Applied Numerical Mathematics 103 (2016) 27–47

Contents lists available at ScienceDirect

Applied Numerical Mathematics

www.elsevier.com/locate/apnum

Piecewise constant policy approximations to Hamilton–Jacobi–Bellman equations

C. Reisinger [a], P.A. Forsyth [b,*]

[a] *Mathematical Institute, Andrew Wiles Building, University of Oxford, Woodstock Road, Oxford, OX2 6GG, United Kingdom*
[b] *Cheriton School of Computer Science, University of Waterloo, Waterloo, ON N2L 3G1, Canada*

ARTICLE INFO

Article history:
Received 19 March 2015
Received in revised form 7 October 2015
Accepted 5 January 2016
Available online 8 January 2016

Keywords:
Fully nonlinear PDEs
Monotone approximation schemes
Piecewise constant policy time stepping
Viscosity solutions
Uncertain volatility model
Mean variance

ABSTRACT

An advantageous feature of piecewise constant policy timestepping for Hamilton–Jacobi–Bellman (HJB) equations is that different linear approximation schemes, and indeed different meshes, can be used for the resulting linear equations for different control parameters. Standard convergence analysis suggests that monotone (i.e., linear) interpolation must be used to transfer data between meshes. Using the equivalence to a switching system and an adaptation of the usual arguments based on consistency, stability and monotonicity, we show that if limited, potentially higher order interpolation is used for the mesh transfer, convergence is guaranteed. We provide numerical tests for the mean-variance optimal investment problem and the uncertain volatility option pricing model, and compare the results to published test cases.

1. Introduction

This article is concerned with the numerical approximation of fully nonlinear second order partial differential equations of the form

$$0 = F(\mathbf{x}, V, DV, D^2V) = \begin{cases} V_\tau - \sup_{q \in Q} L_q V, & \mathbf{x} \in \mathbb{R}^d \times (0, T], \\ V(\mathbf{x}) - \mathcal{G}(\mathbf{x}), & \mathbf{x} \in \mathbb{R}^d \times \{0\}, \end{cases} \tag{1.1}$$

where $\mathbf{x} = (S, \tau)$ contains both 'spatial' coordinates $S \in \mathbb{R}^d$ and *backwards time* τ. For fixed q in a control set Q, L_q is the linear differential operator

$$L_q V = \mathrm{tr}\left(\sigma_q \sigma_q^T D^2 V\right) + \mu_q^T DV - r_q V + f_q, \tag{1.2}$$

where $\sigma_q \in \mathbb{R}^{d \times d}$, $\mu_q \in \mathbb{R}^d$ and $r_q, f_q \in \mathbb{R}$ are functions of the control as well as possibly of $\mathbf{x}$. An initial (in backwards time) condition $V(0, \cdot) = \mathcal{G}(\cdot)$ is also specified.

These equations arise naturally from stochastic optimization problems. By *dynamic programming*, the value function satisfies an HJB equation of the form (1.1). Since dynamic programming works backwards in time from a terminal time T to today $t = 0$, it is conventional to write PDE (1.1) in terms of backwards time $\tau = T - t$, with T being the terminal time, and t being forward time.

* Corresponding author.
E-mail addresses: reisinge@maths.ox.ac.uk (C. Reisinger), paforsyt@uwaterloo.ca (P.A. Forsyth).

http://dx.doi.org/10.1016/j.apnum.2016.01.001
0168-9274/

Many examples of equations of the type (1.1) are found in financial mathematics, including the following: optimal investment problems [32]; transaction cost problems [17]; optimal trade execution problems [1]; values of American options [25]; models for financial derivatives under uncertain volatilities [2,30]; utility indifference pricing of financial derivatives [15]. More recently, enhanced oversight of the financial system has resulted in reporting requirements which include Credit Value Adjustment (CVA) and Funding Value Adjustment (FVA), which lead to nonlinear control problems of the form (1.1) [12,31,13].

If the solution has sufficient regularity, specifically for Cordes coefficients, it has recently been demonstrated that higher order discontinuous Galerkin solutions are possible [36]. Generally, however, these problems have solutions only in the viscosity sense of [16].

A general framework for the convergence analysis of discretization schemes for strongly nonlinear degenerate elliptic equations of type (1.1) is introduced in [7], and has since been refined to give error bounds and convergence orders, see, e.g., [4–6]. The key requirements that ensure convergence are consistency, stability and monotonicity of the discretization.

The standard approach to solve (1.1) by finite difference schemes is to "discretize, then optimize", i.e., to discretize the derivatives in (1.2) and to solve the resulting finite-dimensional control problem. The nonlinear discretized equations are then often solved using variants of policy iteration [20], also known as Howard's algorithm and equivalent to Newton's iteration under common conditions [9].

At each step of policy iteration, it is necessary to find the globally optimal policy (control) at each computational node. The PDE coefficients may be sufficiently complicated functions of the control variable q such that the global optimum cannot be found either analytically or by standard optimization algorithms. Then, often the only way to guarantee convergence of the algorithm is to discretize the admissible control set and determine the optimal control at each node by exhaustive search, i.e., Q is approximated by finite subset $Q_H = \{q_1, \dots q_J\} \subset Q$. This step is the most computationally time intensive part of the entire algorithm. Convergence to the exact solution is obtained by refining Q_H.

Of course, in many practical problems, the admissible set is known to be of *bang–bang* type, i.e., the optimal controls are a finite subset of the admissible set. Then the true admissible set is already a discrete set of the form Q_H.

In both cases, if we use backward Euler timestepping, an approximation to V^{n+1} at time τ^{n+1} is obtained from

$$\frac{V^{n+1} - V^n}{\Delta\tau} - \max_{q_j \in Q_H} L^h_{q_j} V^{n+1} = 0, \tag{1.3}$$

where we have a spatial discretization $L^h_{q_j}$, with h a mesh size and $\Delta\tau$ the timestep.

1.1. Objectives

It is our experience that many industrial practitioners find it difficult and time consuming to implement a solution of equation (1.3). As pointed out in [34], many plausible discretization schemes for HJB equations can generate incorrect solutions. Ensuring that the discrete equations are monotone, especially if accurate *central differencing as much as possible* schemes are used, is non-trivial [38]. Policy iteration is known to converge when the underlying discretization operator for a fixed control is monotone (i.e., an M-matrix) [9]. Seemingly innocent approximations may violate the M-matrix condition, and cause the policy iteration to fail to converge.

A convergent iterative scheme for a finite element approximation with quasi-optimal convergence rate to the solution of a strictly elliptic switching system is proposed and analyzed in [10]. Here, we are concerned with parabolic equations and exploit the fact that approximations of the continuous-time control processes by those piecewise constant in time and attaining only a discrete set of values, lead to accurate approximations of the value function.

A technique which seems to be not commonly used (at least in the finance community) is based on piecewise constant policy time stepping (PCPT) [28,6]. In this method, given a discrete control set $Q_H = \{q_1, \dots q_J\}$, J independent PDEs are solved at each timestep. Each of the J PDEs has a constant control q_j. At the end of the timestep, the maximum value at each computational node is determined, and this value is the initial value for all J PDEs at the next step.

Convergence of an approximation in the timestep has been analyzed in [27] using purely probabilistic techniques, which shows that under mild regularity assumptions a convergence order of $1/6$ in the timestep can be proven. In this and other works [26,28], applications to fully discrete schemes are given and their convergence is deduced. These estimates seem somewhat pessimistic, in that we typically observe (experimentally) first order convergence.

Note that this technique has the following advantages:

- No policy iteration is required.
- Each of the J PDEs has a constant policy, and hence it is straightforward to guarantee a monotone, unconditionally stable discretization.
- Since the PCPT algorithm reduces the solution of a nonlinear HJB equation to the solution of a sequence of linear PDEs (at each timestep), followed by a simple max or min operation, it is straightforward to extend existing (linear) PDE software to handle the HJB case.
- Each of the J PDEs can be advanced in time independently. Hence this algorithm is an ideal candidate for efficient parallel implementation.