Error estimation and adaptive discretization for the discrete stochastic Hamilton–Jacobi–Bellman equation

Lars Grüne
Mathematisches Institut
Universität Bayreuth
95440 Bayreuth, Germany
lars.gruene@uni-bayreuth.de

March 4, 2003

Abstract: Generalizing an idea from deterministic optimal control, we construct a posteriori error estimates for the spatial discretization error of the stochastic dynamic programming method based on a discrete Hamilton–Jacobi–Bellman equation. These error estimates are shown to be efficient and reliable, furthermore, a priori bounds on the estimates depending on the regularity of the approximate solution are derived. Based on these error estimates we propose an adaptive space discretization scheme whose performance is illustrated by two numerical examples.

AMS Classification: 93E20, 65N50, 49L20, 49M25, 65N15

Keywords: optimal control, dynamic programming, numerical solution, adaptive discretization

1 Introduction

The dynamic programming method is a popular and reliable technique for the numerical solution of optimal control problems. In contrast to trajectorywise approaches like direct discretizations or methods based on Pontryagin's maximum principle it gives the full global information about the optimal value function which in turn gives access to the optimal control law in feedback form. This global nature together with the robust convergence behavior makes this very classical method appealing for the solution of optimal control problems in many mathematical application areas, like for instance in technical applications (see, e.g., [16, 28]) or in mathematical economy (see, e.g., [23, 24, 34] or [30] and the references therein).

On the other hand, the global nature of this method already implies the huge computational cost in higher dimensions. Typically, the numerical cost grows exponentially in the space dimension of the problem, a phenomenon which is known as the "curse of dimensionality". In recent years, several attempts have been made to improve the efficiency of dynamic programming algorithms in order to increase the applicability of this method to

a broader class of problems. Techniques like higher order approximations [12, 14, 29, 37], randomization [31] or adaptive space discretization methods have been proposed.

In this paper we focus on this last technique, i.e., on adaptive space discretization. Our main reason for choosing this method is the fact that the alternative techniques rely on the smoothness of the optimal value functions, which are the objects to be approximated by this algorithm. However, such optimal value functions often fail to be smooth. For continuous time problems, this was a severe theoretical problem until the early 80's of the last century, when the method of viscosity solutions was introduced by Crandall and Lions [11, 26]. Numerically, higher order approximations usually lose their efficiency when applied to non–smooth problems, while adaptive methods still work sufficiently well. Of course, for a number of optimal control problems, especially stochastic problems in mathematical economy, smoothness of the solutions can be rigorously proved, in which case high order approximations are typically advantageous, but in general smoothness cannot be expected a priori. Furthermore, Example 7.1, below, shows that also in the smooth case an adaptive method may perform very well. Finally, our proposed technique is by no means restricted to low–order approximations: the basic idea as outlined in Section 3 immediately carries over to arbitrary approximation methods, though, of course, the subsequent analysis in Sections 4 and 5 heavily relies on the type of approximation.

In the literature, several adaptive methods for dynamic programming have been proposed like, e.g., [37] for stochastic problems, where a fixed grid is used but only an adaptively chosen subset of the grid nodes enters into the solution. Here we do not want to use a basic fixed grid (whose size is a priori limited due to memory restrictions) but construct a truly adaptive discretization using a suitable control mechanism. Several of such mechanisms are described in [28] in a deterministic setting, however, all of them use heuristic arguments in one or the other way, and though the numerical experiments reported in this reference show good performance they do not allow the derivation of rigorous mathematical error bounds. In our approach we follow a methodology borrowed from the numerical analysis of PDEs, in the sense that we define efficient and reliable a posteriori error estimates which indicate where the spatial grid should be locally refined and which allow mathematically justified error bounds for the resulting approximations. For deterministic optimal control problems this was done in [18] and this method has turned out to be very efficient in a number of applications, for instance in stability analysis and stabilization of controlled and perturbed systems, see [7, 10, 19], and in the global dynamical analysis of deterministic economic optimal control problems, see [23, 24], where non–smoothness of the optimal value function plays a crucial role.

The main contribution of the present paper is twofold: first, we carry over the technique and results from [18] to the stochastic case. As in the deterministic case we can expect this method to be particularly efficient if the approximated solution is non–smooth. Second, since smoothness of the optimal value function can often be observed in the stochastic context, we investigate theoretically and numerically how our method behaves in the smooth case, providing a posteriori bounds for approximation of the first derivative and asymptotic estimates for the error estimates for vanishing element size.

This paper is organized as follows. In Section 2 we describe the problem, indicate its relation to optimal control problems, describe our basic numerical approximation method and give some preliminary results. In Section 3 we define our a posteriori error estimates

and show that they are efficient and reliable in the sense of numerical approximations to PDEs. In Section 4 we show that the error estimates also allow to derive a bound on the numerical error for the approximation of derivatives. In Section 5 we investigate the asymptotic behavior of our error estimates with respect to the size of the corresponding grid elements. Finally, we turn to numerical examples. We first discuss some implementational aspects including the precise formulation of the adaptive refinement iteration in Section 6 and then illustrate the performance of the method by two examples in Section 7.

2 Problem setup and preliminaries

We consider the discrete stochastic Hamilton–Jacobi–Bellman (or dynamic programming) equation

$$V(x) = \max_{u\in U} \mathbb{E}\{g(x,u,z) + \beta(x)V(\varphi(x,u,z))\}. \tag{2.1}$$

Here $x \in \Omega \subset \mathbb{R}^n$, $U \subset \mathbb{R}^m$, Ω and U are compact sets and z is a random variable with values in $\mathbb{R}^p$. The mappings $\varphi : \Omega \times U \times \mathbb{R}^p \to \mathbb{R}^n$ and $g : \Omega \times U \times \mathbb{R}^p \to \mathbb{R}$ are supposed to be continuous and we assume the existence of constants $L_g > 0$ and $L_\varphi > 0$ such that the inequalities

$$\sup_{u\in U} \mathbb{E}\{|g(x_1,u,z) - g(x_2,u,z)|\} \le L_g\|x_1 - x_2\| \tag{2.2}$$

and

$$\sup_{u\in U} \mathbb{E}\{|W(\varphi(x_1,u,z)) - W(\varphi(x_1,u,z))|\} \le L_\varphi L_W\|x_1 - x_2\| \tag{2.3}$$

hold for any two points x_1, $x_2 \in \Omega$ and any function $W \in C(\Omega,\mathbb{R})$ which is Lipschitz with constant L_W. Furthermore, we assume that either $\varphi(x,u,z) \in \Omega$ almost surely for all $x \in \Omega$ and all $u \in U$, or that suitable boundary values $V(x)$ for $x \notin \Omega$ are specified, such that the right hand side of (2.1) is well defined for all $x \in \Omega$. The value $\beta(x)$ is the (possibly state dependent) discount factor which we assume to be Lipschitz with constant L_β and we assume that there exists $\beta_0 \in (0,1)$ such that $\beta(x) \in (0,\beta_0)$ holds for all $x \in \Omega$.

There are two interpretations of Equation (2.1). First, (2.1) characterizes the optimal value function of the discrete time stochastic optimal control problem

$$V(x) = \max_{u\in U^\infty} \mathbb{E}\left\{\sum_{i=0}^{\infty}\left(\prod_{k=0}^{i}\beta(x_k)\right) g(x_i,u_i,z_i)\right\},$$

where U^∞ denotes the sequences $(u_i)_{i\in\mathbb{N}}$ with $u_i \in U$, the z_i are independent and equidistributed random variables with values in $\mathbb{R}^p$ and x_i solves

$$x_0 = x, \qquad x_{i+1} = \varphi(x_i,u_i,z_i).$$

Optimal control problems of this type occur in many mathematical applications. For comprehensive information on these problems and on the dynamic programming method see, e.g., the monograph [3].

Secondly, (2.1) appears as the semi–discretization of the Hamilton–Jacobi–Bellman equation of continuous time stochastic optimal control problems. This type of semi–discrete

approximation was introduced in [27] and extended to viscosity solutions in [4], a detailed convergence analysis is given in [2]. See Example 7.2, below, for an application. In this context, Equation (2.1) is a discrete time approximation to the continuous time Hamilton–Jacobi–Bellman equation which motivates the name "discrete Hamilton–Jacobi–Bellman equation". For more information on viscosity solutions of Hamilton–Jacobi equations and stochastic optimal control we refer to [15].

Associated to (2.1) we define the dynamic programming operator

$$T : C(\Omega, \mathbb{R}) \to C(\Omega, \mathbb{R})$$

given by

$$T(W)(x) := \max_{u \in U} \mathbb{E}\{g(x, u, z) + \beta(x) W(\varphi(x, u, z))\}. \tag{2.4}$$

The solution V of (2.1) is then the unique fixed point of (2.4), i.e.,

$$T(V) = V. \tag{2.5}$$

For the numerical solution of (2.5) we use a discretization method that goes back to Falcone [13] in the deterministic case (see also [1, Appendix A]) and was applied to stochastic problems in Santos and Vigo–Aguiar [33] (see also [34]). Here we use unstructured cuboidal grids like, e.g., in [22]: We assume that $\Omega \subset \mathbb{R}^n$ is a cuboid and consider a grid Γ covering Ω with cuboid elements Q_l and nodes x_j and the space of continuous and piecewise multilinear functions

$$\mathcal{W}_\Gamma := \{W \in C(\Omega, \mathbb{R}) \,|\, W(x + \alpha e_j) \text{ is linear in } \alpha \text{ on each } Q_l \text{ for each } j = 1, \ldots, n\}$$

where the e_j, $j = 1, \ldots, n$ denote the standard basis vectors of the $\mathbb{R}^n$. If the grid Γ is locally refined, then there exist nodes which do not belong to each adjacent element Q_j. These nodes are called *hanging* or *non–conforming* nodes and in order to ensure continuity the values of W in these hanging nodes have to be computed by interpolation along the corresponding element edges. All other nodes are called *conforming*, in these nodes a function $W \in \mathcal{W}_\Gamma$ can assume arbitrary values. With $\pi_\Gamma : C(\Omega, \mathbb{R}) \to \mathcal{W}_\Gamma$ we denote a projection of an arbitrary continuous function to $\mathcal{W}_\Gamma$, i.e.,

$$\pi_\Gamma(W)(x_j) = W(x_j) \text{ for all conforming nodes } x_j \text{ of the grid } \Gamma.$$

For locally refined grids in dimensions ≥ 3 the projection π_Γ might not be unique because there might be several possible choices for the values in the hanging nodes; in order to avoid technicalities on this level we simply fix one possible projection for each grid Γ. We define the minimal and maximal diameter $\underline{k}_l$ and $\overline{k}_l$ of an element Q_l with nodes $x_{j_1}, \ldots, x_{j_d}$ as

$$\underline{k}_l := \min_{1 \leq i < k \leq d} \|x_{j_i} - x_{j_k}\|, \qquad \overline{k}_l := \max_{x,y \in Q_l} \|x - y\|.$$

In this paper we use an arbitrary norm $\|\cdot\|$ on $\mathbb{R}^n$, for the analysis of certain properties of functions in $\mathcal{W}_\Gamma$, however, special norms like the ∞–norm or the 1–norm turn out to be advantageous, see, e.g., [20].

We now define the discrete dynamic programming operator by

$$T_\Gamma : C(\Omega, R) \to \mathcal{W}_\Gamma, \quad T_\Gamma = \pi_\Gamma \circ T \tag{2.6}$$

with T from (2.4). Then the discrete fixed point equation

$$T_\Gamma(V_\Gamma) = V_\Gamma. \tag{2.7}$$

has a unique solution $V_\Gamma \in \mathcal{W}_\Gamma$ which converges to V if the size of the elements Q_l tends to zero. The convergence is linear if V is Lipschitz on Ω, see [13] or [1, Appendix A] and quadratic if V is C^2, see [33, Theorem 3.5].

There are several reason for restricting ourselves to multilinear, i.e., first order approximation. The main reason is that we want our algorithm to be reliable even if the exact solution V to (2.1) is non–smooth. In general one cannot expect V to (2.1) to be smooth and the approximation of non–smooth functions by smooth approximations can lead to severe problems like spurious oscillations and non–convergence. Another reason is the relative ease of both implementation and analysis of such methods. For example, for this scheme it will turn out to be possible to obtain rigorous a priori bounds on the a posteriori error estimates depending on the element size, which are useful also in the smooth case, see the results in Section 5, below. Also, the analysis of regularity properties for the discrete time approximation is facilitated, see, e.g., [18, Theorem 2.9]. Finally, the error estimation and adaptive discretization method we propose in this paper is by no means restricted to first order approximations; the technique and the basic results can be straightforwardly carried over to higher–order approximations methods. For the description of such methods we refer, e.g., to [14] (using high–order interpolation), [12, 37] (using spline–interpolation) or [29] (using regression techniques).

We end this section with a lemma giving some basic results for the operators (2.4) and (2.6) which will be useful in what follows. Here and in the remainder of the paper we use the norm $\|W\|_\infty := \max_{x\in\Omega} |W(x)|$ for $W \in C(\Omega, \mathbb{R})$.

Lemma 2.1 Let $S = T$ or $S = T_\Gamma$. Then for all functions W, W_1, $W_2 \in C(\Omega, \mathbb{R})$ and all points x_1, $x_2 \in \Omega$ the inequalities

$$\begin{aligned} \text{(i)} \quad |T(W)(x_1) - T(W)(x_2)| \;&\leq\; (L_g + L_\beta \|W\|_\infty)\|x_1 - x_2\| \\ &\quad + \beta_0 \sup_{u\in U} \mathbb{E}\{\|W(\varphi(x_1, u, z)) - W(\varphi(x_2, u, z))\|\} \end{aligned}$$

and

$$\text{(ii)} \quad \|S(W_1) - S(W_2)\|_\infty \leq \beta_0 \|W_1 - W_2\|_\infty$$

hold.

Proof: (i) Assume without loss of generality that $T(W)(x_1) \geq T(W)(x_2)$. Then we obtain

$$\begin{aligned} &|T(W)(x_1) - T(W)(x_2)| \\ &\quad = \; T(W)(x_1) - T(W)(x_2) \\ &\quad = \; \max_{u\in U} \mathbb{E}\{g(x_1, u, z) + \beta(x_1) W(\varphi(x_1, u, z))\} \\ &\qquad\quad - \max_{u\in U} \mathbb{E}\{g(x_2, u, z) + \beta(x_2) W(\varphi(x_2, u, z))\} \end{aligned}$$

$$\begin{aligned}
&\leq \max_{u\in U}\Big\{\mathbb{E}\{g(x_1,u,z)-g(x_1,u,z)\}+\mathbb{E}\{\beta(x_1)W(\varphi(x_1,u,z))-\beta(x_2)W(\varphi(x_2,u,z))\}\Big\}\\
&\leq L_g\|x_1-x_2\|+L_\beta\|W\|_\infty\|x_1-x_2\|+\mathbb{E}\{\beta(x_2)(W(\varphi(x_1,u,z))-W(\varphi(x_2,u,z)))\}\\
&\leq (L_g+L_\beta\|W\|_\infty)\|x_1-x_2\|+\beta_0\sup_{u\in U}\mathbb{E}\{\|W(\varphi(x_1,u,z))-W(\varphi(x_2,u,z))\|\}
\end{aligned}$$

which shows (i).

(ii) We first prove (ii) for $S=T$. Consider an arbitrary $x\in\Omega$ and assume without loss of generality $T(W_1)(x)\geq T(W_2)(x)$. Then

$$\begin{aligned}
&|T(W_1)(x)-T(W_2)(x)|\\
&\quad= T(W_1)(x)-T(W_2)(x)\\
&\quad= \max_{u\in U}\mathbb{E}\{g(x,u,z)+\beta(x)W_1(\varphi(x,u,z))\}-\max_{u\in U}\mathbb{E}\{g(x,u,z)+\beta(x)W_2(\varphi(x,u,z))\}\\
&\quad\leq \max_{u\in U}\Big\{\mathbb{E}\{g(x,u,z)-g(x,u,z)\}+\mathbb{E}\{\beta(x)W_1(\varphi(x,u,z))-\beta(x)W_2(\varphi(x,u,z))\}\Big\}\\
&\quad\leq \beta_0\|W_1-W_2\|_\infty
\end{aligned}$$

which shows (ii) for $S=T$ since $x\in\Omega$ was arbitrary.

For $S=T_\Gamma$ inequality (ii) follows from the multilinearity of the functions in $\mathcal{W}_\Gamma$ which implies

$$|\pi_\Gamma(W_1)(x)-\pi_\Gamma(W_2)(x)|=\left|\sum_{j=1}^{d}\mu_j(W_1(x_j)-W_2(x_j))\right|\leq\|W_1-W_2\|_\infty$$

because the coefficients μ_j are non–negative and sum up to one. Hence (ii) for $S=T_\Gamma$ follows from (ii) for $S=T$. □

Remark 2.2 If the projection π_Γ preserves the Lipschitz constants of Lipschitz continuous functions $W\in C(\Omega,\mathbb{R})$ then (i) also holds for T_Γ. Note, however, that this Lipschitz preserving property in general only holds for suitable norms on $\mathbb{R}^n$ and under suitable regularity properties on the grid Γ. □

3 A posteriori error estimates

In this section we develop suitable a posteriori error estimates for measuring the discretizations error, i.e., the difference between the functions V_Γ and V. The idea for these error estimates was developed for deterministic equations of type (2.1) in [18], some implementational details were discussed in [22]. The concept was subsequently extended to evolutive deterministic equations in [20] and applied to the shape–from–shading problem in image analysis in [32] (see also [5]).

The basic idea of the error estimates simply relies in evaluating the residual of the operator T applied to V_Γ, as made precise in the following definition. Here for any subset $B\subset\Omega$ and any function $W\in C(\Omega,\mathbb{R})$ we use

$$\|W\|_{\infty,B}:=\max_{x\in B}|W|.$$

Definition 3.1 Consider a grid Γ and the corresponding solution $V_\Gamma \in \mathcal{W}_\Gamma$ of (2.7).

(i) We define the *a posteriori error estimate* η as a continuous function $\eta \in C(\Omega, \mathbb{R})$ by

$$\eta(x) := |T(V_\Gamma)(x) - V_\Gamma(x)|.$$

(ii) For any element Q_l of the grid Γ we define the *elementwise error estimate*

$$\eta_l := \|\eta\|_{\infty, Q_l}$$

(iii) We define the *global error estimate* $\eta_{\max}$ by

$$\eta_{\max} := \max_l \eta_l = \|\eta\|_\infty.$$

□

The following theorem shows the important properties of $\eta_{\max}$.

Theorem 3.2 Consider the solutions V of (2.5) and V_Γ of (2.7). Then the inequality

$$\frac{\eta_{\max}}{1+\beta_0} \le \|V - V_\Gamma\|_\infty \le \frac{\eta_{\max}}{1-\beta_0}$$

holds for the global error estimate $\eta_{\max}$ from Definition 3.1 (iii).

Proof: From Lemma 2.1(ii) for any two functions W_1, $W_2 \in C(\Omega, \mathbb{R})$ and any $x \in \Omega$ we obtain

$$|T(W_1)(x) - T(W_2)(x)| \le \|T(W_1) - T(W_2)\|_\infty \le \beta_0 \|W_1 - W_2\|_\infty. \tag{3.1}$$

Since $T(V) = V$ it follows for all $x \in \Omega$ that

$$\begin{aligned} |V_\Gamma(x) - T(V_\Gamma)(x)| &= |V_\Gamma(x) - V(x) + T(V)(x) - T(V_\Gamma)(x)| \\ &\le |V_\Gamma(x) - V(x)| + |T(V)(x) - T(V_\Gamma)(x)| \\ &\le (1+\beta_0)\|V_\Gamma - V\|_\infty \end{aligned}$$

where the last inequality follows by (3.1). Thus

$$\eta_{\max} = \|V_\Gamma - T(V_\Gamma)\|_\infty \le (1+\beta_0)\|V_\Gamma - V\|_\infty$$

which shows the first inequality.

Conversely for all $x \in \Omega$ we have

$$\begin{aligned} |V(x) - V_\Gamma(x)| &= |T(V_\Gamma)(x) - V_\Gamma(x) + T(V)(x) - T(V_\Gamma)(x)| \\ &\le |T(V_\Gamma)(x) - V_\Gamma(x)| + |T(V)(x) - T(V_\Gamma)(x)| \\ &\le \eta_{\max} + \beta_0 \|V - V_\Gamma\|_\infty \end{aligned}$$

using again (3.1) for the last inequality. This implies $\|V - V_\Gamma\|_\infty \le \eta_{\max} + \beta_0 \|V - V_\Gamma\|_\infty$ and thus the second inequality. □

Remark 3.3 An error estimate satisfying these two inequalities is called *reliable* and *efficient*. The upper bound shows that small error estimates guarantee small error (they are *reliable*), the lower bound shows that the error estimates do not overestimate the real error (they are *efficient*). Note that here only the upper bound really depends on the problem data via the constant $1-\beta_0$ in the denominator, which might become small. For very degenerate problems, i.e., with discount factor β_0 close to 1 this might render the upper bound practically useless. In contrast to this, the lower bound is bounded from below independent of β_0 by $\eta_{\max}/2$, which is the reason why the adaptive gridding strategy shows satisfactory performance also in the degenerate case $\beta_0 \approx 1$, see Example 7.2, below. □

Remark 3.4 The proof of the first part of Theorem 3.2 in fact shows the "semi–local" inequality

$$\frac{\eta(x)}{1+\beta_0} \le \|V - V_\Gamma\|_{\infty,B(x)},$$

where $B(x) := \{x\} \cup \varphi(x,U,Z)$ with Z denoting the set of all possible values of the random variable z. Thus, a large error estimate $\eta(x)$ indicates that the real error is large either at the point x itself or at some point in the image of φ. In the second case, refining the element Q_l containing x will not immediately reduce the global error. However, the fact that small elements Q_l carry small error estimates η_l — which will be shown in Section 5 below — ensures that after a number of grid adaptation iterations the elements corresponding to large real errors will eventually be identified and also refined.

The upper inequality can not be localized in this way; in fact the error in a point x depends on all the error estimates in the elements Q_l that can be reached from this point, where the reachability can be formally expressed via the finite state Markov chain associated to the space discretization. This type of error propagation (or error influence) analysis can also be used for the development of adaptive grid refinement strategies, see [28]. □

4 Estimates for the first derivative

As already mentioned, for stochastic optimal control problems it is not too optimistic to expect that the exact solution V is smooth. This leads to the question, whether the error estimates introduced in the last section can also be used in order to give bounds on the approximation error of the derivatives. Note that since our approximation V_Γ is piecewise multi–linear, the derivative $D_x V_\Gamma(x)$ will in general not be continuous, however it is still well defined on the interior $\operatorname{int} Q_l$ for each element; furthermore the directional derivative $D_x V(x)z$ exists and is well defined for each point $x \in \operatorname{int}\Omega$ and each vector $z \in \mathbb{R}^n$, hence these values can be used as a numerical approximation for the derivative $D_x V(x)$, provided that V is smooth.

In this section we provide two results on error bounds for the difference between $D_x V$ and $D_x V_\Gamma$ based on the error estimate $\eta_{\max}$. Note that we cannot expect that $\eta_{\max}$ gives a lower bound for this difference, since if, e.g., V_Γ happens to be simply $V + c$ for some constant $c \in \mathbb{R}$, then the error in the derivatives is 0 but — according to the upper bound in Theorem 3.2 — the inequality $\eta_{\max} \ge (1-\beta_0)c > 0$ must hold. What we can still obtain is an upper bound for this difference.

Proposition 4.1 Assume that V is two times continuously differentiable. Then for each element Q_l with minimal diameter $\underline{k}_l$, each $x \in Q_l$ and each coordinate direction e_i there exists $\xi = x + \lambda e_i \in Q_l$ such that the inequality

$$|D_x V(\xi)e_i - D_x V_\Gamma(\xi)e_i| \leq 4\frac{\eta_{\max}}{\underline{k}_l(1-\beta_0)}$$

holds.

Proof: Let $g(\lambda) := V(x + \lambda e_i)$ and consider the (unique) real values $\lambda_1 < \lambda_2$ with the property that $x_1 := x + \lambda_1 e_i \in \partial Q_l$ and $x_2 := x + \lambda_2 e_i \in \partial Q_l$. From the mean value theorem applied to g we obtain the existence of $\lambda \in (\lambda_1, \lambda_2)$ such that

$$\frac{V(x_2) - V(x_1)}{\lambda_2 - \lambda_1} = \frac{g(\lambda_2) - g(\lambda_1)}{\lambda_2 - \lambda_1} = g'(\lambda) = D_x V(\xi)e_i$$

where $\xi = x + \lambda e_i$. On the other hand, since V_Γ is linear in each coordinate direction we obtain

$$D_x V_\Gamma(\xi)e_i = \frac{V_\Gamma(x_2) - V_\Gamma(x_1)}{\lambda_2 - \lambda_1}.$$

Thus using Theorem 3.2 and the inequality $\lambda_2 - \lambda_1 \geq \underline{k}_l$ we obtain

$$|D_x V(\xi)e_i - D_x V_\Gamma(\xi)e_i| \leq \frac{2\|V - V_\Gamma\|_\infty}{\lambda_2 - \lambda_1} \leq \frac{4\eta_{\max}}{\underline{k}_l(1-\beta_0)}$$

which shows the claim. □

In the following theorem we show how this result implies an error bound for arbitrary points. As above, for subsets $B \subseteq \Omega$ we use the notations $\|D_x V\|_{\infty,B} := \max_{x\in B} \|D_x V(x)\|$ and $\|D_{xx} V\|_{\infty,B} := \max_{x\in B} \|D_{xx} V(x)\|$ where in the second case $D_{xx}V$ denotes the second derivative of V and we denote both the $\mathbb{R}^n$–norm and the induced operator norm by $\|\cdot\|$.

Theorem 4.2 Assume that V is two times continuously differentiable. Then for each element Q_l the estimate

$$\|D_x V - D_x V_\Gamma\|_{\infty,Q_l} \leq 4\frac{\eta_{\max}}{\underline{k}_l(1-\beta_0)} + \overline{k}_l \|D_{xx} V\|_{\infty,Q_l}$$

holds.

Proof: This follows immediately from Proposition 4.1, since by Taylor's theorem $D_x V(x)e_i$ is Lipschitz in x on Q_l with constant $\|D_{xx} V\|_{\infty,Q_l}$. □

5 Bounds on the error estimates

An adaptive grid refinement strategy relies on the idea that local refinement will eventually reduce the local and consequently the global approximation error. While this property is what one would expect intuitively, it is by no means clear that it can be rigorously proved; furthermore it is helpful to know more precisely what order of convergence of η_l to 0 can

be expected depending on the size of the element and which properties are needed in order to ensure this behavior. In this section we provide a couple of estimates on the value of η_l depending on the size of the corresponding element and on certain regularity properties of the discrete approximation. Note that regularity properties of the discrete approximations V_Γ can in principle be proved rigorously, see, e.g., [18, Theorem 2.9] or [20, Proposition 9.1]. Here, in order to keep the presentation simple, we simply impose suitable and reasonable regularity properties.

We would like to emphasize that the error bounds in Sections 3 and 4 do not depend on the results in this section. Rather, the following results indicate the performance that can be expected using an adaptive local refinement strategy. We start with the first estimate which simply assumes Lipschitz continuity of V_Γ.

Theorem 5.1 Assume that V_Γ satisfies $\|V_\Gamma\|_\infty \le M_\Gamma$ and is Lipschitz with constant L_Γ. Then the error estimates $\eta(x)$ are Lipschitz with the constant $L_\eta = L_g + L_\beta M_\Gamma + (1+L_\varphi)L_\Gamma$. In particular, the estimate

$$\eta_l \le L_\eta \overline{k}_l$$

holds.

Proof: Consider x_1, $x_2 \in \Omega$. Assume without loss of generality that $\eta(x_1) \ge \eta(x_2)$. Then from the definition of η, Lemma 2.1(i) and (2.3) we obtain

$$\begin{aligned}
|\eta(x_1) - \eta(x_2)| &= \eta(x_1) - \eta(x_2) \\
&= V_\Gamma(x_1) - T(V_\Gamma)(x_1) - V_\Gamma(x_2) + T(V_\Gamma)(x_2) \\
&\le |V_\Gamma(x_1) - V_\Gamma(x_2)| + |T(V_\Gamma)(x_1) - T(V_\Gamma)(x_2)| \\
&\le L_\Gamma \|x_1 - x_2\| + (L_g + L_\beta \|V_\Gamma\|_\infty)\|x_1 - x_2\| \\
&\quad + \sup_{u\in U} \mathbb{E}\{|V_\Gamma(\varphi(x_1,u,z)) - V_\Gamma(\varphi(x_2,u,z))|\} \\
&\le (L_\Gamma + L_g + L_\beta M_\Gamma + L_\Gamma L_\varphi)\|x_1 - x_2\|
\end{aligned}$$

which shows the Lipschitz property of $\eta(x)$.

The second assertion follows since $\eta(x_k) = 0$ holds for all conforming nodes of Γ. □

Note that the Lipschitz property assumed here could be relaxed if desired, e.g., to Hölder continuity, cf. [18, Corollary 2.10].

Remark 5.2 Apart from showing that the local error indeed tends to 0 as $\overline{k}_l \to 0$, Theorem 5.1 also justifies the numerical evaluation of $\eta(x)$ on an element Q_l in finitely many test points $x_T \in X_T \subset Q_l$ only. The continuity of $\eta(x)$ then implies that $\tilde{\eta}_l := \max_{x_T \in X_T} \eta(x_T)$ gives an approximation for $\eta_l = \max_{x\in Q_l} \eta(x)$. □

As shown in Proposition 4.1 and Theorem 4.2, one can also obtain an error bound on the derivative provided that the exact solution V is C^2. This bound, however, is scaled by

$1/\underline{k}_l$, hence convergence to zero for vanishing element size in an adaptive gridding strategy can only be expected if $\eta_l = o(\underline{k}_l)$ holds, i.e., if we have superlinear convergence of η_l to 0.

Superlinear — more precisely quadratic — convergence for approximations on equidistant grids can be expected if V is C^2, see e.g. [34]. Therefore, it seems natural to impose appropriate "smoothness" properties of our discrete approximations V_Γ if we want to ensure quadratic convergence of the η_l to 0.

The remainder of this section is devoted to giving sufficient conditions under which one can show that the error estimates on an element Q_l are bounded by $\overline{k}_l^2$, i.e., they converge to 0 quadratically with the size of the element. If the used grids satisfy $\overline{k}_l \leq \mu \underline{k}_l$ for some $\mu > 0$, then one can indeed expect $\eta_l = O(\underline{k}_l^2)$, i.e., the error bound $\max_l \eta_l/\underline{k}_l$ for the derivatives tends to zero with $O(\underline{k}_l)$ for $\underline{k}_l \to 0$.

The first result in this direction assumes that the continuous operator T has a smoothing effect.

Theorem 5.3 Assume that $T(V_\Gamma) \in C^2(\Omega, \mathbb{R})$ with $\|D_{xx} T(V_\Gamma)\|_{\infty, Q_l} \leq C$ for some constant $C > 0$. Then the inequality

$$\eta_l \leq C \overline{k}_l$$

holds.

Proof: Let $W \in C^2(\Omega, \mathbb{R})$ be a function satisfying $\|D_{xx} W\|_{\infty, Q_l} \leq C$. Then a straightforward application of Taylor's theorem shows for each $x \in Q_l$ the inequality

$$\|\pi_\Gamma(W)(x) - W(x)\| \leq C \overline{k}_l.$$

For $W = T(V_\Gamma)$ this implies

$$\begin{aligned} |V_\Gamma(x) - T(V_\Gamma)(x)| &= |T_\Gamma(V_\Gamma)(x) - T(V_\Gamma)(x)| \\ &= |\pi_\Gamma(T(V_\Gamma))(x) - T(V_\Gamma)(x)| \leq C \overline{k}_l, \end{aligned}$$

i.e., the assertion. □

If the diffusion is suitably regular then the smoothing property assumed in this theorem might well hold. However, the value function V might also be smooth if T does not have this "immediate" smoothing effect, e.g., for systems with degenerate diffusion or purely deterministic systems, see, e.g., [34, Theorem 3.1] for a smoothness result for a family of such optimal control problems.

Thus, in the next theorem we replace the assumption on T by a suitable "discrete second differentiability" condition on the discrete approximation. This property is most conveniently expressed by the second difference quotient, which for a function $W : \mathbb{R}^n \to \mathbb{R}$ and three points x_1, x_2, y with $y = \mu x_1 + (1-\mu) x_2$ for some $\mu \in (0,1)$ is defined by

$$\Delta^2 W(x_1, x_2, y) := \frac{\mu W(x_1) - W(y) + (1-\mu) W(x_2)}{\|x_1 - x_2\|^2 \mu (1-\mu)}.$$

Note that for a C^2 function W by Taylor's theorem the inequality

$$\Delta^2 W(x_1, x_2, y) \leq \max_{\xi = \mu x_1 + (1-\mu) x_2, \mu \in [0,1]} \|D_{xx} W(\xi)\|$$

holds, hence in particular if W is C^2 then $\Delta^2 W$ is bounded on compact and convex domains. The value $\Delta^2 W$ can be seen as a measure for the "discrete curvature" of W in y along the line connecting x_1 and x_2.

Unlike the previous theorem, here it is not assumed that the approximation V_Γ has sufficient regularity. Instead, we need to assume that the solution on a suitably refined grid has the respective property.

Theorem 5.4 Consider a grid Γ, an element Q_l in Γ and a point $x \in \operatorname{int} Q_l$. Let Γ_x be the grid obtained from Γ by inserting x as a new node. Assume that the second difference quotient satisfies

$$|\Delta^2 V_{\Gamma_x}(x_1, x_2, x)| \le \frac{C}{\mu(1-\mu)} \tag{5.1}$$

for the solution V_{Γ_x} on Γ_x, some $C > 0$, some standard basis vector e_i with $i \in \{1, \dots, n\}$ and the unique points $x_1, x_2 \in \partial Q_l$ with $x = \mu x_1 + (1-\mu)x_2$ for some $\mu \in (0,1)$ and $x_2 - x_1 = \lambda e_i$ for some some $\lambda \in \mathbb{R} \setminus \{0\}$. Then the inequality

$$\eta(x) \le C \frac{1+\beta_0}{1-\beta_0} \overline{k}_l^2$$

holds.

In particular, if (5.1) holds for all $x \in \operatorname{int} Q_l$ then

$$\eta_l \le C \frac{1+\beta_0}{1-\beta_0} \overline{k}_l^2$$

holds.

Proof: Fix $x \in \operatorname{int} Q_l$ and assume that (5.1) holds.

The proof is split into three steps. First we show the inequality

$$\|T_\Gamma(V_{\Gamma_x}) - V_{\Gamma_x}\|_\infty \le C \overline{k}_l^2. \tag{5.2}$$

In order to show (5.2), observe that since Γ_x contains all the nodes of Γ, we obtain $T_\Gamma = \pi_\Gamma \circ T_{\Gamma_x}$ and thus

$$T_\Gamma(V_{\Gamma_x}) = \pi_\Gamma(T_{\Gamma_x}(V_{\Gamma_x})) = \pi_\Gamma(V_{\Gamma_x}).$$

Since Γ_x contains exactly one more conforming node x than Γ, which lies in $\operatorname{int} Q_l$ we obtain

$$\pi_\Gamma(V_{\Gamma_x})|_{\Omega \setminus \operatorname{int} Q_l} = V_{\Gamma_x}|_{\Omega \setminus \operatorname{int} Q_l}.$$

The piecewise multilinearity of V_{Γ_x} implies

$$\|\pi_\Gamma(V_{\Gamma_x}) - V_{\Gamma_x}\|_\infty = |\pi_\Gamma(V_{\Gamma_x})(x) - V_{\Gamma_x}(x)|.$$

Now consider the points $x_1, x_2 \in \partial Q_l$ and $\mu \in (0,1)$ for which (5.1) holds. Then the linearity of $\pi_\Gamma(V_{\Gamma_x})$ on Q_l in direction e_i implies

$$\pi_\Gamma(V_{\Gamma_x})(x) = \mu \pi_\Gamma(V_{\Gamma_x})(x_1) + (1-\mu)\pi_\Gamma(V_{\Gamma_x})(x_2) = \mu V_{\Gamma_x}(x_1) + (1-\mu) V_{\Gamma_x}(x_2).$$

Now the definition of Δ^2 implies

$$|\mu V_{\Gamma_x}(x_1) + (1-\mu)V_{\Gamma_x}(x_2) - V_{\Gamma_x}(x)| = \mu(1-\mu)|\Delta(x_1,x_2,x)|\,\|x_1-x_2\|^2.$$

Combining all these considerations and (5.1) we can conclude

$$\begin{aligned}\|T_\Gamma(V_{\Gamma_x}) - V_{\Gamma_x}\|_\infty &= |\pi_\Gamma(V_{\Gamma_x})(x) - V_{\Gamma_x}(x)| \\ &= |\mu V_{\Gamma_x}(x_1) + (1-\mu)V_{\Gamma_x}(x_2) - V_{\Gamma_x}(x)| \\ &= \mu(1-\mu)|\Delta(x_1,x_2,x)|\,\|x_1-x_2\|^2 \;\le\; C\overline{k}_l^2,\end{aligned}$$

which shows (5.2).

In the second step we show the inequality

$$\|V_{\Gamma_x} - V_\Gamma\|_\infty \le \frac{C}{1-\beta_0}\overline{k}_l^2. \tag{5.3}$$

The proof of (5.3) follows from

$$\begin{aligned}\|V_{\Gamma_x} - V_\Gamma\|_\infty &\le \|V_{\Gamma_x} - T_\Gamma(V_{\Gamma_x})\|_\infty + \|T_\Gamma(V_{\Gamma_x}) - V_\Gamma\|_\infty \\ &\le C\overline{k}_l^2 + \|T_\Gamma(V_{\Gamma_x}) - \underbrace{T_\Gamma(V_\Gamma)}_{=V_\Gamma}\|_\infty \;\le\; C\overline{k}_l^2 + \beta_0\|V_{\Gamma_x} - V_\Gamma\|_\infty,\end{aligned}$$

where we used (5.2) in the second last and Lemma 2.1(ii) in the last inequality. This implies

$$(1-\beta_0)\|V_{\Gamma_x} - V_\Gamma\|_\infty \le C\overline{k}_l^2.$$

and thus (5.3).

Finally, we show the assertion of the theorem. The construction of Γ_x implies

$$T(V_\Gamma)(x) = \pi_{\Gamma_x}(T(V_\Gamma))(x) = T_{\Gamma_x}(V_\Gamma)(x).$$

Thus by the definition of $\eta(x)$ we obtain

$$\begin{aligned}\eta(x) &= |T(V_\Gamma)(x) - V_\Gamma(x)| \\ &= |T_{\Gamma_x}(V_\Gamma)(x) - V_\Gamma(x)| \\ &\le |T_{\Gamma_x}(V_\Gamma)(x) - T_{\Gamma_x}(V_{\Gamma_x})(x)| + |T_{\Gamma_x}(V_{\Gamma_x})(x) - V_\Gamma(x)| \\ &= |T_{\Gamma_x}(V_\Gamma)(x) - T_{\Gamma_x}(V_{\Gamma_x})(x)| + |V_{\Gamma_x}(x) - V_\Gamma(x)| \\ &\le \beta_0|V_{\Gamma_x}(x) - V_\Gamma(x)| + |V_{\Gamma_x}(x) - V_\Gamma(x)| \;\le\; (1+\beta_0)\frac{C}{1-\beta_0}\overline{k}_l^2\end{aligned}$$

where we used Lemma 2.1(ii) in the second last and (5.3) in the last inequality. This shows the claim. □

For a single approximation V_Γ the assumption of this theorem is in general difficult to check, since it requires the knowledge of V_{Γ_x} which is not available. However, it should be noted that the assumption of this theorem holds if *each* approximation V_Γ on *any* admissible grid Γ satisfies (5.1). Thus, we can expect quadratic convergence of the error estimates if the family of discrete approximations satisfies a "discrete C^2 condition". If the approximated function V itself is C^2 then one could expect such a property, though a rigorous proof might be tedious and is beyond the scope of this paper. In any case, the second difference quotient Δ^2 can be evaluated on the nodes of each discrete approximation V_Γ which can be used in order to obtain some "heuristic evidence" for the validity of (5.1).

In numerical examples with smooth V a decreasing of $\eta_l/\underline{k}_l$ could be observed for not too fine grids; after a certain number of refinements, however, the convergence $\eta_l/\underline{k}_l \to 0$ stops (cf. the example in Section 7.1). Numerical experiments indicate that this behavior is due to the effect of additional numerical errors (introduced, e.g., in the solution of the fixed point equation $T_\Gamma(V_\Gamma) = V_\Gamma$) which affect the smoothness of V_Γ, because the higher the numerical accuracy in the solution of $T_\Gamma(V_\Gamma) = V_\Gamma$ was chosen chosen in our numerical experiments the longer the decrease of $\eta_l/\underline{k}_l$ could be observed.

We want to end this discussion by remarking once again that our piecewise multilinear approach is particularly suitable when the exact optimal value function V is *not* smooth. If smoothness of the optimal value function can be guaranteed by some a priori analysis, then higher order approximations may be more efficient. However, our basic a posteriori error estimation technique can be rather straightforwardly carried over to other approximations and we assume that the estimates proved in this section indicate the lines along which an analysis of higher–order schemes should be carried out. This might become more technically involved, but some things might also simplify, e.g., for C^2 approximations the second derivatives of the numerical solution could be computed directly, which should lead to a less technical formulation of the assumption of Theorem 5.4.

6 Implementational aspects

In this section we briefly discuss implementational aspects of our method; in particular we explain some details we have used in order to compute the numerical examples in the Section 7.

6.1 Numerical evaluation of T_Γ and solution of (2.7)

For the solution of (2.7) as well as for the computation of $\eta(x)$ we need to evaluate the operator T_Γ. More precisely, we need to evaluate

$$\max_{u\in U} \mathbb{E}\{g(x,u,z) + \beta(x)W(\varphi(x,u,z))\}.$$

for all conforming nodes $x = x_j$ of Γ and all points $x = x_T$ in which we want to compute $\eta(x_T)$.

This first includes the numerical evaluation of the expectation $\mathbb{E}$. If z is a finite random variable then this is straightforward, if z is a continuous random variable then the

corresponding integral

$$\int_Z (g(x,u,z) + \beta(x)W(\varphi(x,u,z)))f(z)dz$$

has to be computed, where f is the probability density of z which is assumed to be known. In our implementation we approximated this integral by a trapezoidal rule.

The second difficulty in the numerical evaluation of T lies in the maximization over u. In our implementation we simply used a discrete approximation of the set U with equidistant points and then maximized by comparing the corresponding values. For low to medium accurate evaluation of T_Γ this is a feasible method, for high accuracy, however, this method becomes inefficient; in this case other methods like, e.g., Brent's algorithm have been reported in the literature to give good results, see [9, 34].

For the solution of the fixed point equation (2.7) we use the Gauss–Seidel type value space iteration which is described in [18, Section 3] (under the name "increasing coordinate algorithm"), where we subsequently compute $V_{i+1} = S_\Gamma(V_i)$ with S_Γ being a Gauss–Seidel type iteration operator (including the maximization over u) obtained from T_Γ. This iteration is coupled with a policy space iteration as described in [17, 35]: Once a prescribed percentage of the maximizing u–values in the nodes remains constant from one iteration to another we fix all control values and compute the associated value function by solving a linear system of equations using the CGS or BICGSTAB method (in our examples the CGS method turned out to show more reliable convergence behavior). After convergence of this method we continue with the value space iteration using S_Γ until the control values again converge, switch to the linear solver and so on. This combined policy–value space iteration turns out to be much more efficient (often more than 90 percent faster) than the plain Gauss–Seidel value space iteration using S_Γ, which in turn is considerably faster than the Banach iteration $V_{i+1} = T_\Gamma(V_i)$.

6.2 Numerical evaluation of η_l

Clearly, in general the values $\eta_l = \max_{x\in Q_l} \eta(x)$ can not be evaluated exactly since the maximization has to be performed over infinitely many points $x \in Q_l$. Instead, we approximate η_l by

$$\tilde{\eta}_l = \max_{x_T \in X_T(Q_l)} \eta(x_T),$$

where $X_T(Q_l)$ is a set of test points. The continuity of $\eta(x)$ in x implies that $\tilde{\eta}_l$ is indeed an approximation of η_l, cf. Remark 5.2. In our numerical experiments we have used the test points indicated in Figure 6.1.

6.3 Adaptive refinement and coarsening

The adaptive grid itself was implemented on a tree data structure in the programming language C, cf. [22] for details.

The adaptive refinement follows the standard practice in numerical schemes. The iterative construction of adaptive grids works as follows:

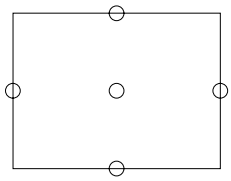

Figure 6.1: Test points $X_T(Q_l)$ for a 2d element Q_l

(0) Choose an initial grid Γ_0, set $i = 0$, fix a refinement threshold $\theta \in (0, 1)$

(1) Compute V_{Γ_i} and the (approximated) error estimates $\tilde{\eta}_l$ and $\tilde{\eta}_{\max}$. If a desired accuracy or a maximally allowed number of nodes is reached, then stop

(2) Refine all elements Q_l with $\tilde{\eta}_l \geq \theta \tilde{\eta}_{\max}$, denote the new grid by Γ_{i+1}

(3) Set $i := i + 1$ and go to (1)

Here for the solution of V_{Γ_i} for $i \geq 1$ we use the previous solution $V_{\Gamma_{i-1}}$ as the initial value for the iteration described in Section 6.1, which turns out to be very efficient.

During the adaptation routine it might happen that the error estimate causes refinements in regions which later turn out to be very regular. It is therefore advisable to include a coarsening mechanism in the above iteration. This mechanism can, e.g., be controlled by comparing the approximation V_{Γ_i} with its projection $\pi_{\widetilde{\Gamma}_i} V_{\Gamma_i}$ onto the grid $\widetilde{\Gamma}_i$ which is obtained from Γ_i by coarsening each element once. Using a specified coarsening tolerance $tol \geq 0$ one can add the following step after Step (2).

(2a) Coarsen all elements Q_l with $\tilde{\eta}_l < \theta \tilde{\eta}_{\max}$ and $\|V_{\Gamma_i} - \pi_{\widetilde{\Gamma}_i} V_{\Gamma_i}\|_{\infty, Q_l} \leq tol$.

This procedure also allows to start from rather fine initial grids Γ_0, which have the advantage of yielding a good approximation $\tilde{\eta}_l$ of η_l. Unnecessarily fine elements in the initial grids will this way be coarsened afterwards.

In addition, it might be desirable to add additional refinements in order to avoid large differences in size between adjacent elements, e.g., to avoid degeneracies or large numbers of hanging nodes. Such regularization steps could be included as a step (2b) after the error based refinement and coarsening has been performed. In the implementation used in Section 7, below, such a criterion was used; there the difference in refinement levels between two adjacent elements was restricted in such a way that at most one hanging node can appear on each edge in the grid.

6.4 Anisotropic grids

Finally, we want to discuss the possibility of anisotropic refinement, i.e., the refinement of a cuboid element Q_l not in all but only in selected coordinate directions, thus allowing to use "flat" or "stretched" elements. For some problems, anisotropic refinement can be controlled by suitably designed error estimates (see, e.g., [36]). Here we apply a heuristic

strategy based on the evaluation of $\eta(x)$ in a suitable selection of test points. The basic idea is to place new nodes only at those points x where the error estimate $\eta(x)$ gives large values.

More precisely, consider an element Q of Γ (we drop the indeces for notational convenience) and let $X_{new,i}$ be the set of potential new nodes which would be added to Γ if the element Q_l was refined in coordinate direction e_i. Figure 6.2 shows these points in 2d.

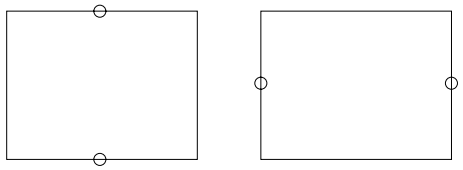

Figure 6.2: Potential new nodes $X_{new,1}$ (left) and $X_{new,2}$ (right) for a 2d element Q

Define the error estimate in these nodes for each coordinate direction e_i by $\eta_{dir,i} := \max_{x \in X_{new,i}} \eta(x)$ and define the overall error measured in these potential new nodes by $\eta_{dir} := \max_{i=1,\dots,n} \eta_{dir,i}$. Note that $\eta_{dir} \leq \eta_l$ always holds. If we include all the points in $X_{new} := \bigcup_{i=1,\dots,n} X_{new,i}$ in our set of test points $X_T(Q)$ (which is reasonable because in order to compute $\eta_{dir,i}$ we have to evaluate $\eta(x)$ for $x \in X_{new}$, anyway) then we can also ensure $\eta_{dir} \leq \tilde{\eta}_l$.

Now we introduce an anisotropy factor $\gamma \in [0, 1]$ and make a two step decision:

(1) We allow anisotropic refinement if $\eta_{dir} \geq (1 - \gamma)\tilde{\eta}_l$, i.e., if the error estimate in the potential new nodes X_{new} is sufficiently large compared to the overall estimated error $\tilde{\eta}_l$.

(2) If anisotropic refinement is allowed according to (1), then we refine the grid in each direction e_i, $i = 1, \dots, n$ for which $\eta_{dir,i} \geq \gamma\eta_{dir}$ holds, i.e., we refine in each direction e_i in which the error in the potential new nodes $X_{new,i}$ is sufficiently large compared to η_{dir}.

For $\gamma = 0$ this implies that anisotropic refinement will never happen, i.e., each element is always refined in each coordinate direction. Conversely, for $\gamma = 1$ in each refinement step each element will only be refined anisotropically, i.e., in one coordinate direction only. The specific structure of this two step decision mechanism has no theoretical foundation; we use it because it shows good results in practice.

Anisotropic refinement can considerably increase the efficiency of the adaptive gridding strategy, in particular if the solution V has certain anisotropic properties, e.g., if V is linear or almost linear in one coordinate direction. On the other hand, a very anisotropic grid Γ can cause degeneracy of the function V_Γ like, e.g., large Lipschitz constants or large discrete curvature even if V is regular, hence the assumptions made in the theorems in Section 5 might be violated or might only hold with large constants. However, according to our numerical experience the positive effects of anisotropic grids are usually predominant.

7 Numerical examples

In this section we describe two test problems in order to illustrate the performance of the proposed numerical method.

7.1 Example: a stochastic economic growth model

Our first example is a one–sector stochastic economic growth model taken from [33][1]. The model is two–dimensional with one–dimensional control and stochastic component. The dynamics and return function are given by

$$\varphi(x,u,z) = \begin{pmatrix} Ax_2x_1^\alpha - u \\ \rho x_2 + z \end{pmatrix} \qquad g(x,u,z) = \log u,$$

where ρ, α and A are positive real constants and $\beta(x) \equiv \beta_0$ which we specified as $\rho = 0.9$, $\alpha = 0.34$, $A = 5$ and $\beta_0 = 0.95$. The random variable is Gaussian distributed with zero mean and standard deviation 0.008.

This simple model has several desirable features as a test problem, because

(i) the exact solution is known. It is given by

$$V(x) = B + C \ln x_1 + Dx_2,$$

where

$$B = \frac{\ln((1-\beta\alpha)A) + \frac{\beta\alpha}{1-\beta\alpha}\ln(\beta\alpha A)}{1-\beta}, \quad C = \frac{\alpha}{1-\alpha\beta}, \quad D = \frac{1}{(1-\alpha\beta)(1-\rho\beta)}$$

(ii) the exact solution is smooth for $x_1 \neq 0$, but the operator T is not smoothing, i.e., we can test the convergence of gradients and the results from Theorem 5.4.

(iii) the exact solution has large curvature for $x_1 \approx 0$, hence the numerical approximation is not too trivial, and it is linear in x_2–direction which makes this a good test problem for the anisotropic refinement strategy.

We have computed the solution to this problem on the domain $\Omega = [0.1, 10] \times [-0.32, 0.32]$. The Gaussian variable z was approximated by a trapezoidal rule with 11 discrete values equidistributed in the interval $[-0.032, 0.032]$ which ensures $\varphi(x,u,z) \in \Omega$ for $x \in \Omega$ and suitable $u \in U = [0.5, 10.5]$. For evaluating the maximum in T the set U was discretized with 161 points. Table 7.1 shows the results of the resulting adaptive gridding scheme applied with refinement threshold $\theta = 0.1$, coarsening tolerance $tol = 0.001$ and anisotropy factor $\gamma = 0.8$. The given CPU time is the accumulated time for all iterations up to the current one including the time needed for error estimation and grid refinement. Figure 7.1 shows the resulting optimal value function and adapted grid.

The error was measured in the L_∞ norm $\|V - V_\Gamma\|_\infty$ while the error in the gradients was measured according to Proposition 4.1 by

$$\max_{l,i} \min_{\xi \in Q_l} |D_x V(\xi) e_i - D_x V_\Gamma(\xi) e_i|$$

in order to rule out effects caused by the size of the elements. As expected, both errors decrease during the refinement iteration, as well as the corresponding upper error bounds

[1]Here we use a different notation. The variables used in [33] are $k = x_1$, $z = e^{x_2}$, $c = u$, $\varepsilon = z$.

iteration	nodes	$\tilde{\eta}_{\max}$	error	$\max \tilde{\eta}_l/\underline{k}_l$	deriv. error	total CPU time (s)
1	49	0.797717	1.357262	7.478598	1.608344	0.94
2	56	0.348198	0.535164	3.264356	0.784972	1.95
3	65	0.170612	0.293198	1.599490	0.570023	3.02
4	109	0.082990	0.133594	1.385438	0.344867	5.16
5	154	0.034092	0.055225	0.598638	0.100389	8.54
6	327	0.012166	0.021920	0.430123	0.103138	14.35
7	889	0.003674	0.009643	0.272726	0.108210	34.84
8	2977	0.001604	0.004342	0.229363	0.089012	132.01

Table 7.1: Numerical results for Example 7.1, adaptive discretization

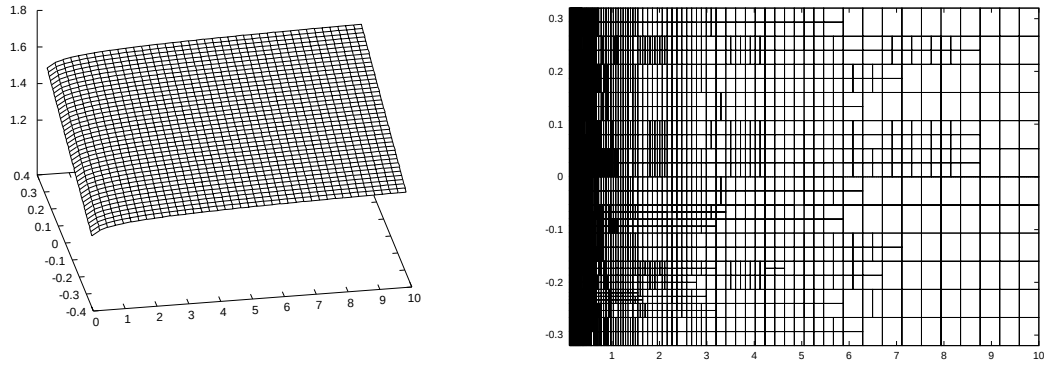

Figure 7.1: Approximated value function and final adaptive grid for Example 7.1

from Theorem 3.2 and Proposition 4.1. Also, these theoretical bounds turn out to be reliable here; recall that $\tilde{\eta}_l$ is only an approximation for η_l which could be a source of numerical errors.

One also observes, that $\max \tilde{\eta}_l/\underline{k}_l$ almost stops decreasing after the seventh iteration (also in further iterations), which indicates that the condition from Theorem 5.4 might be violated here. In fact, additional tests revealed that the number of iterations for which this value decreases strongly depends on the accuracy in the evaluation of T_Γ: the more accurate T_Γ is evaluated, the longer the value $\max \tilde{\eta}_l/\underline{k}_l$ decreases.

In [33, Table 2(a)] on equidistant grids with $143 \times 9 = 1287$ and $500 \times 33 = 16500$ nodes, L_∞–errors of 0.21 and 0.0148, respectively, were reported. In our adaptive iteration these accuracies could be obtained with 109 and 889 nodes, respectively; thus we obtain a reduction in the number of nodes of more than 90% in the first and almost 95% in the second case.

7.2 Example: Zubov's equation

In order to illustrate the performance of our method for discretized continuous time problems, as a second example we consider the numerical approximation of the stochastic Zubov equation

$$\begin{cases} -\frac{1}{2}\mathrm{Tr}\left(a(x)D^2V(x)\right) - b(x)DV(x) - \delta(G(x) - V(x)G(x)) = 0 & x \in \mathbb{R}^n \setminus A, \\ V(x) = 0 & x \in A \end{cases} \tag{7.1}$$

This equation was introduced in [8] and further investigated in [6]. For $\delta \to 0$ it gives the attraction probability of the solutions of the underlying stochastic differential equation

$$\begin{cases} dX(t) &= b(X(t))\,dt + \sigma(X(t))\,dW(t) \\ X(0) &= x. \end{cases} \tag{7.2}$$

to the set $A \subset \mathbb{R}^n$ which is assumed to be a.s. locally exponentially stable. Note that here no control variable is present, though it could be added if desired, see [8].

In order to apply the discretization method from [4] we first need to regularize (7.1) according to [8] replacing the second $G(x)$–term by $G_\varepsilon(x) = \max\{\varepsilon, G(x)\}$. In the computation, below, we used $\varepsilon = 0.001$. Then the semi–discrete approximation of (7.1) with time–step $h > 0$ is of type (2.1) with

$$\varphi(x,z) = x + hb(x) + z\sigma(x), \;\; g(x) = h\delta G_\varepsilon(x) \;\text{ and }\; \beta(x) = 1 - h\delta G_\varepsilon(x),$$

where z is a two–point distributed random variable which assumes the values $\pm\sqrt{h}$ with probability $1/2$ (φ is the discretization of (7.2) using the simplified weak Euler scheme, cf. [25, Section 14.1]).

We have solved this equation for the SDE

$$\begin{aligned} dX_1(t) &= (-3 + \cos X_2(t))X_1(t) + X_1(t)^3 dt \\ dX_2(t) &= \sigma dW(t) \end{aligned}$$

with $\sigma = 5$ on the domain $\Omega = [-2.5, 2.5] \times \mathbb{R}$ with $A = \{0\} \times \mathbb{R}$ and boundary condition $V \equiv 1$ outside Ω. Due to the 2π–periodicity of the equation in the second component the actual computation only needs to be carried out on $\Omega = [-2.5, 2.5] \times [0, 2\pi]$, with periodic boundary conditions $V((x_1, 0)^T) = V((x_1, 2\pi)^T)$.

The numerical solution shown in Table 7.2 and Figure 7.2 was computed with $h = 0.05$, $\delta = 0.0001$, $g(x) = |x_1|$, refinement threshold $\theta = 0.1$ and coarsening tolerance $tol = 0.001$. Anisotropic refinement was not used in this computation.

For this problem we expect the solution of (7.1) to be non–smooth (the solution to (7.1) has to interpreted in the viscosity sense), and indeed the values $\max \tilde{\eta}_l / \underline{k}_l$ remain almost constant during the iterative refinement which supports this conjecture.

It should be noted that here $\beta_0 = 1 - \delta\varepsilon h = 1 - 5 \cdot 10^{-9} \approx 1$ holds, thus the upper error bound from Theorem 3.2 does not give useful information. Nevertheless, the lower error bound in this theorem still holds which implies that large error estimates still indicate large real errors, which explains why the resulting adaptive discretization routine provides

iteration	nodes	$\tilde{\eta}_{\max}$	$\max \tilde{\eta}_l / \underline{k}_l$	total CPU time (s)
1	2401	0.050458	0.484395	0.17
2	2569	0.028417	0.545600	0.51
3	7011	0.012574	0.482858	1.25
4	19092	0.005579	0.428443	3.54
5	54091	0.003665	0.436802	10.53

Table 7.2: Numerical results for Example 7.2

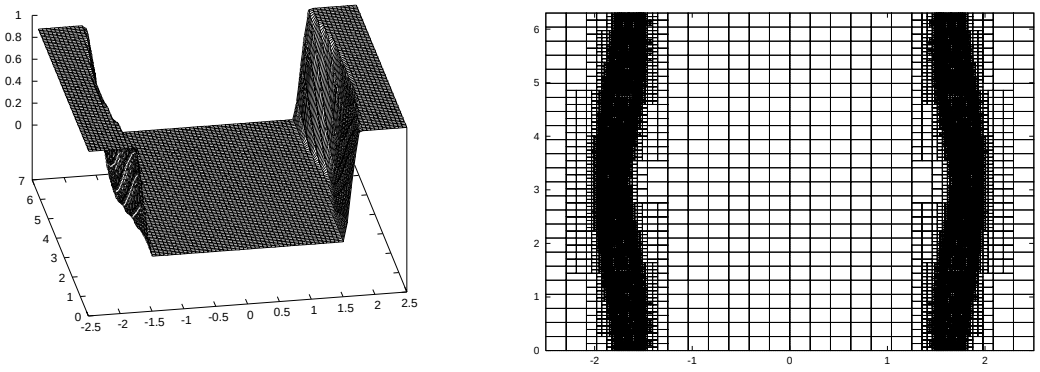

Figure 7.2: Approximated value function and final adaptive grid for Example 7.2

a very reasonable final discretization also for this highly degenerate problem (cf. also the deterministic numerical example in [18] which exhibits a similar degeneracy).

In fact, for this problem the convergence of the scheme does not rely on the contraction property of T but on additional dynamical stability properties of the solution, which for the analogous deterministic problem have been investigated in [21, Section 7.6] using ideas from numerical dynamics. A thorough convergence analysis of this scheme in the stochastic setting along the lines of [21] is currently under investigation.

Acknowledgement: This research was supported by the Center for Empirical Macroeconomics, University of Bielefeld. The support is gratefully acknowledged.

References

[1] M. Bardi and I. Capuzzo Dolcetta, *Optimal Control and Viscosity Solutions of Hamilton-Jacobi-Bellman equations*, Birkhäuser, Boston, 1997.

[2] G. Barles and E. R. Jakobsen, *On the convergence rate of approximation schemes for Hamilton–Jacobi–Bellman equations*, M^{2}AN, Math. Model. Numer. Anal., 36 (2002), pp. 33–54.

[3] D. P. Bertsekas, *Dynamic Programming and Optimal Control. Vol. 1 and 2.*, Athena Scientific, Belmont, MA, 1995.

[4] F. Camilli and M. Falcone, *An approximation scheme for the optimal control of diffusion processes*, RAIRO, Modélisation Math. Anal. Numér., 29 (1995), pp. 97–122.

[5] F. Camilli and L. Grüne, *Numerical approximation of the maximal solutions for a class of degenerate Hamilton–Jacobi equations*, SIAM J. Numer. Anal., 38 (2000), pp. 1540–1560.

[6] ——, *Characterizing attraction probabilities via the stochastic Zubov equation*, Discrete Contin. Dyn. Syst. Ser. B, (2003). To appear.

[7] F. Camilli, L. Grüne, and F. Wirth, *A regularization of Zubov's equation for robust domains of attraction*, in Nonlinear Control in the Year 2000, Volume 1, A. Isidori, F. Lamnabhi-Lagarrigue, and W. Respondek, eds., Lecture Notes in Control and Information Sciences 258, NCN, Springer Verlag, London, 2000, pp. 277–290.

[8] F. Camilli and P. Loreti, *A characterization of the domain of attraction for a locally exponentially stable stochastic system.* Quaderno IAC, n. 25, 2000.

[9] E. Carlini, M. Falcone, and R. Ferretti, *An efficient algorithm for Hamilton–Jacobi equations in high dimensions.* Report 2002/7, Università di Roma "La Sapienza", Dipartimento di Matematica "Guido Castelnuovo", 2002. Submitted.

[10] F. Colonius and W. Kliemann, *The Dynamics of Control*, Birkhäuser, Boston, 2000.

[11] M. G. Crandall and P. L. Lions, *Viscosity solutions of Hamilton–Jacobi equations*, Trans. Amer. Math. Soc., 277 (1983), pp. 1–42.

[12] J. W. Daniel, *Splines and efficiency in dynamic programming*, J. Math. Anal. Appl., 54 (1976), pp. 402–407.

[13] M. Falcone, *A numerical approach to the infinite horizon problem of deterministic control theory*, Appl. Math. Optim., 15 (1987), pp. 1–13. *Corrigenda*, ibid., 23 (1991), 213–214.

[14] M. Falcone and R. Ferretti, *Fully discrete high–order schemes for viscosity solutions of Hamilton–Jacobi–Bellman equations.* Preprint, Università di Roma "La Sapienza", 2000.

[15] W. H. Fleming and M. H. Soner, *Controlled Markov processes and viscosity solutions*, Springer–Verlag, New York, 1993.

[16] R. Franke, P. Terwiesch, and M. Meyer, *Development of an algorithm for the optimal control of trains*, in Proceedings of the 39th IEEE Conference on Decision and Control, Sydney, Australia, 2000, pp. 2123–2128.

[17] R. L. V. González and C. A. Sagastizábal, *Un algorithme pour la résolution rapide d'équations discrètes de Hamilton-Jacobi-Bellman*, C. R. Acad. Sci., Paris, Sér. I, 311 (1990), pp. 45–50.

[18] L. Grüne, *An adaptive grid scheme for the discrete Hamilton–Jacobi–Bellman equation*, Numer. Math., 75 (1997), pp. 319–337.

[19] ——, *Homogeneous state feedback stabilization of homogeneous systems*, SIAM J. Control Optim., 38 (2000), pp. 1288–1314.

[20] ——, *Adaptive grid generation for evolutive Hamilton–Jacobi–Bellman equations*, in Numerical methods for viscosity solutions and applications, M. Falcone and C. Makridakis, eds., vol. 59 of Ser. Adv. Math. Appl. Sci., World Scientific, Singapore, 2001.

[21] ——, *Asymptotic Behavior of Dynamical and Control Systems under Perturbation and Discretization*, Lecture Notes in Mathematics, Vol. 1783, Springer–Verlag, 2002.

[22] L. Grüne, M. Metscher, and M. Ohlberger, *On numerical algorithm and interactive visualization for optimal control problems*, Comput. Vis. Sci., 1 (1999), pp. 221–229.

[23] L. Grüne and W. Semmler, *Using dynamic programming with adaptive grid scheme for optimal control problems in economics*. Working Paper No. 38, Center for Empirical Macroeconomics, University of Bielefeld, 2002. Submitted.

[24] L. Grüne, W. Semmler, and M. Sieveking, *Thresholds in a credit market model with multiple equilibria*. Discussion Paper Nr. 482, Department of Economics, University of Bielefeld, 2001. Submitted.

[25] P. E. Kloeden and E. Platen, *Numerical Solution of Stochastic Differential Equations*, Springer–Verlag, Heidelberg, 1992. (3rd revised and updated printing, 1999).

[26] P. L. Lions, *Generalized solutions of Hamilton-Jacobi equations*, Pitman, London, 1982.

[27] J. L. Menaldi, *Some estimates for finite difference approximations*, SIAM J. Control Optim., 27 (1989), pp. 579–607.

[28] R. Munos and A. Moore, *Variable resolution discretization in optimal control*, Mach. Learn., (2002), pp. 291–323.

[29] M. Reiter, *Solving higher–dimensional continuous–time stochastic control problems by value function regression*, J. Econom. Dynam. Control, 23 (1999), pp. 1329–1353.

[30] J. Rust, *Numerical dynamic programming in economics*, in Handbook of Computational Economics, H. M. Amman, D. A. Kendrick, and J. Rust, eds., Elsevier, Amsterdam, 1996.

[31] ——, *Using randomization to break the curse of dimensionality*, Econometrica, 65 (1997), pp. 487–516.

[32] M. Sagona and A. Seghini, *An adaptive scheme on unstructured grids for the shape–from–shading problem*, in Numerical methods for viscosity solutions and applications, M. Falcone and C. Makridakis, eds., vol. 59 of Ser. Adv. Math. Appl. Sci., World Scientific, Singapore, 2001.

[33] M. S. Santos and J. Vigo-Aguiar, *Accuracy estimates for a numerical approach to stochastic growth models.* Discussion Paper 107, Institute for Empirical Macroeconomics, Federal Reserve Bank of Minneapolis, December 1995.

[34] ——, *Analysis of a numerical dynamic programming algorithm applied to economic models*, Econometrica, 66 (1998), pp. 409–426.

[35] A. Seeck, *Iterative Lösungen der Hamilton-Jacobi-Bellman-Gleichung bei unendlichem Zeithorizont.* Diplomarbeit, Universität Kiel, 1997.

[36] K. G. Siebert, *An a posteriori error estimator for anisotropic refinement*, Numer. Math., 73 (1996), pp. 373–398.

[37] M. A. Trick and S. E. Zin, *Spline approximations to value functions: a linear programming approach*, Macroecon. Dyn., 1 (1997), pp. 255–277.