CSE 538 Notes

Roch Guérin
guerin@wustl.edu

January 2, 2014

Chapter 2

A Brief Introduction to Fluid-Flow Models

This chapter describes a model for evaluating the performance of systems where the *arrival of work* is proportional to time. Such models arise commonly in communication networks where work is often in the form of *packets* that require transmission on network links. Hence, work is in proportion to the size of the packet to be transmitted, with larger packets taking longer to arrive on the system's input link(s). It is this *correlation* between service times and inter-arrival times that we seek to capture. The introduction of correlation typically makes for a much more complex analysis, *e.g.,* see [1] for an example, and the goal is to capture the main feature of correlation, without adding too much complexity. Fluid-flow models offer such an intermediate alternative.

2.1 Limitations of Point Processes

Before describing the main characteristics of fluid-flow models, we highlight the impact that failure to account for correlation (between arrivals and services times) can have.

Specifically, consider the traditional point process based approach used by, say, M/M/1 or M/G/1 queues, to represent the arrival of work (packets) at a communication link/multiplexer. Those models do not account for the fact that the finite speed of (input) links couple inter-arrival times and service times. This is not unreasonable when considering a multiplexer with many input links, each of higher speed than the output link, where this coupling is weak, but can result in overly pessimistic outcomes in other cases. Consider for example the configuration of Fig. 2.1, where a single high-speed link, *e.g.,* a 1 Gbps Ethernet link, feeds packets to a buffer connected to a slower 100 Mbps Ethernet wide-area uplink. The presence of a single, finite speed input link imposes a tight upper bound on the maximum amount of data that can arrive to the buffer during any period of time. This constraint on the arrival process in turn limits the amount of buffering required[1].

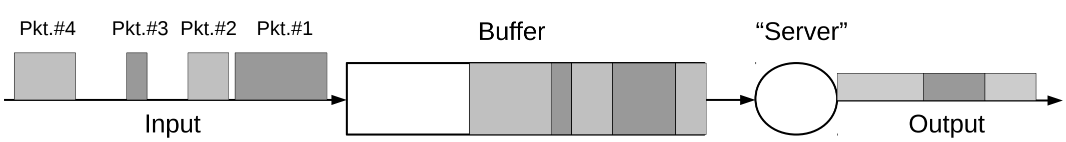

Figure 2.1: A simple link model.

For illustration purposes, assume that the 1 Gbps Ethernet input link of Figure 2.1 is active (receiving packets) only 10 % of the time, *i.e.,* the average incoming bit rate is 100 Mbps or exactly the speed of the output link, *i.e.,* for a resulting load of $\rho = 1$. If we further assume that the size of incoming packets follows (approximately) an

[1] *E.g.,* when input and output links speeds are equal and the output link is work conserving, the buffer never stores more than one packet.

exponential distribution with mean $2,000$ bits, we could rely on an M/M/1 queue to model the system (the packet arrival rate is $\lambda = 50,000$ packets/sec). Under an M/M/1 model, a load of $\rho = 1$ implies an infinite average waiting time for incoming packets (assuming an infinite buffer as well). This is, however, a gross over-estimate, in part because it assumes arrival patterns that are not all feasible. For example, under a Poisson arrival process of rate $\lambda = 50,000$ packets/sec, the probability that two consecutive arrivals are less than 1μsec apart is about 0.05, and given exponentially distributed packet sizes with mean $2,000$ bits, the probability that the second packet is larger than the average packet size of $2,000$ bits is about 0.37. Hence, an M/M/1 model allows two consecutive packets that are less than 1μsec apart with the second one being larger than $2,000$ bits with a probability of about 0.018. However, we know that this is physically not possible, *i.e.,* has zero probability, since in 1μsec no more than $1,000$ bits can arrive on a 1 Gbps link. Our goal is to develop models that can capture this constraint.

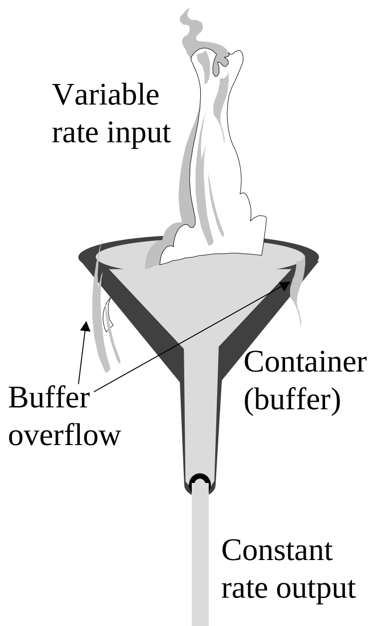

Figure 2.2: A fluid model analogy.

2.2 Fluid-Flow Models Overview

The basic premise of a fluid-flow model is that the arrival of work (packets) on a link is *progressive* with a rate (in bits/sec) that is a function of the input link speed. Packet transmissions (service times) are also progressive, but at a rate that depends on the speed of the output link. The main simplification that makes fluid-flow models tractable is that they ignore packet boundaries. This is best understood through an analogy that further illustrates the similarity to fluid systems. In the context of a fluid-flow model for the system of Fig. 2.1, the input link is a *pipe* through which bits *flow* at a maximum rate of 1 Gbps, the buffer is a *container* that stores the bits as they arrive, and the output link is another *pipe* that drains bits out of the container at a maximum rate of 100 Mbps. The flow of bits on the input pipe can be turned on and off (see Fig. 2.3) as packets transmissions start and stop. Conversely, any bit is deemed available for transmission over the output pipe, as soon as it has been received. This is an approximation of the common behavior of packet networks, which typically operate in a store-and-forward fashion. This primarily affects packets that arrive to find an empty buffer or a low level of buffer occupancy[2]. Hence, the impact should be small unless the (output) link load is very low.

In the rest of this chapter, we describe and analyze this basic model, with a focus on the probability distribution of the buffer content, *e.g.,* $P(B \leq x)$, where B is a random variable denoting the number of bits in the buffer, and x is in bits.

2.3 Buffer Content Distribution

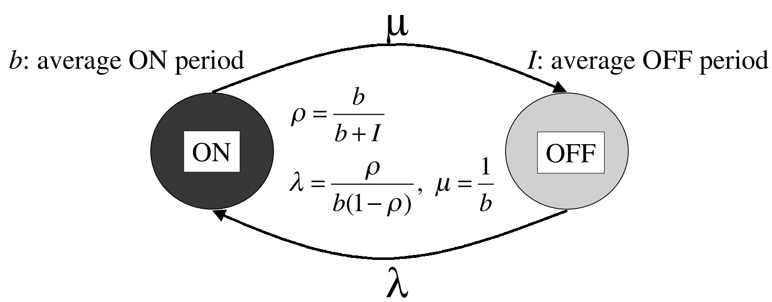

Figure 2.3: Source model.

2.3.1 Notation and System Model

We assume that the input to the system is characterized by a source that alternates between active and idle states, *i.e.,* an ON-OFF source. In the ON state, the source transmits data (bits) at its peak rate R. In the OFF state, the source does not transmit any data. The source is further characterized by the distributions of its ON and OFF periods that are both take to be exponential. The average duration of an ON period is denoted as b, and the average duration of an OFF period is denoted as I. The *source utilization* (probability π_1 that the source is in the ON state) is denoted as ρ, and is readily seen to be equal to the ratio of the average ON period to the sum of the average ON and OFF periods, *i.e.,* the duty cycle of the source or $\rho = b/(b+I)$. Note that conversely, the probability that the source is OFF is $\pi_0 = 1 - \rho$.

Under the assumption of exponentially distributed ON and OFF states, the source is essentially a two-state (ON and OFF) continuous time Markov chain as shown in Fig. 2.3. The transition rate λ out of the OFF state and the transition rate μ out of the ON state can be expressed as functions of b and ρ, namely

$$\begin{aligned} \lambda &= \frac{\rho}{b(1-\rho)} \\ \mu &= \frac{1}{b} \end{aligned} \tag{2.1}$$

The source, when active, feeds data into a buffer that is being emptied (served) by a link (server) of speed C. Unless stated otherwise, we assume that the buffer is infinite. When considering finite buffer systems, X will be used to denote the buffer size. Our goal is to obtain an expression for the buffer content distribution $P(B \leq x), x \geq 0$, where B is the random variable denoting the buffer content, and we proceed with this derivation next.

[2]Note that as long as the input link speed is higher than that of the output, packet transmissions on the output still proceed continuously once started even if their transmission starts before the arrival of the last bit of the packet.

2.3.2 System State and Evolution

Let the tuple (t, x) denote the state of the buffer at time t, *i.e.,* the buffer content is $\leq x$ at time t, and let $P_i(t, x)$ be the probability that at time t the source is in state i ($i = 0$ if the source is idle and $i = 1$ if the source is active) and the buffer content is at most x. Our first step is to express the evolution over time of the quantities $P_i(t, x)$, $i = 0, 1$, based on the characteristics of the source and the server. Specifically, the evolution of the system can be described by the following set of equations:

$$\begin{aligned} P_0(t+dt, x) &= (1-\lambda dt)P_0(t, x+Cdt) + \mu dt P_1(t, x-(R-C)dt) \\ P_1(t+dt, x) &= \lambda dt P_0(t, x+Cdt) + (1-\mu dt)P_1(t, x-(R-C)dt). \end{aligned} \tag{2.2}$$

Let us consider first the equation giving $P_0(t+dt, x)$. There are two possible ways for the buffer content to be below x at time $t+dt$ with the source being in state $i=0$. One possibility is if the source was in state $i=0$ at time t and did not make a transition to state $i=1$ in the interval $t, t+dt)$ (this has probability $(1-\lambda dt)$), and the buffer content was below $x+Cdt$ at time t (the link empties Cdt during the interval $[t, t+dt)$). The second possibility is that the source was in state $i=1$ at time t and that it switched to state $i=0$ in the interval $t, t+dt)$ (this has probability μdt), and the buffer content at time t was below $x-(R-C)dt$ ($(R-C)$ is the rate at which the buffer fills when the source is active). A similar reasoning can be carried out for $P_1(t+dt, x)$ to yield the second equation.

Our next step is to rewrite Eq. (2.2) so as to ultimately be able to express it in terms of differential equations. Specifically, we can rewrite the first expression in Eq. (2.2) as follows:

$$\begin{aligned} & P_0(t+dt, x) - P_0(t, x+Cdt) = \\ & [-\lambda P_0(t, x+Cdt) + \mu P_1(t, x-(R-C)dt)]\, dt \\ \Rightarrow \quad & \frac{P_0(t+dt, x) - P_0(t, x)}{dt} + C\frac{P_0(t, x) - P_0(t, x+Cdt)}{Cdt} = \\ - \quad & \lambda P_0(t, x+Cdt) + \mu P_1(t, x-(R-C)dt). \end{aligned}$$

Letting $dt \to 0$ in the above expression gives

$$\frac{\partial P_0(t, x)}{\partial t} - C\frac{\partial P_0(t, x)}{\partial x} = -\lambda P_0(t, x) + \mu P_1(t, x) \tag{2.3}$$

Similarly, we can rewrite the second expression in Eq. (2.2) as follows:

$$\begin{aligned} & P_1(t+dt, x) - P_1(t, x-(R-C)dt) = \\ & [\lambda P_1(t, x+Cdt) - \mu P_1(t, x-(R-C)dt)]\, dt \\ \Rightarrow \quad & \frac{P_1(t+dt, x) - P_1(t, x)}{dt} + (R-C)\frac{P_1(t, x) - P_1(t, x-(R-C)dt)}{(R-C)dt} = \\ & \lambda P_1(t, x+Cdt) - \mu P_1(t, x-(R-C)dt). \end{aligned}$$

Letting again $dt \to 0$ in the above expression gives

$$\frac{\partial P_1(t, x)}{\partial t} + (R-C)\frac{\partial P_1(t, x)}{\partial x} = \lambda P_0(t, x) - \mu P_1(t, x) \tag{2.4}$$

2.3.3 A Matrix Differential Equation for the Stationary Distribution

We have so far obtained a set of partial differential equations that characterize the evolution of the buffer content over time. However, we are primarily interested in the stationary behavior, and assuming the existence of a stationary distribution independent of t, we get

$$\begin{aligned} -CF_0'(x) &= -\lambda F_0(x) + \mu F_1(x) \\ (R-C)F_1'(x) &= \lambda F_0(x) - \mu F_1(x), \end{aligned} \tag{2.5}$$

where $F_i(x) = \lim_{t\to\infty} P_i(t,x)$, $i = 0, 1$, gives the stationary buffer content distribution in state i, and where $f'(x)$ denotes the derivative of $f(x)$ with respect to x. Using the vector notation $F(x) = [F_0(x), F_1(x)]^T$, Eq. (2.5) can be expressed in matrix form as follows:

$$\begin{bmatrix} -C & 0 \\ 0 & -(C-R) \end{bmatrix} F'(x) = \begin{bmatrix} -\lambda & \mu \\ \lambda & -\mu \end{bmatrix} F(x),$$

where again $F'(x)$ is the derivative of $F(x)$ with respect to x.

The next step is to solve the above matrix differential equation. It can be rewritten as follows:

$$F'(x) = \begin{bmatrix} \frac{\lambda}{C} & -\frac{\mu}{C} \\ \frac{\lambda}{R-C} & -\frac{\mu}{R-C} \end{bmatrix} F(x) \tag{2.6}$$

In order to solve the above matrix differential equation, we first "guess" that a solution is of the form $F_i(x) = \beta_i + \gamma_i e^{\delta x}$. Using this expression in Eq. (2.6) gives:

$$\begin{aligned} \gamma_0 \delta e^{\delta x} &= \frac{\lambda}{C}\beta_0 - \frac{\mu}{C}\beta_1 + e^{\delta x}\left(\frac{\lambda}{C}\gamma_0 - \frac{\mu}{C}\gamma_1\right) \\ \gamma_1 \delta e^{\delta x} &= \frac{\lambda}{R-C}\beta_0 - \frac{\mu}{R-C}\beta_1 + e^{\delta x}\left(\frac{\lambda}{R-C}\gamma_0 - \frac{\mu}{R-C}\gamma_1\right) \end{aligned}$$

Concentrating first on the constant term, we easily find $\lambda\beta_0 = \mu\beta_1$, from which we deduce:

$$\beta_0 = \alpha_0\mu \text{ and } \beta_1 = \alpha_0\lambda \tag{2.7}$$

Focusing next on the exponential term, we get

$$\begin{aligned} \gamma_0\delta &= \frac{\lambda}{C}\gamma_0 - \frac{\mu}{C}\gamma_1 \\ \gamma_1\delta &= \frac{\lambda}{R-C}\gamma_0 - \frac{\mu}{R-C}\gamma_1 \end{aligned}$$

Eliminating δ yields

$$\frac{\gamma_1}{\gamma_0} = \frac{C}{R-C},$$

which gives

$$\gamma_0 = \alpha_1\frac{R-C}{C} \text{ and } \gamma_1 = \alpha_1 \tag{2.8}$$

It now remains to obtain an expression for δ. From the above expression for the ratio γ_1/γ_0 and the previous equations, we get:

$$\delta = \frac{\lambda}{C} - \frac{\mu}{R-C} \tag{2.9}$$

Combining Eqs. (2.7) and (2.8), and (2.9) with Eq. (2.1), we get the following expression for $F(x)$:

$$F(x) = \left[\begin{array}{c} F_0(x) \\ F_1(x) \end{array} \right] = \alpha_0 \left[\begin{array}{c} 1/b \\ \frac{\rho}{b(1-\rho)} \end{array} \right] + \alpha_1 \left[\begin{array}{c} \frac{R-C}{C} \\ 1 \end{array} \right] \times e^{-\left[\frac{x(C-\rho R)}{b(1-\rho)(R-C)C}\right]}, \tag{2.10}$$

where the coefficients α_0 and α_1 can be determined from boundary conditions. However, before we proceed with deriving expressions for α_0 and α_1, we first obtain an expression for the buffer content distribution based on Eq. (2.10).

Recall that $F_i(x)$ corresponds to the stationary probability that the buffer content is less than or equal to x and the source feeding it is in state i. Since states 0 and 1 are complementary (the source is in either one of them), we readily obtain:

$$P(B \leq x) = F_0(x) + F_1(x).$$

This gives

$$P(B \leq x) = \frac{\alpha_0}{b(1-\rho)} + \alpha_1 \frac{R}{C} e^{-\left[\frac{x(C-\rho R)}{b(1-\rho)(R-C)C}\right]} \tag{2.11}$$

Boundary conditions depend on the type of system being considered, and we first derive expressions for α_0 and α_1 assuming an infinite buffer system.

Infinite Buffer Case

There are two boundary conditions we can identify. The first is obtained by noticing that the buffer cannot be empty when the source is active, *i.e.,* we have $F_1^{\infty}(0) = 0$. This first relation gives

$$\alpha_0 \frac{\rho}{b(1-\rho)} + \alpha_1 = 0 \tag{2.12}$$

Assuming a stable system so that a stationary distribution exists, we get a second boundary condition from letting $x \to \infty$ in either Eq. (2.10) or Eq. (2.11), which yields

$$\alpha_0 = b(1-\rho) \tag{2.13}$$

Combining Eqs. (2.12) and (2.13) in Eq. (2.11) gives

$$P_{\infty}(B \leq x) = 1 - \frac{\rho R}{C} e^{-\left[\frac{x(C-\rho R)}{b(1-\rho)(R-C)C}\right]} = 1 - \frac{\rho R}{C} e^{-\left[\frac{x(1-\rho R/C)}{b(1-\rho)(R-C)}\right]}, \tag{2.14}$$

where ρR is the average rate generated by the source, so that $\frac{\rho R}{C}$ corresponds to the system's load.

Finite Buffer Case

In this section, we assume that the available buffer size is finite and limited to X. Our goal is again to derive an explicit expression for the buffer content distribution $P_X(B \leq x)$ by using boundary conditions to determine the two unknown quantities α_0 and α_1.

The boundary condition $F_1(0) = 0$ still holds even assuming a finite buffer, so that Eq. (2.12) remains valid. The second boundary condition, *i.e.,* for $x \to \infty$ is, however, not available anymore because of the finite buffer assumption. Nevertheless, a similar boundary condition can be obtained by focusing on the system state when the queue is full. A full queue corresponds to an event with a non-zero probability mass at $x = X$. This then gives rise to the equivalent of Eq. (2.13)

$$\pi_i = P_X^{i,ovflw} + F_i(X^-), \tag{2.15}$$

where $P_X^{i,ovflw}$ is the probability mass associated with an overflow event when the source is in state i. Overflow is obviously not possible when the source is in state $i = 0$, so that $P_X^{0,ovflw} = 0$. Using this in Eq. (2.15) gives

$$\frac{\alpha_0}{b} + \alpha_1 \frac{(R-C)}{C} e^{-\left[\frac{X(C-\rho R)}{b(1-\rho)(R-C)C}\right]} = 1 - \rho \tag{2.16}$$

From combining Eqs. (2.12) and (2.16), we get:

$$\begin{aligned} \alpha_0 &= \frac{b(1-\rho)^2 C}{\Delta} \\ \alpha_1 &= -\frac{\rho(1-\rho)C}{\Delta}, \end{aligned} \tag{2.17}$$

where

$$\Delta = (1-\rho)C - \rho(R-C)e^{-\frac{X(C-\rho R)}{b(1-\rho)(R-C)C}}.$$

Combining this with Eq. (2.11) gives the following expression for the buffer content distribution:

$$P_X(B \leq x) = \begin{cases} \frac{(1-\rho)C}{\Delta} - \frac{\rho(1-\rho)R}{\Delta} e^{-\frac{x(C-\rho R)}{b(1-\rho)(R-C)C}} & \text{for } 0 \leq x < X, \\ 1 & \text{for } x \geq X. \end{cases} \tag{2.18}$$

The buffer overflow probability, P_X^{ovflw}, can then be obtained from Eq. (2.15) or directly from the above expression by computing the "missing" probability mass between the values X^- and X. Using either approach, one obtains:

$$P_X^{ovflw} = P_X^{1,ovflw} = \frac{\rho(C-\rho R)e^{-\frac{X(C-\rho R)}{b(1-\rho)(R-C)C}}}{(1-\rho)C - \rho(R-C)e^{-\frac{X(C-\rho R)}{b(1-\rho)(R-C)C}}} \tag{2.19}$$

Note, however, that in practice we are interested in the overflow probability *given that the source is active*. This probability is readily obtained by dividing the above expression by the probability $\pi_1 = \rho$ that the source is active.

$$\overline{P}_X^{ovflw} = \frac{(C-\rho R)e^{-\frac{X(C-\rho R)}{b(1-\rho)(R-C)C}}}{(1-\rho)C - \rho(R-C)e^{-\frac{X(C-\rho R)}{b(1-\rho)(R-C)C}}} \tag{2.20}$$

Eq. (2.20) can be inverted to obtain the buffer size $\tilde{X}$ that will ensure a (conditional) buffer overflow probability below a desired value ϵ. Setting $\overline{P}_X^{ovflw} = \epsilon$ in Eq. (2.20) and solving for $\tilde{X}$ gives:

$$\tilde{X} = \ln\left[\frac{(C-\rho R) + \epsilon\rho(R-C)}{\epsilon(1-\rho)C}\right] \times \frac{b(1-\rho)(R-C)C}{C-\rho R}. \tag{2.21}$$

2.4 Effective Bandwidth of a Fluid ON-OFF Source

A symmetric quantity of interest to the buffer size $\tilde{X}$ needed to ensure a target (conditional) overflow probability ϵ for a given link capacity C, is the required link capacity $\widehat{c}$ (often called effective bandwidth or equivalent capacity) to ensure a (conditional) overflow probability ϵ given a buffer size of X.

Letting $\overline{P}_X^{ovflw} = \epsilon$ in Eq. (2.20) can be rewritten as

$$\epsilon = \beta e^{-\delta X}, \tag{2.22}$$

where

$$\beta = \frac{(c - \rho R) + \epsilon\rho(R - c)}{(1 - \rho)c} \qquad \text{and} \qquad \delta = \frac{c - \rho R}{b(1 - \rho)(R - c)c},$$

Computing the effective bandwidth $\widehat{c}$ then calls for solving the transcendental equation Eq. (2.22) for c, which can only be done numerically. However, a reasonable approximation is available by assuming $\beta = 1$, which can be shown to hold in most cases (actually we always have $\beta \leq 1$). With this simplification, $\widehat{c}$ can be obtained simply by solving a second order equation.

$$\begin{aligned} \alpha &= \frac{X(c - \rho R)}{b(1 - \rho)(R - c)c} \\ \Rightarrow 0 &= c^2\alpha b(1 - \rho) + c[X - b\alpha(1 - \rho)R] - \rho X R\,, \end{aligned} \tag{2.23}$$

where $\alpha = \ln 1/\epsilon$. Solving equation (2.23) gives the following expression for the effective bandwidth $\widehat{c}$ for the flow,

$$\widehat{c} = \frac{\alpha b(1 - \rho)R - X + \sqrt{[\alpha b(1 - \rho)R - X]^2 + 4x\alpha b\rho(1 - \rho)R}}{2\alpha b(1 - \rho)} \tag{2.24}$$

Bibliography

[1] I. Cidon, R. Guerin, A. Khamisy, and M. Sidi, "Analysis of a correlated queue in a communication system," *IEEE Trans. Infor. Theory*, vol. 39, no. 2, March 1993.