

A TRANSIENT ANALYSIS OF A DATA NETWORK WITH A PROCESSOR-SHARING SWITCH

Alan A. Weiss and Debasis Mitra

Alan A. Weiss is a member of technical staff and Debasis Mitra is head of the Mathematics of Networks and Systems Department of AT&T Bell Laboratories in Murray Hill, New Jersey. Mr. Weiss joined the company in 1981 and is working on applications of the theory of large deviations to computers and communications. He has a B.S. in physics and a B.S. in mathematics from Case Western Reserve University, Cleveland, Ohio; and an M.A. and a Ph.D. in mathematics from the Courant Institute of Mathematical Sciences, New York University. Mr. Mitra joined AT&T in 1968 and is responsible for research on networks and systems in communications, computing, and manufacturing. He has a B.Sc. and a Ph.D. in electrical engineering from the University of London.

This paper gives a simple, accurate, first-order asymptotic analysis of the transient behavior of a data network. The network is based on a switch, such as the Datakit® virtual circuit switch, which receives packetized traffic from many message sources of diverse classes and multiplexes the traffic over a trunk line. The network has a regulatory mechanism based on acknowledgments that produces a closed model. Assuming that the switch has a processor-sharing discipline, and assuming an asymptotic regime of high loading of the switch and high-capacity trunks, we derive the explicit, first-order transient behavior of the means of queue lengths—and hence response times. We then give a simple procedure for finding the time constants (eigenvalues) that govern the approach to steady state. Our numerical experiments show that the analysis is quite accurate.

Introduction

In a model of a data network based on a switch, such as the Datakit virtual circuit switch, packetized traffic flows from many message sources to the switch, which multiplexes the traffic over a trunk line.¹ How can one calculate the transient response of such a network? Even a simple M/M/1 model has complications. When there are several classes of message sources, a closed network, and a processor-sharing discipline in the switch, the problem would seem to be analytically insurmountable. Yet it is important to know how such a switch responds, especially to heavy traffic. Analytic formulae are desirable because they show dependence on parameters. As Fritz John of the Courant Institute once remarked, "One asymptotic (analytic) expression is worth 1000 pages of printout."

For a simple model of a switch, asymptotic expressions can, in fact, be derived. Our asymptotic regime has heavy traffic and high-

capacity trunks, where a fluid approximation may be expected to hold.^{2,3} To show the accuracy of our analysis, we compare some exact numerical calculations with our asymptotics.

These fluid limits (or fluid models, as we shall also call them) are systems of ordinary differential equations. In our case, the system of first-order equations is nonlinear and has as many dimensions as there are traffic classes, for which a typical value is three.⁴ On the other hand, in the exact description of the Markovian system, the number of equations is exponential in the number of classes and is typically in the hundreds of thousands.

In its primary use, the fluid model is numerically integrated to yield values of the aggregated state of the switch as a function of time. We also use it here to obtain the time constants of the network. First, we get an explicit representation of the *linear* system of differential equations that describes the behavior of the fluid model near the equilibrium state. From this, we obtain its eigenvalues.

The work reported here originated in questions relating to the behavior of the Datakit virtual circuit switch. The service disciplines that have been studied for this and related switches are generally based on the round-robin discipline, with other features, such as priority, added on.^{4,5} The reason for depending on round-robin is because it provides a service in which the mean delay seen by sources conditioned on message length is approximately proportional to the message length. (This relation is exact for processor sharing.⁶) This property is highly desirable in a data communications environment where diverse classes of sources exist. For example, interactive terminals have messages of only a few characters, while host-to-host file transfers have considerably longer messages.⁴ Processor sharing is the discipline obtained in the limit as the quantum of service in round-robin goes to zero.⁷ Because packet size actually defines the quantum of service, processor sharing is only an approximation, but one with a long history of being reasonably accurate and robust.

There are many potential applications of the knowledge of time constants and transient behavior of the network, although none is pursued in any depth in this paper. Such knowledge would be especially important in the design of network flow-control schemes. For example,

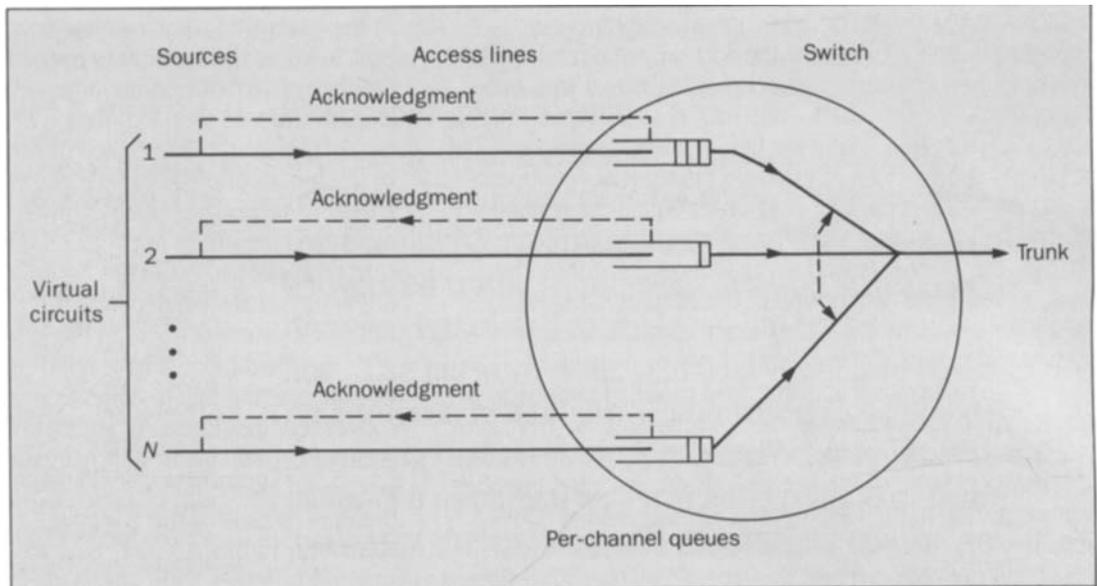
in the control of very high bandwidth sources, it is important to know the sensitivity of the time constants to the number of virtual circuits in that class. Also, it is known that, in the real-time estimation of traffic parameters from measurements, the variance of sample means as a function of the length of observations requires knowledge of the transient response.⁸ Such estimation procedures are an essential element in network control schemes. Yet another related application is in the design of simulations: Samples are approximately independent if they are separated in time by an amount equal to the system time constant.

Fluid models have been considered both as effective approximations to complex systems and as limits of convergent sequences of approximating systems in learning theory,⁹ stochastic approximations,¹⁰ other recursive stochastic algorithms,^{11,12} and queueing theory.¹³ Note that the fluid models of this paper are stripped of all stochastic elements. In this respect, they are distinct from, and considerably simpler than, *stochastic* fluid models. (See, for example, Reference 14 and accompanying references.)

The closed, product-form queueing network of this paper has been extensively studied in connection with the distributions of queue lengths¹⁵ and the waiting-time moments and distributions¹⁶ for various asymptotic regimes, including the present one, but only with the system in equilibrium. Even though the analysis in Reference 16 is quite different from this paper's, it is worth noting that the waiting-time distributions are also obtained by solving a system of a like number of ordinary differential equations. An earlier paper of ours considered the transient behavior of Erlang's model, which is the model considered here specialized to one class of message sources.¹⁷ However, this earlier paper considered various asymptotic regimes and also developed multiple-term asymptotic expansions and a large deviation theory.

In this paper, we first describe a system with one class of customers, and analyze this system. We then describe and analyze systems with more than one class of customers. The final section gives extracts from the numerical investigations. We examine three networks with parameters selected from the data traffic model of Fraser and Morgan⁴ and contrast the "exact" transient solutions with the fluid approximation.

Figure 1. Physical system of a single-class model.



6

The Single-Class Model

Figure 1 shows the physical system, which is similar to that of Fraser and Morgan.⁴ Each of N statistically homogeneous, asynchronous sources of virtual circuits sends a message over relatively low-speed access lines, and waits for an acknowledgment from the switch. It waits until the acknowledgment is received, at which point it starts a new cycle by transmitting another message. The length of the time interval required to transmit the message over access lines is the "think time."

The switch is constantly taking nibbles from the occupied per-channel queues and transmitting these nibbles over the high-speed trunk. The switch discipline is approximated by processor sharing.⁷ When a per-channel queue empties, the switch sends an acknowledgment to the matching source.

Let

$$r \triangleq \text{mean think rate} = \frac{\text{access line speed}}{\text{mean message length}} \quad (1a)$$

and

$$\lambda \triangleq \text{mean service rate at switch} = \frac{\text{trunk speed}}{\text{mean message length}}$$

If the mean acknowledgment time is significant, it is simply added to the mean think time. The think time and the service time for each message are assumed to be independent, exponentially distributed, random variables. The following dimensionless parameter is an important indicator of traffic conditions:

$$\gamma = \frac{1}{N} \frac{\lambda}{r} \quad (1b)$$

We will make the following assumption of heavy usage:

$$\gamma < 1 \quad (2)$$

For example, if the trunk operates at the T1 rate of 1.544 megabits per second (Mb/s), the access lines at the rate of 9.6 kb/s, and the number of virtual circuits, N , is 512, then γ is about 1/3 and the heavy usage condition is satisfied by a comfortable margin.

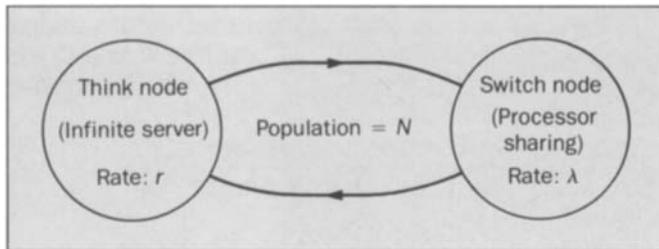


Figure 2. Queueing network representation of the physical system.

Figure 2 is a conventional queueing-network representation of the same system. The network is a closed, product-form network, and the single class shown has a population of N "jobs." A snapshot at time t would reveal that a certain number of jobs, $m(t)$, are in the think node and the remainder, $N - m(t) = n(t)$, are in the switch node. Translated to Figure 1, this corresponds to $m(t)$ virtual circuits in the course of transmitting messages to the switch and $n(t)$ per-channel queues being serviced in the switch.

It is quite possible that specific modeling contexts may require that one modify the narrow and specific definition of think time used here. For example, think time may be extended to include independent, exponentially distributed, random time periods taken by the source between receipt of an acknowledgment and transmission of the next message.

We draw the reader's attention to a feature of real data networks that is not reflected in the simplified model considered here. In some applications (including the Datakit switch), the switch does not wait until an entire message is received before beginning service; instead, service begins as soon as the header arrives. This form of pipelining is not modeled here. However, we expect that our results are reasonably accurate even for pipelined systems for the following reasons.

- In heavy traffic where congestion effects dominate, pipelining will be a relatively small factor.
- In other comparative studies between fluid models, which capture pipelining effects, and packet models, performance differences are not large.¹⁸

Results for the Single-Class Model. Let $\bar{n}(t) \triangleq$ mean number of occupied per-channel queues in the switch at time t , and $u(t) \triangleq$ utilization of the switch at time $t = \Pr(n(t) > 0)$.

It is straightforward from the birth-and-death equations describing the queueing network in Figure 2 that

$$\left. \begin{aligned} \frac{d}{dt} \bar{n}(t) &= -r\bar{n}(t) + Nr(1 - \gamma u(t)) \\ \bar{n}(0) &= 0 \end{aligned} \right\} (t \geq 0) \quad (3)$$

for the initial condition of an empty switch. This equation, while exact and slightly tedious to derive, is quite intuitive.

We cannot solve equation (3) because the utilization function $u(t)$ is not known a priori. However, we obtained the following facts for the condition of heavy usage.¹⁷ As $t \rightarrow \infty$,

$$1 - u(t) \rightarrow \frac{1}{\sqrt{2\pi N}} (\gamma e^{1-\gamma})^N \left[1 + O\left(\frac{1}{N}\right) \right]$$

Because $\gamma < 1$, $(\gamma e^{1-\gamma}) < 1$ and thus the steady-state value is exponentially small in N . We will therefore take the steady-state value of $1 - u(t)$ to be 0.

Secondly, $u(t)$ goes from 0 to its steady-state value extremely quickly, with a time constant that is $O(1/N)$. This is easily seen from the following observation. When $n(t)$ is small, it behaves very nearly as an M/M/1 queue-length process with arrival rate Nr and service rate $Nr\gamma$. The time constant of this system is proportional to $1/N$. A consequence of these two facts is that $u(t) \approx 1$, $t > 0$. Substituting in equation (3), we obtain the following solution:

$$\bar{n}(t) = N(1 - \gamma)(1 - e^{-rt}), \quad \text{for } t \geq 0 \quad (4)$$

which corresponds to a time constant of $1/r$.

Figure 3. Transient response for a single-class model. (a) $N = 200$; (b) $N = 512$.

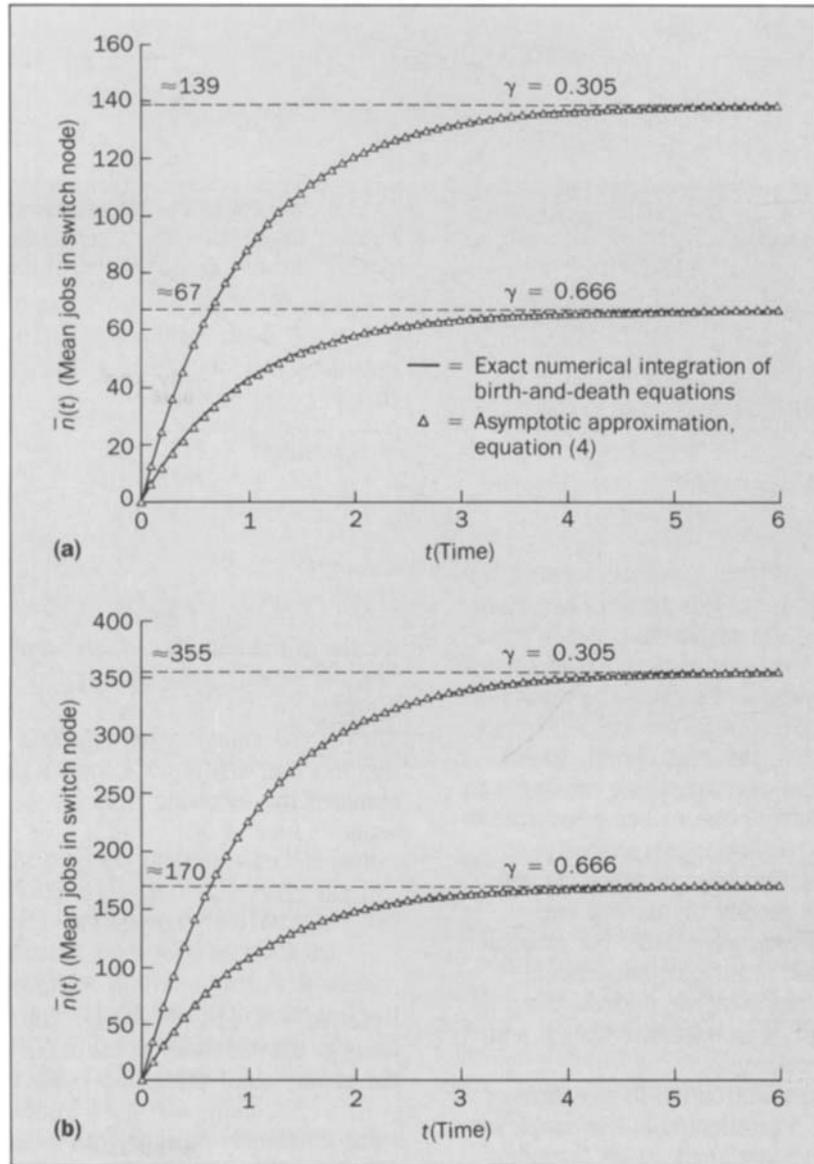


Figure 3 shows plots of $\bar{n}(t)$ as obtained from equation (4). The figure also contains the exact values obtained by numerically integrating the system of $(N + 1)$ first-order differential equations, which constitute the birth-and-death equations. For these plots, $1/r$ has been the unit of time chosen. In each case considered in Figure 3, the approximation is indistinguishable from the exact values. Our integration routine for obtaining the exact values is based on an adaptive Runge-Kutta scheme and is quite time-consuming.

Multiple Classes

Now consider a model such as one used by Fraser and Morgan with multiple types of virtual circuits or classes of jobs.⁴ Virtual circuits may be lumped into different classes by the following considerations:

- Access line speed
- Message length
- Service rate at the switch.

The classes are indexed $1, 2, \dots, p$, with K_j virtual circuits (jobs) of class j , $1 \leq j \leq p$. Figure 2 still

applies, except that there is a think rate, r_j , and a service rate at the switch, λ_j , for each class j . Based on equation (1), let

$$r_j \triangleq \frac{\text{access line speed for class } j}{\text{mean message length for class } j} \quad (5a)$$

$$\lambda_j \triangleq \frac{\text{trunk speed}}{\text{mean message length for class } j} \quad (5b)$$

For the asymptotic analysis, we introduce N as a large parameter. In our software, N is the total number of jobs over all classes (i.e., $\sum K_j$). We define

$$\beta_j \triangleq \frac{1}{N} K_j, \quad \text{for } 1 \leq j \leq p \quad (6a)$$

and, in correspondence to equation (1b),

$$\gamma_j \triangleq \frac{1}{N} \frac{\lambda_j}{r_j}, \quad \text{for } 1 \leq j \leq p \quad (6b)$$

We assume that $\{\beta_j\}$ and $\{\gamma_j\}$ are $O(1)$ as $N \rightarrow \infty$.

The state of the system is a p -dimensional vector $\mathbf{z}_N(t) \triangleq \frac{1}{N} (n_1(t), \dots, n_p(t))$, where $n_j(t)$ is the number of jobs of class j in the switch node at time t . $\mathbf{z}_N(t)$ is a Markov process with generator¹⁹

$$L_N f(\mathbf{z}) = \sum_{i=1}^p \left[N(\beta_i - z_i) r_i \{f(\mathbf{z} + \frac{1}{N} \mathbf{e}_i) - f(\mathbf{z})\} + N r_i \gamma_i \frac{z_i}{|\mathbf{z}|} \{f(\mathbf{z} - \frac{1}{N} \mathbf{e}_i) - f(\mathbf{z})\} \right]$$

where, as in the rest of the paper, \mathbf{e}_i is the unit vector in \mathbb{R}^p in the i -direction and $|\mathbf{z}| = \sum_{i=1}^p |z_i|$ for $\mathbf{z} \in \mathbb{R}^p$.

The factor $z_i/|\mathbf{z}|$ is characteristic of the processor-sharing discipline.

We will now show that $\mathbf{z}_N(t)$ is very nearly a

deterministic process. Consider the following system of ordinary differential equations, $t \geq 0$

$$\frac{d}{dt} y_i(t) = r_i(\beta_i - y_i(t)) - r_i \gamma_i \frac{y_i(t)}{|\mathbf{y}(t)|}, \quad \text{for } 1 \leq i \leq p,$$

$$\mathbf{y}(0) = \mathbf{z}_N(0) \quad (7)$$

The following property of any solution of equation (7) is noted:

$$0 \leq y_i(0) \leq \beta_i \implies 0 < y_i(t) < \beta_i \quad (1 \leq i \leq p) \quad (t > 0) \quad (8)$$

The proof follows from observing that, if $\mathbf{y}(t) \geq \mathbf{0}$ (the case of $\mathbf{y}(t) = \mathbf{0}$ is considered in the section on initial behavior), then

$$y_i(t) = 0 \implies y_i'(t) > 0$$

$$y_i(t) = \beta_i \implies y_i'(t) < 0$$

where $y_i' = dy_i/dt$.

The ordinary differential equation (7) is singular at $\mathbf{y} = \mathbf{0}$. In the section on initial behavior, we show how to obtain a unique solution of equation (7) for $\mathbf{z}_N(0) = \mathbf{0}$.

We have the following result:

THEOREM 1. *For any $T > 0$ and $\epsilon > 0$, there exist positive numbers c_1 and c_2 such that*

$$P(\sup_{0 \leq t \leq T} |\mathbf{z}_N(t) - \mathbf{y}(t)| > \epsilon) < c_1 e^{-N c_2} \quad (9)$$

Because $\mathbf{z}_N(t)$ is bounded, we obtain the following simple corollary:

COROLLARY. *For all $t \geq 0$*

$$\lim_{N \rightarrow \infty} E(\mathbf{z}_N(t)) = \mathbf{y}(t) \quad (10)$$

The theorem and corollary are mathematical statements of the physically plausible "fluid approximation"

for $\mathbf{z}_N(t)$. The flow $\mathbf{y}(t)$ is like a mean field, or continuous-flow approximation to the random (but nearly deterministic) $\mathbf{z}_N(t)$. The rare excursions of $\mathbf{z}_N(t)$ far from $\mathbf{y}(t)$ are in the province of the theory of large deviations and will not be examined here. (See References 3, 20, 21, and 22 for some related results.)

The theorem is an extension of Kurtz's theorem^{2,3} and is proven in Reference 23. Kurtz's theorem does not quite apply to our situation because the vector field for $d\mathbf{y}/dt$ in equation (7) is not continuous at $\mathbf{y} = \mathbf{0}$, let alone Lipschitz-continuous as required. The extension is nontrivial because Kurtz's theorem is known not to hold in some cases where the vector field is non-Lipschitz. Ordinary differential equations with non-Lipschitz vector fields may have non-unique solutions and this may be reflected in non-unique behavior of the random process and consequent breakdown of Kurtz's theorem. In Reference 23, the reader may find a more detailed description of the solutions to equation (7).

Our analysis of $E(\mathbf{z}_N(t))$ has been reduced to the analysis of $\mathbf{y}(t)$. There are three parts to the latter:

1. Steady-state behavior
2. Initial behavior
3. Approach to steady state (i.e., time constants, exponential rates).

Steady-State Behavior. To find the stationary solutions of equation (7), set $\frac{d}{dt} y_i = 0$ to obtain

$$0 = r_i(\beta_i - s_i) - r_i \gamma_i \frac{s_i}{|\mathbf{s}|}, \quad \text{for } 1 \leq i \leq p \quad (11)$$

where we have let s_i be the steady-state value of $y_i(t)$. To solve, set $L \triangleq \sum_{i=1}^p s_i$, and note that $0 < L < \infty$. Then

$$s_i = \frac{\beta_i}{1 + \frac{\gamma_i}{L}}, \quad \text{for } 1 \leq i \leq p \quad (12)$$

and on summing over i ,

$$1 = \sum_{i=1}^p \frac{\beta_i}{L + \gamma_i} \quad (13)$$

The right-hand side of equation (13) is decreasing in L and has a positive solution if and only if

$$\sum_{i=1}^p \frac{\beta_i}{\gamma_i} = \sum_{i=1}^p \frac{K_i \gamma_i}{\lambda_i} > 1 \quad (14)$$

Equation (14) generalizes equation (2) and represents our "heavy-usage condition." We sometimes refer to the quantity on the left as the "usage parameter." Assuming equation (14) holds, it is simple to solve equation (13) numerically for L and so calculate \mathbf{s} by substituting L into equation (12).

We now establish the stability of the fluid approximation of equation (7) by showing that solutions approach \mathbf{s} exponentially. The proof consists of showing the existence of a Lyapunov function.

PROPOSITION 1. *Let*

$$R(t) \triangleq \sum_{i=1}^p \frac{|y_i(t) - s_i|}{r_i \gamma_i} \quad (15)$$

Then

$$R(t) \leq e^{-\hat{r}t} R(0) \quad (t \geq 0) \quad (16)$$

where $\hat{r} = \min_i r_i$.

Proof. Define $x_i(t) \triangleq y_i(t) - s_i$ ($1 \leq i \leq p$). Subtract equation (11) from equation (7) to obtain

$$\frac{d}{dt} x_i(t) = -r_i x_i(t) - \frac{r_i \gamma_i}{|\mathbf{y}|} \left\{ x_i - \frac{s_i \sum x_j}{|\mathbf{s}|} \right\} \quad (17)$$

Now from equation (15),

$$\frac{d^+}{dt} R(t) = \sum_i \frac{\text{sgn } x_i}{r_i \gamma_i} \frac{d}{dt} x_i$$

where d^+/dt is the right derivative.²⁴ By substituting equation (17),

$$\begin{aligned} \frac{d^+}{dt} R(t) &= -\sum_i r_i \frac{|x_i|}{r_i \gamma_i} - \frac{1}{|\mathbf{y}|} \left\{ \sum_i |x_i| \right. \\ &\quad \left. - \frac{1}{|\mathbf{s}|} \sum_i \{s_i \operatorname{sgn} x_i\} \sum_j x_j \right\} \\ &\leq -\hat{r}R(t) - \frac{1}{|\mathbf{y}|} \left\{ \sum_i |x_i| - \left| \sum_j x_j \right| \right\} \\ &\leq -\hat{r}R(t) \end{aligned}$$

Hence,

$$R(t) \leq e^{-\hat{r}t} R(0), \quad \text{for } t \geq 0$$

The above proposition has established an important qualitative property that holds for all time for the fluid approximation. On the other hand, Theorem 1 only shows that equation (7) is a valid approximation to the Markov process $\mathbf{z}_N(t)$ over finite intervals (i.e., finite T). This situation is corrected in the following theorem, whose proof is also in Reference 23.

THEOREM 2. *For every $\epsilon > 0$ there are positive constants c_1 and c_2 such that*

$$\lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t P(|\mathbf{z}_N(u) - \mathbf{s}| > \epsilon) du \leq c_1 e^{-Nc_2}$$

This theorem shows that the steady-state distribution of $\mathbf{z}_N(\bullet)$ is very nearly a delta distribution at \mathbf{s} .

As mentioned in the introduction, the complete asymptotic expansion of the steady-state distribution of the particular network being considered here is available in Reference 15. It is interesting to check the correspondence between these results and the steady-state values of equation (7); the latter are exactly the leading terms of the

asymptotic expansions for the heavy usage regime given in Reference 15.

Initial Behavior. Equation (7) is singular at $\mathbf{y} = \mathbf{0}$; we do not know a priori $\lim_{t \rightarrow 0} \{y_i(t) / |\mathbf{y}(t)|\}$. We overcome this ambiguity in the following manner: for $t \sim 0$, set $y_i = t y'_i(0)$, and obtain

$$y'_i(0) = r_i \beta_i - r_i \gamma_i \frac{y'_i(0)}{|\mathbf{y}'(0)|}$$

Now let $M \triangleq |\mathbf{y}'(0)| = \sum_1^p y'_i(0)$, and we have

$$y'_i(0) = \frac{r_i \beta_i}{1 + \frac{r_i \gamma_i}{M}}, \quad \text{for } 1 \leq i \leq p \quad (18)$$

Note the similarity to equation (12). Summing over i as before,

$$1 = \sum_{i=1}^p \frac{r_i \beta_i}{M + r_i \gamma_i} \quad (19)$$

which will have a unique positive solution, M , assuming condition (14). Then $\mathbf{y}'(0)$ comes from equation (18). The justification of this procedure forms part of the proof of Theorem 1.²³

We have incorporated in our software the above procedure for calculating $\mathbf{y}'(0)$ and we start the numerical integration of equation (7) using these calculated values. The singularity of equation (7) at $\mathbf{y} = \mathbf{0}$ still makes itself felt sometimes in the delicate way we must integrate out of the corner.

Approach to Steady State. A particularly interesting quantity is, of course, the rate at which $\mathbf{z}_N(t)$ approaches \mathbf{s} , the steady state. We analyze this by seeing how quickly $\mathbf{y}(t)$ approaches \mathbf{s} . We do this by linearizing, i.e., by constructing a linear system of differential equations that describes the behavior of the solution of equation (7) in the neighborhood of \mathbf{s} . To proceed, define the error vector

$$\epsilon(t) = \mathbf{y}(t) - \mathbf{s} \quad (20)$$

Then, by equation (7), we have

$$\frac{d}{dt} \epsilon_i(t) = r_i(\beta_i - y_i(t)) - \frac{r_i \gamma_i y_i(t)}{|\mathbf{y}(t)|} \quad (21)$$

But

$$0 = r_i(\beta_i - s_i) - \frac{r_i \gamma_i s_i}{|\mathbf{s}|} \quad (22)$$

Letting $L(t) \triangleq |\mathbf{y}(t)|$, and recalling that $L = |\mathbf{s}|$, we have

$$L(t) = L + \sum \epsilon_j(t)$$

and so, subtracting equation (22) from (21),

$$\begin{aligned} \frac{d}{dt} \epsilon_i(t) &= -r_i \epsilon_i(t) - r_i \gamma_i \left[\frac{L \epsilon_i(t) - s_i \sum \epsilon_j(t)}{L^2 + L \sum \epsilon_j(t)} \right] \\ &= -r_i \epsilon_i(t) - \frac{r_i \gamma_i}{L} \epsilon_i(t) + \frac{r_i \gamma_i s_i}{L^2} \sum \epsilon_j(t) + O(|\epsilon(t)|^2) \end{aligned}$$

Thus, to first order, as $|\epsilon(t)| \rightarrow 0$, i.e., as steady state is approached,

$$\frac{d}{dt} \epsilon(t) = -\mathbf{B} \epsilon(t) \quad (23)$$

where

$$B_{ij} \triangleq \begin{cases} -\frac{r_i \gamma_i s_i}{L L} & i \neq j \\ r_i + \frac{r_i \gamma_i}{L} - \frac{r_i \gamma_i s_i}{L L} & i = j \end{cases}$$

That is, \mathbf{B} is the sum of a diagonal matrix and a rank 1 dyad.

Clearly, the eigenvalues of \mathbf{B} govern the exponential rates at which the solutions of equation (23) approach $\mathbf{0}$, and this gives the time constants of the original system of equation (7). Let us now calculate the eigenvalues of \mathbf{B} . Our result is the following equation that immediately gives a simple numerical procedure for finding all eigenvalues.

PROPOSITION 2. *The eigenvalues of \mathbf{B} solve the following equation in z :*

$$\frac{1}{L^2} \sum_{i=1}^p \frac{r_i \gamma_i s_i}{r_i (1 + \frac{\gamma_i}{L}) - z} = 1 \quad (24)$$

COROLLARY. *With the convention*

$$r_1 \left(1 + \frac{\gamma_1}{L}\right) < r_2 \left(1 + \frac{\gamma_2}{L}\right) < \dots < r_p \left(1 + \frac{\gamma_p}{L}\right)$$

the eigenvalues $\{z_i\}$ occur in the following intervals:

$$0 < z_1 < r_1 \left(1 + \frac{\gamma_1}{L}\right) < z_2 < r_2 \left(1 + \frac{\gamma_2}{L}\right) < \dots < z_p < r_p \left(1 + \frac{\gamma_p}{L}\right) \quad (25)$$

Because the function in equation (24) is monotone, this makes it quite trivial to find the eigenvalues numerically. *Proof of Proposition 2.* If z is an eigenvalue of \mathbf{B} with associated eigenvector $\mathbf{c} = (c_1, \dots, c_p)$, then

$$z c_i = -\sum_{j=1}^p \frac{r_j \gamma_j s_j}{L^2} c_j + r_i \left(1 + \frac{\gamma_i}{L}\right) c_i$$

This reduces to

Table I. System Parameters

| Parameter | Example 1 | Example 2 | Example 3 |
|--|-------------------------------|--------------------------|--------------------------|
| K, numbers of virtual circuits | (200,20,10) | (120,12,6) | (40,160,35) |
| r, rates in the "think node" | (1.0, 0.14583, 0.01139) | Same as for Example 1 | (1.0, 0.025, 0.01139) |
| λ, rates in the "switch node" | (156.25, 3.90625, 0.30518) | Same as for Example 1 | Same as for Example 1 |
| $\sum \frac{K_i r_i}{\lambda_i}$, usage parameter (>1 for "heavy usage") | 2.4 | 1.44 | 2.59 |

$$\frac{r_i \gamma_i \delta_i}{L^2 \left(r_i \left(1 + \frac{\gamma_i}{L} \right) - z \right)} = \frac{c_i}{\sum c_j}$$

and summing on i yields equation (24).

This proposition shows that any eigenvalue must satisfy equation (24). The corollary shows that there are p distinct solutions to equation (24) and each one is an eigenvalue (because it satisfies the characteristic equation); so the set of solutions to equation (24) is exactly the set of eigenvalues.

We have seen that the eigenvalues are all real, positive, and easy to calculate. The time constant of the system is $1/z_j$.

Numerical Investigations

From our rather extensive numerical investigations, we extract a sample of three systems that we shall call Examples 1, 2, and 3. In probing the limits of the fluid model, we find these examples interesting. For each system, we will present the "exact" transient response obtained by numerically integrating the multidimensional birth-and-death equations together with the transient response of the fluid approximation obtained by numerically integrating equation (7). In each example, there are

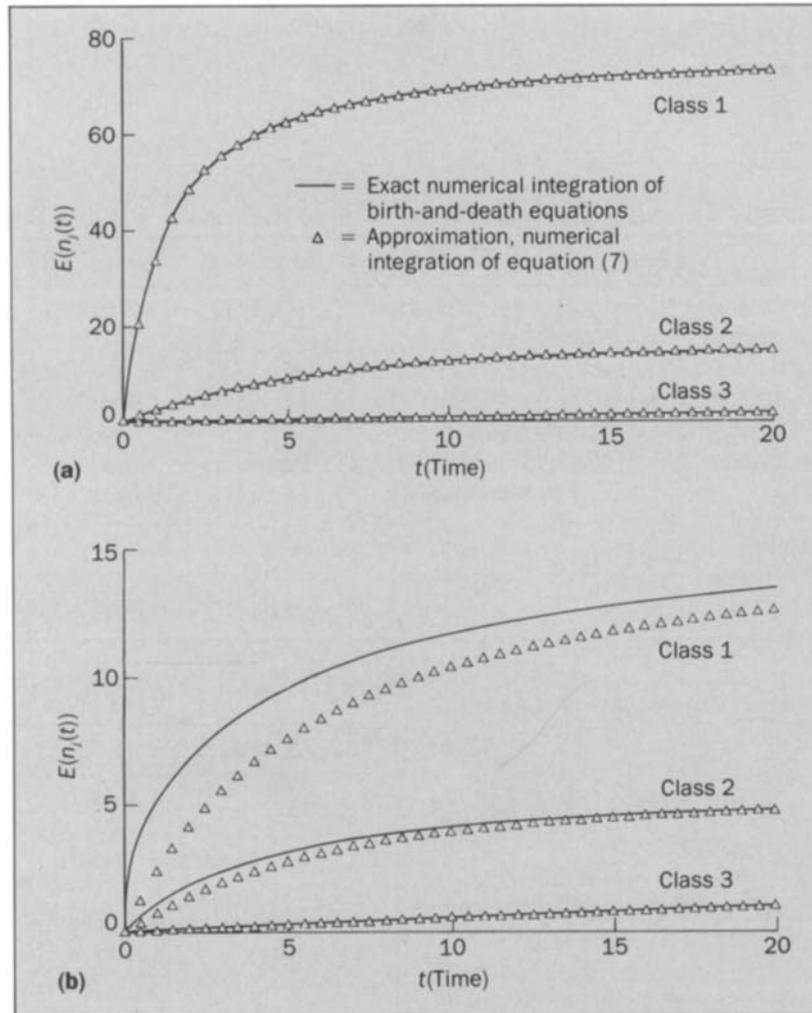
Table II. Parameters of the Fluid Model

| Computed parameter | Example 1 | Example 2 | Example 3 |
|---------------------------------------|--------------------------------|--------------------------------|--------------------------------|
| M | 0.23777 | 0.02633 | 0.011178 |
| L | 0.45021 | 0.18799 | 0.47451 |
| Steady state utilization of switch | (0.7698, 0.1534, 0.0767) | (0.6586, 0.2276, 0.1138) | (0.1494, 0.5976, 0.2530) |
| {z} | (0.014, 0.172, 1.353) | (0.019, 0.216, 3.105) | (0.013, 0.038, 2.194) |
| Time constant (ms) | 60.0 | 44.2 | 65.8 |

three classes, ($p = 3$); and hence, the fluid approximation is a system of three first-order differential equations. In contrast, the exact models are approximately of dimension 4.6×10^4 , 1.1×10^4 , and 2.4×10^5 , respectively. The transient responses presented are for the initial condition of the empty switch.

In the choice of the number of classes and other parameters, we have been guided by Fraser and Morgan's data traffic model.⁴ We let the trunk have the T1 rate of

Figure 4. Transient responses for (a) Example 1, (b) Example 2, (c) Example 3, and (d) an asymptotic approximation for Example 3 with 100 time units. j = class index.



14

1.5×10^6 b/s, and average message length (characters) = (1, 40, 512). Our convention is that, in any parameter vector (p_1, p_2, p_3) , p_i is the parameter for class i where $1 \leq i \leq 3$. In contrast to the above, the access line speeds vary from example to example:

| Access line speeds (kb/s) | |
|---------------------------|----------------|
| Example 1 | (9.6, 56, 56) |
| Example 2 | (9.6, 56, 56) |
| Example 3 | (9.6, 9.6, 56) |

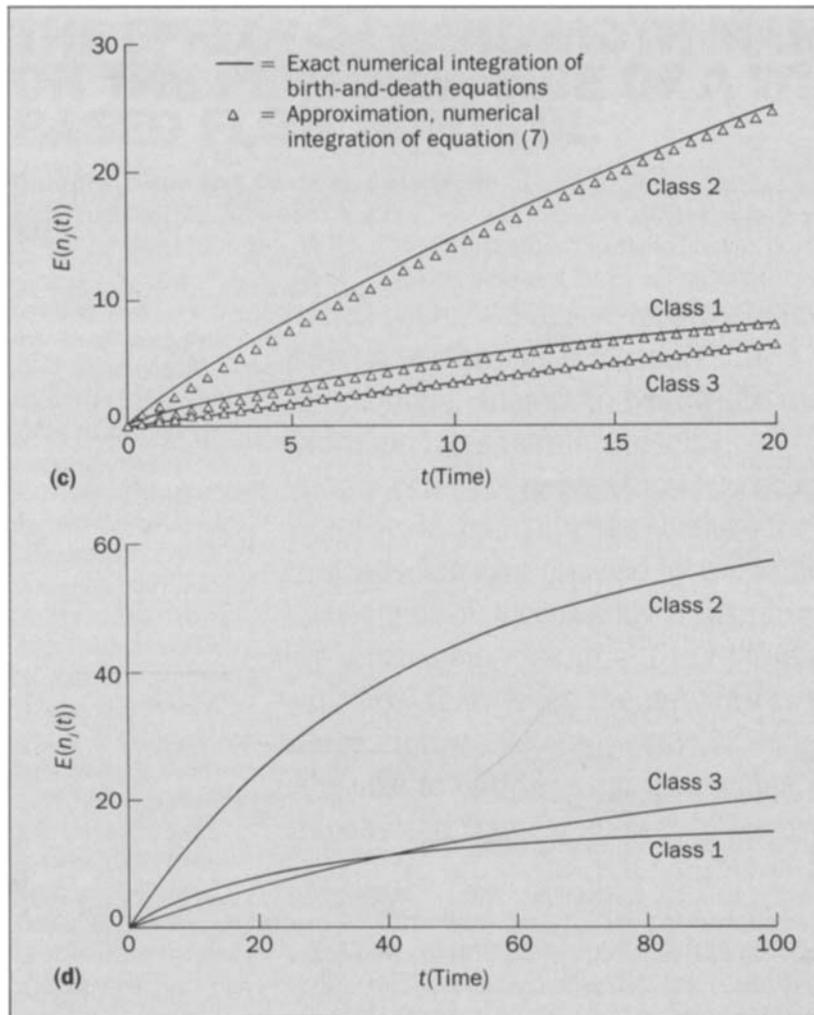
We normalize time so that unity is the amount of time it takes a class 1 source to send a message of average length over its access lines; thus, $r_1 = 1.0$ [see equation (5a)]. It turns out that for all examples, the unit of time is 0.833

milliseconds (ms). From these data, it is possible to calculate r , λ . These, together with the remaining parameters K , are shown in Table I.

Table II shows the important parameters of the fluid model:

- M , norm of the vector field at $\mathbf{0}$. See equation (19).
- L , the mean number of jobs of all classes in the switch as a fraction of the total, for the network in equilibrium. See equation (13).
- The steady-state utilization of the switch by message class, which is obtained from the steady-state mean values and Little's Law.
- The eigenvalues of the linearized fluid model. See equation (25).
- The time constant $1/z_1$.

Notice that to leading order, the estimated total utilization



of the switch in steady state is unity for all the examples.

Figures 4a, b, and c give the transient responses for Examples 1, 2, and 3, respectively. We should mention that the "exact" transient response for Example 3 consumed about 4000 minutes (CPU time) of a VAX™ 8550 computer, while the "elapsed time" was about a week! (VAX is a trademark of Digital Equipment Corporation.) That is why Figure 4c shows only 20 units of time. Figure 4d displays the response of the fluid approximation for the same network but for 100 time units. In every case, the elapsed time for solving the approximation was less than 30 seconds.

The value for the usage parameter in Example 2 is, in our estimation, about the lowest that will give an acceptable quality of approximation for N of the same order as considered here. Of course, if N is larger, we can

allow smaller usage parameters. The usage parameters in Examples 1 and 3 are more typical of networks in which we would expect the fluid approximation to work well.

Example 2 is a difficult case for another reason, namely, the population of class 3 (i.e., 6) is rather small to be approximated by a fluid model. Yet Figure 4b shows that the quality of the approximation is good.

Examples 2 and 3 are also difficult cases because M is quite small. The fluid model underestimates the derivatives at $\mathbf{0}$, the corner of the positive orthant, and this misestimation is more pronounced for smaller M . Figures 4b and 4c show this.

Two trajectories cross in Figure 4d. This was a surprise at first, but on reflection, we realized that it is not an artifact of the approximation. The initial directions $y'_i(0)$ [see equation (18)] are simply ordered differently from the

steady-state values s_i [see equation (12)].

We have investigated several other networks and, in each case, the agreement between the exact transient solution and its fluid approximation has been, if anything, better than in the examples above and for explicable reasons.

Acknowledgments

We are especially grateful to S. P. Morgan for proposing the problem and for the benefit of discussions. Our thanks also to P. E. Wright, with whom we have had several lengthy and useful discussions.

References

1. A. G. Fraser, "Towards a Universal Data Transport System," *IEEE Journal of Selected Areas in Communication*, Vol. SAC-1, No. 5, November 1983, pp. 803-816.
2. T. G. Kurtz, "Strong Approximation Theorems for Density Dependent Markov Chains," *Stochastic Processes and Their Applications*, No. 6, 1978, pp. 223-240.
3. A. Weiss, monograph in preparation; completion anticipated in 1989.
4. A. G. Fraser and S. P. Morgan, "Queueing and Framing Disciplines for a Mixture of Data Traffic Types," *AT&T Bell Laboratories Technical Journal*, Vol. 63, No. 6, July-August 1984, pp. 1061-1087.
5. C.-Y. Lo, "Performance Analysis and Applications of a Two-Priority Packet Queue," *AT&T Technical Journal*, Vol. 66, No. 3, May/June 1987, pp. 82-99.
6. P. Tsoucas and J. Walrand, "A Note on the Processor-Sharing Queue in a Quasi-Reversible Network," *Advances in Applied Probability*, Vol. 15, 1983, pp. 468-469.
7. L. Kleinrock, *Queueing Systems, Volume 1: Theory*, John Wiley & Sons, New York, 1975.
8. V. E. Beneš, *Mathematical Theory of Connecting Networks and Telephone Traffic*, Chapter 6, Academic Press, New York, 1965.
9. M. F. Norman, "A Central Limit Theorem For Markov Processes That Move By Small Steps," *Annals of Probability*, Vol. 2, No. 6, 1974, pp. 1065-1074.
10. H. J. Kushner and D. S. Clark, *Stochastic Approximation Methods for Constrained and Unconstrained Systems*, Springer-Verlag, New York, 1978.
11. L. Ljung, "Analysis of Recursive Stochastic Algorithms," *IEEE Transactions on Automatic Control*, Vol. AC-22, 1977, pp. 551-575.
12. P. R. Kumar and P. Varaiya, *Stochastic Systems: Estimation, Identification and Adaptive Control*, Prentice-Hall, Englewood Cliffs, New Jersey, 1986.
13. A. Weiss, "A New Technique for Analyzing Large Traffic Systems," *Advances in Applied Probability*, Vol. 18, 1986, pp. 506-532.
14. D. Mitra, "Stochastic Fluid Models," in *Performance '87*, P.-J. Courtois and G. Latouche, eds., Elsevier Science Publishers, New York, 1988, pp. 39-51.
15. J. McKenna, D. Mitra, and K. G. Ramakrishnan, "A Class of Closed Markovian Queueing Networks: Integral Representation, Asymptotic Expansions, and Generalizations," *The Bell System Technical Journal*, Vol. 60, No. 5, May 1981, pp. 599-641.
16. J. A. Morrison and D. Mitra, "Heavy-Usage Asymptotic Expansions for the Waiting Time in Closed Processor-Sharing Systems with Multiple Classes," *Advances in Applied Probability*, Vol. 17, 1985, pp. 163-185.
17. D. Mitra and A. Weiss, "The Transient Behavior in Erlang's Model for Large Trunk Groups and Various Traffic Conditions," *Proceedings, 12th International Teletraffic Congress*, Turin, Italy, June 1988, pp. 5.1B4.1-5.1B4.8.
18. H. Hagirahim, AT&T Bell Laboratories, and D. Mitra, private communication, 1988.
19. S. N. Ethier and T. G. Kurtz, *Markov Processes: Characterization and Convergence*, John Wiley & Sons, New York, 1986.
20. M. I. Freidlin and A. D. Wentzell, *Random Perturbations of Dynamical Systems*, Springer-Verlag, New York, 1984.
21. S. R. S. Varadhan, *Large Deviations and Applications*, SIAM, Society for Industrial and Applied Mathematics, Philadelphia, 1984.
22. R. Ellis, *Entropy, Large Deviations, and Statistical Mechanics*, Springer-Verlag, New York, 1985.
23. A. Weiss and D. Mitra, "The Fluid Limit of a Markovian Closed Network with a Discriminatory Processor-Sharing Server," to be submitted for publication, 1988; preprints available from authors.
24. W. A. Coppel, *Stability and Asymptotic Behavior of Differential Equations*, D. C. Heath & Co., Boston, 1965.

(Manuscript received August 18, 1988)