Chapter 3

Examples

> *... Don't apply any model until you understand the simplifying assumptions on which it is based, and you can test their validity. Catch phrase: use only as directed. Don't limit yourself to a single model: More than one model may be useful for understanding different aspects of the same phenomenon. Catch phrase: legalize polygamy."*
>
> Saul Golomb in his 1970 paper "Mathematical Models—Uses and Limitations" [Gol70].

In this chapter we present a collection of examples spanning many different fields of science and engineering. These examples will be used throughout the text and in exercises to illustrate different concepts. First time readers may wish to focus only on a few examples with which they have the most prior experience or insight to understand the concepts of state, input, output, and dynamics in a familiar setting.

3.1 Cruise Control

The cruise control system of a car is one of the most common control systems encountered in everyday life. The system attempts to keep the speed of the car constant in spite of disturbances caused by changes in the slope of a road and variations in the wind and road surface. The controller compensates for these unknowns by measuring the speed of the car and adjusting the throttle appropriately.

To model the complete system we start with the block diagram in Figure 3.1. Let v be the speed of the car and v_r the desired (reference) speed. The controller, which typically is of the proportional-integral (PI) type described briefly in Chapter 1, receives the signals v and v_r and generates a

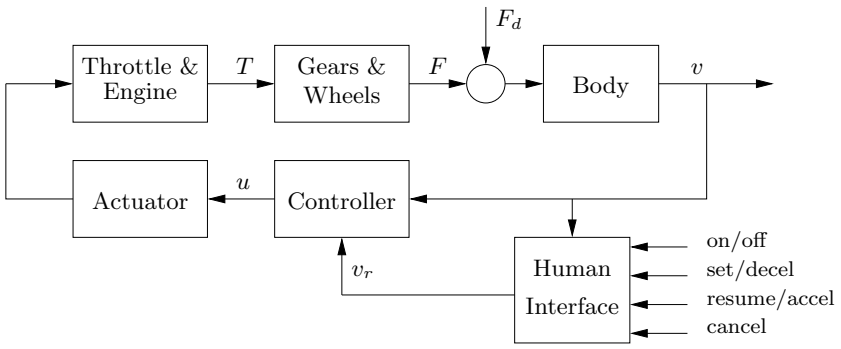

Figure 3.1: Block diagram of a cruise control system for an automobile.

control signal u that is sent to an actuator that controls throttle position. The throttle in turn controls the torque T delivered by the engine, which is then transmitted through gears and the wheels, generating a force F that moves the car. There are disturbance forces F_d due to variations in the slope of the road, the effects of rolling resistance and aerodynamic forces. The cruise controller also has a man-machine interface that allows the driver to set and modify the desired speed. There are also functions that disconnects cruise control when the brake is touched as well as functions to resume cruise control operation.

The system has many individual components—actuator, engine, transmission, wheels and car body—and a detailed model can be very complicated. In spite of this, the model required to design the cruise controller can be quite simple. In essence the model should describe how the car's speed is influenced by the slope of the road and the control signal u that drives the throttle actuator.

To model the system, it is natural to start with a momentum balance for the car body. Let v be the speed measured in m/s, m the total mass of the car in kg (including passengers), F the force generated by the contact of the wheels with the road, and F_d the disturbance force due to gravity and friction. The equation of motion of the car is simply

$$m\frac{dv}{dt} = F - F_d. \tag{3.1}$$

The force F is generated by the engine, whose torque is proportional to the rate of fuel injection, which is itself proportional to the control signal $0 \leq u \leq 1$ that controls throttle position. The torque also depends on engine speed ω. A simple representation of the torque at full throttle is given by

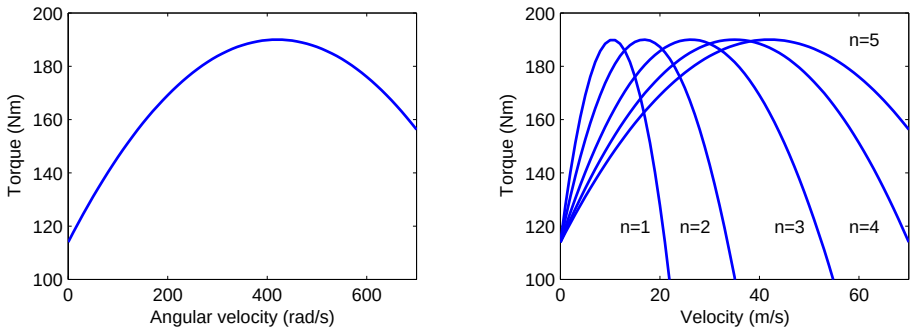

(a) (b)

Figure 3.2: Torque curves for typical car engine: (a) torque as a function of the angular velocity of the engine and (b) torque as a function of car speed for different gears.

the torque curve

$$T(\omega) = T_m \left(1 - \beta\left(\frac{\omega}{\omega_m} - 1\right)^2\right), \tag{3.2}$$

where the maximum torque T_m is obtained at engine speed ω_m. Typical parameters are $T_m = 190$ Nm, $\omega_m = 420$ rad/sec (about 4000 RPM) and $\beta = 0.4$.

Let n be the gear ratio and r the wheel radius. The engine speed is related to the velocity through the expression

$$\omega = \frac{n}{r} v =: \alpha_n v,$$

and the driving force can be written as

$$F = \frac{nu}{r} T(\omega) = \alpha_n u T(\alpha_n v).$$

Typical values of α_n for gears 1 through 5 are $\alpha_1 = 40$, $\alpha_2 = 25$, $\alpha_3 = 16$, $\alpha_4 = 12$ and $\alpha_5 = 10$. The inverse of α_n has physical interpretation as the *effective wheel radius*. Figure 3.2 shows the torque as function of engine speed and vehicle speed. The figure shows that the effect of the gear is to "flatten" the torque curve so that a torque close to maximum can be obtained almost over the full speed range.

The disturbance force F_d has three major components: F_g, the forces due to gravity; F_r, the forces due to rolling friction; and F_a, the aerodynamic

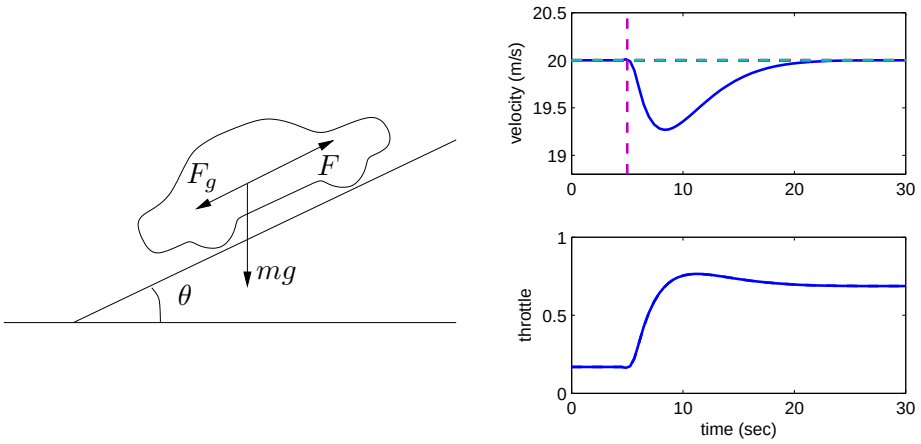

(a) (b)

Figure 3.3: Car with cruise control encountering a sloping road: a schematic diagram is shown in (a) and (b) shows the response in speed and throttle when a slope of 4° is encountered. The hill is modeled as a net change in hill angle, θ, of 4 degrees, with a linear change in the angle between $t = 5$ and $t = 6$. The PI controller has proportional gain is $k_p = 0.5$ and the integral gain is $k_i = 0.1$.

drag, Letting the slope of the road be θ, gravity gives the retarding force $F_g = mg \sin\theta$, as illustrated in Figure 3.3a, where $g = 9.8$ m/sec^2 is the gravitational constant. A simple model of rolling friction is

$$F_r = mgC_r$$

where C_r is the coefficient of rolling friction; a typical value is $C_r = 0.01$. Finally, the aerodynamic drag is proportional to the square of the speed:

$$F_a = \frac{1}{2}\rho C_d A v^2,$$

where ρ is the density of air, C_d is the shape-dependent aerodynamic drag coefficient and A is the frontal area of the car. Typical parameters are $\rho = 1.3$ kg/m^3, $C_d = 0.32$ and $A = 2.4$ m^2.

Summarizing, we find that the car can be modeled by

$$m\frac{dv}{dt} = \alpha_n u T(\alpha_n v) - mgC_r - \tfrac{1}{2}\rho C_v A v^2 - mg\sin\theta, \tag{3.3}$$

where the function T is given by equation (3.2). The model (3.3) is a dynamical system of first order. The state is the car velocity v, which is also

the output. The input is the signal u that controls the throttle position, and the disturbance is the force F_d, which depends on the slope of the road. The system is nonlinear because of the torque curve and the nonlinear character of the aerodynamic drag. There can also be variations in the parameters, e.g. the mass of the car depends on the number of passengers and the load being carried in the car.

We add to this model a feedback controller that attempts to regulate the speed of the car in the presence of disturbances. We shall use a PI (proportional-integral) controller, which has the form

$$u(t) = k_p e(t) + k_i \int_0^t e(\tau)\, d\tau.$$

This controller can itself be realized as an input/output dynamical system by defining a controller state z and implementing the differential equation

$$\frac{dz}{dt} = v_r - v \qquad u = k_p(v_r - v) + k_i z, \tag{3.4}$$

where v_r is the desired (reference) speed. As discussed briefly in the introduction, the integrator (represented by the state z) insures that in steady state the error will be driven to zero, even when there are disturbances or modeling errors. (The design of PI controllers is the subject of Chapter 10.) Figure 3.3b shows the response of the closed loop system, consisting of equations (3.3) and (3.4), when it encounters a hill. The figure shows that even if the hill is so steep so that the throttle changes from 0.17 to almost full throttle, the largest speed error is less than 1 m/s, and the desired velocity is recovered after 20s.

The model (3.3) is essentially a momentum balance for the car. Many approximations were made when deriving it. It may be surprising that such a seemingly complicated system can be described by the simple model (3.3). As we shall see in later chapters, the reason for this is the inherent robustness of feedback systems: even if the model is not perfectly accurate, we can use it to design a controller and make use of the feedback in the controller to manage the uncertainty in the system.

The cruise control system also has a human-machine interface (HMI) that allows the driver to communicate with the system. There are many different ways to implement this system; one version is illustrated in Figure 3.4. The system has four buttons: on-off, set/decelerate, resume/accelerate and cancel. The operation of the system is governed a finite state system and which controls the modes of the PI controller and the reference generator.

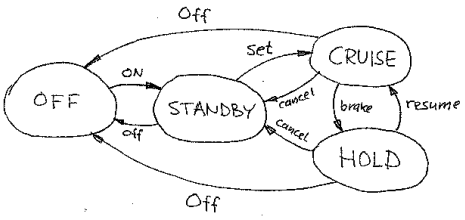

Figure 3.4: Finite state machine for cruise control system.

The controller can operate in two ways: in the normal cruise control mode and in a tracking mode, where the integral is adjusted to match given process inputs and outputs. The tracking mode is introduced to avoid switching transients when the system is controlled manually. The generator for the reference signal has three modes: a normal control mode when the output is controlled by the set/accelerate and resume/decelerate buttons, a tracking mode and a hold mode where the reference is held constant.

To control the overall operation of the controller and reference generator, we use a finite state machine with four states: off, standby, cruise and hold. The states of the controller and the reference generator in the different modes are given in Figure 3.4. The cruise mode is the normal operating mode where the speed can be then be decreased by pushing set/decelerate and increased by pushing the resume/accelerate. When the system is switched on it goes to standby mode. The cruise mode is activated by pushing the set/accelerate button. If the brake is touched or if the gear is changed, the system goes into hold mode and the current velocity is stored in the reference generator. The controller is then switched to tracking mode and the reference generator is switched to hold mode, where it holds the current velocity. Touching the resume button then switches the system to cruise mode. The system can be switched to standby mode from any state by pressing the cancel button.

The PI controller should be designed to have good regulation properties and to give good transient performance when switching between resume and control modes. Implementation of controllers and reference generators will be discussed more fully in Chapter 10. A popular description of cruise control system can be found on the companion web site. Many automotive applications are discussed in detail in [BP96] and [KN00].

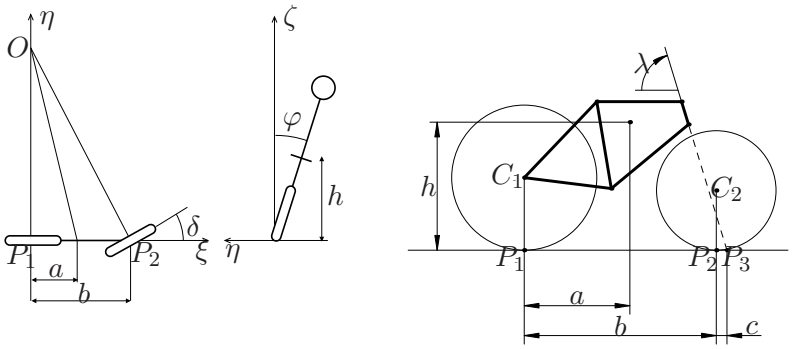

Figure 3.5: Schematic top (left), rear (middle), and side (right) views of a bicycle. The steering angle is δ, the roll angle is φ. The center of mass has height h and on the distance a from a vertical through the contact point P_1 of the rear wheel. The wheel base is b and the trail is c.

3.2 Bicycle Dynamics

The bicycle is an interesting dynamical system system with the feature that one of its key properties is due to a feedback mechanism that is created by a clever design of the front fork. A detailed model of a bicycle is complex because the system has many degrees of freedom and the geometry is complicated. However, a great deal of insight can be obtained from simple models.

To derive the equations of motion we assume that the bicycle rolls on the horizontal xy plane. Introduce a coordinate system that is fixed to the bicycle with the ξ-axis through the contact points of the wheels with the ground, the η-axis horizontal and the ζ-axis vertical, as shown in Figure 3.5. Let v_0 be the velocity of the bicycle at the rear wheel, b the wheel base, φ the tilt angle and δ the steering angle. The coordinate system rotates around the point O with the angular velocity $\omega = v_0\delta/b$, and an observer fixed to the bicycle experiences forces due to the motion of the coordinate system.

The tilting motion of the bicycle is similar to an inverted pendulum, as shown in the rear view in Figure 3.5b. To model the tilt, consider the rigid body obtained when the wheels, the rider and the front fork assembly are fixed to the rear frame. Let m be the total mass of the system, J the moment of inertia of this body with respect to the ξ-axis, and D the product of inertia with respect to the $\xi\zeta$ axes. Furthermore, let the ξ and ζ coordinates of the center of mass be a and h, respectively. We have $J \approx mh^2$ and $D = mah$.

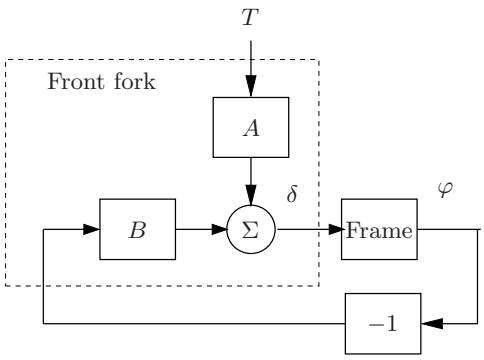

Figure 3.6: Block diagram of the bicycle with a front fork. The steering torque applied to the handlebars is T, the roll angle is φ, and the steering angle δ. Notice that the front fork creates a feedback from the roll angle φ to the steering angle δ that under certain conditions can stabilize the system.

The torques acting on the system are due to gravity and centripetal action. Assuming that the steering angle δ is small, the equation of motion becomes

$$J\frac{d^2\varphi}{dt^2} - \frac{Dv_0}{b}\frac{d\delta}{dt} = mgh\sin\varphi + \frac{mv_0^2h}{b}\delta, \tag{3.5}$$

The term $mgh\sin\varphi$ is the torque generated by gravity. The terms containing δ and its derivative are the torques generated by steering, with the term $(Dv_0/b)\,d\delta/dt$ due to inertial forces and the term $(mv_0^2h/b)\,\delta$ due to centripetal forces.

The steering angle is influenced by the torque the rider applies to the handle bar. Because of the tilt of the steering axis and the shape of the front fork, the contact point of the front wheel with the road P_2 is behind the axis of rotation of the front wheel assembly, as shown in Figure 3.5. The distance c between the contact point of the front wheel P_2 and the projection of the axis of rotation of the front fork assembly P_3 is called the *trail*. The steering properties of a bicycle depend critically on the trail. A large trail increases stability but make the steering less agile.

A consequence of the design of the front fork is that the steering angle δ is influence both by steering torque T and by the tilt of the frame φ. This means that the bicycle with a front fork is a *feedback system* as illustrated by the block diagram in Figure 3.6. The steering angle δ influences the tilt angle φ and the tilt angle influences the steering angle giving rise to the circular causality that is characteristic for reasoning about feedback. For a front fork with positive trail, the bicycle will steer into the lean creating

a centrifugal force that attempts to diminish the lean. The effect can be verified experimentally by biking on a straight path, creating a lean by tilting the body and observing the steering torque required to keep the bicycle on a straight path when leaning. Under certain conditions, the feedback can actually stabilize the bicycle. A crude empirical model is obtained by assuming that the blocks A and B are static gains k_1 and k_2 respectively:

$$\delta = k_1 T - k_2 \varphi. \tag{3.6}$$

This model neglects the dynamics of the front fork, the tire-road interaction and the fact that the parameters depend on the velocity. A more accurate model is obtained by the rigid body dynamics of the front fork and the frame. Assuming small angles this model becomes

$$M \begin{pmatrix} \ddot{\varphi} \\ \ddot{\delta} \end{pmatrix} + C v_0 \begin{pmatrix} \dot{\varphi} \\ \dot{\delta} \end{pmatrix} + (K_0 + K_2 v_0^2) \begin{pmatrix} \varphi \\ \delta \end{pmatrix} = \begin{pmatrix} 0 \\ T \end{pmatrix}, \tag{3.7}$$

where the elements of the 2×2 matrices M, C, K_0 and K_2 depend on the geometry and the mass distribution of the bicycle. Even this model is inaccurate because the interaction between tire and road are neglected. Taking this into account requires two additional state variables.

Interesting presentations of the development of the bicycle are given in the books by D. Wilson [Wil04] and Herlihy [Her04]. More details on bicycle modeling is given in the paper [ÅKL05], which has many references. The model (3.7) was presented in a paper by Whipple in 1899 [Whi99].

3.3 Operational Amplifier

The operational amplifier (op amp) is a modern implementation of Black's feedback amplifier. It is a universal component that is widely used for for instrumentation, control and communication. It is also a key element in analog computing.

Schematic diagrams of the operational amplifier are shown in Figure 3.7. The amplifier has one inverting input (v_-), one non-inverting input (v_+), and one output (v_{out}). There are also connections for the supply voltages, e_- and e_+ and a zero adjustment (offset null). A simple model is obtained by assuming that the input currents i_- and i_+ are zero and that the output is given by the static relation

$$v_{\text{out}} = \text{sat}_{(v_{\min}, v_{\max})}\big(k(v_+ - v_-)\big), \tag{3.8}$$

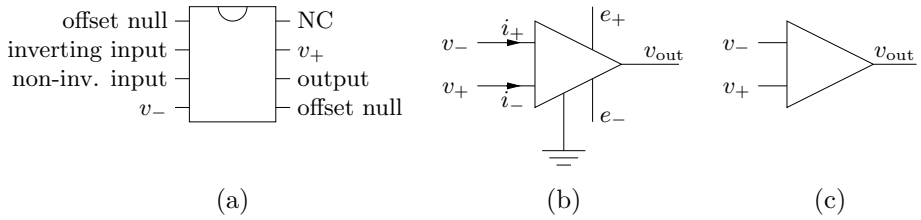

Figure 3.7: An operational amplifier and two schematic diagrams. The figure on the left shows the amplifier pin connections on an integrated circuit chip, the middle figure shows a schematic with all connections, and the diagram on the right shows only the signal connections.

where sat denotes the saturation function

$$\text{sat}_{(a,b)}(x) = \begin{cases} a & \text{if } x < a \\ x & \text{if } a \leq x \leq b \\ b & \text{if } x > b. \end{cases} \tag{3.9}$$

We assume that the gain k is very large, in the range of 10^6–10^8, and the voltages $v_{\min}$ and $v_{\max}$ satisfy

$$e_- \leq v_{\min} < v_{\max} \leq e_+$$

and hence are in the range of the supply voltages. More accurate models are obtained by replacing the saturation function with a smooth function as shown in Figure 3.8. For small input signals the amplifier characteristic (3.8) is linear

$$v_{\text{out}} = k(v_+ - v_-) =: -kv. \tag{3.10}$$

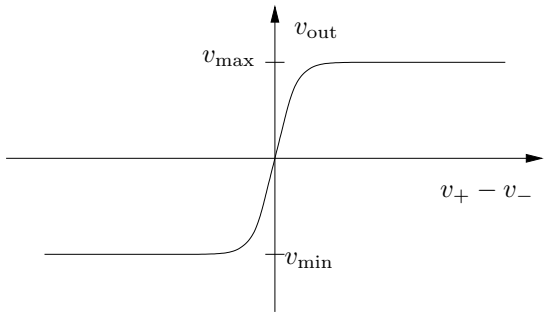

Figure 3.8: Input-output characteristics of an operational amplifier.

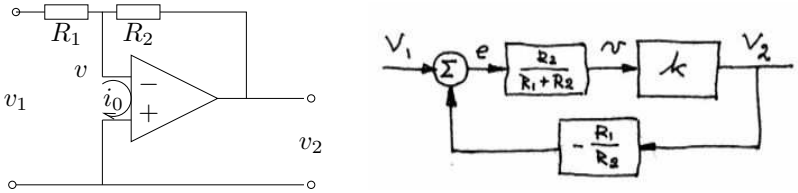

(a) (b)

Figure 3.9: Circuit diagram of a stable amplifier based on negative feedback around an operational amplifier (a) and the corresponding block diagram (b).

Since the open loop gain k is very large, the range of input signals where the system is linear is very small.

A simple amplifier is obtained by arranging feedback around the basic operational amplifier as shown in Figure 3.9a. To model the feedback amplifier in the linear range, we assume that the current $i_0 = i_- + i_+$ is zero, and that the gain of the amplifier is so large that the voltage $v = v_- - v_+$ is also zero. It follows from Ohm's law that the currents through resistors R_1 and R_2 are given by

$$\frac{v_1}{R_1} = -\frac{v_2}{R_2}$$

and hence

$$\frac{v_2}{v_1} = -k_{\text{cl}} \quad \text{where} \quad k_{\text{cl}} = \frac{R_2}{R_1} \tag{3.11}$$

is the closed loop gain of the amplifier.

A more accurate model is obtained by neglecting the current i_0 but assuming that the voltage v is small but not negligible. The current balance then becomes

$$\frac{v_1 - v}{R_1} = \frac{v - v_2}{R_2}. \tag{3.12}$$

Assuming that the amplifier operates in the linear range and using equation (3.10) the gain of the closed loop system becomes

$$k_{\text{cl}} = -\frac{v_2}{v_1} = \frac{R_2}{R_1} \frac{1}{1 + \frac{1}{k}\left(1 + \frac{R_2}{R_1}\right)} \tag{3.13}$$

If the open loop gain k of the operational amplifier is large, the closed loop gain k_{cl} is the same as in the simple model given by equation (3.11). Notice that the closed loop gain only depends on the passive components, and that variations in k only have a marginal effect on the closed loop gain. For example if $k = 10^6$ and $R_2/R_1 = 100$, a variation of k by 100% only

gives a variation of 0.01% in the closed loop gain. The drastic reduction in sensitivity is a nice illustration of how feedback can be used to make good systems from bad components. In this particular case, feedback is used to trade high gain and low robustness for low gain and high robustness. Equation (3.13) was the formula that inspired Black when he invented the feedback amplifier.

It is instructive to develop a block diagram for the feedback amplifier in Figure 3.9a. To do this we will represent the pure amplifier with input v and output v_2 as one block. To complete the block diagram we must describe how v depends on v_1 and v_2. Solving equation (3.12) for v gives

$$v = \frac{R_2}{R_1 + R_2} v_1 + \frac{R_1}{R_1 + R_2} v_2 = \frac{R_2}{R_1 + R_2} \Big(v_1 + \frac{R_1}{R_2} \Big),$$

and we obtain the block diagram shown in Figure 3.9b. The diagram clearly shows that the system has feedback and that the gain from v_2 to v is $R_1/(R_1 + R_2)$, which can also be read from the circuit diagram in Figure 3.9a. If the loop is stable and if gain of the amplifier is large it follows that the error e is small and then we find that $v_2 = -(R_2/R_1)v_1$. Notice that the resistor R_1 appears in two blocks in the block diagram. This situation is typical in electrical circuits and it is one reason why block diagrams are not always well suited for some types of physical modeling.

The simple model of the amplifier given by equation (3.10) gives qualitative insight but it neglects the fact that the amplifier is a dynamical system. A more realistic model is

$$\frac{dv_{\text{out}}}{dt} = -av_{\text{out}} - bv. \tag{3.14}$$

The parameter b which has dimensions of frequency is called the gain-bandwidth product of the amplifier.

The operational amplifier is very versatile and many different systems can be built by combining it with resistors and capacitors. Figure 3.10 shows the circuit diagram for analog PI (proportional-integral) controller. To develop a simple model for the circuit we assume that the current i_0 is zero and that the open loop gain k is so large that the input voltage v is negligible. The current i through the capacitor is $i = Cdv_c/dt$, where v_c is the voltage across the capacitor. Since the same current goes through the resistor R_1 we get

$$i = \frac{v_1}{R_1} = C\frac{dv_c}{dt},$$

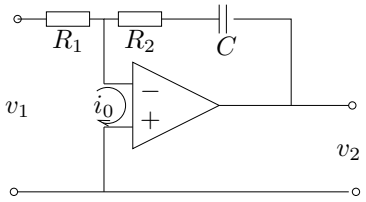

Figure 3.10: Circuit diagram of a PI controller obtained by feedback around an operational amplifier.

which implies that

$$v_c(t) = \frac{1}{C}\int i(t)\,dt = \frac{1}{R_1 C}\int_0^t v_1(\tau)d\tau.$$

The output voltage is thus given by

$$v_2(t) = -R_2 i - v_c = -\frac{R_2}{R_1}v_1(t) - \frac{1}{R_1 C}\int_0^t v_1(\tau)d\tau,$$

which is the input/output relation for a PI controller.

The development of operational amplifiers is based on the work of Philbrick [Lun05, Phi48] and their usage is described in many textbooks (e.g. [CD75]). Very good information is also available from suppliers [Jun02, Man02].

3.4 Web Server Control

Control is important to ensure proper functioning of web servers, which are key components of the Internet. A schematic picture of a server is shown in Figure 3.11. Requests are arriving, queued and processed by the server, typically on a first-come-first-serve basis. There are typically large variations in arrival rates and service rates. The queue length builds up when the

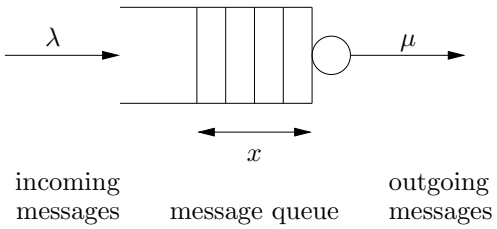

Figure 3.11: Schematic diagram of a web server.

arrival rate is larger than the service rate. When the queue becomes too large, service is denied using some admission control policy.

The system can be modeled in many different ways. One way is to model each incoming request, which leads to an event-based model, where the state is an integer that represents the queue length. The queue changes when a request arrived or a request is served. A discrete time model that captures these dynamics is given by the difference equation

$$x_{k+1} = x_k + u_i - u_o, \quad x \in I$$

where u_i and u_o are *random variables* representing incoming and outgoing requests on the queue. These variables take on the values 0 or 1 with some probability at each time instant. To capture the statistics of the arrival and servicing of messages, we model each of these as a *Poisson process* in which the number of events occurring in a fixed time has a given rate, with the specific timing of events independent of the time since the last event. (The details of random processes are beyond the scope of this text, but can be found in standard texts such as [**?**].)

The system can also described using a *flow model* by approximating the requests and services by continuous flows and the queue length by a continuous variable. A flow model can be obtained by making probabilistic assumptions on arrival and service rates and computing the average queue length. For example, assuming that the arrival and service rates are Poisson processes with intensities λ and μ it can be shown that the average queue length x is described by the first-order differential equation

$$\frac{dx}{dt} = \lambda u - \mu \frac{x}{x+1}. \tag{3.15}$$

The control variable $0 \leq u \leq 1$ is the fraction of incoming requests that are serviced, giving an effective arrival rate of $u\mu$. The average time to serve a request is

$$T_s = \frac{x}{\lambda}.$$

If μ, λ and u are constants with $\mu > u\lambda$, the queue length x approaches the steady state value

$$x_{ss} = \frac{u\lambda}{\mu - u\lambda}. \tag{3.16}$$

Figure 3.12a shows the steady state queue length as a function of $\mu - u\lambda$, the effective service rate excess. Notice that the queue length increases rapidly as $\mu - u\lambda$ approaches zero. To have a queue length less than 20 requires $\mu > u\lambda + 0.05$.

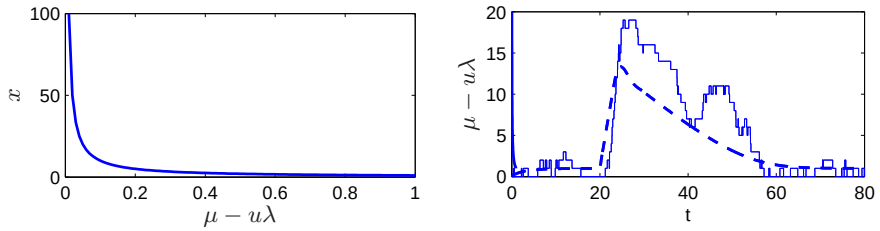

(a) (b)

Figure 3.12: The figure on the left shows steady state queue length as a function of $u\lambda - \mu$, and the figure on the right shows the behavior of the queue length when there is a temporary overload in the system. The full line shows a realization of an event based simulation and the dashed line shows the behavior of the flow model (3.15).

Figure 3.12b illustrates the behavior of the server in a typical overload situation. The service rate is $\mu = 1$, while the arrival rate starts at $\lambda = 0.5$. The arrival rate is increased to $\lambda = 4$ at time 20, and it returns to $\lambda = 0.5$ at time 25. The figure shows that the queue builds up quickly and clears very slowly. Since the response time is proportional to queue length, it means that the quality of service is poor for a long period after an overload. The behavior illustrated in Figure 3.12b, which is called the *rush-hour effect*, has been observed in web servers and in many other queuing systems like automobile traffic. Congestion avoidance is a main reason for controlling queues.

The dashed line in Figure 3.12b shows the behavior of the flow model, which describes the average queue length. The simple model captures behavior qualitatively, but since the queue length is short there is significant variability from sample to sample. The behavior shown in Figure 3.12b can be explained quantitatively by observing that the queue length increases at constant rate over large time intervals. It follows from equation (3.15) that the rate of change is approximately 3 messages/second when the queue length builds up at time $t = 20$, and approximately 0.5 messages/second when the queue length decreases after the build up. The time to return to normal is thus approximately 6 times the overload time.

Admission Control

The long delays created by temporary overloads can be reduced by access control. The queue length can be controlled by only admitting a fraction of the incoming requests. Figure 3.13 shows what happens when a simple ad-

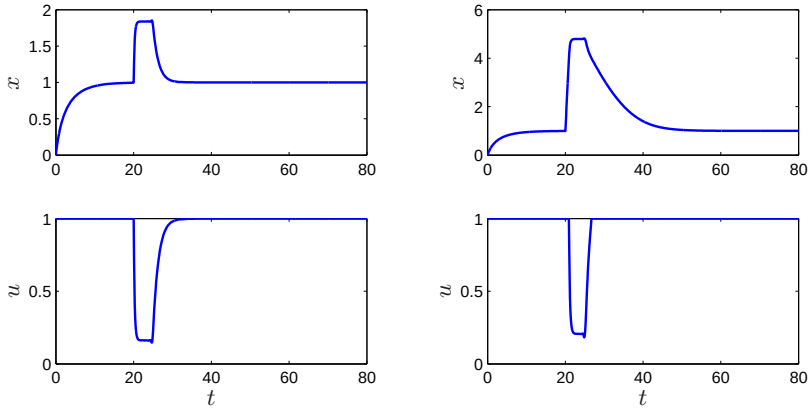

(a) $r = 2$ (b) $r = 5$

Figure 3.13: Behavior of queue length for a server with admission control when there is a temporary overload in the system. The figure on the left has $r = 2$ and the right figure has $r = 5$, with $k = 1$ in both cases. Compare with a simulation of the uncontrolled system in Figure 3.12b.

mission control strategy is introduced. The feedback used in the simulation is a simple proportional control with saturation described by

$$u = \text{sat}_{(0,1)}(k(r - x)), \tag{3.17}$$

where $\text{sat}_{(a,b)}$ is defined in equation (3.9) and r is the desired (reference) queue length. The feedback gain is $k = 1$, and the saturation ensures that the control variable is in the interval $0 \leq u \leq 1$. Comparing Figures 3.12b and 3.13, we see that simple access control works very well in comparison with the uncontrolled server. The control law ensures that the access is restricted when overload occurs.

The maximum queue length is determined by the reference value r. A low value of r gives a short queue length and the service delay is short, as is clearly seen in Figure 3.13a. A number of customers are, however, denied service. The simulation also indicates that the control problem is not too difficult and that a simple control strategy works quite well. It allows all requests arrived to be serviced if the arrival rate is slow and it restricts admission when the system is overloaded. Admission control is activated when the queue length approaches the value r. Since service time is proportional to queue length, r is a measure of service time.

Notice that the web server control problem we have discussed is not a conventional regulation problem where we wish to keep a constant queue

length. The problem is instead to make sure that the queue length does not become too large when there are many service requests. The key trade-off is to find a good reference value r. A large value gives few rejections but long service time after an overload; a small value guarantees a short service time but more messages will be rejected. The simulation of the simple flow model indicates that the simple admission control strategy works well.

To execute admission control in a real queue, where arrival and departure from the queue are discrete events, we argue as follows. Figure 3.13 shows that all requests are serviced ($u = 1$) except when the system is overloaded, at which point service is reduced significantly. A simple strategy that mimics this for event-based systems is to admit customers as long as the queue length is less than r and deny service for requests if the queue length is greater than r.

Delay Control

An alternative to admission control is *delay control*, where the goal is to keep the delay for serving individual requests constant. An advantage of this approach is that all requests are treated fairly. A block diagram of such a system, with a controller combining feedback u_{fb} and feedforward u_{ff}, is shown in Figure 3.14a. The server delay is estimated based on arriving server requests and queue waiting times of requests that have not been serviced. Feedforward control requires good models and the simple model (3.15) that captures the average behavior of the system is not sufficient.

The control variable is the processing speed u, which can be varied by the changing the number of servers and their processing capacity. It is assumed that u can be regarded as a continuous variable. The delays in serving the requests is the output of the system. An average of past service is easily obtained, but this information is only available with a time delay.

To obtain a better model we consider the situation in Figure 3.14b. A request has just been serviced at time $t = t_k$ and N requests are waiting to be serviced. The average delay of the N requests that are waiting to be serviced is d_k^-, which is a measurable quantity. To predict the additional time required to serve these request we assume that they require the same service time C/u where u is the service rate. The average additional service time for the requests that are processed is then $d_k^+ = (N + 1)C/(2u)$, as indicated in Figure 3.14b. Combining this with the measurable quantity d_k^- we obtain the following estimate of the average service time for the N

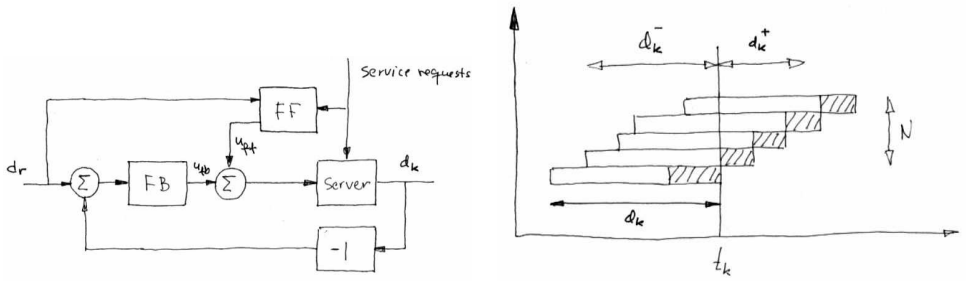

(a) (b)

Figure 3.14: The left figure (a) shows a block diagram of a web server system with a controller based on a combination of feedback and feedforward. The right figure (b) shows the history of arrivals and departures of requests. The dashed square indicates the time used to service the requests. The true delay of the request serviced at time t_k is d_k, $d_k^- + d_k^+$ is an estimate of future delays used to calculate the service rate.

requests that are waiting to be serviced

$$d_k = d_k^- + d_k^+ = d_k^- + \frac{(N+1)C}{2u}.$$

Requiring that d_k is equal to the desired delay time d_r, we find that the service rate at instant k should be chosen as

$$u_k = \frac{(N+1)C}{2(d_r - d_k^-)}, \tag{3.18}$$

which is the formula used to calculate the feedforward control signal at time t_k. The control action can recalculated at each time instant, resulting in a control strategy called *receding horizon control*. The choice of recalculation rate is a compromise because frequent recalculations improves control quality but it also consumes computer resources.

The feedforward is complemented with a feedback controller in the form of a PI controller based on the measured delay at event k. Since the queue dynamics varies with the delay time it is useful to let the parameters of the PI controller depend on the desired delay d_r, an example of *gain scheduling*.

The control algorithm has been tested experimentally on a testbed of PC's connected via Ethernet. One PC was was assigned to run the web server, and the others were generating a synthetic workload. The goal of the system was to provide the delay guarantee for that class with as few resources as possible. The input load patterns generated by the clients are

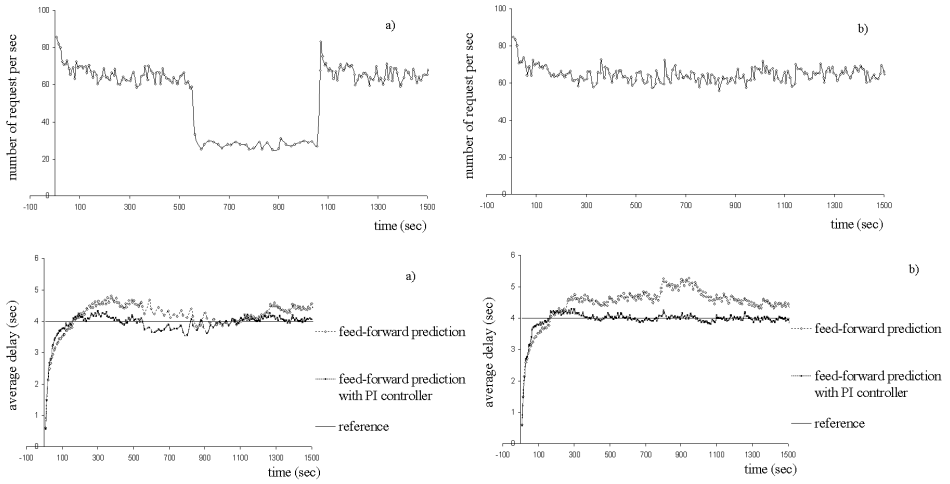

Figure 3.15: Arrival rate (top) and average service delay (bottom) for an experiment with web server control (from [HLA04]).

shown in Figure 3.15. The desired delay for the class was set to $d_r = 4s$ in all experiments. The figure shows that the control algorithm keeps the service time reasonably constant and that the PI controller reduces the variations in delay compared with a pure feedforward controller.

This example illustrates that simple models can give good insight and that nonlinear control strategies are useful. The example also illustrates that continuous time models can be useful for phenomena that are basically discrete. There are also converse examples. Therefore it is a good idea to keep an open mind and master both discrete and continuous time modeling.

The book by Hellerstein et al. [HDPT04] gives many examples of use of feedback in computer systems. The example on delay control is based on the work of Henriksson [HLA04, Hen06].

3.5 Atomic Force Microscope

The 1986 Nobel Prize in Physics was shared by Gerd Binnig and Heinrich Rohrer for their design of the scanning tunneling microscope (SCM). The idea of an SCM is to bring an atomically sharp tip so close to a conducting surface that tunneling occurs. An image is obtained by traversing the tip and measuring the tunneling current as a function of tip position. The image reflects the electron structure of the upper atom-layers of the sample.

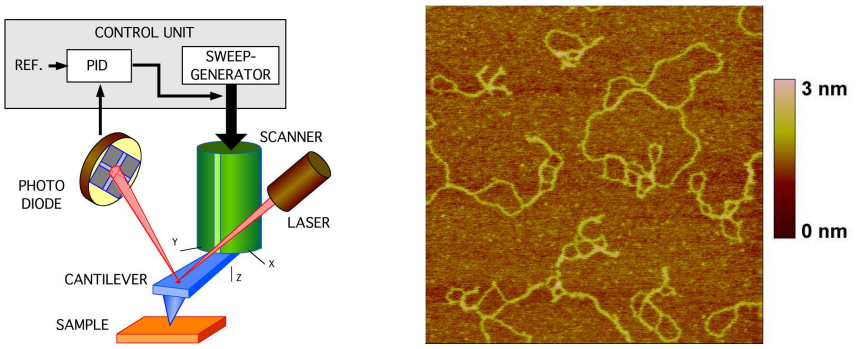

(a) (b)

Figure 3.16: Schematic diagram of an atomic force microscope and a sample AFM image of DNA.

This invention has stimulated development of a family of instruments that permit visualization of surface structure at the nanometer scale, including the atomic force microscope (AFM). These instruments are now standard tools for exploring nanoscale structures.

In the atomic force microscope, a sample is probed by a tip on a cantilever which is controlled to exert a constant force on the sample. The control system is essential because it has a direct influence on picture quality and scanning rate. Since the dynamic behavior of the system changes with the properties of the sample, it is necessary to tune the feedback loop, which is currently done manually by adjusting parameters of a PI controller. There are interesting possibilities to make the systems easier to use by introducing automatic tuning and adaptation.

A schematic picture of an atomic force microscope is shown in Figure 3.16a. A micro-cantilever with a tip having a radius of the order of 10 nm is placed close to the sample. The tip can be moved vertically and horizontally using a piezoelectric scanner. Atomic forces bend the cantilever and the cantilever tilt is measured by sensing the deflection of the beam using a photo diode. The signal from the photo diode is amplified and sent to a controller that drives the amplifier for the vertical deflection of the cantilever. By controlling the piezo scanner so that the deflection of the cantilever is constant, the signal driving the vertical deflection of the scanner is a measure of the atomic forces between the cantilever tip and the atoms of the sample. An image of the surface is obtained by scanning the cantilever along the sample. The resolution makes it possible to see the structure of the sample on the atomic scale, as illustrated in Figure 3.16b,

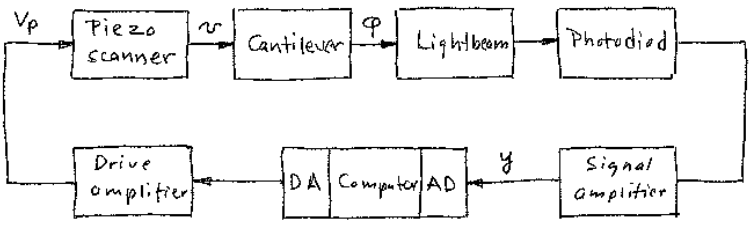

Figure 3.17: Block diagram of the system for vertical positioning of the cantilever.

which shows an AFM image of DNA.

To model the system, we start with the block diagram shown in Figure 3.17, which shows the major components. Signals that are easily accessible are: the voltage V_p that drives the piezo scanner, the input voltage u to its power amplifier and the output voltage y of the signal amplifier for the photo diode. The controller is a PI controller implemented by a computer, which is connected to the system by A/D and D/A converters. The deflection of the cantilever, φ, is also shown.

For a more detailed model we will start with the cantilever, which is at the heart of the system. The micro-cantilever is modeled as a spring-mass-damper system. Let z be the distance from the tip of the cantilever to the sample and let v be the position of the cantilever base. Furthermore let m, k and c be the effective values of mass, spring and damping coefficients. The equation of motion of the cantilever is then

$$m\frac{d^2z}{dt^2} + c\frac{dz}{dt} + k(z - v) = F, \tag{3.19}$$

where F is the atomic force between the sample and the cantilever tip.

Neutral atoms and molecules are subject to two forces, an attractive van der Waals force, and a repulsion force due to the Pauli exclusion principle. The force between two atoms can be approximately described by the Lennard-Jones potential given by

$$V_{LJ}(z) = A\left(\left(\frac{\sigma}{z}\right)^{12} - \left(\frac{\sigma}{z}\right)^{6}\right),$$

where σ is the atom radius and r the distance between the atoms. Approximating the cantilever tip by a sphere with radius R and the sample by a flat surface then integrating the Lennard-Jones potential, the interaction

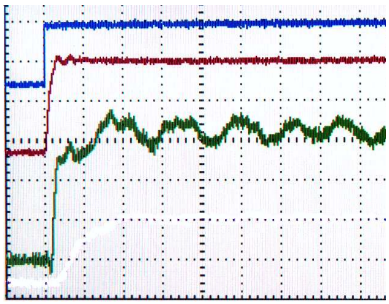

(a)

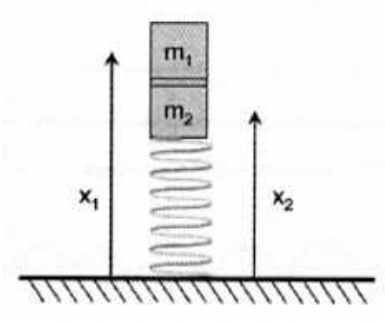

(b)

Figure 3.18: Measured step response and model of piezo scanner. The left figure shows a measured step response. The blue signal shows the input is the voltage applied to the drive amplifier (50 mV/div), the red curve is the output of the power amplifier (500 mV/div) and the red curve is the output of the signal amplifier (500 mV/div). The time scale is 25 μs/div. The right figure is a simple mechanical model for the vertical positioner and the piezo crystal.

between the cantilever and the sample can be described by the following potential

$$V(z) = \frac{HR}{6\sigma}\left(\frac{1}{120}\left(\frac{\sigma}{z}\right)^7 - \frac{\sigma}{z}\right),$$

where $H \approx 10^{-19}$ J is the Hamaker constant, and a typical atom radius is $\sigma = 0.4$ nm. The potential has a minimum where the distance between the tip is less than an atom size from the sample and the tip is essentially clamped at the minimum by the atomic forces. The natural frequency of the clamped cantilever is so high that the dynamics of the cantilever can be neglected and we can model the cantilever as a static system. For small deviations, the bending φ of the cantilever is then proportional to the vertical translation of the cantilever.

The piezo scanner gives a deflection that is proportional to the applied voltage, but the system and the amplifiers also have dynamics. Figure 3.18a shows a step response of a scanner from the input voltage u to the drive amplifier to the output voltage y of the signal amplifier for the photo diode. A schematic mechanical representation of the vertical motion of the scanner is shown in Figure 3.18b. The figure shows that the system responds quickly but that there is a poorly damped oscillatory mode caused by the dynamics of the scanner. The instrument designer has two choices, either to accept the oscillation and to have a slow response time or else to design a control

system that can damp the oscillations which gives a faster response and a faster imaging. Damping the oscillations is a significant challenge because there are many oscillatory modes and they can change depending on how the instrument is used. An instrument designer also has the choice to redesign the mechanics so that the resonances occur at higher frequencies.

The book by Sarid [Sar91] gives a broad coverage of atomic force microscopes. The interaction of atoms close to surfaces is fundamental to solid state physics. A good source is Kittel [Kit95] where the Lennard-Jones potential is discussed. Modeling and control of atomic force microscopes are discussed by Schitter [Sch01].

3.6 Drug Administration

The phrase "take two pills three times a day" is a recommendation that we are all familiar with. Behind this recommendation is a solution of an open loop control problem. The key issue is to make sure that the concentration of a medicine in a part of our bodies will be sufficiently high to be effective but not so high that it will cause undesirable side effects. The control action is quantized, *take two pills*, and sampled, *every 8 hours*. The prescriptions can be based on very simple models in terms of empirical tables where the dosage is based on the age and weight of the patient. A more sophisticated administration of medicine is used to keep concentration of insulin and glucose at a right level. In this case the substances are controlled by continuous measurement and injection, and the control schemes are often model based.

Drug administration is clearly a control problem. To do it properly it is necessary to understand how a drug spreads in the body after it is administered. This topic, called *pharmacokinetics*, is now its own discipline and the models used are called *compartment models*. They go back to 1920 when Widmark modeled propagation of alcohol in the body [WT24]. Pharmacokinetics describes how drugs are distributed in different organs of the body. Compartment models are now important for screening of all drugs used by humans. The schematic diagram in Figure 3.19 illustrates the idea of a compartment model. Compartment models are also used in many other fields such as environmental science.

One-Compartment Model

The simplest dynamic model is obtained by assuming that the body behaves like a single compartment: that the drug is spread evenly in the body after

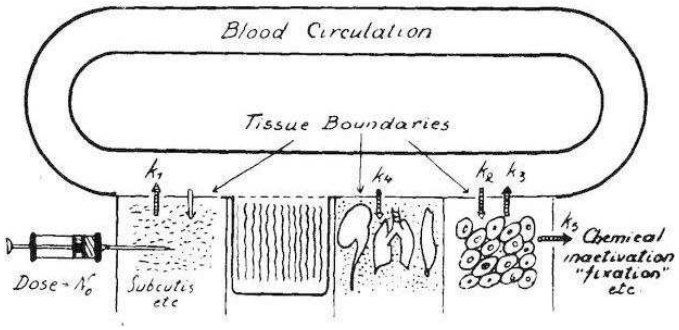

Figure 3.19: Schematic diagram of the circulation system (from Teorell [Teo37]).

it has been administered, and that it is then removed at a rate proportional to the concentration. Let c be the concentration, V the volume and q the outflow rate or the clearance. Converting the description of the system into differential equations, the model becomes

$$V\frac{dc}{dt} = -qc. \tag{3.20}$$

This equation has the solution

$$c(t) = c_0 e^{-qt/V} = c_0 e^{-kt},$$

which shows that the concentration decays exponentially after an injection. The input is introduced implicitly as an initial condition in the model (3.20). The way the input enters the model depends on how the drug is administered. The input can be represented as a mass flow into the compartment where the drug is injected. A pill that is dissolved can also be interpreted as an input in terms of a mass flow rate.

The model (3.20) is called a a *one-compartment model* or a *single pool model*. The parameter q/V is called the elimination rate constant. The simple model is often used in studies where the concentration is measured in the blood plasma. By measuring the concentration at a few times, the initial concentration can be obtained by extrapolation. If the total amount of injected substance is known, the volume V can then be determined as $V = m/c_0$; this volume is called the *the apparent volume of distribution*. This volume is larger than the real volume if the concentration in the plasma is lower than in other parts of the body. The model (3.20) is very simple and there are large individual variations in the parameters. The parameters V and q are often normalized by dividing with the weight of the person.

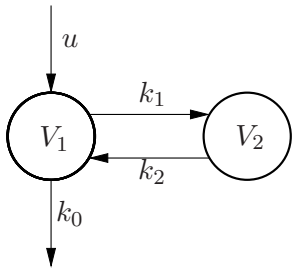

Figure 3.20: Schematic diagram of a model with two compartments.

Typical parameters for aspirin are $V = 0.2$ l/kg and $q = 0.01$ l/h/kg. These numbers can be compared with a blood volume of 0.07 l/kg, a plasma volume of 0.05 l/kg and intracellular fluid volume of 0.4 l/kg.

The simple one compartment model gives the gross behavior but it is based on strong simplifications. Improved models can be obtained by considering the body as composed of several compartments. We will work out the details for a system with two compartments.

Two-Compartment Model

Consider the system shown in Figure 3.20, where the compartments are represented as circles and the flows by arrows. We assume that there is perfect mixing in each compartment and that the transport between the compartments are driven by concentration differences. We further assume that a drug with concentration c_0 is injected in compartment 1 at a volume flow rate of u and that the concentration in compartment 2 is the output.

Let x_1 and x_2 be the total mass of the drug in the compartments and let V_1 and V_2 be the volumes of the compartments. A mass balance for the system gives

$$\begin{aligned}
\frac{dx_1}{dt} &= q(c_2 - c_1) - q_0 c_1 + c_0 u = q\Big(\frac{x_2}{V_2} - \frac{x_1}{V_1}\Big) - \frac{q_0}{V_1} c_1 + c_0 u \\
&= -(k_1 + k_0)x_1 + k_2 x_2 + c_0 u \\
\frac{dx_2}{dt} &= q(c_1 - c_2) = q\Big(\frac{x_1}{V_1} - \frac{x_2}{V_2}\Big) = k_1 x_1 - k_2 x_2 \\
y &= c_2 = \frac{1}{V_2} x_2,
\end{aligned}$$

where $k_0 = q_0/V_1$, $k_1 = q/V_1$ and $k_2 = q/V_2$. Introducing matrices, this

model can be written as

$$\begin{aligned} \frac{dx}{dt} &= \begin{pmatrix} -k_0 - k_1 & k_2 \\ k_1 & -k_2 \end{pmatrix} x + \begin{pmatrix} c_0 \\ 0 \end{pmatrix} u \\ y &= \begin{pmatrix} 0 & 1/V_2 \end{pmatrix} x. \end{aligned} \tag{3.21}$$

In this model we have used the total mass of the drug in each compartment as state variables. If we instead choose to use the concentrations as state variables, the model becomes

$$\begin{aligned} \frac{dc}{dt} &= \begin{pmatrix} -k_0 - k_1 & k_1 \\ k_2 & -k_2 \end{pmatrix} c + \begin{pmatrix} b_0 \\ 0 \end{pmatrix} u \\ y &= \begin{pmatrix} 0 & 1 \end{pmatrix} x, \end{aligned} \tag{3.22}$$

where $b_0 = c_0/V_1$. Mass is called an *extensive variable* and concentration is called an *intensive variable*.

The papers by Widmark and Tandberg [WT24] and Teorell [Teo37] are classics. Pharmacokinetics is now an established discipline with many textbooks [Dos68, Jac72, GP82]. Because of its medical importance pharmacokinetics is now an essential component of drug development. Compartment models are also used in other branches of medicine and in ecology. The problem of determining rate coefficients from experimental data is discussed in [BÅ70] and [God83].

3.7 Population Dynamics

Population growth is a complex dynamic process that involves the interaction of one or more species with their environment and the larger ecosystem. The dynamics of population groups are interesting and important in many different areas of social and environmental policy. There are examples where new species have been introduced in new habitats, sometimes with disastrous results. There are also been attempts to control population growth both through incentives and through legislation. In this section we describe some of the models that can be used to understand how populations evolve with time and as a function of their environment.

Simple Growth Model

Let x the the population of a species at time t. A simple model is to assume that the birth and death rates are proportional to the total population. This

gives the linear model

$$\frac{dx}{dt} = bx - dx = (b-d)x = rx \tag{3.23}$$

where birth rate b and death rate d are parameters. The model gives an exponential increase if $b > d$ or an exponential decrease if $B < d$. A more realistic model is to assume that the birth rate decreases when the population is large. The following modification of the model (3.23) has this property:

$$\frac{dx}{dt} = rx(1 - \frac{x}{x_c}) = f(x), \tag{3.24}$$

where x_c is the *carrying capacity* of the environment. The model (3.24) is called the *logistic growth* model.

Predator Prey Models

A more sophisticated model of population dynamics includes the effects of competing populations, where one species may feed on another. This situation, referred to as the *predator prey problem*, was already introduced in Example 2.3, where we developed a discrete time model that captured some of the features of historical records of lynx and hare populations.

In this section, we replace the difference equation model used there with a more sophisticated differential equation model. Let $H(t)$ represent the number of hares (prey) and $L(t)$ represent the number of lynxes (predator). The dynamics of the system are modeled as

$$\begin{aligned} \frac{dH}{dt} &= r_h H\left(1 - \frac{H}{K}\right) - \frac{aHL}{1 + aHT_h} \quad H \geq 0 \\ \frac{dL}{dt} &= r_l L\left(1 - \frac{L}{kH}\right) \qquad\qquad\qquad\quad L \geq 0. \end{aligned}$$

In the first equation, r_h represents the growth rate of the hares, K represents the maximum population of hares (in the absence of lynxes), a represents the interaction term that describes how the hares are diminished as a function of the lynx population, and T_h depends is a time constant for prey consumption. In the second equation, r_l represents the growth rate of the lynxes and k represents the fraction of hares versus lynxes at equilibrium. Note that both the hare and lynx dynamics include terms that resemble the logistic growth model (3.24).

Of particular interest are the values at which the population values remain constant, called *equilibrium points*. The equilibrium points for this

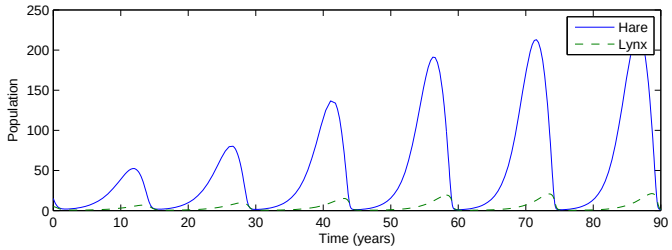

Figure 3.21: A simulation of the predator prey model with $r_h = 0.02$, $K = 500$, $a = 0.03$, $T_h = 5$, $r_l = 0.01$, $k = 0.2$ and time scale chosen to correspond to weeks.

system can be determined by setting the right hand side of the above equations to zero. Letting H_e and L_e represent the equilibrium state, from the second equation we have

$$L_e = kH_e.$$

Substituting this into the first equation, we must solve

$$r_h H_e \left(1 - \frac{H_e}{K}\right) - \frac{akH_e^2}{1 + aH_eT_h} = 0.$$

Multiplying through by the denominator, we get

$$\begin{aligned} 0 &= H_e \cdot \left(r_h \left(1 - \frac{H_e}{K}\right)(1 + aH_eT_h) - akH_e\right) \\ &= H_e \cdot \left(\frac{r_h aT_h}{K} H_e^2 + (ak + r_h/K - r_h aT_h)H_e - r_h\right). \end{aligned}$$

This gives one solution at $H_e = 0$ and a second that can be solved analytically or numerically.

Figure 3.21 shows a simulation of the dynamics starting from a set of population values near the nonzero equilibrium values. We see that for this choice of parameters, the simulation predicts an oscillatory population count for each species, reminiscent of the data shown in Figure 2.6 (page 48).

Fisheries Management

We end this section by discussing a control problem that has had significant impact on international legislation for fishing.

The dynamics of a commercial fishery can be described by the following simple model

$$\frac{dx}{dt} = f(x) - h(x, u), \tag{3.25}$$

where x be the total biomass, $f(x)$ the growth rate and $h(x,u)$ the harvesting rate. The logistic function (3.24) is a simple model for the growth rate and the harvesting can be modeled by

$$h(x,u) = axu, \tag{3.26}$$

where the control variable u is the harvesting effort, and a is a constant. The rate of revenue is

$$g(x,u) = bh(x,u) - cu, \tag{3.27}$$

where b and c are constants representing the price of fish and the cost of fishing. Using equations (3.26) and (3.27) we find that the rate of revenue is

$$g(x,u) = (abx - c)u.$$

In a situation where there are many fishermen and no concern for the environment, it is economic to fish as long as $abx > c$ and there will then be an equilibrium where the biomass is

$$x_\infty = \frac{c}{ab}, \tag{3.28}$$

which is the equilibrium with unrestricted fishing.

Assume that the population is initially at equilibrium at $x(0) = x_c$. The revenue rate with unrestricted fishing is then $(abx_c - c)u$, which can be very large. The fishing effort then naturally increases until the equilibrium (3.28), where the revenue rate is zero.

We can contrast unrestricted fishing with the situation for a single fishery. A typical case is when a country has all fishing rights in a large area. In such a case it is natural to maximize the rate of *sustainable revenue*. This can be accomplished by adding the constraint that the biomass x in equation (3.25) is constant, which implies that

$$f(x) = h(x,u).$$

Solving this equation for u gives

$$u = u_d(x) = \frac{f(x)}{ax}.$$

Inserting the value of u into equation (3.27) gives the following rate of revenue

$$\begin{aligned} g(x) &= bh(x,u_d) - cu_d(x) = \Big(b - \frac{c}{ax}\Big) f(x) \\ &= rx\Big(b - \frac{c}{ax}\Big)\Big(1 - \frac{x}{x_c}\Big) = \frac{r}{x_c}\Big(-abx^2 + (c + abx_c)x - cx_c\Big). \end{aligned} \tag{3.29}$$

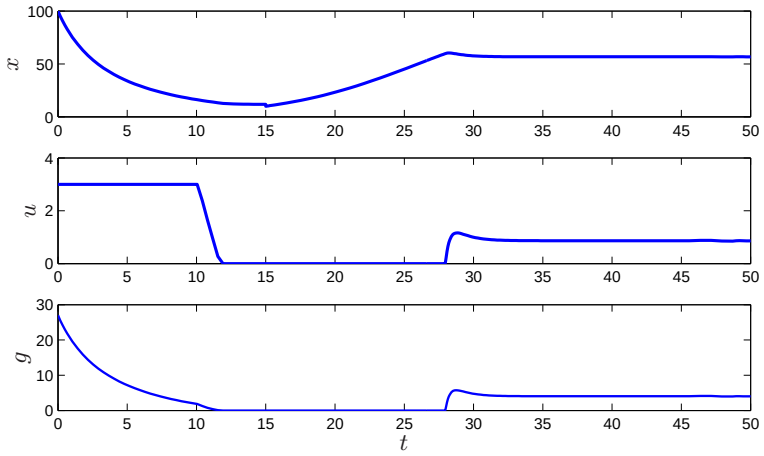

Figure 3.22: Simulation of a fishery. The curves show total biomass x, harvesting rate u and revenue rate g as a function of time t. The fishery is modeled by equations (3.25), (3.26), (3.27) with parameters $x_c = 100$, $a = 0.1$, $b = 1$ and $c = 1$. Initially fishing is unrestricted at rate $u = 3$, at time $t = 15$ fishing is changed to harvesting at a sustainable rate, accomplished by a PI controller with parameters $k = 0.5$ and $k_i = 0.5$.

The rate of revenue has a maximum

$$r_0 = \frac{r(c - abx_c)^2}{4abx_c}, \tag{3.30}$$

for

$$x_0 = \frac{x_c}{2} + \frac{c}{2ab}. \tag{3.31}$$

Figure 3.22 shows a simulation of a fishery. The system is initially in equilibrium with $x = 100$. Fishing begins with constant harvesting rate $u = 3$ at time $t = 0$. The initial revenue rate is large, but it drops rapidly as the population decreases. At time $t = 12$ the revenue rate is practically zero. The fishing policy is changed to a sustainable strategy at time $t = 15$. This is accomplished by using a PI controller where the reference is the optimal sustainable population size $x_0 = 55$, given by equation (3.31). The feedback stops harvesting for a period but the biomass increases rapidly. At time $t = 28$ the harvesting rate increases rapidly and a sustainable steady state is reached in a short time.

Volume I of the two volume set by J. Murray [Mur04] give a broad coverage of population dynamics. Maintaining a sustainable fish population

is a global problem that has created many controversies and conflicts. A detailed mathematical treatment is given in [**?**]. The mathematical analyses has influenced international agreement on fishing.

3.8 Exercises

1. Consider the cruise control example described in Section 3.1. Build a simulation that recreates the response to a hill shown in Figure 3.3b and show the effects of increasing and decreasing the mass of the car by 25%. Redesign the controller (using trail and error is fine) so that it returns to the within 10% of the desired speed within 3 seconds of encountering the beginning of the hill.

2. Consider the inverted pendulum model of the bicycle given in Figure 3.6. Assume that the block labeled body is modeled by equation (3.5) and that the front fork is modeled by (3.6). Derive the equations for the closed loop. Show that when $T = 0$ the equation is the same as for a mass spring damper system. Also show that the spring coefficient is negative for low velocities but positive if the velocity is sufficiently large.

3. Show that the dynamics of a bicycle frame given by equation (3.5) can be written in state space form as

$$\frac{d}{dt}\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} 0 & mgh/J \\ 1 & 0 \end{pmatrix}\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} + \begin{pmatrix} 1 \\ 0 \end{pmatrix} u$$
$$y = \begin{pmatrix} \frac{Dv_0}{bJ} & \frac{mv_0^2 h}{bJ} \end{pmatrix} x,$$

   where the input u is the torque applied to the handle bars and the output y is the title angle φ. What do the states x_1 and x_2 represent?

4. Combine the bicycle model given by equation (3.5) and the model for steering kinematics in Example 2.8 to obtain a model that describes the path of the center of mass of the bicycle.

5. Consider the op amp circuit shown below:

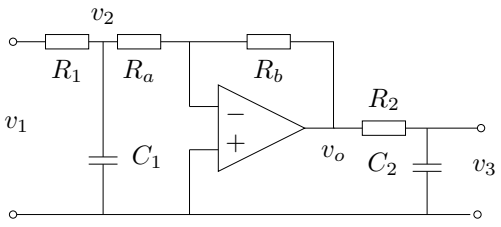

Show that the dynamics can be written in state space form as

$$\frac{dx}{dt} = \begin{pmatrix} -\frac{1}{R_1C_1} - \frac{1}{R_aC_1} & 0 \\ \frac{R_b}{R_a}\frac{1}{R_2C_2} & -\frac{1}{R_2C_2} \end{pmatrix} x + \begin{pmatrix} \frac{1}{R_1C_1} \\ 0 \end{pmatrix} u$$
$$y = \begin{pmatrix} 0 & 1 \end{pmatrix} x$$

where $u = v_1$ and $y = v_3$. (Hint: Use v_2 and v_3 as your state variables.)

6. (Atomic force microscope) A simple model for the vertical motion of the scanner is shown in Figure 3.18b, where the system is approximated with two masses. The mass m_1 is half of the piezo crystal and the mass m_2 is the other half of the piezo crystal and the mass of the support. A simple model is obtained by assuming that the piezo crystal generates a force F between the masses and that there is a damping c in the spring. Let the positions of the center of the masses be x_1 and x_2, and let the elongation of the piezo stack is $u = x_1 - x_2$. A momentum balance gives the following model for the system.

$$m_1\frac{d^2x_1}{dt^2} = F$$
$$m_2\frac{d^2x_2}{dt^2} = -c\frac{dx_2}{dt} - kx_2 - F$$
$$u = x_1 - x_2.$$

Review the assumptions made in the simplified model. Let the elongation u of the piezo stack be the control variable and the height of the sample x_1 be the output. Show that the relation between x_1 and u is given by

$$(m_2 - m_1)\frac{d^2x_1}{dt^2} + c\frac{dx_1}{dt} + kx_1 = m_1\frac{d^2u}{dt^2} + c\frac{du}{dt} + ku.$$

Simulate the system and show that the response is qualitatively the same as the one shown in Figure 3.18a. Can the parameters of the model be determined from a step response experiment of the type shown in Figure 3.18a?

7. (Drug administration) Consider the compartment model in Figure 3.20. Assume that there is no outflux, i.e. $k_0 = 0$. Compare the models where the states are masses and concentrations. Compute the steady state solutions for the different cases. Give a physical interpretation of the results.

8. (Drug administration) Show that the model represented by the schematic diagram in Figure 3.19 can be represented by the compartment model shown below:

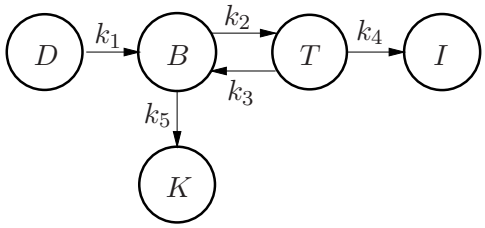

where compartment D represents the issue where the drug is injected, compartment B represents the blood, compartment T represents tissue where the drug should be active, compartment K the kidney where the drug is eliminated, and I a part of the body where the drug is inactive.

Write a simulation for the system and explore how the amount of the drug in the different compartments develops over time. Relate you observations to your physical intuition and the schematic diagram above. Modify your program so that you can investigate what happens if the drug is injected directly to the blood stream, compartment B, instead of in compartment D.

9. (Drug administration) The metabolism of alcohol in the body has can be modeled by the nonlinear compartment model

$$V_b \frac{dc_b}{dt} = q(c_l - c_b) + q_{iv}$$
$$V_l \frac{dc_l}{dt} = q(c_b - c_l) - q_{max} \frac{c_l}{c_0 + c_l} + q_{gi}$$

where $V_b = 48$ l and $V_l = 0.6$ l are the effective distribution volume of body water and liver water, c_b and c_l the corresponding concentrations of alcohol, q_{iv} and q_{gi} are the injection rates for intravenously and gastrointestinal intake, $q = 1.5$ l/min is the total hepatic blood flow,

$q_{max} = 2.75$ mmol/min and $k_m = 0.1$ mmol. Simulate the system and compute the concentration in the blood for oral and intravenous doses of 12 g and 40 g of alcohol.

10. (Population dynamics) Consider the model for logistic growth given by equation (3.24). Show that the maximum growth rate occurs when the size of the population is half of the steady state value.

11. (Population dynamics) Verify the curves in Figure 3.21 by creating a program that integrates the differential equations.