for a control problem is like taking your car in the morning and trying to drive to work with your eyes closed, assuming you have perfectly memorized the itinerary. For most purpose, you'll probably be better off designing a reasonable even suboptimal closed-loop policy...

Remark. More generally, the choice of control can be *randomized*, that is, u_k is chosen according to a probability distribution depending on I_k and k, although in this course we will usually not need this extra generality. It becomes necessary to consider such randomized control laws when dealing with constrained Markov decision processes for example [Alt99], or minimax control criteria.

The next example illustrates the ideas seen so far, and their equivalent for continuous-time systems. Do not worry too much about the technical details associated with continuous-time noise models, we won't deal with them in the course. It also shows one way of obtaining a discrete-time linear state space model from a continuous-time one by sampling (the so-called *step-invariant transformation* [CF96]).

Example 1.1.2 (discretization of a continuous-time stochastic linear time-invariant system). A vehicle moves on a one-dimensional line, with its position $\xi \in \mathbb{R}$ evolving in continuous time according the differential equation

$$\ddot{\xi}(t) = u(t) + w(t),$$

where $w(t)$ is a zero mean white Gaussian noise with power spectral density $\hat{w}$. Hence $E[w(t)] = 0$, and the autocorrelation of the process is $E[w(t)w(\tau)] = \hat{w}\,\delta(t-\tau)$, where $\delta(\cdot)$ is the Dirac delta function [2]. The deterministic part of the equation corresponds to Newton's law. That is, there is a force $u(t)$ available for the controller to modify the acceleration of the vehicle. However, the acceleration is also subject to a perturbation $w(t)$ modeled as a random noise. What is the state of this system? Ignoring the stochastic component $w(t)$ at least for the moment, and by analogy with our definition in the discrete-time domain, elementary notions of differential equations tell us that the information necessary at time t to characterize the future evolution of the system consists of both the position and velocity of the vehicle, that is, the state is the two-dimensional vector $x = [\xi, \dot{\xi}]^T$. We can rewrite the dynamics as

$$\dot{x} = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} x + \begin{bmatrix} 0 \\ 1 \end{bmatrix} u + \begin{bmatrix} 0 \\ 1 \end{bmatrix} w. \tag{1.3}$$

More generally, the state-space representation of a *continuous-time linear Gaussian time-invariant system* can be written as

$$\dot{x}(t) = Ax(t) + Bu(t) + Qw(t), \tag{1.4}$$

[2] we will not spend time in this course on building rigorous foundations in stochastic calculus, because most of the time we will work directly with a discrete-time model. When needed, we will use continuous-time "white noise" freely, as does most of the engineering literature, even though it is usually more convenient from a mathematical point of view to work with integrals of the noise. For a gentle introduction to some of the mathematics involved in a more rigorous presentation, see for example [Oks07].

where x is the state vector of dimension n_x, u is the input vector of dimension n_u, w is a zero mean white Gaussian *process noise* vector of dimension n_w and power spectral density matrix $E[w(t)w(\tau)^T] = W\,\delta(t-\tau)$, and A, B, Q, W are known matrices of appropriate dimensions. The matrix A is called the system matrix, B is the input gain and Q is the noise gain. The output of the system (i.e., the available measurements) is assumed to be a vector of dimension n_y of the form

$$y(t) = Cx(t) + Du(t) + Rv(t). \tag{1.5}$$

Here v is a zero mean white Gaussian *measurement noise* vector of dimension n_v and power spectral density matrix $E[v(tv(\tau)] = V\,\delta(t-\tau)$, and C is called the measurement matrix. We assume that $RV^{1/2}$ is invertible. The general solution of (1.4) can be written explicitely as

$$x(t) = e^{A(t-t_0)}x(t_0) + \int_{t_0}^{t} e^{A(t-\tau)}[Bu(\tau) + Qw(\tau)]d\tau. \tag{1.6}$$

The fact that $w(t)$ is a white noise translates into the fact that $x(t)$ satisfies the properties necessary to represent the state of the system. In this probabilistic model, it means that the distribution of $x(t)$ at some time t conditioned on its values up to an earlier time t' depends only on the last value $x(t')$:

$$P(x(t)|x_{[-\infty,t']}, u_{[t',t]}) = P(x(t)|x(t'), u_{[t',t]}), \text{ for } t' < t.$$

In probabilistic terms, $x(t)$ is a *Markov process* when the system is driven by white noise (this would not necessarily be the case if w were not white, because states prior to t' could be used to predict $w_{[t',t]}$ and $x(t)$ in some way).

Now suppose that we sample the system (1.4) periodically with period T. It is assumed that between samples, the input $u(t)$ is kept constant

$$u(t) = u(kT),\ \forall\, kT \le t < (k+1)T.$$

Let us write $x(k) := x(kT)$, and similarly for the other signals. Then from (1.6) we deduce the following linear time invariant difference equation[3] for $x(k)$

$$x(k+1) = A_d\, x(k) + B_d\, u(k) + \tilde{w}(k), \tag{1.7}$$

where

$$A_d := e^{AT},\ B_d := e^{AT}\left(\int_0^T e^{-A\,s}ds\right)B,$$

$$\tilde{w}(k) = \int_{kT}^{(k+1)T} e^{A((k+1)T-\tau)}Qw(\tau)d\tau.$$

[3]this discretization step is exact for linear systems with such piecewise constant inputs, that is, $x(k)$ represents exactly the value of the signal $x(t)$ at the sampling times, with no approximation involved so far. It is one of the basic techniques in digital control system design, worth remembering.

From the assumption that $w(t)$ is Gaussian, zero mean white and stationary (i.e., W independent of t), it follows that $\{\tilde{w}(k)\}_{k\geq 0}$ is i.i.d. Gaussian with mean $E[\tilde{w}(k)] = 0$ and covariance matrix $E[\tilde{w}(j)\tilde{w}(k)^T] = W_d\,\delta_{jk}$ where

$$\begin{aligned} W_d &= \int_{kT}^{(k+1)T} e^{A((k+1)T-\tau)} Q W Q^T e^{A^T((k+1)T-\tau)} d\tau \\ &= \int_0^T e^{A(T-u)} Q W Q^T e^{A^T(T-u)} du. \end{aligned}$$

Hence after discretization the system at the sampling times is described by (1.7), which is a discrete time controlled random walk (CRW) of the form (1.1). The process $x(k)$ is a Markov chain, evolving over the continuous state space $\mathbb{R}^{n_x}$. The CRW is driven by a discrete time zero mean Gaussian white noise $\tilde{w}$ with covariance matrix Q_d. Note that a first order approximation (as $T \to 0$) of W_d is

$$W_d \approx (QWQ^T)\,T$$

Exercise 1. Rederive the expression of W_d in more details, starting from the definition of the covariance $E[\tilde{w}(j)\tilde{w}(k)^T]$. You can proceed formally, using properties of the Dirac delta, and not worry about the theory.

We now proceed to discretize the output of the sensor, i.e., the measurement signal $y(t)$ which can be used to design the control $u(k)$. Here the mathematical issues involved in dealing with continuous-time white observation noise appear more clearly. We cannot usually work directly with the samples $v(kT)$ of the observation noise, notably in the very important case of white Gaussian noise, because this would not be well defined mathematically. Instead, we *define* the discrete-time measurement equation to be

$$\tilde{y}(k) = Cx(k) + Du(k) + \tilde{v}(k), \tag{1.8}$$

where $\tilde{v}$ is a discrete-time zero-mean white noise with covariance matrix

$$E[\tilde{v}(j)\tilde{v}(k)^T] =: V_d \delta_{jk}.$$

This measurement equation is of the form (1.2) introduced earlier, except for a time indexing convention for u_k which is not important for the development of the theory. It is possible to establish more formal relationships between the continuous-time and discrete-time measurement equations (1.5) and (1.8), however. Typically, we need to introduce an integration step, which we did not have for the process noise because the system dynamics (1.4) already involve a differential equation with w on the right hand side. It turns out that the covariance of the corresponding discrete-time noise in (1.8) should be chosen as

$$V_d := \frac{RVR^T}{T}. \tag{1.9}$$

In (1.8), we used the notation $\tilde{y}(k)$ instead of $y(k)$ because $\tilde{y}(k)$ is not equal to $y(kT)$. Instead, as in the derivation presented below, $\tilde{y}(k)$ can be thought of as the time average of the signal $y(t)$ over the period T

$$\tilde{y}(k) = \frac{1}{T}\int_{kT}^{(k+1)T} y(\tau)d\tau.$$

This is compatible with real-world measurement devices, which cannot sample continuous-time values perfectly but perform a "short-term" integration. Moreover, in contrast to the discretization of the system dynamics, (1.8) *involves an approximation*, namely, it assumes that the state $x(t)$ is constant over the interval $kT \leq t < (k+1)T$. The derivation of (1.8) below is only given for completeness and can be skipped. It serves to justify the choice of the discrete-time covariance matrix (1.9), which is useful to pass from continuous-time to discrete-time filters for example. But for all practical purposes, in this course equation (1.8) can simply be taken as the definition of the discrete-time values returned by a digital sensor.

Derivation of the measurement equation (1.8): with the definition of $\tilde{y}(k)$ above, and approximating $x(t)$ by the constant $x(k)$ over the interval $kT \leq t < (k+1)T$, we see that

$$\tilde{v}(k) = \frac{1}{T}\int_{kT}^{(k+1)T} R\,v(\tau)d\tau. \tag{1.10}$$

Now by definition the Gaussian white noise $v(t)$ is the formal derivative of a Brownian motion B_t, with mean zero and covariance V. In stochastic calculus, equation (1.10) would be written

$$\tilde{v}(k) = \frac{1}{T}\int_{kT}^{(k+1)T} R\,dB_\tau.$$

The property $E[\tilde{v}(k))] = 0$ is then a standard property of the stochastic integral, and the discrete-time covariance matrix is

$$\begin{aligned} E[\tilde{v}(j)\tilde{v}(k)^T] &= \delta_{jk}\frac{1}{T^2}R\,E\left[\left(\int_{kT}^{(k+1)T} dB_\tau\right)\left(\int_{kT}^{(k+1)T} dB_\tau\right)^T\right]R^T \\ &= \delta_{jk}\frac{R}{T^2}\left(\int_{kT}^{(k+1)T} V\,d\tau\right)R^T \\ &= \delta_{jk}\frac{RVR^T}{T}. \end{aligned}$$

Here the first line (introduction of δ_{jk}) follows from basic properties of the Brownian motion (the independence property of the increments and their zero-mean distribution), and the second equality is called the Itô isometry, see e.g. [Oks07].

Objective to Optimize

So far we have discussed the dynamical systems of interest in this course, as described in discrete-time by the state space equations (1.1) and (1.2). The specification of a problem involves choosing an appropriate state, describing the available measurements, the process and measurement disturbances, and determining what the available controls or decisions are and how they influence the dynamics of the system and the measurements. Control theory studies