13

Finite Alphabet Controllers and Estimators

Contributed by Daniel Quevedo

13.1 Introduction

In this chapter we address the issue of control and estimation when the decision variables must satisfy a finite set constraint. We will distinguish between finite alphabet control and estimation problems. As in the case of convex constraints, the essential difference is whether or not the initial state is given or can be considered a decision variable.

Finite alphabet control occurs in many practical situations including: on-off control, relay control, control where quantisation effects are important (in principle this covers all digital control systems and control systems over digital communication networks), and switching control of the type found in power electronics.

Exactly the same design methodologies can be applied in other areas; for example, the following problems can be directly formulated as finite alphabet control problems:

- quantisation of audio signals for compact disc production;
- design of filters where the coefficients are restricted to belong to a finite set (it is common in digital signal processing to use coefficients that are powers of two to facilitate implementation issues);
- design of digital-to-analog [D/A] and analog-to-digital [A/D] converters.

Finite alphabet estimation problems are also frequently encountered in practice. Common examples are:

- estimation of transmitted signals in digital communication systems where the signals are known to belong to a finite alphabet (say ± 1);
- state estimation problems where a disturbance is known to take only a finite set of values (for example, either "on" or "off").

In this chapter we show how these problems can be formulated in the same general framework as described in earlier chapters. However, special care is needed to address the finite set nature of the constraints. In particular, this restriction gives rise to a hard combinatorial optimisation problem, which is exponential in the dimension of the problem. Thus, various approximation techniques are necessary to deal with optimisation problems in which the horizon is large. Commonly used strategies are variants of well-known branch and bound algorithms (Land and Doig 1960, Bertsekas 1998).

We will show how the receding horizon principle can be used in these problems. A key observation in this context is the fact that often the "first" decision variable is insensitive to increasing the optimisation horizon beyond some modest value (typically 3 to 10 in many real world problems).

Also, a *closed form* expression for the control law is derived by exploiting the geometry of the underlying optimisation problem. The solution can also be characterised by means of a partition of the state space, which is closely related to the partition induced by the *interval-constrained* solution,[1] as developed in Chapter 6. As a consequence, the controller can be implemented without relying upon on-line numerical optimisation. Furthermore, the insight obtained from this viewpoint into the nature of the control law can be used to study the dynamic behaviour of the closed loop system.

13.2 Finite Alphabet Control

Consider a linear system having a scalar input u_k and state vector $x_k \in \mathbb{R}^n$ described by

$$x_{k+1} = Ax_k + Bu_k. \tag{13.1}$$

(Here we treat only the scalar input case, but the extension to multiple inputs presents no additional conceptual difficulties.) A key consideration here is that the input is restricted to belong to the finite set

$$\mathbb{U} = \{s_1, s_2, \ldots, s_{n_{\mathbb{U}}}\}, \tag{13.2}$$

where $s_i \in \mathbb{R}$ and $s_i < s_{i+1}$ for $i = 1, 2, \ldots, n_{\mathbb{U}} - 1$.

We will formulate the input design problem as a receding horizon quadratic regulator problem with finite set constraints. Thus, given the state $x_k = x$, we seek the optimising sequence of present and future control inputs:

[1] The explicit solution in problems with "interval-type" constraints of the form $|u| \leq \Delta$.

$$\mathbf{u}^{\mathrm{OPT}}(x) \triangleq \arg\min_{\mathbf{u}_k \in \mathbb{U}^N} V_N(x, \mathbf{u}_k), \tag{13.3}$$

where

$$\mathbf{u}_k \triangleq \begin{bmatrix} u_k \\ u_{k+1} \\ \vdots \\ u_{k+N-1} \end{bmatrix}, \quad \mathbb{U}^N \triangleq \mathbb{U} \times \cdots \times \mathbb{U}. \tag{13.4}$$

As in previous chapters, in (13.3) V_N is the finite horizon quadratic objective function[2]

$$V_N(x, \mathbf{u}_k) \triangleq \|x_{k+N}\|_P^2 + \sum_{t=k}^{k+N-1} (\|x_t\|_Q^2 + \|u_t\|_R^2), \tag{13.5}$$

with $Q = Q^{\mathrm{T}} > 0$, $P = P^{\mathrm{T}} > 0$, $R = R^{\mathrm{T}} > 0$ and where $x_k = x$. Note that, as usual, the formulation of $V_N(\,\cdot\,,\,\cdot\,)$, uses predictions of future plant states.

Whilst we concentrate here upon plant state deviations from the origin, nonzero references can also be encompassed within this framework. In order to accomplish this, the objective function (13.5) needs to be modified by considering *shifted coordinates* as is common when dealing with nonzero constant references in standard receding horizon control schemes (see Chapter 5).

The minimisation of (13.5) subject to the finite set constraint on $\mathbf{u}_k$ and the plant dynamics expressed in (13.1) yields the optimal sequence $\mathbf{u}^{\mathrm{OPT}}(x)$. It is a function only of the current state value $x_k = x$.

Following the usual receding horizon principle (see Chapter 4), only the first control action, namely

$$u^{\mathrm{OPT}}(x) \triangleq \begin{bmatrix} 1 & 0 & \cdots & 0 \end{bmatrix} \mathbf{u}^{\mathrm{OPT}}(x), \tag{13.6}$$

is applied. At the next time instant, the optimisation is repeated with a new *initial* state and the finite horizon window shifted by one.

In the next section we present a closed form expression for $\mathbf{u}^{\mathrm{OPT}}(x)$. This is directly analogous to the geometric interpretation of the constrained solution developed in Chapter 6. This result will allow us to characterise the control law as a partition of the state space and provide a tool for studying the dynamic behaviour of the resulting closed loop system.

13.3 Nearest Neighbour Characterisation of the Solution

Since the constraint set $\mathbb{U}^N$ is finite, the optimisation problem (13.3) is non-convex. Indeed, it is a hard combinatorial optimisation problem whose solution requires a computation time that is exponential in the horizon length. Thus,

[2] $\|\nu\|_S^2$ denotes $\nu^{\mathrm{T}} S \nu$, where ν is any vector and S is a matrix.

one needs either to use a relatively small horizon or to resort to approximate solutions. We will adopt the former strategy based on the premise that, due to the receding horizon technique, the *first* decision variable is all that is of interest. Moreover, it is a practical observation that this first decision variable is often insensitive to increasing the horizon length beyond some relative modest value. To proceed, it is useful to *vectorise* the objective function (13.5) as follows:

Define

$$\mathbf{x}_k \triangleq \begin{bmatrix} x_{k+1} \\ x_{k+2} \\ \vdots \\ x_{k+N} \end{bmatrix}, \quad \Phi \triangleq \begin{bmatrix} B & 0 & \dots & 0 & 0 \\ AB & B & \dots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ A^{N-1}B & A^{N-2}B & \dots & AB & B \end{bmatrix}, \quad \Lambda \triangleq \begin{bmatrix} A \\ A^2 \\ \vdots \\ A^N \end{bmatrix}, \tag{13.7}$$

so that, given $x_k = x$ and by iterating (13.1), the predictor $\mathbf{x}_k$ satisfies

$$\mathbf{x}_k = \Phi \mathbf{u}_k + \Lambda x. \tag{13.8}$$

Hence, the objective function (13.5) can be re-written as

$$V_N(x, \mathbf{u}_k) = \bar{V}_N(x) + \mathbf{u}_k^{\mathrm{T}} H \mathbf{u}_k + 2\mathbf{u}_k^{\mathrm{T}} F x, \tag{13.9}$$

where

$$H \triangleq \Phi^{\mathrm{T}} \mathbf{Q} \Phi + \mathbf{R} \in \mathbb{R}^{N \times N}, \quad F \triangleq \Phi^{\mathrm{T}} \mathbf{Q} \Lambda \in \mathbb{R}^{N \times n},$$
$$\mathbf{Q} \triangleq \mathrm{diag}\{Q, \dots, Q, P\} \in \mathbb{R}^{Nn \times Nn}, \quad \mathbf{R} \triangleq \mathrm{diag}\{R, \dots, R\} \in \mathbb{R}^{N \times N},$$

and $\bar{V}_N(x)$ does not depend upon $\mathbf{u}_k$.

By direct calculation, it follows that the minimiser to (13.9), without taking into account any constraints on $\mathbf{u}_k$, is

$$\mathbf{u}_{\mathrm{UC}}^{\mathrm{OPT}}(x) = -H^{-1} F x. \tag{13.10}$$

Our subsequent development will utilise a nearest neighbour vector quantiser in order to characterise the constrained optimiser. This is defined as follows:

Definition 13.3.1 (Nearest Neighbour Vector Quantiser) *Given a countable (not necessarily finite) set of nonequal vectors $\mathbb{B} = \{b_1, b_2, \dots\} \subset \mathbb{R}^{n_{\mathbb{B}}}$, the nearest neighbour quantiser is defined as a mapping $q_{\mathbb{B}} \colon \mathbb{R}^{n_{\mathbb{B}}} \to \mathbb{B}$ that assigns to each vector $c \in \mathbb{R}^{n_{\mathbb{B}}}$ the closest element of $\mathbb{B}$ (as measured by the Euclidean norm), that is, $q_{\mathbb{B}}(c) = b_i \in \mathbb{B}$ if and only if c belongs to the region*

$$\begin{aligned} &\left\{ c \in \mathbb{R}^{n_B} \colon \|c - b_i\|^2 \le \|c - b_j\|^2 \text{ for all } b_j \ne b_i,\ b_j \in \mathbb{B} \right\} \\ &\quad \setminus \left\{ c \in \mathbb{R}^{n_B} \colon \text{there exists } j < i \text{ such that } \|c - b_i\|^2 = \|c - b_j\|^2 \right\}. \end{aligned} \tag{13.11}$$

$\circ$

Note that in the special case, when $n_{\mathbb{B}} = 1$, this vector quantiser reduces to the standard scalar quantiser.

In the above definition, the zero measure set of points that satisfy (13.11) with equality have been arbitrarily assigned to the element having the smallest index. This is done in order to avoid ambiguity in the case of *frontier points*, that is, points equidistant to two or more elements of $\mathbb{B}$. If this aspect does not matter, then expression (13.11) can be simplified to

$$\left\{c \in \mathbb{R}^{n_B} : \|c - b_i\|^2 \leq \|c - b_j\|^2 \text{ for all } b_j \neq b_i,\ b_j \in \mathbb{B}\right\}. \tag{13.12}$$

Given Definition 13.3.1, we can now restate the solution to (13.3). This leads to:

Theorem 13.3.1 (Closed Form Solution) *Let $\mathbb{U}^N = \{v_1, v_2, \ldots, v_r\}$, where $r = (n_{\mathbb{U}})^N$. Then the optimiser $\mathbf{u}^{\text{OPT}}(x)$ in (13.3) is given by*

$$\mathbf{u}^{\text{OPT}}(x) = H^{-1/2} q_{\tilde{\mathbb{U}}^N}(-H^{-\text{T}/2} F x), \tag{13.13}$$

where the nearest neighbour quantiser $q_{\tilde{\mathbb{U}}^N}(\cdot)$ maps $\mathbb{R}^N$ to $\tilde{\mathbb{U}}^N$, defined as

$$\tilde{\mathbb{U}}^N \triangleq \{\tilde{v}_1,\ \tilde{v}_2, \ldots, \tilde{v}_r\}, \quad \tilde{v}_i = H^{1/2} v_i, \quad v_i \in \mathbb{U}^N. \tag{13.14}$$

Proof. For fixed x, the level sets of the objective function (13.9) are ellipsoids in the input sequence space $\mathbb{R}^N$. These are centred at the point $\mathbf{u}_{\text{UC}}^{\text{OPT}}(x)$ defined in (13.10). Thus, the optimisation problem (13.3) can be geometrically interpreted as the one where we find the point $\mathbf{u}_k \in \mathbb{U}^N$ that belongs to the smallest ellipsoid defined by (13.9) (that is, the point which provides the smallest objective function value whilst satisfying the constraints).

In order to simplify the problem, we introduce the same coordinate transformation utilised in Chapter 6, that is,

$$\tilde{\mathbf{u}}_k = H^{1/2} \mathbf{u}_k, \tag{13.15}$$

which transforms the constraint set $\mathbb{U}^N$ into $\tilde{\mathbb{U}}^N$ defined in (13.14). The optimiser $\mathbf{u}^{\text{OPT}}(x)$ can be defined in terms of this auxiliary variable as

$$\mathbf{u}^{\text{OPT}}(x) = H^{-1/2} \arg\min_{\tilde{\mathbf{u}}_k \in \tilde{\mathbb{U}}^N} J_N(x, \tilde{\mathbf{u}}_k), \tag{13.16}$$

where

$$J_N(x, \tilde{\mathbf{u}}_k) \triangleq \tilde{\mathbf{u}}_k^{\text{T}} \tilde{\mathbf{u}}_k + 2\tilde{\mathbf{u}}_k^{\text{T}} H^{-\text{T}/2} F x. \tag{13.17}$$

The level sets of J_N are spheres in $\mathbb{R}^N$, centred at

$$\tilde{\mathbf{u}}_{\text{UC}}^{\text{OPT}}(x) \triangleq -H^{-\text{T}/2} F x. \tag{13.18}$$

Hence, the constrained optimiser (13.3) is given by the nearest neighbour to $\tilde{\mathbf{u}}_{\text{UC}}^{\text{OPT}}(x)$, namely

$$\arg\min_{\tilde{\mathbf{u}}_k \in \tilde{\mathbb{U}}^N} J_N(x, \tilde{\mathbf{u}}_k) = q_{\tilde{\mathbb{U}}^N}(-H^{-\text{T}/2} F x). \tag{13.19}$$

The result (13.13) follows by substituting (13.19) into (13.16). □

We observe that, with $N > 1$, the optimiser $\mathbf{u}^{\mathrm{OPT}}(x)$ provided in Theorem 13.3.1 is, in general, different to the sequence obtained by direct quantisation of the unconstrained minimum (13.10), namely, $q_{\mathbb{U}^N}(\mathbf{u}_{\mathrm{UC}}^{\mathrm{OPT}}(x))$.

As a consequence of Theorem 13.3.1, the receding horizon controller (13.6) satisfies

$$u^{\mathrm{OPT}}(x) = \begin{bmatrix} 1 & 0 & \cdots & 0 \end{bmatrix} H^{-1/2} q_{\tilde{\mathbb{U}}^N}(-H^{-\mathrm{T}/2} F x). \tag{13.20}$$

This solution can be illustrated as the composition of the following transformations:

$$x \in \mathbb{R}^n \xrightarrow{-H^{-\frac{\mathrm{T}}{2}} F} \tilde{\mathbf{u}}_{\mathrm{UC}}^{\mathrm{OPT}} \in \mathbb{R}^N \xrightarrow{H^{-\frac{1}{2}} q_{\tilde{\mathbb{U}}^N}(\cdot)} \mathbf{u}^{\mathrm{OPT}} \in \mathbb{U}^N \xrightarrow{[1\,0\,\cdots\,0]} u^{\mathrm{OPT}} \in \mathbb{U}\,. \tag{13.21}$$

It is worth noticing that $q_{\tilde{\mathbb{U}}^N}(\cdot)$ is a *memoryless nonlinearity*, so that (13.20) corresponds to a time-invariant nonlinear state feedback law. In a direct implementation, at each time step, the quantiser needs to perform $r - 1$ comparisons. However, in some cases, it is possible to exploit the nature of the problem to obtain more efficient search algorithms (Quevedo and Goodwin 2003a).

13.4 State Space Partition

Expression (13.11) partitions the domain of the quantiser into polyhedra, called *Voronoi partition*. Since the constrained optimiser $\mathbf{u}^{\mathrm{OPT}}(x)$ in (13.13) (see also (13.21)) is defined in terms of $q_{\tilde{\mathbb{U}}^N}(\cdot)$, an equivalent partition of the state space can be derived, as shown next.

Theorem 13.4.1 *The constrained optimising sequence* $\mathbf{u}^{\mathrm{OPT}}(x)$ *in* (13.13) *can be characterised as*

$$\mathbf{u}^{\mathrm{OPT}}(x) = v_i \Longleftrightarrow x \in \mathcal{R}_i,$$

where

$$\begin{aligned} \mathcal{R}_i \triangleq \Big\{ z \in \mathbb{R}^n \colon 2(v_i - v_j)^{\mathrm{T}} F z \leq \|v_j\|_H^2 - \|v_i\|_H^2 \text{ for all } v_j \neq v_i,\, v_j \in \mathbb{U}^N \Big\} \\ \setminus \Big\{ z \in \mathbb{R}^n \colon \text{there exists } j < i \text{ such that } 2(v_i - v_j)^{\mathrm{T}} F z = \|v_j\|_H^2 - \|v_i\|_H^2 \Big\}. \end{aligned} \tag{13.22}$$

Proof. From expressions (13.13) and (13.14) it follows that $\mathbf{u}^{\mathrm{OPT}}(x) = v_i$ if and only if $q_{\tilde{\mathbb{U}}^N}(-H^{-\mathrm{T}/2} F x) = \tilde{v}_i$. On the other hand,

$$\| - H^{-\mathrm{T}/2} F x - \tilde{v}_i \|^2 = \|H^{-\mathrm{T}/2} F x\|^2 + \|\tilde{v}_i\|^2 + 2\tilde{v}_i^{\mathrm{T}} H^{-\mathrm{T}/2} F x,$$

so that

$$\| - H^{-\mathrm{T}/2} F x - \tilde{v}_i \|^2 \leq \| - H^{-\mathrm{T}/2} F x - \tilde{v}_j \|^2$$

holds if and only if

$$2(\tilde{v}_i - \tilde{v}_j)^{\mathrm{T}} H^{-\mathrm{T}/2} F x \leq \|\tilde{v}_j\|^2 - \|\tilde{v}_i\|^2.$$

This inequality together with expressions (13.14) and (13.11) shows that

$$q_{\tilde{\mathbb{U}}^N}(-H^{-\mathrm{T}/2} F x) = \tilde{v}_i$$

if and only if x belongs to the region $\mathcal{R}_i$ defined in (13.22). This fact completes the proof. □

The $n_{\mathbb{U}}^N$ regions $\mathcal{R}_i$ defined in (13.22) are polyhedra. Without taking into account constraint borders, we can write these in a compact form as

$$\mathcal{R}_i = \left\{ x \in \mathbb{R}^n \colon D_i x \leq h_i \right\},$$

where the rows of D_i are equal to all terms $2(v_i - v_j)^{\mathrm{T}} F$ as required, whilst the vector h_i contains the scalars $\|v_j\|_H^2 - \|v_i\|_H^2$.

Some of the inequalities in (13.22) may be redundant. In these cases, the corresponding regions do not share a common edge, that is, they are not adjacent. This phenomenon is illustrated in Figure 13.4 of Example 13.8.1, where the regions $\mathcal{R}_1$ and $\mathcal{R}_4$ are not adjacent. The inequality separating them is redundant.

Also, depending upon the matrix $H^{-\mathrm{T}/2} F$, some of the regions $\mathcal{R}_i$ may be empty. This might happen, in particular, if $N > n$. In this case, the rank of F is smaller than N and the transformation $H^{-\mathrm{T}/2} F$ does not span the entire space $\mathbb{R}^N$. Figure 13.1 illustrates this for the case $n = 1$, $n_{\mathbb{U}} = 2$ and $N = 2$. As can be seen from this figure, depending on the *unconstrained optimum locus* given by the (dashed) line $-H^{-\mathrm{T}/2} F x,\ x \in \mathbb{R}$, there exist situations in which some sequences $\tilde{v}_j$ will never be optimal, thus yielding empty regions in the state space.

On the other hand, if the pair (A, B) is completely controllable and A is invertible, then the rank of F is equal to $\min(N, n)$. In this case, if $n \geq N$, then $H^{-\mathrm{T}/2} F$ is onto, so that for every $\tilde{v}_j \in \tilde{\mathbb{U}}^N$ there exists at least one x such that $q_{\tilde{\mathbb{U}}^N}(-H^{-\mathrm{T}/2} F x) = \tilde{v}_j$ and none of the regions $\mathcal{R}_i$ are empty.

13.5 The Receding Horizon Case

In the receding horizon law (13.20), only $n_{\mathbb{U}}$ instead of (at most) $(n_{\mathbb{U}})^N$ regions are needed to characterise the control law. Each of these $n_{\mathbb{U}}$ regions is given by the union of all regions $\mathcal{R}_i$ that correspond to vertices v_i having the same first element. The appropriate extension of Theorem 13.4.1 is presented below. This result follows directly from Theorem 13.4.1.

Corollary 13.5.1 (State Space Partition) *Let the constraint set* $\mathbb{U}$ *be given in* (13.2) *and consider the partition into equivalence classes*

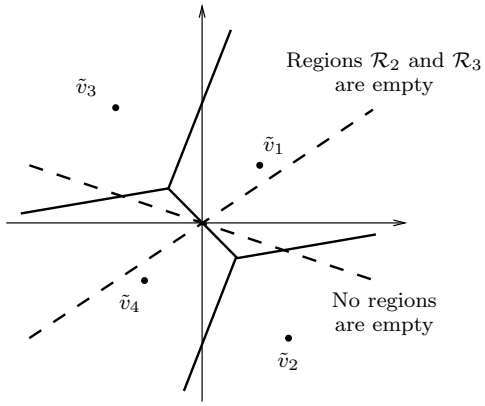

Figure 13.1. Partition of the transformed input sequence space with $N = 2$ (solid lines) and two examples of $-H^{-\mathrm{T}/2}Fx,\ x \in \mathbb{R}$ (dashed lines).

$$\mathbb{U}^N = \bigcup_{i=1,\ldots,n_{\mathbb{U}}} \mathbb{U}_i^N,$$

where

$$\mathbb{U}_i^N \triangleq \left\{ v \in \mathbb{U}^N \colon \begin{bmatrix} 1 & 0 & \cdots & 0 \end{bmatrix} v = s_i \right\}.$$

Then, the receding horizon control law (13.20) *is equivalent to*

$$u^{\mathrm{OPT}}(x) = s_i, \quad \textit{if } x \in \mathcal{X}_i, \quad i = 1, 2, \ldots n_{\mathbb{U}}. \tag{13.23}$$

Here, the polyhedra $\mathcal{X}_i$ *are given by*

$$\mathcal{X}_i \triangleq \bigcup_{j\,\colon\, v_j \in \mathbb{U}_i^N} \mathcal{X}_{ij},$$

where

$$\begin{aligned} \mathcal{X}_{ij} \triangleq & \left\{ z \in \mathbb{R}^n \colon 2(v_j - v_k)^{\mathrm{T}} F z \leq \|v_k\|_H^2 - \|v_j\|_H^2 \textit{ for all } v_k \in \mathbb{U}^N \backslash \mathbb{U}_i^N \right\} \\ & \backslash \Big\{ z \in \mathbb{R}^n \colon \textit{there exists } v_k \in \mathbb{U}^N \backslash \mathbb{U}_i^N,\ k < j, \textit{ such that} \\ & 2(v_j - v_k)^{\mathrm{T}} F z = \|v_k\|_H^2 - \|v_j\|_H^2 \Big\}. \end{aligned}$$

It should be emphasised that this description requires evaluation of less inequalities than the direct calculation of the union of all $\mathcal{R}_j$ (as defined in (13.22)) with $v_j \in \mathbb{U}_i^N$, since inequalities corresponding to *internal borders* are not evaluated. Moreover, the definition of $\mathcal{X}_i$ (and of $\mathcal{R}_i$) can be simplified if the *ambiguity problem* is not addressed.

The state space partition obtained can be calculated off-line so that on-line computational burden can be reduced. The partition induced is related to the partition that characterises the interval-constrained case, as detailed in the following section.

13.6 Relation to Interval Constraints

If in the setup described above, the input is not constrained to belong to a finite set $\mathbb{U}$, but instead, needs to satisfy the interval-type constraint

$$-\Delta \leq u_k \leq \Delta \quad \text{for all } k, \tag{13.24}$$

where $\Delta \in \mathbb{R}$ is fixed, then a convex optimisation problem is obtained.

For this case, as shown in Chapter 6, the control law can be finitely parameterised and calculated off-line. The state space is partitioned into polytopes in which the receding horizon controller is piecewise affine in the state.

The partition of the $\tilde{\mathbf{u}}$-space using the transformation (13.15) and a geometric argument similar to the one used in the proof of Theorem 13.3.1 is sketched in Figure 13.2 for the case $N = 2$ and the restriction (13.24). In Figure 13.2, the polytope Θ_0 is obtained by applying the transformation (13.15) to the region in which the constraints are not active. It is the *allowed set*. The regions denoted as Θ_{si} are adjacent to a face of Θ_0.

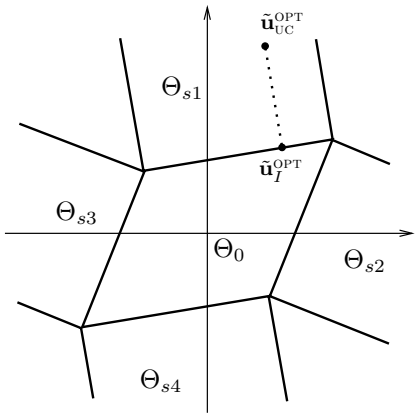

Figure 13.2. Partitions of $\tilde{\mathbf{u}}$-space with the interval constraint set (13.24).

As shown in Section 13.3, in the finite set-constrained case, the constrained solution $\mathbf{u}^{\mathrm{OPT}}$ is related to $\tilde{\mathbf{u}}_{\mathrm{UC}}^{\mathrm{OPT}}$ by means of a nearest neighbour quantiser as stated in (13.13). (For ease of notation, the dependence on x of this and other vectors to follow has not been explicitly included.) A similar result holds in the interval-constrained case. Given (13.24), the constrained optimiser, denoted here as $\tilde{\mathbf{u}}_{I}^{\mathrm{OPT}}$, is related to $\tilde{\mathbf{u}}_{\mathrm{UC}}^{\mathrm{OPT}}$ via a minimum Euclidean distance projection to the allowed set. This result was shown earlier in Chapter 6 and can be summarised as follows:

Remark 13.6.1. (Projection in the Interval-Constrained Case) If $\tilde{\mathbf{u}}_{\mathrm{UC}}^{\mathrm{OPT}}$ lies inside of Θ_0, then it holds that $\tilde{\mathbf{u}}_{I}^{\mathrm{OPT}} = \tilde{\mathbf{u}}_{\mathrm{UC}}^{\mathrm{OPT}}$. On the other hand, if $\tilde{\mathbf{u}}_{\mathrm{UC}}^{\mathrm{OPT}} \notin$

Θ_0, then the constrained solution is obtained by the minimum Euclidean distance projection (see, for example, (10.15) in Chapter 10) onto the border of Θ_0. In particular, if the unconstrained solution lies in any of the regions adjacent to a face of Θ_0, then $\tilde{\mathbf{u}}_I^{\text{OPT}}$ is obtained by an orthogonal projection onto the nearest face (as illustrated in Figure 13.2 by means of a dotted line). ∘

As a consequence of the foregoing discussion, we obtain the following theorem, which establishes a connection between the partition of the $\tilde{\mathbf{u}}$-space in the interval-constrained case and the Voronoi partition of the quantiser defining the solution with a special finite set constraint.

Theorem 13.6.1 (Relationship Between the Binary and Interval-Constrained Cases) *Consider the binary constraint set $\mathbb{U} = \{-\Delta, \Delta\}$ and the region outside of Θ_0. Then, the borders of the Voronoi partition of the quantiser in* (13.13) *are parallel and equidistant to the borders of those regions of the interval-constrained case, which are adjacent to an $(N-1)$-dimensional face of Θ_0. (These regions are denoted in Figure 13.2 as Θ_{si}.)*

Proof. From (13.18) it follows that the solution (13.13) can be stated alternatively as $\mathbf{u}^{\text{OPT}}(x) = H^{-1/2} q_{\tilde{\mathbb{U}}^N}(\tilde{\mathbf{u}}_{\text{UC}}^{\text{OPT}}(x))$. The result is a consequence of the fact that, as can be seen in Figure 13.3, the borders of the regions Θ_{si} are formed by orthogonal projections to $\tilde{v}_i$, and that the Voronoi partition is formed by equidistant hyperplanes, which are also orthogonal to the corresponding $(N-1)$-dimensional face of Θ_0. □

This result is illustrated in Figure 13.3, where the Voronoi partition is depicted via dashed lines. Due to linearity of the mapping $H^{-\text{T}/2}F$ in (13.21), the induced partition of the state space given the constraint (13.24) and the partition defined in (13.22) are similarly related.

13.7 Stability

In Chapters 4 and 5, we found that, for the case of interval-type constraints, one could utilise the value function of the optimal control problems as a candidate Lyapunov function to establish stability. The situation in the finite alphabet case is more difficult. Indeed, asymptotic stability is, in general, too strong a requirement for the finite alphabet problem. In this section we explore various stability issues associated with this case.

The closed loop that results when controlling the plant (13.1) with the receding horizon law (13.23) is described via the following piecewise-affine map, which follows from Corollary 13.5.1:

$$\begin{aligned} &x_{k+1} = g(x_k), \\ &g(x_k) \triangleq Ax_k + Bs_i, \quad \text{if } x_k \in \mathcal{X}_i, \quad i = 1, 2, \ldots n_{\mathbb{U}}. \end{aligned} \tag{13.25}$$

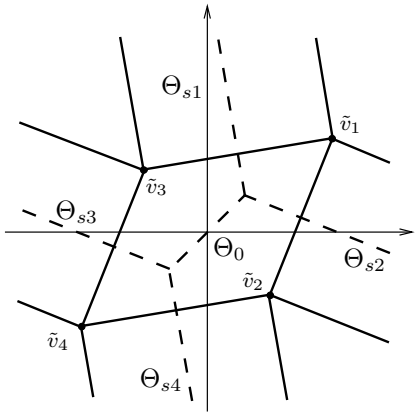

Figure 13.3. Relationship between partitions induced by binary constraints (dashed line) and interval-type constraints (solid line).

Piecewise-affine maps are *mixed mappings* and also form a special class of hybrid systems with underlying discrete time dynamics (see, for example, Bemporad, Ferrari-Trecate and Morari (2000), Heemels, De Schutter and Bemporad (2001) and the references therein). They also appear in connection with some signal processing problems, namely arithmetic overflow of digital filters (Chua and Lin 1988) and $\Sigma\Delta$-modulators (Feely 1997, Norsworthy, Schreier and Temes 1997) and have also been studied in a more theoretical mathematical context (see, for example, Adler, Kitchens and Tresser 2001, Wu and Chua 1994).

Since there exist fundamental differences in the dynamic behaviour of (13.25), depending on whether the plant (13.1) is open loop stable or unstable, that is, on whether the matrix A is Hurwitz or not, it is convenient to divide the discussion that follows accordingly.

13.7.1 Stable Plants

If the plant (13.1) is stable, then its states are always bounded when controlled by means of any finite set constraint law. This follows directly from the fact that $\mathbb{U}$ is always bounded.

Moreover, it can also be shown that all state trajectories[3] of (13.25) either converge towards a fixed point or towards a limit cycle (see, for example, Wu and Chua 1994, Ramadge 1990).

The properties stated so far apply to general systems described by (13.25), where $\mathcal{X}_i$ defines any partition of the state space. In contrast, the following theorem is more specific. It utilises the fact that the control law $u^{\mathrm{OPT}}(x)$ is optimising in a receding horizon sense in order to establish a stronger result.

[3] Exceptions are limited to trajectories that emanate from initial conditions belonging to a zero-measure set.

Theorem 13.7.1 (Asymptotic Stability) *If A is Hurwitz, $0 \in \mathbb{U}$ and $P = P^{\mathrm{T}} > 0$ satisfies the Lyapunov equation $A^{\mathrm{T}}PA + Q = P$, then the closed loop* (13.25) *has a globally attractive, locally asymptotically stable, equilibrium point at the origin.*

Proof. The proof follows standard techniques used in the receding horizon control framework as summarised in Chapter 4 (see also Section 5.6.1 in Chapter 5). In particular, we will use Theorem 4.4.2 of Chapter 4. We choose $\mathbb{X}_f = \mathbb{R}^n$ and $\mathcal{K}_f(x) = 0$ for all $x \in \mathbb{X}_f$. Clearly conditions **B1**, **B3**, **B4** and **B5** hold and $\mathbb{S}_N = \mathbb{R}^n$.

Direct calculation yields that $F(x) = x^{\mathrm{T}}Px$ satisfies

$$\begin{aligned} F(f(x,\mathcal{K}_f(x))) - F(x) + L(x,\mathcal{K}_f(x)) &= (Ax + B\mathcal{K}_f(x))^{\mathrm{T}}P(Ax + B\mathcal{K}_f(x)) \\ &\quad - x^{\mathrm{T}}Px + x^{\mathrm{T}}Qx + (\mathcal{K}_f(x))^{\mathrm{T}}R\mathcal{K}_f(x) \\ &= x^{\mathrm{T}}(A^{\mathrm{T}}PA + Q - P)x \\ &= 0 \quad \text{for all } x \in \mathbb{X}_f, \end{aligned}$$

so that condition **B2** is also satisfied. Global attractivity of the origin then follows from Theorem 4.4.2.

Next, note that there exists a region containing an open neighbourhood of the origin where $u^{\mathrm{OPT}}(x) = 0$, hence local asymptotic stability of the origin follows since A is Hurwitz. □

As can be seen, if the conditions of this theorem are satisfied, then the receding horizon law (13.6) ensures that the origin is not only a fixed point, but also that it has region of attraction $\mathbb{R}^n$.

It should be emphasised here that, in a similar manner, it can be shown that a finite alphabet control law can steer the plant state asymptotically to any point $x^\star$, such that there exist $s_i \in \mathbb{U}$ that allow one to write $x^\star = (I - A)^{-1}Bs_i$.

13.7.2 Unstable Plants

In case of strictly unstable plants (13.1), the situation becomes more involved. Although fixed points and periodic sequences may be admissible, they are basically nonattractive.

Moreover, with control signals that are limited in magnitude, as is the case with finite set constraints (and also with interval constraints), there always exists an unbounded region, such that initial states contained in it lead to unbounded state trajectories. This does not mean that every state trajectory of (13.25) is unbounded. Despite the fact that the unstable open loop dynamics (as expressed in A) makes neighbouring trajectories diverge locally, under certain circumstances the control law may keep the state trajectory bounded.

As a consequence of the highly nonlinear (non-Lipschitz) dynamics resulting from the quantiser defining the control law (13.13), in the bounded case

the resulting closed loop trajectories may be quite complex. In order to analyse them without exploring their fine geometrical structure, it is useful that we relax the usual notion of asymptotic stability of the origin. A more useful characterisation here is that of *ultimate boundedness* of state trajectories. This notion refers to convergence towards a bounded region of $\mathbb{R}^n$, instead of to a point or a specific periodic orbit (Blanchini 1999). (Ultimate boundedness has also been considered in Li and Soh (1999), and by several other authors in the context of *practical stability.*) We refer the reader to the literature, especially Quevedo, De Doná and Goodwin (2002), where these more detailed issues are discussed and analysed for the case of finite alphabet receding horizon control of unstable open loop plants.

13.8 Examples

13.8.1 Open Loop Stable Plant

Consider an open loop stable plant described by

$$x_{k+1} = \begin{bmatrix} 0.1 & 2 \\ 0 & 0.8 \end{bmatrix} x_k + \begin{bmatrix} 0.1 \\ 0.1 \end{bmatrix} u_k, \tag{13.26}$$

and the binary constraint set $\mathbb{U} = \{-1, 1\}$. The receding horizon control law with $R = 0$ and

$$P = Q = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \tag{13.27}$$

partitions the state space into the regions depicted in Figure 13.4, for constraint horizons $N = 2$ and $N = 3$. In this figure x_k^1 and x_k^2 denote the two components of the state vector x_k.

The receding horizon control law is

$$u^{\mathrm{OPT}}(x) = \begin{cases} -1 & \text{if } x \in \mathcal{X}_1, \\ 1 & \text{if } x \in \mathcal{X}_2, \end{cases}$$

where

$$\mathcal{X}_1 = \bigcup_{i=2^{N-1}+1, 2^{N-1}+2, \ldots, 2^N} \mathcal{R}_i, \qquad \mathcal{X}_2 = \bigcup_{i=1,2,\ldots,2^{N-1}} \mathcal{R}_i.$$

13.8.2 Open Loop Unstable Plant

We next analyse a situation when the plant is open loop unstable. For that purpose, consider

$$x_{k+1} = \begin{bmatrix} 1.02 & 2 \\ 0 & 1.05 \end{bmatrix} x_k + \begin{bmatrix} 0.1 \\ 0.1 \end{bmatrix} u_k, \tag{13.28}$$

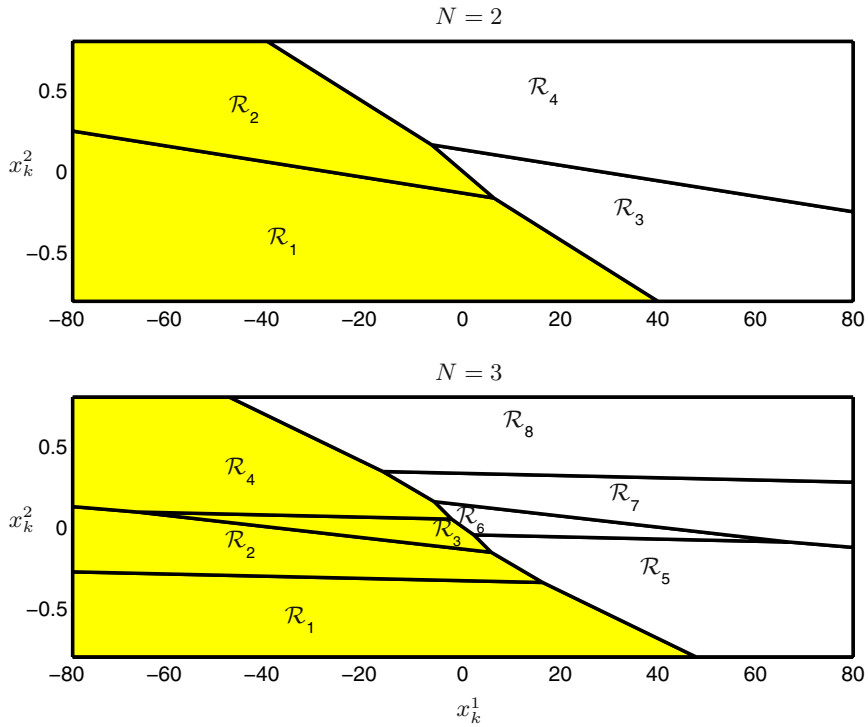

Figure 13.4. State space partition for the plant (13.26).

controlled with a receding horizon controller with parameters $\mathbb{U}$, P, Q and R as in Example 13.8.1 above. The constraint horizon is chosen to be $N = 2$.

Figure 13.5 illustrates the induced state space partition and a closed loop trajectory, which starts at $x = [-10 \quad 0]^{\mathrm{T}}$. As can be seen, due to the limited control action available, the trajectory becomes unbounded.

The situation is entirely different when the initial condition is chosen as $x = [0.7 \quad 0.2]^{\mathrm{T}}$. As depicted in Figure 13.6, the closed loop trajectory now converges to a bounded region, which contains the origin in its interior. Within that region, the behaviour is not periodic, but appears to be random, despite the fact that the system is deterministic. Neighbouring trajectories diverge due to the action of the unstable poles of the plant. However, the control law manifests itself by maintaining the plant state *ultimately bounded*. As already mentioned in Section 13.7.2, the *complex* dynamic behaviour obtained is a consequence of the insertion of a nonsmooth nonlinearity in the feedback loop.

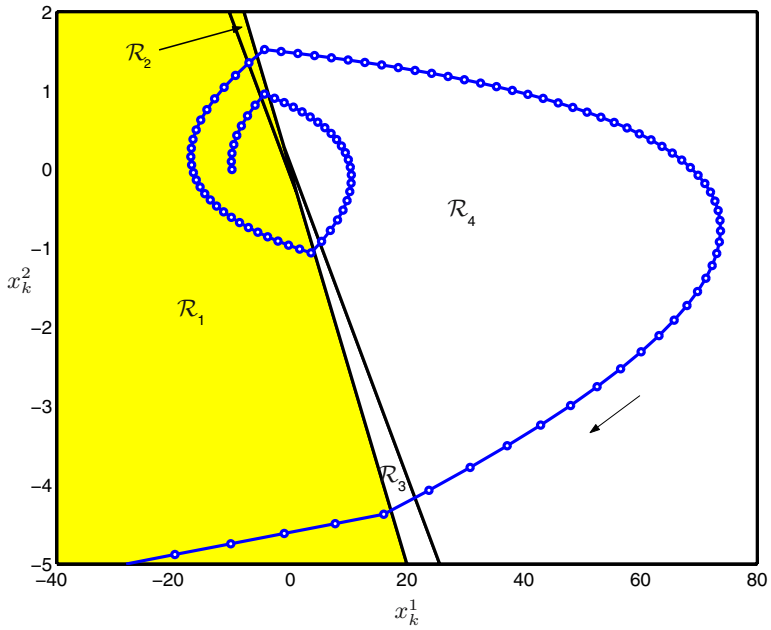

Figure 13.5. State trajectories of the controlled plant (13.28) with initial condition $x = [-10 \quad 0]^{\mathrm{T}}$.

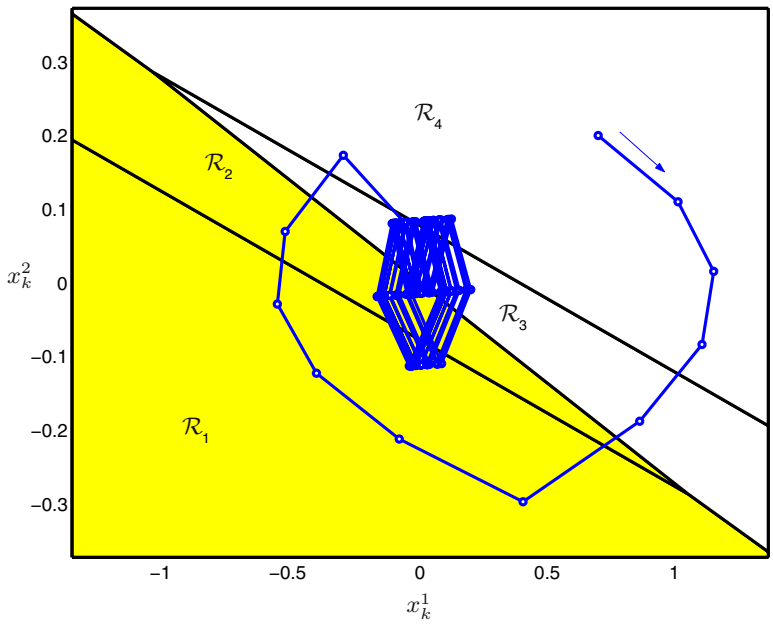

Figure 13.6. State trajectories of the controlled plant (13.28) with initial condition $x = [0.7 \quad 0.2]^{\mathrm{T}}$.

13.9 Finite Alphabet Estimation

As we have seen in Chapter 9, the problems of constrained control and estimation are very similar (differing essentially only with respect to the nature of the boundary conditions). Here we give a brief description of finite alphabet estimation. To fix ideas, we refer to the specific problem of estimating a signal drawn from a given finite alphabet that has been transmitted over a noisy dispersive communication channel.

This problem, which is commonly referred to as one of *channel equalisation*, can be formulated as a fixed-delay maximum likelihood detection problem. The resultant detector estimates each symbol based upon the entire sequence received to a point in time and hence constitutes, in principle, a growing memory structure. In the case of finite impulse response [FIR] channels, the Viterbi algorithm can be used for solving the resultant optimisation problem. However, for more general infinite impulse response [IIR] channels, the complexity of the Viterbi Algorithm is infinite. This is a direct consequence of the requirement to take into account the finite alphabet nature of the transmitted signal, which makes this a hard combinatorial optimisation problem.

In order to address this problem, various simplified detectors of fixed memory and complexity have been proposed. The simplest such scheme is the decision feedback equaliser [DFE] (Qureshi 1985), which is a symbol-by-symbol detector. It basically corresponds to the scheme depicted in Figure 1.10 in Chapter 1. It is a feedback loop comprising linear filters and a scalar quantiser. The DFE is extended and outperformed by more complex multistep detector structures, which estimate channel inputs based upon blocks of sampled outputs of fixed size (see, for example, Williamson, Kennedy and Pulford 1992, Duel-Hallen and Heegard 1989).

In these schemes, decision feedback (also called *genie-aided feedback*) is used to overcome the growing memory problem. The information contained in the sampled outputs received before the block where the constraints are taken into account explicitly is summarised by means of an estimate of the channel state. This estimate is based upon previous decisions, which are assumed to be correct. Not taking into account that various decisions may contain errors can lead to error propagation problems (see also Cantoni and Butler 1976).

Here we show how the idea of the "benevolent genie" can be extended by means of an *a priori* state estimate and a measure of its degree of belief.

13.10 Maximum Likelihood Detection Utilising An A Priori State Estimate

Consider a linear channel (which may include a whitening matched filter and any other pre-filter) with scalar input u_k drawn from a finite alphabet $\mathbb{U}$. The channel output y_k is scalar and is assumed to be perturbed by zero-mean

additive white Gaussian noise n_k of variance r, denoted by $n_k \sim N(0, r)$, yielding the state space model

$$\begin{aligned} x_{k+1} &= Ax_k + Bu_k, \\ y_k &= Cx_k + Du_k + n_k, \end{aligned} \tag{13.29}$$

where $x_k \in \mathbb{R}^n$. The above model may equivalently be expressed in transfer function form as

$$y_k = H(\rho)u_k + n_k, \quad H(\rho) = D + C(\rho I - A)^{-1}B = h_0 + \sum_{i=1}^{\infty} h_i \rho^{-i},$$

where[4]

$$h_0 = D, \quad h_i = CA^{i-1}B, \quad i = 1, 2, \ldots. \tag{13.30}$$

We incorporate an a priori state estimate into the problem formulation. This is achieved as follows:

As described in Section 9.9 of Chapter 9, we fix integers $L_1 \geq 0$, $L_2 \geq 1$ and suppose, for the moment, that

$$x_{k-L_1} \sim N(z_{k-L_1}, P), \tag{13.31}$$

that is, z_{k-L_1} is a given a priori estimate for x_{k-L_1} which has a Gaussian distribution. The matrix P^{-1} reflects the degree of belief in this a priori state estimate. Absence of prior knowledge of x_{k-L_1} can be accommodated by using $P^{-1} = 0$, and decision feedback is achieved by taking $P = 0$, which effectively locks x_{k-L_1} at z_{k-L_1}.

Additionally, we define the vectors

$$\begin{aligned} \mathbf{u}_k &\triangleq \left[u_{k-L_1}\ u_{k-L_1+1} \cdots u_{k+L_2-1}\right]^{\mathrm{T}}, \\ \mathbf{y}_k &\triangleq \left[y_{k-L_1}\ y_{k-L_1+1} \cdots y_{k+L_2-1}\right]^{\mathrm{T}}. \end{aligned}$$

The vector $\mathbf{y}_k$ gathers time samples of the channel output and $\mathbf{u}_k$ contains channel inputs, which are the decision variables of the estimation problem considered here.

The *maximum a posteriori* [MAP] sequence detector, which at time $t = k$ provides an estimate of $\mathbf{u}_k$ and x_{k-L_1} based upon the received data contained in $\mathbf{y}_k$, maximises the probability density function (see Chapter 9 for further discussion)[5]

$$p\left(\begin{bmatrix} \mathbf{u}_k \\ x_{k-L_1} \end{bmatrix} \Big| \mathbf{y}_k\right) = \frac{p\left(\mathbf{y}_k \Big| \begin{bmatrix} \mathbf{u}_k \\ x_{k-L_1} \end{bmatrix}\right) p\left(\begin{bmatrix} \mathbf{u}_k \\ x_{k-L_1} \end{bmatrix}\right)}{p\left(\mathbf{y}_k\right)}, \tag{13.32}$$

[4] ρ denotes the forward shift operator, $\rho v_k = v_{k+1}$, where $\{v_k\}$ is any sequence.

[5] For ease of notation, in what follows we will denote all (conditional) probability density functions by p. The specific function referred to will be clear from the context.

where we have utilised Bayes' rule.

Note that only the numerator of this expression influences the maximisation. Assuming that $\mathbf{u}_k$ and x_{k-L_1} are independent (which is a consequence of (13.29) if u_k is white), it follows that

$$p\left(\begin{bmatrix}\mathbf{u}_k\\ x_{k-L_1}\end{bmatrix}\right) = p\left(x_{k-L_1}\right) p\left(\mathbf{u}_k\right).$$

Hence, if all finite alphabet-constrained symbol sequences $\mathbf{u}_k$ are equally likely (an assumption that we make in what follows), then the MAP detector that maximises (13.32) is equivalent to the following *maximum likelihood* sequence detector

$$\begin{bmatrix}\hat{\mathbf{u}}_k\\ \hat{x}_{k-L_1}\end{bmatrix} \triangleq \underset{\mathbf{u}_k, x_{k-L_1}}{\arg\max} \left\{ p\left(\mathbf{y}_k \middle| \begin{bmatrix}\mathbf{u}_k\\ x_{k-L_1}\end{bmatrix}\right) p\left(x_{k-L_1}\right)\right\}. \tag{13.33}$$

Here,

$$\hat{\mathbf{u}}_k \triangleq \left[\hat{u}_{k-L_1}\ \hat{u}_{k-L_1+1} \cdots \hat{u}_k \cdots \hat{u}_{k+L_2-1}\right]^{\mathrm{T}}, \tag{13.34}$$

and $\mathbf{u}_k$ needs to satisfy the constraint

$$\mathbf{u}_k \in \mathbb{U}^N, \quad \mathbb{U}^N \triangleq \mathbb{U} \times \cdots \times \mathbb{U}, \qquad N \triangleq L_1 + L_2, \tag{13.35}$$

in accordance with the restriction $u_k \in \mathbb{U}$. Our working assumption (see (13.31)) is that the initial channel state x_{k-L_1} has a Gaussian probability density function

$$p\left(x_{k-L_1}\right) = \frac{1}{(2\pi)^{n/2}(\det P)^{1/2}} \exp\left\{\frac{-\|x_{k-L_1} - z_{k-L_1}\|^2_{P^{-1}}}{2}\right\}. \tag{13.36}$$

In order to derive analytic expressions for the other probability density functions in (13.33), we rewrite the channel model (13.29) at time instants $t = k - L_1, k - L_1 + 1, \ldots, k + L_2 - 1$ in block form as

$$\mathbf{y}_k = \Psi \mathbf{u}_k + \Gamma x_{k-L_1} + \mathbf{n}_k.$$

Here,

$$\mathbf{n}_k \triangleq \begin{bmatrix} n_{k-L_1}\\ n_{k-L_1+1}\\ \vdots\\ n_{k+L_2-1}\end{bmatrix}, \quad \Gamma \triangleq \begin{bmatrix} C\\ CA\\ \vdots\\ CA^{N-1}\end{bmatrix}, \quad \Psi \triangleq \begin{bmatrix} h_0 & 0 & \ldots & 0\\ h_1 & h_0 & \ddots & \vdots\\ \vdots & \ddots & \ddots & 0\\ h_{N-1} & \ldots & h_1 & h_0\end{bmatrix}.$$

The entries of Ψ obey (13.30), that is, its columns contain truncated impulse responses of the model (13.29).

Since the noise n_k is assumed Gaussian with variance r, it follows that

$$p\left(\mathbf{y}_k \middle| \begin{bmatrix} \mathbf{u}_k \\ x_{k-L_1} \end{bmatrix}\right) = \frac{1}{(2\pi)^{N/2}(\det R)^{1/2}} \exp\left\{\frac{-\|\mathbf{y}_k - \Psi\mathbf{u}_k - \Gamma x_{k-L_1}\|^2_{R^{-1}}}{2}\right\}, \tag{13.37}$$

where the matrix $R \triangleq \text{diag}\{r, \ldots, r\} \in \mathbb{R}^{N\times N}$.

After substituting expressions (13.36) and (13.37) into (13.33) and applying the natural logarithm, one obtains the sequence detector

$$\begin{bmatrix} \hat{\mathbf{u}}_k \\ \hat{x}_{k-L_1} \end{bmatrix} = \underset{\mathbf{u}_k, x_{k-L_1}}{\arg\min}\, V(\mathbf{u}_k, x_{k-L_1}), \tag{13.38}$$

subject to the constraint (13.35). In (13.38), the objective function V is defined as

$$\begin{aligned} V(\mathbf{u}_k, x_{k-L_1}) &\triangleq \|x_{k-L_1} - z_{k-L_1}\|^2_{P^{-1}} + \|\mathbf{y}_k - \Psi\mathbf{u}_k - \Gamma x_{k-L_1}\|^2_{R^{-1}} \\ &= \|x_{k-L_1} - z_{k-L_1}\|^2_{P^{-1}} + r^{-1}\sum_{j=k-L_1}^{k+L_2-1}(y_j - C\check{x}_j - Du_j)^2, \end{aligned} \tag{13.39}$$

and the vectors $\check{x}_j$ denote predictions of the channel states x_j. They satisfy (13.29), that is,

$$\begin{aligned} \check{x}_{j+1} &= A\check{x}_j + Bu_j \quad \text{for } j = k-L_1, \ldots, k+L_2-1, \\ \check{x}_{k-L_1} &= x_{k-L_1}. \end{aligned} \tag{13.40}$$

Remark 13.10.1. (Notation) Since $\hat{\mathbf{u}}_k$ and $\hat{x}_{k-L_1}$ in (13.33) are calculated using data up to time $t = k + L_2 - 1$, they could perhaps be more insightfully denoted as $\hat{\mathbf{u}}_{k|k+L_2-1}$ and $\hat{x}_{k-L_1|k+L_2-1}$, respectively (see Chapter 9). However, in order to keep the notation simple, we will here avoid double indexing, in anticipation that the context will always allow for correct interpretation. ∘

As a consequence of considering the joint probability density function (13.32), the objective function (13.39) includes a term which allows one to obtain an a posteriori state estimate $\hat{x}_{k-L_1}$ which differs from the a priori estimate z_{k-L_1} as permitted by the confidence matrix P^{-1}.

13.11 Information Propagation

Having set up the fixed horizon estimator as the finite alphabet optimiser (13.38)–(13.40), in this section we show how this information can be utilised as part of a moving horizon scheme. Here we essentially follow the methodology outlined in Section 9.9 of Chapter 9.

13.11.1 Moving Horizon Implementation

Minimisation of the objective function V in (13.39) yields the entire optimising sequence $\hat{\mathbf{u}}_k$ defined in (13.38). However, following our usual procedure, we will utilise a moving horizon approach in which only the *present* value[6]

$$\hat{u}_k^{\text{OPT}} \triangleq \begin{bmatrix} 0_{L_1} & 1 & 0_{L_2-1} \end{bmatrix} \hat{\mathbf{u}}_k, \tag{13.41}$$

will be delivered at the output of the detector.

At the next time instant the optimisation is repeated, providing $\hat{u}_{k+1}^{\text{OPT}}$ and so on. Thus, the data window "slides" (or moves) forward in time. The scheme previews $L_2 - 1$ samples, hence has a decision-delay of $L_2 - 1$ time units.

The *window length* $N = L_1 + L_2$ fixes the complexity of the computations needed in order to minimise (13.39). It is intuitively clear that good performance of the detector can be ensured if N is sufficiently large. However, in practice, there is a strong incentive to use small values for L_1 and L_2, since large values give rise to high complexity in the associated computations to be performed at each time step.

13.11.2 Decision-Directed Feedback

The provision of an a priori estimate, z_{k-L_1}, together with an associated degree of belief via the term $\|x_{k-L_1} - z_{k-L_1}\|_{P^{-1}}^2$ in (13.39) provides a means of propagating the information contained in the data received before $t = k - L_1$. Consequently, an *information horizon* of growing length is effectively obtained in which the computational effort is fixed by means of the window length N.

One possible approach to choose the a priori state estimate is as follows: Each optimisation step provides estimates for the channel state and input sequence (see (13.38)). These decisions can be re-utilised in order to formulate a priori estimates for the channel state x_k. We propose that the estimates be propagated in blocks according to[7]

$$z_k = A^N \hat{x}_{k-N} + M \hat{\mathbf{u}}_{k-L_2},$$

where $M \triangleq \begin{bmatrix} A^{N-1}B & A^{N-2}B & \dots & AB & B \end{bmatrix}$. In this way, the estimate obtained in the previous block is rolled forward. Indeed, in order to operate in a moving horizon manner, it is necessary to store N a priori estimates. This is depicted graphically in Figure 13.7.

13.11.3 The Matrix P as a Design Parameter

Since channel states depend on the finite alphabet input, one may well question the assumption made in Section 13.10 that x_{k-L_1} is Gaussian. (This

[6] The row vector $0_m \in \mathbb{R}^{1\times m}$ contains only zeros.

[7] Since z_k is based upon channel outputs up to time $k-1$, it could alternatively be denoted as $\hat{x}_{k|k-1}$; see also Remark 13.10.1.

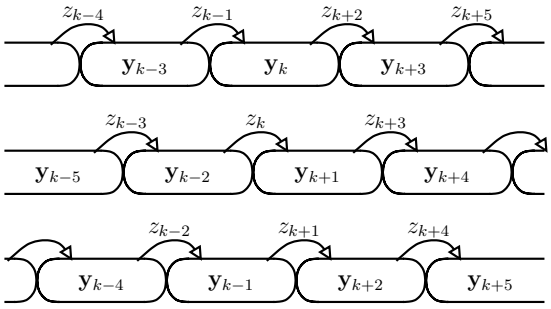

Figure 13.7. Information propagation with parameters $L_1 = 1$ and $L_2 = 2$.

situation is similar to that of other detectors that utilise Gaussian approximations; (see, for example, Lawrence and Kaufman 1971, Thielecke 1997, Baccarelli, Fasano and Zucchi 2000). However, we could always use this structure by interpreting the matrix P in (13.39) as a design parameter.

As a guide for tuning P, we recall that in the unconstrained case, where the channel input and initial state are Gaussian, that is, $u_k \sim N(0, Q)$ and $x_0 \sim N(\mu_0, P_0)$, the Kalman filter provides the minimum variance estimate for x_{k-L_1} (see, for example, Anderson and Moore 1979). Its covariance matrix P_{k-L_1} obeys the Riccati difference equation (see Chapter 9),

$$P_{k+1} = AP_kA^{\mathrm{T}} - K_k(CP_kC^{\mathrm{T}} + r + DQD^{\mathrm{T}})K_k^{\mathrm{T}} + BQB^{\mathrm{T}}, \qquad k \geq 0, \quad (13.42)$$

where $K_k \triangleq (AP_kC^{\mathrm{T}} + BQD^{\mathrm{T}})(CP_kC^{\mathrm{T}} + r + DQD^{\mathrm{T}})^{-1}$.

A further simplification occurs if we replace the recursion (13.42) by its steady state equivalent. In particular, it is well-known (Goodwin and Sin 1984) that, under reasonable assumptions, P_k converges to a steady state value P as $k \to \infty$. The matrix P satisfies the following algebraic Riccati equation:

$$P = APA^{\mathrm{T}} - K(CPC^{\mathrm{T}} + r + DQD^{\mathrm{T}})K^{\mathrm{T}} + BQB^{\mathrm{T}}, \quad (13.43)$$

where $K = (APC^{\mathrm{T}} + BQD^{\mathrm{T}})(CPC^{\mathrm{T}} + r + DQD^{\mathrm{T}})^{-1}$. Of course, the Gaussian assumption on u_k is not valid in the constrained case. However, the choice (13.43) may still provide good performance. Alternatively, one may simply use P as a design parameter and test different choices via simulation studies.

13.12 Closed Loop Implementation of the Finite Alphabet Estimator

Here we follow similar arguments to those used with respect to finite alphabet control in Section 13.3 to obtain a closed form expression for the solution to

the finite alphabet estimation problem. This closed form expression utilises a vector quantiser as defined earlier in Definition 13.3.1.

For general recursive channels, it is useful to assume that, whilst the input is always constrained to a finite alphabet, the channel state x_k in (13.29) is left unconstrained. In this case, the optimisers (13.38) are characterised as follows:

Lemma 13.12.1 (Closed Form Solution) *The optimisers corresponding to* (13.38) *given the constraint* $\mathbf{u}_k \in \mathbb{U}^N$ *are given by*

$$\hat{\mathbf{u}}_k = \Omega^{-1/2} q_{\tilde{\mathbb{U}}^N}\big(\Omega^{-1/2}(\Lambda_1 \mathbf{y}_k - \Lambda_2 z_{k-L_1})\big), \tag{13.44}$$

$$\hat{x}_{k-L_1} = \Upsilon\left(P^{-1} z_{k-L_1} + \Gamma^{\mathrm{T}} R^{-1} \mathbf{y}_k - \Gamma^{\mathrm{T}} R^{-1} \Psi \hat{\mathbf{u}}_k\right), \tag{13.45}$$

where

$$\begin{aligned} \Omega &= \Psi^{\mathrm{T}}\left(R^{-1} - R^{-1}\Gamma\Upsilon\Gamma^{\mathrm{T}}R^{-1}\right)\Psi, \qquad \Omega^{\mathrm{T}/2}\Omega^{1/2} = \Omega, \\ \Upsilon &= (P^{-1} + \Gamma^{\mathrm{T}} R^{-1}\Gamma)^{-1}, \\ \Lambda_1 &= \Psi^{\mathrm{T}}\left(R^{-1} - R^{-1}\Gamma\Upsilon\Gamma^{\mathrm{T}}R^{-1}\right), \\ \Lambda_2 &= \Psi^{\mathrm{T}} R^{-1}\Gamma\Upsilon P^{-1}. \end{aligned} \tag{13.46}$$

The nonlinear function $q_{\tilde{\mathbb{U}}^N}(\cdot)$ *is the nearest neighbour vector quantiser described in Definition 13.3.1. The image of this mapping is the set*

$$\tilde{\mathbb{U}}^N = \Omega^{1/2}\mathbb{U}^N \triangleq \{\tilde{v}_1,\ \tilde{v}_2, \ldots, \tilde{v}_r\} \subset \mathbb{R}^N,\ \textit{with}\ \tilde{v}_i = \Omega^{1/2} v_i,\ v_i \in \mathbb{U}^N. \tag{13.47}$$

Proof. The objective function (13.39) can be expanded as

$$\begin{aligned} V(\mathbf{u}_k, x_{k-L_1}) = {} & \|x_{k-L_1}\|^2_{\Upsilon^{-1}} + \|z_{k-L_1}\|^2_{P^{-1}} + \|\mathbf{y}_k\|^2_{R^{-1}} \\ & + \|\mathbf{u}_k\|^2_{\Psi^{\mathrm{T}} R^{-1}\Psi} + \mathbf{u}_k^{\mathrm{T}}\Psi^{\mathrm{T}}R^{-1}\Gamma x_{k-L_1} + x_{k-L_1}^{\mathrm{T}}\Gamma^{\mathrm{T}}R^{-1}\Psi\mathbf{u}_k \\ & - 2\big[\mathbf{u}_k^{\mathrm{T}}\Psi^{\mathrm{T}}R^{-1}\mathbf{y}_k + x_{k-L_1}^{\mathrm{T}}\big(P^{-1}z_{k-L_1} + \Gamma^{\mathrm{T}}R^{-1}\mathbf{y}_k\big)\big], \end{aligned} \tag{13.48}$$

with Υ defined in (13.46). This expression can be written as

$$\begin{aligned} V(\mathbf{u}_k, x_{k-L_1}) = {} & \alpha(\mathbf{u}_k, \mathbf{y}_k, z_{k-L_1}) + \|x_{k-L_1}\|^2_{\Upsilon^{-1}} \\ & - 2x_{k-L_1}^{\mathrm{T}}\left[P^{-1}z_{k-L_1} + \Gamma^{\mathrm{T}}R^{-1}\mathbf{y}_k - \Gamma^{\mathrm{T}}R^{-1}\Psi\mathbf{u}_k\right], \end{aligned}$$

where $\alpha(\mathbf{u}_k, \mathbf{y}_k, z_{k-L_1})$ does not depend upon x_{k-L_1}.

Since x_{k-L_1} is assumed unconstrained, it follows that, for every fixed value of $\mathbf{u}_k$, the objective function is minimised by means of $x_{\mathrm{UC}}^{\mathrm{OPT}} = \Upsilon\left(P^{-1}z_{k-L_1} + \Gamma^{\mathrm{T}}R^{-1}\mathbf{y}_k - \Gamma^{\mathrm{T}}R^{-1}\Psi\mathbf{u}_k\right)$ from where (13.45) follows.

In order to obtain the constrained optimiser $\hat{\mathbf{u}}_k \in \mathbb{U}^N$, observe that

$$\hat{\mathbf{u}}_k = \arg\min_{\mathbf{u}_k \in \mathbb{U}^N} J(\mathbf{u}_k), \tag{13.49}$$

where $J(\mathbf{u}_k) \triangleq V(\mathbf{u}_k, x_{\mathrm{UC}}^{\mathrm{OPT}})$. Substitution of (13.45) into (13.48) yields

$$J(\mathbf{u}_k) = \beta(\mathbf{y}_k, z_{k-L_1}) + \mathbf{u}_k^{\mathrm{T}}\Omega\mathbf{u}_k - 2\mathbf{u}_k^{\mathrm{T}}(\Lambda_1\mathbf{y}_k - \Lambda_2 z_{k-L_1}), \tag{13.50}$$

where Ω, Λ_1 and Λ_2 are defined in (13.46) and $\beta(\mathbf{y}_k, z_{k-L_1})$ does not depend upon $\mathbf{u}_k$.

As in the proof of Theorem 13.3.1, it is useful to introduce the coordinate transformation $\tilde{\mathbf{u}}_k \triangleq \Omega^{1/2}\mathbf{u}_k$. This transforms $\mathbb{U}^N$ into $\tilde{\mathbb{U}}^N$ defined in (13.47). Equation (13.50) then allows one to rewrite (13.49) as

$$\hat{\mathbf{u}}_k = \Omega^{-1/2} \arg\min_{\tilde{\mathbf{u}}_k \in \tilde{\mathbb{U}}^N} \tilde{J}(\tilde{\mathbf{u}}_k), \tag{13.51}$$

with $\tilde{J}(\tilde{\mathbf{u}}_k) \triangleq \tilde{\mathbf{u}}_k^{\mathrm{T}}\tilde{\mathbf{u}}_k - 2\tilde{\mathbf{u}}_k^{\mathrm{T}}\Omega^{-1/2}(\Lambda_1\mathbf{y}_k - \Lambda_2 z_{k-L_1})$. The level sets of $\tilde{J}$ are spheres in $\mathbb{R}^N$, centred at $\Omega^{-1/2}(\Lambda_1\mathbf{y}_k - \Lambda_2 z_{k-L_1})$. Hence,

$$\arg\min_{\tilde{\mathbf{u}}_k \in \tilde{\mathbb{U}}^N} \tilde{J}(\tilde{\mathbf{u}}_k) = q_{\tilde{\mathbb{U}}^N}\big(\Omega^{-1/2}(\Lambda_1\mathbf{y}_k - \Lambda_2 z_{k-L_1})\big),$$

which, after substituting into (13.51) yields (13.44). □

13.13 Example

Consider an FIR channel described by

$$H(z) = 1 + 2z^{-1} + 2z^{-2}. \tag{13.52}$$

In order to illustrate the performance of the multistep optimal equaliser presented, we carry out simulations of this channel with an input consisting of 10000 independent and equiprobable binary digits drawn from the alphabet $\mathbb{U} = \{-1, 1\}$. The system is affected by Gaussian noise with different variances. The following detection architectures are used: direct quantisation of the channel output, decision feedback equalisation and moving horizon estimation, with parameters $(L_1, L_2) = (1, 2)$ and also with $(L_1, L_2) = (2, 3)$.

Figure 13.8 documents the results. It contains the empirical probabilities of symbol errors obtained at several noise levels. It can be appreciated how moving horizon estimation clearly outperforms direct quantisation of the channel output and also the DFE for this example.

13.14 Conclusions

In this chapter we have presented an approach that addresses control and estimation problems where the decision variables are constrained to belong to a finite alphabet.

It turns out that concepts introduced in previous chapters, namely receding horizon optimisation, exploration of the geometry of the underlying

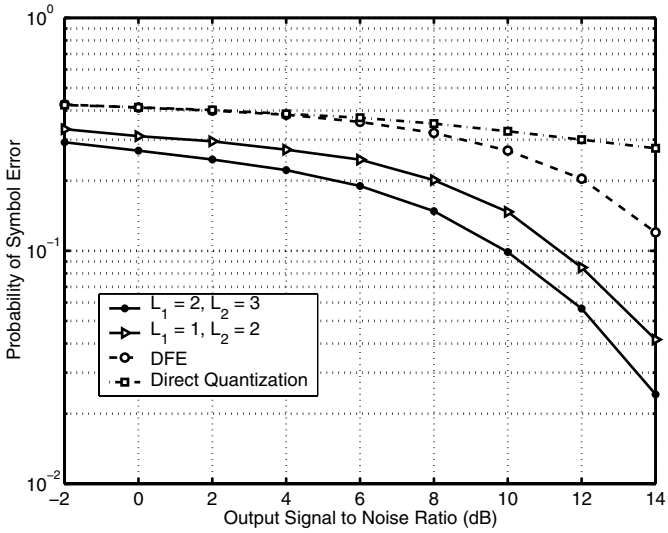

Figure 13.8. Bit error rates of the communication systems simulated.

optimisation problem and information propagation can be readily utilised. It is also apparent that some aspects, such as dynamics and stability of the closed loop, demand for other, more specialised, tools.

Bearing in mind the wide range of applications that can be cast as finite alphabet-constrained control and estimation problems, we invite the reader to apply the acquired expertise in these *nontraditional* areas. The cross-fertilisation of ideas gives new insight and may lead to improved design methodologies in various realms of application.

13.15 Further Reading

General

A more detailed presentation of the ideas outlined in this chapter, including several application studies can be found in Quevedo and Goodwin (2004c). More information on computational complexity of combinatorial optimisation problems can be found in Garey and Johnson (1979). Vector quantisers and Voronoi partitions are described thoroughly in Gersho and Gray (1992), Gray and Neuhoff (1998).

Finite Set Control

In relation to finite set-constrained control problems, Quevedo, Goodwin and De Doná (2004) forms the basis of our presentation. Alternative views are

given, for example, in Brockett and Liberzon (2000), Ishii and Francis (2003), Richter and Misawa (2003), Sznaier and Damborg (1989), and Bicchi, Marigo and Piccoli (2002).

Channel Equalisation

Channel equalisation is an important problem in digital communications. It is described in standard textbooks, such as Proakis (1995) and also in the survey papers Qureshi (1985), Tugnait et al. (2000). The presentation given in this chapter follows basically Quevedo, Goodwin and De Doná (2003) and is also related to the multistep estimation schemes described in Williamson et al. (1992), Duel-Hallen and Heegard (1989).

Other Application Areas

Design of networked control systems based upon the ideas presented in this chapter can be found, for example, in Quevedo, Goodwin and Welsh (2003), Goodwin, Haimovich, Quevedo and Welsh (2004), Kiihtelys (2003) and also in Chapter 16. Other interesting references include Bushnell (ed.) (2001), Wong and Brocket (1999), Ishii and Francis (2002), Zhivoglyadov and Middleton (2003), Hristu and Morgansen (1999), Elia and Mitter (2001). The related problem of state estimation with quantised measurements has also been treated in Curry (1970) and Delchamps (1989), Haimovich, Goodwin and Quevedo (2003). (See also Chapter 16.)

Applications to the design of FIR filters with finite set constrained coefficients can be found in Quevedo and Goodwin (2003b), Goodwin, Quevedo and De Doná (2003). These problems have been studied extensively in the signal-processing literature, see, for example, Evangelista (2002), Lim and Parker (1983b), Kodek (1980), Lim and Parker (1983a).

Audio quantisation and A/D conversion can be dealt with in a way similar to finite alphabet-constrained control problems, as detailed in Goodwin, Quevedo and McGrath (2003), Quevedo and Goodwin (2003a), Quevedo and Goodwin (2004b). Other references include Norsworthy et al. (1997), Lipschitz, Vanderkooy and Wannamaker (1991) and the collection Candy and Temes (1992).

Applications of finite set-constrained control in power electronics abound. One particular case resides in the design of the switching signal for switch-mode power supplies, as described in Quevedo and Goodwin (2004a). The book Rashid (1993) is a good introductory level textbook on power electronics.

Part III

Case Studies