a sufficient condition for existence of a forecast horizon for a very general class of infinite horizon optimization problems in [1]. However, although it is believed that real world problems generally satisfy this condition, we typically do not know how to check for this condition in a specific instance [12]. Indeed, if a forecast horizon does not exist (a situation we will call failure to be well-posed), the problem cannot be solved. In [6], authors give an example of such a problem whereby no matter how many periods of data are forecasted, the data not yet seen can make any decision either optimal or not optimal, thus rendering the task of determining the next best decision impossible. Such a problem is not solvable by forecasts of finite data sets, no matter how large.

* Corresponding author.
E-mail address: dolira@northwestern.edu (I.S. Dolinskaya).

http://dx.doi.org/10.1016/j.orl.2015.07.003

This paper formalizes these notions of well-posedness and solvability for a very general class of infinite horizon problems and introduces the new notion of *coalescence* of optimal states, which we show to be equivalent to being well-posed and solvable. We give necessary and sufficient conditions for finite detection of an optimal initial decision under a broad set of assumptions involving reachability between states. We show that under these assumptions, coalescence and well-posedness are equivalent, and give a solution procedure that is guaranteed to solve any well-posed problem with these properties. In particular, we provide a finite algorithm that is guaranteed to solve every solvable problem within this class of infinite horizon, deterministic nonstationary problems. We show that these reachability conditions are satisfied in infinite horizon production planning. An additional example of equipment replacement problem can be found in [10].

1.1. Related work

There is a substantial literature both in establishing general conditions under which a forecast horizon exists, and in developing solution procedures that will yield an infinite horizon optimal initial decision for an infinite horizon nonstationary optimization problem. An excellent classified bibliography of research endeavors in both of these directions over a broad spectrum of applications and theory can be found in [5]. In light of that article, we will only mention other research as it directly pertains to our work.

The problems studied in this paper all satisfy the condition that there exists a feasible solution with finite total discounted cost, and our optimality criterion will be that of total discounted cost. Under such conditions, one can bound the maximum deviation from optimal total cost by solving to optimality a finite horizon version of the problem. That is, one can obtain optimal solutions over a finite horizon together with bounds on cost-to-go that ensure total cost within error ϵ of the optimal infinite horizon cost for any $\epsilon > 0$. This paper adopts the more challenging objective of *seeking convergence by looking at incrementally longer finite horizon problems not just in terms of cost error, but also in terms of policy error.*

It is often the case that there exists a sequence of selections of finite horizon optimal initial decisions that agrees, in finite time, with an infinite horizon optimal initial decision. Algorithms to determine such selections typically require the existence of forecast horizons. In general deterministic optimization, Bean and Smith [1] show that a forecast horizon exists for a very general class of problems when the optimal solution is unique. Later Bean and Smith [2] extending an algorithm in [8] show that a weak reachability condition is necessary and sufficient for finite discovery of the optimal initial decision whenever it is unique. They also give a solution procedure that will detect an optimal initial decision in finite time when the uniqueness and weak reachability conditions are met, and stopping sets of states are appropriately chosen. In subsequent work, Bes and Sethi [3] adopted a more abstract approach to exploring general conditions for the existence of forecast horizons although these remained difficult to check in practice. However, the assumption of uniqueness of the optimal initial decision is only a sufficient condition for the solvability of a particular problem instance. Since uniqueness is not a necessary condition for discovery of an optimal initial decision, we will depart from [1,2,8], in that *we explicitly drop the assumption that the optimal initial decision (or any optimal decision, for that matter) is unique, allowing for the presence of multiple optima.* Moreover, we establish coalescence properties that are not only sufficient but, unlike uniqueness and weak reachability, are also necessary to solve an infinite horizon optimization problem. We follow [6], which seeks a finite algorithm that can solve every well-posed Markov Decision Process problem (a problem is well-posed if it is solvable). Such an algorithm is presented in [6] for solving well-posed problems whose optimal policies are monotone in horizon. In this paper, we show well-posed problems are those satisfying a coalescence property for optimal states and we provide an algorithm that solves every such problem satisfying some mild structural properties.

1.2. Paper outline

The rest of the paper is organized as follows. Section 2 presents the problem statement, including the notation used, necessary assumptions and an illustrative application. In Section 3, we establish definitions of the terminology used throughout the paper, including that of a well-posed problem, forecast horizon and coalescence, and establish the key relationship between problem *solvability*, *coalescence* and *well-posed* properties. The solution procedure for our class of problems is described in Section 4. We come back to the motivating example in Section 5 and demonstrate application of our key results to infinite horizon single-item production planning.

2. Notation and problem statement

We consider a class of infinite horizon, deterministic, discrete time, non-stationary problems parameterized by data *forecast* $\phi \in \Phi$, where $\phi = (\phi_1, \phi_2, \ldots)$, ϕ_i is a finite set of problem data associated with period i for $i = 1, 2, \ldots,$ and Φ is the set of all possible forecasts ϕ. We will often refer to Problem ϕ by which we mean the problem with forecast $\phi \in \Phi$. Let $\phi^n = (\phi_1, \ldots, \phi_n)$, for $n \geq 1$, denote the problem data for the first n periods from a forecast ϕ in Φ. We call ϕ^n a *truncated forecast at* n, and Φ^n the class of all truncated forecasts ϕ^n from $\phi \in \Phi$. Then, $\Phi(\phi^n)$ is the set of all forecasts $\theta \in \Phi$ whose data for the first n periods match that of ϕ^n, that is, $\theta^n = \phi^n$.

For a given forecast $\phi \in \Phi$, the underlying system is observed at the beginning of periods $n = 1, 2, \ldots$ to be in state $s_n \in S_n(\phi^{n-1})$, where $S_n(\phi^{n-1})$ is the set of all feasible states associated with truncated forecast ϕ^{n-1} at the beginning of period n. Then, an action $y_n \in Y_n(s_n; \phi^n)$ is chosen and a cost $c_n(y_n, s_n; \phi^n)$ is incurred, where $Y_n(s_n; \phi^n)$ is the *feasible decision space at period n*, corresponding to the set of all feasible actions when the system is in state s_n beginning period n under truncated forecast ϕ^n. The system then transitions to state s_{n+1} at the end of period n, following the state equation:

$$s_{n+1} = f_n(s_n, y_n; \phi^n), \quad \phi^n \in \Phi^n, \; y_n \in Y_n(s_n; \phi^n), \; \forall n \geq 1, \tag{1}$$

where f_n is the given state transition function for period n.

Then,

$$S_n(\phi^{n-1}) = \{f_{n-1}(s_{n-1}, y_{n-1}; \phi^{n-1}) : s_{n-1} \in S_{n-1}(\phi^{n-2}), \\ y_{n-1} \in Y_{n-1}(s_{n-1}; \phi^{n-1})\}$$

represents all the *feasible states* in period n under truncated forecast ϕ^{n-1}. Finally, s_1 denotes the initial state at the beginning of period 1.

Note that implicit from the notation provided above, we assume the following.

Remark 2.1. $S_n(\phi^{n-1})$ only depends on the data parameters for the first $n-1$ periods from a forecast ϕ. Therefore, for any forecast ϕ, if state s_n is feasible in period n (i.e., $s_n \in S_n(\phi^{n-1})$), then $s_n \in S_n(\theta^{n-1})$ for all $\theta \in \Phi(\phi^{n-1})$.

Remark 2.2. Feasible decision space $Y_n(s_n; \phi^n)$ is nonempty for any $\phi^n \in \Phi^n$, $s_n \in S_n(\phi^{n-1})$, and all n, so that any finite horizon feasible state or decision sequence can be feasibly extended arbitrarily far beyond period n, for any forecast in $\Phi(\phi^n)$.