Revisiting the Separation Principle in Stochastic Control

Tryphon T. Georgiou and Anders Lindquist

***Abstract*— The separation principle is the statement that under suitable conditions the design of stochastic control can be divided into two separate problems, one of optimal control with state information and one of filtering. The literature over the past 50 years contains several derivations where subtle difficulties are overlooked and inadmissible shortcuts taken. Other contributions that have established the separation principle under various hypotheses require considerable mathematical sophistication, which makes the ideas difficult to include in standard textbooks. The contribution of the present work is a new set of conditions that are in line with basic engineering thinking and ensure that the separation principle holds. The feedback system is required to be well-posed in the sense that it defines a map between sample paths, representing signals rather than stochastic processes *per se*. This approach allows certain generalizations of the separation theorem to a wide class of feedback laws, models and stochastic noise, including martingales with possible jumps.**

I. Introduction

The separation principle of stochastic control — the fact that the problems of optimal control and state estimation can be decoupled in certain cases — was discovered in the early 1960's, and the term was coined in [12], [22]. This is also closely connected to the idea of certainty equivalence; see, e.g., [28]. Since then, a constant stream of accounts have appeared in the literature attempting to identify the most general setting where the principle is valid.

While the separation principle has been established under various conditions, rigorous treatments tend to require considerable mathematical sophistication that are difficult to present in a classroom setting. The purpose of this work is to present a framework for the separation principle which is more in line with basic engineering thinking. In this, signals "travel" around feedback loops and the feedback equations define maps between signal spaces rather than stochastic processes *per se*. Besides a cleaner presentation of the classical separation principle, our approach allows for a considerable generalization to the case where the driving noise is a martingale that does not need to be Gaussian and can have jumps.

The outline is as follows. In Section II we review the standard quadratic regulator problem and prove it for linear control laws. In Section III we point out the challenges and subtleties that arise when one allows the control laws to be nonlinear. In Section IV we review previous treatments in the literature and point out certain shortcomings. In Section V we present our framework and main results.

This research was supported by grants from AFOSR (FA9550-12-1-0319), NSF (1027696), VR, SSF, the Göran Gustafsson Foundation, the Vincentine Hermes-Luh Endowment, and the Thousand Talent Program.

T. T. Georgiou is with the Dept. of Electrical & Computer Engineering, Univ. of Minnesota, Minneapolis, Minnesota; `tryphon@umn.edu`

A. Lindquist is with the Department of Automation, Shanghai Jiao Tong Univerity, Shanghai, China, the Center for Industrial and Applied Mathematics (CIAM), and the ACCESS Linnaeus Center, Royal Institute of Technology, 100 44 Stockholm, Sweden; `alq@kth.se`

II. The separation principle in its basic form

Consider the linear stochastic system

$$\begin{cases} dx = A(t)x(t)dt + B_1(t)u(t)dt + B_2(t)dw \\ dy = C(t)x(t)dt + D(t)dw \end{cases} \tag{1}$$

where x is the state process, y the output process, and u the control, while $w(t)$ is a vector-valued Wiener process and $x(0)$ is a zero-mean Gaussian random vector independent of $w(t)$. Moreover, $y(0) = 0$, and A, B_1, B_2, C, D are continuous matrix-valued functions of compatible sizes and bounded variation. We take DD' to be nonsingular on the interval $[0, T]$. Finally, if we want the noise processes in the state and output equations to be independent, as often is assumed but not required here, we take $B_2D' \equiv 0$.

Next consider the control problem to design an output feedback law

$$\pi \,:\, y \mapsto u \tag{2}$$

over the window $[0, T]$ which specifies the control input u based on the observation process y in a nonanticipatory manner so as to minimize the functional

$$\begin{aligned} J(u) = E\Bigg\{ &\int_0^T x(t)'Q(t)x(t)dt \\ &+ \int_0^T u(t)'R(t)u(t)dt + x(T)'Sx(T) \Bigg\}. \end{aligned} \tag{3}$$

Here, Q and R are continuous matrix-valued functions of bounded variation, $Q(t)$ is positive semi-definite and $R(t)$ is positive definite for all values of t. The separation principle is the statement that, under suitable conditions to be discussed in this paper, the optimal feedback law is linear in the data and given by

$$u(t) = K(t)\hat{x}(t), \tag{4}$$

where $\hat{x}(t)$ is obtained by the Kalman filter

$$\begin{aligned} d\hat{x} &= A(t)\hat{x}(t)dt + B_1(t)u(t)dt \\ &\quad + L(t)(dy - C(t)\hat{x}(t)dt) \\ \hat{x}(0) &= 0 \end{aligned} \tag{5}$$

while the gains

$$K(t) = -R(t)^{-1}B_1(t)'P(t), \text{ and} \tag{6}$$

$$L(t) = \big(\Sigma(t)C(t)' + B_2(t)D(t)'\big)R(t)^{-1} \tag{7}$$

are obtained by solving the pair of dual Riccati equations

$$\begin{cases} \dot{P} = -A'P - PA - Q + PB_1R^{-1}B_1'P \\ P(T) = S \end{cases} \tag{8}$$

$$\begin{cases} \dot{\Sigma} = A\Sigma + \Sigma A' + B_2B_2' \\ \qquad - (\Sigma C' + B_2D')R^{-1}(\Sigma C' + B_2D')' \\ \Sigma(0) = E\{x(0)x(0)'\} \end{cases} \tag{9}$$

A standard approach establishing the separation principle is based on a completion-of-squares argument similar to the one used in deterministic linear-quadratic-regulator theory; see e.g. [1]. We briefly review this construction as this is central to our theme. Itô's differential rule (see, e.g., [13], [21]) gives

$$d(x'Px) = x'\dot{P}xdt + 2x'Pdx + \mathrm{tr}(B_2'PB_2)dt,$$

where $\mathrm{tr}(M)$ denotes the trace of a matrix M. Then, using (1) and (8),

$$d(x'Px) = [-x'Qx - u'Ru + (u - Kx)'R(u - Kx)]dt \\ + \mathrm{tr}(B_2'PB_2)dt + 2x'PB_2dw.$$

By integrating from 0 to T and taking expectation, we obtain

$$J(u) \;=\; E\left\{x(0)'P(0)x(0) + \int_0^T (u - Kx)'R(u - Kx)dt\right\} + \int_0^T \mathrm{tr}(B_2'PB_2)dt. \tag{10}$$

Now, if complete state information is available, (1) is replaced by

$$\begin{cases} dx = A(t)x(t)dt + B_1(t)u(t)dt + B_2(t)dw \\ y = x \end{cases} \tag{11}$$

and, since the last term in (10) does not depend on the control, we immediately conclude that the feedback law

$$u(t) = K(t)x(t) \tag{12}$$

is optimal.

However, in general the control is a function of the observed process $\{y(\tau); 0 \leq \tau \leq t\}$, and we need to carry the analysis one step further. To this end, first note that $u = \pi(y)$ is adapted to the *filtration*

$$\mathcal{Y}_t := \sigma\{y(\tau), \tau \in [0, t]\}, \quad t \in [0, T], \tag{13}$$

generated by the output process; i.e., the family of increasing sigma fields representing the data. Now define

$$\hat{x}(t) := E\{x(t) \mid \mathcal{Y}_t\}, \tag{14}$$

and

$$\tilde{x}(t) := x(t) - \hat{x}(t). \tag{15}$$

The components of $\tilde{x}(t)$ are orthogonal to $\mathcal{Y}_t$-measurable random variables in the sense of the inner product $\langle \xi, \eta \rangle = E\{\xi\eta\}$. Hence, since u and $\hat{x}$ are functions of the data,

$$E\{[u(t) - K(t)\hat{x}(t)]\tilde{x}(t)'\} = 0.$$

Therefore

$$E\int_0^T (u - Kx)'R(u - Kx)dt \tag{16}$$
$$= E\int_0^T [(u - K\hat{x})'R(u - K\hat{x}) + \mathrm{tr}(K'RK\Sigma)]dt,$$

where Σ is the covariance matrix

$$\Sigma(t) = E\{\tilde{x}(t)\tilde{x}(t)'\}. \tag{17}$$

Consequently, in view of (10), it would follow that (4) is the optimal control law provided we can show that Σ is independent of the choice of control law.

The state process can be written as

$$x(t) = x_0(t) + \int_0^t \Phi(t, s)B_1(s)u(s)ds, \tag{18}$$

where x_0 is the state process of the uncontrolled system

$$\begin{cases} dx_0 = A(t)x_0(t)dt + B_2(t)dw \\ dy_0 = C(t)x_0(t)dt + D(t)dw \end{cases} \tag{19}$$

obtained from (1) by setting the control u identically equal to zero, and $\Phi(t, s)$ is the transition matrix function of (1). Then

$$\hat{x}(t) = \hat{x}_0(t) + \int_0^t \Phi(t, s)B_1(s)u(s)ds \tag{20}$$

with

$$\hat{x}_0(t) := E\{x_0(t) \mid \mathcal{Y}_t\}. \tag{21}$$

since u is a function of the data and hence adapted to $\mathcal{Y}_t$. Therefore

$$\tilde{x}(t) = \tilde{x}_0(t) := x_0(t) - \hat{x}_0(t), \tag{22}$$

as the control terms cancel.

We now complete the argument in the *special case* where the admissible control laws are restricted to those in the linear class

$$(\mathcal{L}) \quad u(t) = \bar{u}_0(t) + \int_0^t F(t, \tau)dy, \tag{23}$$

where $\bar{u}$ is a deterministic function and F is an L_2 kernel. In this case, as shown in the appendix,

$$\mathcal{Y}_t = \mathcal{Y}_t^0 = \sigma\{y_0(\tau), \tau \in [0, t]\}, \quad t \in [0, T]. \tag{24}$$

Therefore

$$\hat{x}_0(t) = E\{x_0(t) \mid \mathcal{Y}_t^0\} \tag{25}$$

does not depend on the control law and is generated by the Kalman filter

$$d\hat{x}_0 = A\hat{x}_0dt + L(dy_0 - C\hat{x}_0dt), \quad \hat{x}_0(0) = 0. \tag{26}$$

This together with (20) and (22) yields the Kalman filter (5) for the controlled process as well as

$$d\tilde{x} = (A - LC)\tilde{x}dt + (B_2 - LD)dw, \quad \tilde{x}(0) = x(0),$$

from which we readily derive the Riccati equation (9).

Consequently, the separation principle is valid for all control laws in $\mathcal{L}$ since obviously Σ does not depend on the particular choice. The remainder of the paper considers versions of the separation principle where nonlinear feedback is admissible.

III. Nonlinear control laws: a delicate point

In general the separation principle is stated for control laws that are allowed to be nonlinear. In this generality, a frequent mistake in the literature is to assume *without further investigation* that Σ defined by (17) does not depend on the choice of control. Indeed, if this were the case, it would follow directly that (10) is minimized by choosing the control as (4), and the proof of the separation principle would be immediate as in Section II. This mistake appears to be due to the fact that the control term in (18) cancels when forming (15). As we have seen in Section II such a conclusion would require that $\hat{x}_0$ given by (21) does not depend on the choice of control law. This would follow if the filtrations $\mathcal{Y}_t$ and $\mathcal{Y}_t^0$ were equal as exhibited in (24), and needs to be proven in the case where nonlinear control laws are considered. A detailed discussion of this dilemma can be found in [19].

Given a nonlinear control law $u = \pi(y)$, the output process

$$dy = dy_0 + \int_0^t C(t)\Phi(t,s)B_1(s)u(s)dsdt, \qquad (27)$$

can be written as

$$dy = \varphi(t;\, y(s), 0 \leq s \leq t)dt + dy_0 \qquad (28)$$

where φ is the non-anticipatory map

$$\varphi(t,y) = \int_0^t C(t)\Phi(t,s)B_1(s)\pi(y)(s)ds. \qquad (29)$$

From (28) we readily see that

$$\mathcal{Y}_t^0 \subset \mathcal{Y}_t, \quad t \in [0,T]. \qquad (30)$$

For the separation principle to hold we need (24). The reverse containment needs to be proven. The potential pitfalls are underscored by a celebrated example by Tsirel'son [27] (also see [2], [8, p. 298], [24, Section V.18]) taking as functional

$$\varphi(t,y) = \sum_{k<0} \left\{ \frac{y(t_k) - y(t_{k-1})}{t_k - t_{k-1}} \right\} \mathbf{1}_{(t_k, t_{k+1}]}(t) \qquad (31)$$

in (28) instead, where $(t_k)_{k\leq 0}$ satisfies $t_k < t_{k+1}$, $\lim_{k\to-\infty} t_k = 0$, $\{\lambda\}$ denotes the fractional part $\lambda \in \mathbb{R}$, $\mathbf{1}_S$ is the indicator function of the set S, and y_0 is a Wiener process. Tsirel'son showed that for the functional (31), the stochastic differential equation (28) does not have strong solutions and that (30) is a proper inclusion (see also [24, p. 156, Theorem 18.3]). Tsirel'son's example is not in the form (29) that arises in the feedback problem, but it still underscores the potential subtlety in showing the reverse inclusion in general.

To avoid these problems one might begin by uncoupling the feedback loop and determine an optimal control process in the class of stochastic processes u that are adapted to the family of (uncontrolled) sigma fields $(\mathcal{Y}_t^0)_{t\in[0,T]}$. Such a problem, where one optimizes over the class of all control processes adapted to a fixed filtration, was called a *stochastic open loop (SOL) problem* in [19]. In the literature on the separation principle it is not uncommon to assume from the outset that the control is adapted to $(\mathcal{Y}_t^0)_{t\in[0,T]}$; see, e.g., [4, Section 2.3], [11], [30].

In [19] it was suggested how to embed various SOL classes in a problem-dependent manner, and then construct the corresponding feedback law. More precisely, the class of admissible feedback laws was taken to consist of the nonanticipatory functions $u := \pi(y)$ such that the feedback equations have a unique solution with an output process y_π and $u = \pi(y_\pi)$ adapted to $\{\mathcal{Y}_t^0\}$. In the next section we shall give a few examples.

IV. Historical remarks

We have already discussed the case where the control law is a linear function of the data in Section II. An early treatment can be found in [5] where control laws are restricted to being finite-dimensional compensators. For the more general class of linear control laws $\mathcal{L}$, see [19].

Regarding the case of nonlinear control laws we begin with a construction due to Kushner [15]. Let

$$\hat{\xi}_0(t) := E\{x_0(t) \mid \mathcal{Y}_t^0\},$$

be the Kalman state estimate of the uncontrolled system (19), where we use the notation $\hat{\xi}_0$ to distinguish it from $\hat{x}_0$, defined by (21), which *a priori* might depend on the control. Then the corresponding Kalman filter is

$$d\hat{\xi}_0 = A\hat{\xi}_0(t)dt + L(t)dv_0,\ \hat{\xi}_0(0) = 0$$

where the filtration $(\mathcal{V}_t^0)$ of the innovation process

$$dv_0 = dy_0 - C\hat{\xi}_0(t)dt,\ v_0(0) = 0$$

is the same as that of y_0; i.e., $\mathcal{V}_t^0 = \mathcal{Y}_t^0$ for $t \in [0,T]$. In analogy with (18), we define

$$\hat{\xi}(t) = \hat{\xi}_0(t) + \int_0^t \Phi(t,s)B_1(s)u(s)ds,$$

taking

$$u(t) = \psi(t, \hat{\xi}(t)) \qquad (32)$$

with $\psi(t,x)$ Lipschitz in x and $\hat{\xi}$ the unique strong solution of the stochastic differential equation

$$d\hat{\xi} = \big(A\hat{\xi}(t) + B_1\psi(t,\hat{\xi}(t))\big)dt + L(t)dv_0,\ \hat{\xi}(0) = 0. \qquad (33)$$

Clearly $\hat{\xi}$ is adapted to $(\mathcal{V}_t^0)$ and hence to $(\mathcal{Y}_t^0)$; see, e.g., [13, p. 120]. Hence the selection (32) of control law forces u to be adapted to $(\mathcal{Y}_t^0)$, and hence, due to (27), $\mathcal{Y}_t \subset \mathcal{Y}_t^0$ for $t \in [0,T]$. Since the control-dependent terms cancel,

$$dv_0 = dy_0 - C\hat{\xi}_0(t)dt = dy - C\hat{\xi}(t)dt,$$

which inserted into (33) yields a stochastic differential equation, obeying the appropriate Lipschitz condition, driven by dy and having $\hat{\xi}$ as a strong solution. Therefore, $\hat{\xi}$ is adapted to $\{\mathcal{Y}_t\}$, and hence, by (32), so is u. Consequently, (27) implies that $\mathcal{Y}_t^0 \subset \mathcal{Y}_t$ for $t \in [0,T]$ so that actually $\mathcal{Y}_t = \mathcal{Y}_t^0$. Finally, this implies that $\hat{\xi} = \hat{x}$, and thus u is given by

$$u(t) = \psi(t, \hat{x}(t)). \qquad (34)$$

However, it should be noted that the class of control laws (32) is a subclass of (34) as it has been constructed to make u *a priori* adapted to $\{\mathcal{Y}_t^0\}$. Therefore, the relevance of these results, presented in [15], for the proof in [16, page 348] is unclear. In their popular textbook [14], widely used as a reference source for the validity of the separation principle over a general class of admissible (including nonlinear) controls, Kwakernaak and Sivan prove the separation principle over a class of linear laws but claim with reference to [16], [15] that it holds "without qualification" in general [14, p. 390].

In his pioneering paper [31], Wonham proved the separation theorem for the class of control laws (34) and also for a more general cost functional than (3). However, his proof is far from simple and marred by many technical assumptions. A case in point is the assumption that $C(t)$ is square and has a determinant bounded away from zero, which is a serious restriction. A subsequent proof by Fleming and Rishel [9] for quadratic cost functionals is considerably simpler and applies to control laws $u = \varphi(t; y(\tau), 0 \leq \tau \leq t)$ that are Lipschitz in y.

We note that if there is a delay $\varepsilon > 0$ in the processing of the observed data so that, for each t, $u(t)$ is a function of $\{y(\tau), 0 \leq \tau \leq t-\varepsilon\}$, then $(\mathcal{Y}_t) = (\mathcal{Y}_t^0)$. To see this, let n be a positive integer, and suppose that $\mathcal{Y}_t = \mathcal{Y}_t^0$ for $t \in [0, n\varepsilon]$. Then, in view of (27) and the fact that $u(t)$ is $\mathcal{Y}_{t-\varepsilon}^0$- measurable on $[0,(n+1)\varepsilon]$, $y(t)$ is $\mathcal{Y}_t^0$-measurable on the same interval, and hence $\mathcal{Y}_t = \mathcal{Y}_t^0$ for $t \in [0,(n+1)\varepsilon]$. Since $\mathcal{Y}_t = \mathcal{Y}_t^0$ for $t \in [0,\varepsilon]$, the claim follows by induction. This observation shows why control-dependent sigma fields do not occur in the usual discrete-time formulation. In contrast, careful analysis is needed to rule out that the same is true in continuous-time. This point is overlooked in several textbooks (see, e.g., [26]) where a continuous-time Σ is constructed as limits of finite difference quotients of a discrete-time one, which does not depend on the control and is the solution of a discrete-time matrix Riccati equation. However, we cannot *a priori* conclude that the continuous-time Σ satisfies this Riccati equation. For this to be true, $(\mathcal{Y}_t) = (\mathcal{Y}_t^0)$ is needed, otherwise such an argument is circular.

A popular viewpoint in Duncan and Varaiya [7] and Davis and Varaiya [6] relies on weak solutions. It is tailored to the case of a Brownian input process and utilizes the Girsanov transformation for the purpose of avoiding control dependence of the filtration $(\mathcal{Y}_t)$ [4, Section 2.4]. By an appropriate change of probability measure,

$$d\tilde{w} = B_1 u dt + B_2 dw$$

can be transformed into a (weighted) Brownian motion process, which in the sense of weak solutions [13, page 128] is the same as any other Brownian motion process. In this way, the filtration $(\mathcal{Y}_t)$ can be fixed to be constant with respect to variations in the control. The engineering interpretation of this scheme is unclear.

Yet another approach to the separation principle is based on the fact that although (1) with a nonlinear control is non-Gaussian, the model is conditionally Gaussian given the filtration $(\mathcal{Y}_t)$ [20, Chapters 16.1]. This fact can be used to show that $\hat{x}$ is actually generated by a Kalman filter [20, Chapters 11 and 12]. This last approach requires a lengthy and sophisticated analysis and applies only to the case where the driving noise w is a Wiener process.

In the sequel, we take a viewpoint which is more in line with engineering thinking. We consider control laws that render the feedback equations well posed in the sense that they represent non-anticipatory measurable maps between all sample paths of the various processes.

V. Main results

Following [18], [19] we rewrite the model in (1) in an integrated form which allows similar conclusions for more general classes of linear systems. Setting

$$z(t) = \begin{pmatrix} x(t) \\ y(t) \end{pmatrix},$$

the system (1) can be cast in the form

$$\begin{cases} z(t) = z_0(t) + \int_0^t G(t,\tau)u(\tau)d\tau \\ y(t) = Hz(t), \end{cases} \tag{35}$$

where G is a Volterra kernel. This is shown in Figure 1, where g represents the Volterra operator

$$g \;:\; (t,u) \mapsto \int_0^t G(t,\tau)u(\tau)d\tau, \tag{36}$$

and where H is a constant matrix. As usual, Figure 1 is a graphical representation of the algebraic relationship

$$z = z_0 + g\pi Hz. \tag{37}$$

In the stochastic system (1) $H = [0, I]$, but H could be any matrix or linear system. Setting $z = x$ and $H = I$ we obtain the special case of complete state information.

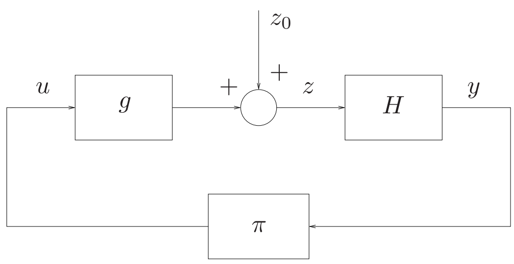

Fig. 1. A feedback interconnection.

In the sequel, we take the point of view that for any natural class of feedback laws (2), the function π should act on sample paths of the stochastic process y rather than on the process itself. Likewise, feedback systems are thought of as maps between sample paths which are *well-posed* in the sense that the feedback equations admit a unique solution which causally depends on the input. This formulation is substantially different from the probabilistic viewpoint which focuses on sample paths in the complement of a zero measure set, whereas we insist instead on the engineering viewpoint that signals (rather than processes) drive a feedback system.

A. Signals and systems

Signals are thought of as sample paths of a stochastic process with possible discontinuities. This is quite natural from several points of view. First, it encompasses the response of a typical nonlinear operation that involves thresholding and switching, and second, it includes sample paths of counting processes and other martingales. More specifically we consider signals to belong to the Skorohod space D; this is defined as the space of functions which are continuous on the right and have a left limit at all points, i.e., the space of càdlàg functions.[1] It contains the space C of continuous functions as a proper subspace.

Systems are thought of as general measurable non-anticipatory maps from $D \to D$, sending sample paths to sample paths so that the output at any given time t is a measurable function of past values of the input and of time. An important class of systems is provided by stochastic differential equations that have strong solutions. Strong solutions induce such maps between corresponding path spaces. In particular, semimartingales have sample paths in D and, under fairly general conditions (see e.g., [23, Chapter V]), stochastic differential equations driven by martingales have strong solutions who are themselves semimartingales.

Besides stochastic differential equations in general, and those in (35) in particular, other nonlinear maps may serve as systems. For instance, discontinuous hystereses nonlinearities with continuous inputs as well as non-Lipschitz static maps are reasonable as systems from an engineering viewpoint. Indeed, these induce maps from $D \to D$ (or from $C \to D$, as in the case of relay hysteresis) and can be considered as components of nonlinear feedback laws. These typify systems that need be considered as an option in control laws when establishing separation theorems.

B. Sample-path well-posedness

The question of well-posedness of feedback systems has been studied from different angles for over forty years (see, e.g., [29]). In our present setting of stochastic control we need a concept of well-posedness which ensures that signals inside a feedback loop are causally dependent on external inputs.

Definition 1: The feedback system depicted in Figure 2 is *(sample-pathwise) well-posed* if the closed-loop maps are themselves systems; i.e., the feedback equation

$$z = z_0 + f(z)$$

has a unique solution z for inputs z_0 and the operator $(1-f)^{-1}$ is itself a system.

Thus, now thinking about z_0 and z in the feedback system in Figure 2 as stochastic processes, well-posedness implies that $\mathcal{Z}_t \subset \mathcal{Z}_t^0$ for $t \in [0,T]$, where $\mathcal{Z}_t$ and $\mathcal{Z}_t^0$ are the sigma-fields generated by z and z_0, respectively. This is a consequence of the fact that $(1-f)^{-1}$ is a system. Likewise, since $(1-f)$ is also a system, $\mathcal{Z}_t^0 \subset \mathcal{Z}_t$ so that in fact

$$\mathcal{Z}_t^0 = \mathcal{Z}_t, \quad t \in [0,T]. \tag{38}$$

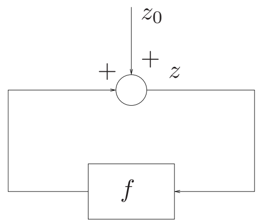

Fig. 2. Basic feedback system.

Next we consider the situation in Figure 1 and the relation between $\mathcal{Y}_t$ and the filtration $\mathcal{Y}_t^0$ of the process $y_0 = Hz_0$. The latter represents the "uncontrolled" output process where the control law π is taken to be identically zero. A key technical lemma for what follows states that the filtrations $\mathcal{Y}_t$ and $\mathcal{Y}_t^0$ are also identical. This is not obvious at first sight, solely on the basis of the linear relationships $y = Hz$ and $y_0 = Hz_0$, as the following simple example demonstrates: the two vector processes $\binom{w}{0}$ and $\binom{0}{w}$ generate the same filtrations while $(1 \quad 0)\binom{w}{0}$ and $(1 \quad 0)\binom{0}{w}$ do not.

Lemma 2: Assume that $g\pi$ is a system, $H = \begin{bmatrix} 0 & I \end{bmatrix}$, and the feedback interconnection in Figure 1 is well-posed. Then $(1-Hg\pi)^{-1}$ is a system and $\mathcal{Y}_t = \mathcal{Y}_t^0$ for $t \in [0,T]$.

The conditions of the lemma can be further relaxed to requiring that H be a linear dynamical system having a right inverse which is itself a system. The idea of the proof, which is given in detail in [10], is to show that provided $(1-g\pi H)^{-1}$ is a system (assumed by well-posedness), then $(1-Hg\pi)^{-1}$ exists and is a system as well. The essence of the lemma is to underscore the equivalence between the configuration in Figure 1 and that in Figure 3. It is this equivalence which accounts for the identity $(\mathcal{Y}_t) = (\mathcal{Y}_t^0)$ between the respective filtrations. An analogous notion of well-posedness was considered by Willems in [30] where however, in contrast, the well-posedness of the feedback configuration in Figure 3, and consequently the validity of $\mathcal{Y}_t = \mathcal{Y}_t^0$, is assumed at the outset.

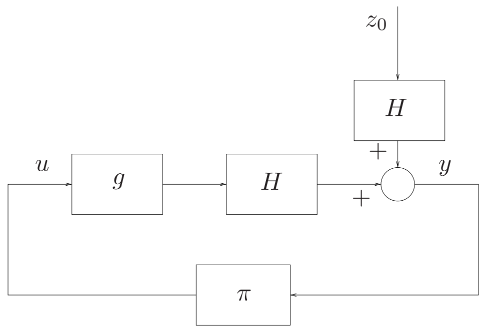

Fig. 3. An equivalent feedback configuration.

[1] "continue à droite, limite à gauche" in French, alternatively RCLL ("right continuous with left limits") in English.

All reasonable feedback laws from an engineering standpoint must render the feedback system well-posed, which we now formally define.

Definition 3: A feedback law π is *well-posed* for the system (35) if $g\pi$ is a system and the feedback loop of Figure 1 is well-posed.

Thus, if the feedback law π is well-posed, then, by Lemma 2, the feedback loop in Figure 3 is also well-posed.

Remark 4: For pedagogical reasons, it is worth considering the case of complete state information, as in (11), i.e., take $H = I$ and $z = x$, with the feedback loop depicted in Figure 2 and z, z_0 replaced by x, x_0, respectively. Then well-posedness (38) amounts to the filtration $(\mathcal{X}_t)$, where $\mathcal{X}_t := \sigma\{x(s);\ s \in [0,T]\}$, being independent of the control. Therefore, Lemma 2 is not needed to resolve the circular control dependence. This is consistent with the analysis leading up to (12).

C. The separation principle

Our first theorem is a very general statement of the separation principle admitting nonlinear control laws for the classical stochastic control problem stated in Section II.

Theorem 5: Given the system (1), consider the problem of minimizing the functional (3) over the class of all feedback laws π that are well-posed. Then the unique optimal control law is given by (4) and $\hat{x}$ is given by the Kalman filter (5).

The proof is based on Lemma 2 and the analysis in Section II. In fact, equality of the filtrations $(\mathcal{Y}_t)$ and $(\mathcal{Y}_t^0)$ implies that (21) does not depend on the choice of control law and hence neither does (17). Therefore, in view of the analysis in Section II, (4) is the unique optimal control provided it defines a well-posed control law. A detailed proof is given in [10], where well-posedness is also proven.

Thus, for a system driven by a Wiener process with Gaussian initial condition, the linear control law defined by (4) and (5) is optimal in the class of all linear and nonlinear control laws for which the feedback system makes sense as a map between path spaces of signals. There is no need for any Lipschitz condition or change of measures.

Interestingly, in this framework, if we dispense with the requirement that $\hat{x}$ is given by a Kalman filter, we can allow x_0 to be non-Gaussian and w to be an arbitrary martingale, even allowing jumps. This is stated next, and the proof is given in [10].

Theorem 6: Given the system (1), where w is a martingale and $x(0)$ is an arbitrary zero-mean random vector independent of w, consider the problem to minimize the functional (3) over the class of all feedback laws π that are well-posed for (1). The control law given by (4) with $\hat{x}$ being the conditional mean (14) is the unique optimal control provided it is well-posed.

Clearly, although the feedback law in Theorem 6 is linear in $\hat{x}$, computing $\hat{x}$ is in general a nonlinear operation that requires a nonlinear filter. Here, well-posedness is not guaranteed by the theorem and needs to be verified separately.

VI. Concluding remarks

The central technical point of the separation principle in continuous time is the dichotomy between the information content of the controlled and the uncontrolled observation processes. In fact, Lemma 2 raises the following question. Given that $z = z_0 + g\pi H z$ has a unique strong solution, does $y = y_0 + Hg\pi y$ admit a unique strong solution as well? If $\mathcal{Y}_t \neq \mathcal{Y}_t^0$, then y is only a weak solution. A celebrated example of Tsirel'son [3] shows that there are not always strong solutions. Hence there may be nonlinear feedback laws that are not admissible for this reason. In our formulation we restrict the class of admissible feedback laws a bit more. Whereas in stochastic theory equations need only be satisfied with probability one, our definition of well-posedness requires that the equations represent (measurable) maps which are well defined for every sample path. Hence, insisting on well-posed feedback loops and on signals as sample functions comes at a certain price.

The present paper is based on [10] which contains further technical details, examples, and a detailed treatment of more general linear systems with possible time-delays, generalizing [18], [19]. During the review process of [10] an anonymous referee has provided us with important insightful suggestions and corrections which have helped us improve the paper.

VII. Appendix: linear control laws

Herein we present the statement that the controlled and uncontrolled filtrations $(\mathcal{Y}_t)$ and $(\mathcal{Y}_t^0)$ coincide for linear control laws, i.e., for $\pi \in \mathcal{L}$. We do this in the general framework of Section V.

Lemma 7: Consider the stochastic system (35), and let $\mathcal{L}$ be defined by (23). Then, $\mathcal{Y}_t$ equals $\mathcal{Y}_t^0$ for any $\pi \in \mathcal{L}$.

Proof: For simplicity and without loss of generality we let $\bar{u}_0 = 0$ in (23). Thus,

$$u(t) = \int_0^t F(t,\tau)dy$$

which we now substitute into (27) to obtain an expression of the form

$$dy = dy_0 + \int_0^t N(t,\tau)dy(\tau)dt. \tag{39}$$

Define the Volterra resolvent which is the unique solution of

$$R(t,\tau) = \int_\tau^t R(t,s)N(s,\tau)ds + N(t,s),$$

see, e.g., [25], [32]. Then, it can be shown that

$$\int_0^t N(t,\tau)dy(\tau) = \int_0^t R(t,\tau)dy_0(\tau),$$

which together with (39) yields

$$dy = dy_0 + \int_0^t R(t,\tau)dy_0(\tau)dt. \tag{40}$$

Then, $\mathcal{Y}_t^0 \subset \mathcal{Y}_t$ follows from (39), and $\mathcal{Y}_t \subset \mathcal{Y}_t^0$ follows from (40). ∎

REFERENCES

[1] K.J. Åström, *Introduction to Stochastic Control Theory*, Academic Press, 1970.

[2] A. Bain and D. Crisan, *Fundamentals of Stochastic Filtering*, Springer, 2007.

[3] V.E. Beneš, "Nonexistence of strong nonanticipating solutions to stochastic DEs: implications for functional DEs, filtering, and control," *Stoch. Proc. Appl.*, **5 (3)**: 243263, 1977.

[4] A. Bensoussan, *Stochastic Control of Partially Observable Systems*, Cambridge University Press, 1992.

[5] M.H.A. Davis, *Linear Estimation and Stochastic Control*, Chapman and Hall, 1977.

[6] M.H.A. Davis and P. Varaiya, "Information states for stochastic systems", *J. Math. Analysis and Applications* **37** (1972), 384–402.

[7] T. Duncan and P. Varaiya, "On the solutions of a stochastic control system, " *SIAM J. Control* **9(3)** (1971), 354-371.

[8] M. Emery and W. Schachermayer, "A remark on Tsirelson's stochastic differential equation," *Séminaire de probabilités XXXIII*, 291–303, 1999.

[9] W.H. Fleming and R.W. Rishel, *Deterministic and Stochastic Optimal Control*, Springer-Verlag, 1975.

[10] T.T. Georgiou and A. Lindquist, "The Separation Principle in Stochastic Control, Redux," http://arxiv.org/abs/1103.3005, submitted for publication.

[11] R. van Handel, *Stochastic Calculus, Filtering, and Stochastic Control*, unpublished lecture notes, Caltech, Spring 2007.

[12] P.D. Joseph and J.T. Tou, "On linear control theory", *AIEE Trans. Applications and Industry* **80** (1961), 193-196.

[13] F.C. Klebaner, *Introduction to Stochastic Calculus with Applications*, Imperial College Press, 1998.

[14] H. Kwakernaak and R. Sivan, *Linear Optimal Control Systems*, John Wiley & Sons, 1972.

[15] H.J. Kushner, *Stochastic Stability and Control*, Academic Press, 1967.

[16] H.J. Kushner, *Introduction to Stochastic Control*, Holt, Rinehart and Winston, 1971.

[17] A. Lindquist, "On optimal stochastic control with smoothed information," *Information Sciences* **1** (1968), 55-85.

[18] A. Lindquist, "Optimal control of linear stochastic systems with applications to time lag systems", *Information Sciences* 5 (January 1973), 81-126.

[19] A. Lindquist, "On feedback control of linear stochastic systems," *SIAM J. Control* **11**: 323-343, 1973.

[20] R.S. Liptser and A.N. Shirayev, *Statistics of Random Processes II, Applications*, Springer-Verlag, New York, Heidelberg, Berlin, 1978.

[21] B. Øksendal, *Stochastic Differential Equations*, Springer, 2007.

[22] J.E. Potter, "A guidance-navigation separation theorem", Rep. Rell, Experimental Astronomy Laboratory, Massachusetts Institute of Technology, Cambridge, 1964.

[23] P.E. Protter, *Stochastic Integration and Differential Equations*, Springer, 2004.

[24] L.C.G. Rogers and D. Williams, *Diffusions, Markov Processes and Martingales, Vol. 2*, Cambridge University Press, 2nd ed., 2000.

[25] F. Smithies, *Integral Equations*, Cambridge University Press, 1958.

[26] R.F. Stengel, *Optimal Control and Estimation*, unabridged and corrected, Dover (1994) republication of *Stochastic Optimal Control: Theory and Applications*, Wiley, 1986.

[27] B.S. Tsirel'son, "An example of a stochastic differential equation which has no strong solution," *Teoriya Veroyatnostei i ee Primeneniya*, **20(2)**: 427–430, 1975.

[28] H. van Warter and J.C. Willems, "The certainty equivalence property in stochastic control theory", *IEEE Trans. Automatic Control* **AC-26**(5) (1981), 1080–1087.

[29] J.C. Willems, *The Analysis of feedback systems*, MIT Press, 1971.

[30] J.C. Willems, "Recursive filtering," *Statistica Neerlandica*, **32(1)** (1978), 1-39.

[31] W.M. Wonham, "On the separation theorem of stochastic control", *SIAM J. Control* **6**(2) (1968), 312–326.

[32] P.P. Zabreyko et. al., *Integral Equations – A Reference Text*, Noordhoff International Publishing, Leyden, 1975.