OPTIMAL ROUTING AND FLOW CONTROL IN NETWORKS WITH REAL-TIME TRAFFIC

C. Cassandras, M. Kallmes and D. Towsley

COINS Technical Report 89-07

January 1989

OPTIMAL ROUTING AND FLOW CONTROL IN NETWORKS WITH REAL-TIME TRAFFIC [(1)]

Christos G. Cassandras and **Michelle Hruby Kallmes**
Department of Electrical and Computer Engineering
University of Massachusetts
Amherst, MA 01003
Tel. 413 - 545 1340

Don Towsley
Department of Computer and Information Science
University of Massachusetts
Amherst, MA 01003

ABSTRACT

We address the problem of flow control and routing of *real-time traffic* in a network, where messages must arrive at their destination within given deadlines; otherwise they are considered lost. Performance in this case is measured in terms of the probability of losing a message. For the case of n parallel links, the problem is formulated as one of optimal flow allocation and solved under general conditions. It is shown that for a FCFS service discipline an admission policy rejecting messages before link assignment is optimal when the load exceeds a critical value. Thus, we take advantage of the fact that if some messages will exceed their deadlines anyway, it is beneficial not to admit them in the first place. An efficient algorithm for explicitly solving the problem is presented and specific examples are analyzed. We also discuss the applicability of on-line algorithms for this problem when modeling assumptions cannot be made.

January 1989

(1) This work is supported in part by the Office of Naval Research under contract N00014-87-K-0304 and by the ROME Air Development Center under contract F30602-88-D-0027.

1. INTRODUCTION

In this paper, we address issues of performance analysis and control for systems with *real-time traffic*. Real-time traffic is characterized by deadlines within which service must be completed, and is encountered in at least two settings: communication networks and multiprocessor computer systems [1],[2]. We will present the problem with a communication network application in mind, but all results can be applied to analogous problems in a computer system environment. Thus, when a message with a real-time constraint is transmitted, we require that its destination be reached by some given *deadline*, characteristic of the message itself or of the traffic class to which it belongs. If this condition is not satisfied, the message is useless and is considered lost. In such problems, the objective of interest is the minimization of the probability of loss, measuring the fraction of messages that exceed their deadlines.

Given this framework, we shall consider the problem of assigning messages to n links connecting a source node to a destination node so as to minimize the fraction of messages that do not reach the destination node by some deadline. There are two variations of this problem: (a) all messages are required to be transmitted, and (b) messages are allowed to be rejected instead of being assigned to a link. We shall consider case (b) only, and accomplish the flow control through an admission policy. A typical workload consists of voice and/or video packets with a time constraint by which they must be received to be useful. In current implementations, the scheduling policies at the link level are very simple and do not allow for deliberate packet rejection. It is, however, possible to include flow control at the source in the form of an admission policy that does reject a fraction of the packets whenever it is beneficial to do so.

A similar situation is encountered in static load balancing for a multiprocessor system, where if we consider messages to be jobs and links to be processors, the same results apply. In past analysis of computer systems, most work has focussed on the problem of balancing loads so as to minimize a metric such as average delay [3],[4]. However, a better objective seems to be minimizing the fraction of jobs whose response times fall above some threshold. Thus, if $E[x]$ is the average processing time, good performance may be determined by a response time that lies below some multiple of E[x]. In this application, it would be beneficial to impose a policy that prohibits admission to some fraction of jobs so as to increase the fraction of jobs that complete processing within their threshold response time.

To illustrate the potential advantage of an admission policy allowing rejection, we present the following simple example. Suppose the system consists of a single link. Upon arrival into the system, a message is either accepted with some probability p and placed in the FCFS queue, or rejected with probability (1-p). The objective is to adjust p so as to minimize the *probability of loss*. A message can be *lost* either at the source by being rejected, or at the destination by violating its deadline. The crucial observation here is that a proper choice of p can result in *rejecting some messages that may not have made their deadlines anyway*. Note that as p is decreased, the link becomes less congested, resulting in a larger fraction of messages making their deadline. On the other hand, more messages are being rejected. Moreover, if we dealt with limited buffering space, then a third result of decreasing p would be a welcomed decrease in the blocking probability. Thus, there is a tradeoff between rejecting a message so that others may make their deadline and serving the message in hopes that it might make its deadline. As we shall show, it is generally beneficial to reject some fraction of the incoming messages. It is only for the case of lightly loaded systems that the optimal admission policy is $p = 1$.

The paper is organized as follows. In section 2, we formulate the model and the optimization problem. We then characterize the optimal admission and link assignment policies under fairly general conditions. In section 3, we present an algorithm for explicitly solving this problem. We also discuss the problem of on-line optimization when knowledge of the system characteristics is limited. Section 4 includes some specific examples. Finally, we summarize our results in section 5, and indicate directions for future work.

2. MODEL DESCRIPTION AND PROBLEM FORMULATION

Consider a stream of messages arriving to n parallel links according to an arbitrary arrival process with rate λ, as shown in Fig. 1. The rate of transmission at link i is $\mu_i > 0$, and all links are independent. We will assume that there is the option of rejecting messages before they are assigned to any link; let λ_0 denote the corresponding flow. In addition, let λ_i be the flow to link i, for $i = 1,\ldots,n$. Finally, let D be a random variable that denotes the deadline associated with a random message.

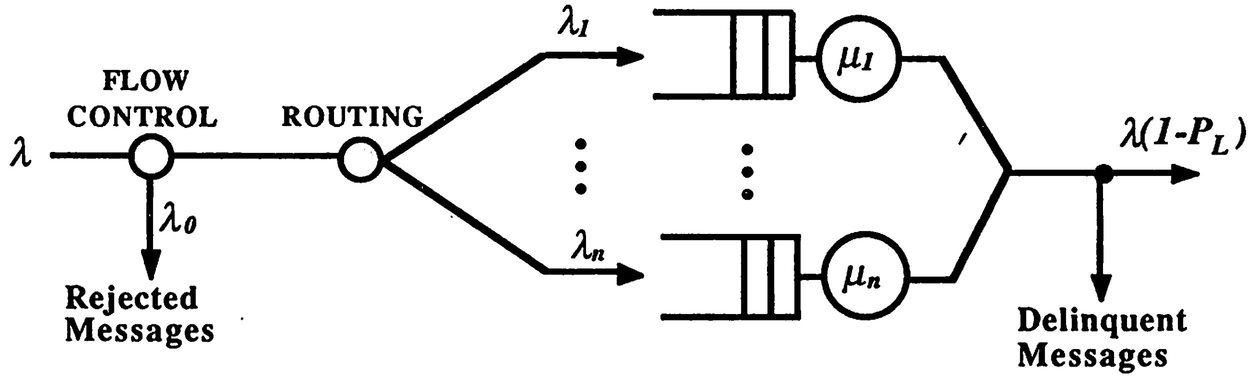

Figure 1: *Link Assignment Model with Admission Policy*

2.1. Problem Formulation

We wish to determine a flow allocation $(\lambda_0,\ldots,\lambda_n)$ so as to minimize the *probability of loss*, where a message is considered *lost* if it is either *rejected* upon arrival or *accepted* but is *delinquent* (i.e., misses its deadline). Let $P^r(\lambda_0)$ be the probability that a random message is rejected, and let $P^d(\lambda_1,\ldots,\lambda_n)$ be the probability that a message is delinquent given that it is accepted. Then the total probability of loss is

$$P_L = [P^r + (1\text{-}P^r)\cdot P^d].$$

We shall assume in what follows that a message is delinquent if its total delay T exceeds the deadline D, i.e. $P^d = \mathrm{P}[T > D]$. However, the analysis also applies to other definitions of delinquency, such as the event that the queueing time of a message exceeds the deadline. For example, see section 4.2.

At this point, we translate the problem from one of minimizing the *probability of loss* to one of minimizing the *loss rate*, *L*, where

$$L = \lambda \cdot [P^r + (1-P^r) \cdot P^d].$$

Clearly, for a fixed load λ, the two problems are equivalent. Note that

$$\lambda \cdot P^r = \lambda_0 , \qquad \lambda \cdot (1-P^r) \cdot P^d = \sum_{i=1}^{n} \lambda_i \cdot P_i^d(\lambda_i) ,$$

where $P_i^d(\lambda_i)$ is the probability that a job is delinquent given that it is routed over link i. Thus, the performance metric can be written as

$$L(\lambda_0, \ldots, \lambda_n) = \lambda_0 + \sum_{i=1}^{n} \lambda_i \cdot P_i^d(\lambda_i).$$

Furthermore, let

$$f_i(\lambda_i) = \lambda_i \cdot P_i^d(\lambda_i), \qquad i = 1, \ldots, n, \tag{1}$$

$$f_0(\lambda_0) = \lambda_0. \tag{2}$$

Then a general statement of the optimal flow allocation problem is the following:

$$\min_{(\lambda_0, \ldots, \lambda_n)} \sum_{i=0}^{n} f_i(\lambda_i)$$

subject to

$$\sum_{i=0}^{n} \lambda_i = \lambda$$

$$\lambda_i \geq 0, \qquad i = 0, \ldots, n.$$

Note that the flows λ_i are not constrained to be less than the corresponding link capacity μ_i; this will naturally emerge as a property of the optimal flow allocation. Also note that by introducing the additional constraint $\lambda_0 = 0$ to the problem above, we can obtain a formulation where no rejections are allowed, as in variation (a) presented in the introduction. The analysis in this case would proceed along the same lines.

2.2. The Optimal Solution

We will make the following assumptions regarding the functions $P_i^d(\cdot)$ and $f_i(\cdot)$, $i=1,\ldots,n$:

A1 $P_i^d(x)$ is increasing for all $x \in [0,\mu_i)$.

A2 $P_i^d(\mu_i) = 1$.

A3 $f_i(x)$ is increasing and strictly convex for all $x \in [0,\mu_i)$.

A4 If $\lambda_i < \mu_i$, then $f_i(\lambda_i) < \lambda_i$.

Most of our subsequent results depend on **A3** only. **A1** simply requires that increasing the load to a link must increase the probability that messages on that link exceed their deadlines; this is natural, since increasing the load causes longer average delays. **A2** requires that waiting times become infinite at capacity. It is satisfied for the FCFS service discipline and possibly by RSS, but it does not hold for a LCFS service discipline. Finally, from (1), **A4** is equivalent to the requirement that if $\lambda_i < \mu_i$, then $P_i^d(\lambda_i) < 1$. This assumption is not critical, since we can always determine some μ_i' $\leq \mu_i$ such that $P_i^d(\mu_i') = 1$, and treat μ_i' as the effective capacity of the link.

Let

$$h_i(\lambda_i) = \partial f_i(\lambda_i)/\partial \lambda_i, \quad i = 0,\ldots,n,$$

where the derivative is defined for $\lambda_i \in [0,\mu_i)$. Then, from **A3**, $h_i(x)$ is monotonically increasing for all $x \in [0,\mu_i)$, $i=0,\ldots,n$. Moreover, $h_i(\lambda_i)$ has the following properties for $i=0,\ldots,n$:

Lemma:

(a) $0 \leq h_i(0) < 1$,

(b) $h_i(\mu_i) > 1$.

Proof: From the definition (1):

$$h_i(\lambda_i) = \frac{\partial f_i(\lambda_i)}{\partial \lambda_i} = P_i^d(\lambda_i) + \lambda_i \frac{\partial P_i^d(\lambda_i)}{\partial \lambda_i}.$$

Hence, $h_i(0) = P_i^d(0)$, where $0 \leq P_i^d(0) \leq 1$. Statement (a) then follows from **A4**. Similarly, $h_i(\mu_i)$ $= 1 + \mu_i \cdot (\partial P_i^d/\partial \mu_i) > 1$, by **A1** and **A2**. **QED**

The following result identifies conditions satisfied by the optimal solution to the flow allocation problem (see also [3]).

Theorem 1: (a) The optimal solution $(\lambda_0^*,\ldots,\lambda_n^*)$ to the above problem satisfies the following necessary and sufficient conditions:

$$h_i(\lambda_i^*) = \alpha \quad \text{if } \lambda_i^* > 0, \tag{3}$$

$$h_i(\lambda_i^*) \geq \alpha \quad \text{if } \lambda_i^* = 0, \tag{4}$$

such that:

$$\sum_{i \text{ s.t. } \lambda_i^* > 0} h_i^{-1}(\alpha) = \lambda, \qquad \alpha > 0. \tag{5}$$

(b) All optimal flows satisfy $\lambda_i^* < \mu_i$, $i = 1,\ldots,n$.

Proof: Using a Lagrange multiplier $\alpha > 0$, we adjoin the flow conservation constraint to $L(\lambda_0,\ldots,\lambda_n)$ to obtain

$$\tilde{L}(\lambda_0,\ldots,\lambda_n,\alpha) = \sum_{i=0}^{n} f_i(\lambda_i) + \alpha \cdot \left(\lambda - \sum_{i=0}^{n} \lambda_i \right).$$

Thus, the necessary (Kuhn-Tucker) conditions require that at $\lambda_i = \lambda_i^*$

$$\frac{\partial \tilde{L}}{\partial \lambda_i} = h_i(\lambda_i) - \alpha = 0, \quad i = 0,\ldots,n$$

provided $\lambda_i^* > 0$, and (3) is obtained. On the other hand, if the constraint $\lambda_i \geq 0$ is active (i.e. $\lambda_i^* = 0$), then at the feasible region boundary, defined by $\lambda_i = 0$, optimality requires that

$$\frac{\partial \tilde{L}}{\partial \lambda_i} \geq 0 \quad \text{at} \quad \lambda_i = 0,$$

which yields (4). Finally, for an optimal allocation $(\lambda_0^*,\ldots,\lambda_n^*)$, the flow conservation constraint becomes

$$\sum_{i=0}^{n} \lambda_i^* = \sum_{i \text{ s.t. } \lambda_i^* > 0} \lambda_i^* = \lambda$$

and since all λ_i^* such that $\lambda_i^* > 0$ must satisfy (3), we immediately obtain (5). Note that $h_i^{-1}(\cdot)$ is well-defined in $(0,h_i(\mu_i))$, and, as we show in part (b), $h_i^{-1}(\alpha) = \lambda_i^* < \mu_i$.

The sufficiency of these conditions follows directly from the convexity assumption A3. This guarantees that a unique α is determined through (5), and hence unique $\lambda_i^* > 0$ are determined from (3).

Part (b) of the Theorem is trivial if $\lambda_i^* = 0$. Thus, we consider only the case $\lambda_i^* > 0$. First, note from (2) that $h_0(\lambda_0) = 1$ for all $\lambda_0 \geq 0$. Since (3) and (4) imply that $h_0(\lambda_0^*) \geq \alpha$, it follows that $\alpha \leq 1$. Thus, from (3), we get

$$h_i(\lambda_i^*) = \alpha \leq 1.$$

From the Lemma above, $h_i(\mu_i) > 1$. Thus, $h_i(\lambda_i^*) < h_i(\mu_i)$. Since $h_i(\lambda_i)$ is monotonically increasing for all $\lambda_i \in [0,\mu_i)$, it follows that $\lambda_i^* < \mu_i$. **QED**

Remark. Since $h_0(\lambda_0) = 1$, it follows from (3) that if $\lambda_0^* > 0$, i.e. rejections do occur, then we have $\alpha = 1$.

An illustration of the Theorem is shown in Fig. 2. For any $i=1,\ldots,n$, as long as $\lambda_i^* > 0$, the optimal flow λ_i^* is determined by the point where $h_i(\lambda_i)$ intersects the constant α line. Otherwise, $h_i(\lambda_i)$ lies above the constant α line and $\lambda_i^* = 0$. Furthermore, as pointed out above, $\alpha = 1$ whenever $\lambda_0^* > 0$. In the example shown in Fig. 2, α takes on a value less than 1 and $h_1(\lambda_1)$ does not intersect the line defined by α. Thus, $\lambda_0^* = \lambda_1^* = 0$, corresponding to a lightly loaded case with link 1 unutilized.

Some interesting properties of the optimal flows are presented below as a Corollary of Theorem 1. First, we define $g_i(y)$ to be the bounded inverse of $h_i(x)$, defined for $0 \leq y < h_i(\mu_i)$. More specifically,

$$g_i(y) = \begin{cases} x & \text{if } y = h_i(x),\ x > 0, \\ 0 & \text{if } y \leq h_i(0). \end{cases}$$

Thus, (5) can be written as $\Sigma_i\, g_i(\alpha) = \lambda$. Observe that since $h_i(\cdot)$ is increasing, $g_i(\cdot)$ is non-decreasing.

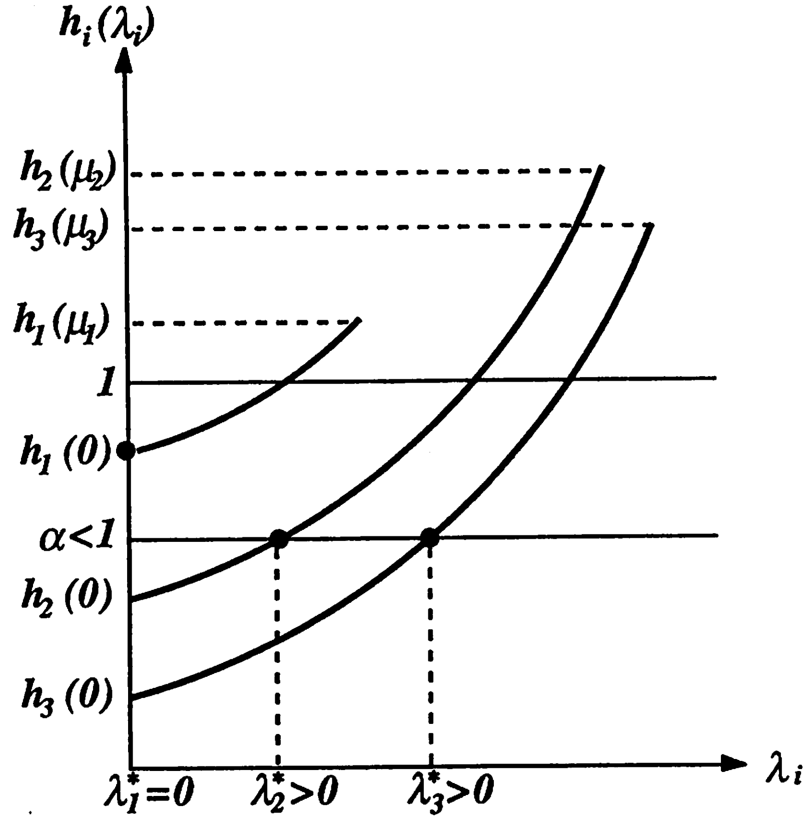

Figure 2: *Characterization of Optimal Flows*

Corollary 1.1: Suppose links are indexed so as to satisfy the following ordering:

$$h_0(0) \geq h_1(0) \geq \ldots \geq h_n(0).$$

Then the optimal flows, $\lambda_0^*,\ldots,\lambda_n^*$, have the following properties:

$$\lambda_i^* > 0 \text{ if and only if } \lambda > \sum_{j>i} h_j^{-1}[h_i(0)],\ i = 0,\ldots,n, \tag{6}$$

$$\lambda_0^* > 0 \text{ if and only if } \lambda > \sum_{i=1}^{n} h_i^{-1}(1), \tag{7}$$

$$\lambda_i^* = 0 \Rightarrow h_i(0) \geq \alpha,\ \ h_i(0) > \alpha \Rightarrow \lambda_i^* = 0,\ \ i = 0,1,\ldots,n. \tag{8}$$

Proof: For any $i = 0,1,\ldots,n$, suppose $\lambda_i^* > 0$. Then (3) implies $\alpha = h_i(\lambda_i^*) \geq h_i(0)$ since $h_i(\cdot)$, $i=1,\ldots,n$ is monotonically increasing by A3 and $h_0(\cdot) = 1$. Thus, from (5)

$$\lambda = \lambda_i^* + \sum_{j \neq i} g_j(\alpha) > \sum_{j > i} h_j^{-1}[h_i(0)],$$

which establishes the "if" part in (6). We prove the converse by contradiction. Thus, assume that the inequality in (6) is satisfied and that $\lambda_i^* = 0$. From (4), this implies that $h_i(\lambda_i^*) \geq \alpha$, i.e. $h_i(0) \geq \alpha$. Because of the ordering we have established, it follows that

$$h_j(0) \geq \alpha \quad \text{for all } j < i.$$

Using this fact, $g_j(\alpha) = 0$ for all $j < i$, and (5) now gives:

$$\lambda = \sum_{j < i} g_j(\alpha) + \sum_{j > i} g_j(\alpha) \leq \sum_{j > i} h_j^{-1}[h_i(0)],$$

which contradicts the inequality in (6). This completes the proof of (6).

Since this proof holds for all $i = 0,1,\ldots,n$, (7) is simply a special case of (6) with $i = 0$, noting that $h_0(0) = 1$ from (2).

To establish (8), first assume that $\lambda_i^* = 0$. Then, by (4), $h_i(\lambda_i^*) \geq \alpha$, i.e. $h_i(0) \geq \alpha$. On the other hand, if $h_i(0) > \alpha$, then, by the monotonicity of $h_i(\cdot)$ we have $h_i(\lambda_i^*) > \alpha$, since $\lambda_i^* \geq 0$. Invoking (4), the result follows and the proof is complete. **QED**

An implication of this Corollary, specifically (7), is that there exists a critical load value, λ_{crit}, given by

$$\lambda_{crit} = \sum_{i=1}^{n} h_i^{-1}(1). \tag{9}$$

The value λ_{crit} defines the *load threshold*, beyond which rejection of messages will always improve performance. This is illustrated in Fig. 3, where the effective throughput, $\lambda \cdot (1-P_L)$, is always higher with $\lambda_0 > 0$ than with $\lambda_0 = 0$ (no admission policy) as long as $\lambda > \lambda_{crit}$. That is, when the load exceeds the critical value, the optimal flow policy will reject some *positive* flow of messages.

Remark. For any $i = 0,\ldots,n$, the values of $h_i^{-1}(1)$ in (9) are independent of the flows $\lambda_0,\ldots,\lambda_n$, and the parameter λ; they depend only on the structure of the system (the distribution of interarrival times, service times, and deadlines), and on the parameters $\mu_1,\ldots,\mu_n$ and E[D]. Hence, provided that an analytical expression for $h_i^{-1}(\cdot)$ can be obtained, one can determine λ_{crit} without explicitly solving the optimization problem .

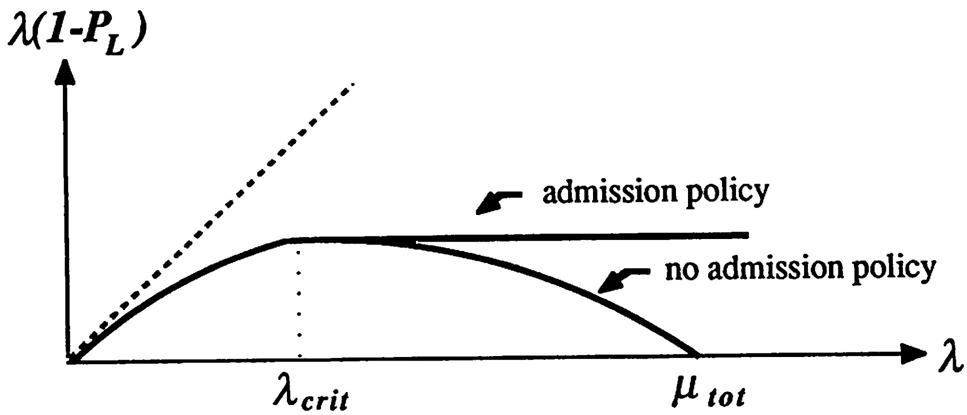

Figure 3: *Load Threshold in Optimal Admission Policy*

Some additional properties of the optimal flow allocation are stated below without proof. In expressing these properties, it is convenient to introduce the following definitions. Let $(\lambda_0,\ldots,\lambda_n)$ denote a feasible solution to our problem. With respect to this solution, there are two classes of links, *active* and *idle*. Link i is considered active if $\lambda_i > 0$; otherwise, it is considered idle. Let A denote the set of links that are active and I the set of links that are idle. We now present four properties of the optimal solution.

Property 1 The optimal flow λ_i^* is a nondecreasing function of λ. Furthermore, if $\lambda_0^* > 0$ for some λ, then as λ increases, λ_i^* remains constant for all $i = 1,\ldots,n$.

Property 2 The set of active links A is a nondecreasing function of λ. In particular, if the nodes are ordered according to the rule specified in Corollary 1.1, then as λ increases, links join A in the reverse order.

Property 3 Let $A' = A - \{j\}$. Then A' is a nonincreasing function of μ_j. Moreover, links transform from active to idle in the same order of $h_i(0)$ as μ_j increases. Furthermore, j always becomes active for sufficiently large μ_j, and remains active thereafter. Last, $A' \rightarrow \emptyset$ as $\mu_j \rightarrow \infty$.

Property 4 If $\lambda > \Sigma_i \mu_i$, then $\lambda_0^* > 0$.

3. OPTIMAL FLOW ALLOCATION ALGORITHMS

In this section, we address the issue of obtaining explicit solutions for the optimization problem presented above. Clearly, this depends on the availability of closed-form expressions for the link loss rate functions $f_i(\cdot)$, as well as the inverse functions $h_i^{-1}(\cdot)$. Such expressions are generally not easy to obtain, even under reasonable stochastic modeling assumptions [5].

In section 3.1, we present an efficient algorithm (of order $n\log n$) for obtaining an explicit optimal flow solution under the assumption that analytical expressions for $h_i(\cdot)$ are available. However, its inverse, $h_i^{-1}(\cdot)$, need not be available in closed form; numerical methods can be employed to evaluate it.

In section 3.2, we briefly address the issue of on-line optimization, when limited knowledge of the overall system is available. In this case, we invoke gradient-based stochastic optimization algorithms, taking advantage of recent developments in stochastic gradient estimation [6]-[8].

3.1. An Efficient Algorithm

The following algorithm can be used to obtain the solution to the problem of minimizing the message loss rate for the system of Fig. 1.

Algorithm.

1. Ordering of links:
   - order links so that $h_0(0) \geq h_1(0) \geq h_2(0) \geq \cdots \geq h_n(0)$.
2. Determination of the set of active links, A:
   - define $\lambda(m) = \Sigma_{i>m}\, g_i(h_m(0))$,
   - use binary search to find m such that $\lambda(m) \leq \lambda < \lambda(m+1)$,
   - $I = \{1,\ldots,m\}$, $A = \{m+1,\ldots,n\}$.
3. Determination of link flows:

• solve equations (3)-(5) to obtain $(\lambda_0^*,\ldots,\lambda_n^*)$.

The algorithm is explained as follows. The first step orders the links so that the properties (6)-(8) will hold. Step 2 applies property (6) to determine which links have positive flow. For links $i \in I$, we can immediately determine $\lambda_i^* = 0$. In step 3, we notice that equation (3) holds for links $i \in A$, so those $(n\text{-}m)$ equations along with (5) uniquely determine λ_i^* for $i \in A$.

The algorithm is most easily understood by referring back to Fig. 2. In the example shown, the ordering of step 1 has already been applied. Step 2 determines the region of the constant α line, that is, m is found such that $h_m(0) \geq \alpha > h_{m+1}(0)$. Then step 3 calculates the intersection point of $h_i(\lambda_i)$ with the constant α line for all active links; λ_i at the intersection is optimal.

3.2. On-Line Optimization

Our analysis thus far relies on the availability of analytical expressions for $f_i(\lambda_i)$, $i=1,\ldots,n$. This means that specific link models can be developed; in section 4, for example, we assume each link to behave like an M/M/1 queueing system. Often, however, accurate models are not available and operating conditions change. In such cases, one is limited to observing actual system data (e.g. message arrival times at a node) and making admission and routing decisions based on that information alone. There are several algorithms suitable for such on-line control schemes, mostly dependent on estimating gradients (sensitivities) of the performance metrics with respect to the parameters of interest (see [9],[10]). Our purpose here is only to make brief mention of the fact that stochastic gradient estimation techniques can be developed and used in the context of our problem, and to illustrate the use of this information for the simple flow control problem corresponding to n=1.

If a single link is present, the critical load value in (9) becomes $\lambda_{crit} = h_1^{-1}(1)$. Recalling the definitions of $f_i(\cdot)$ and $h_i(\cdot)$, this implies that when $\lambda = \lambda_{crit}$, the accepted flow, λ_1^c, is determined from

$$\min_{\lambda_1} [\ (\lambda-\lambda_1) + \lambda_1 \cdot P_1^d(\lambda_1)\ ],$$

yielding

$$\frac{\partial}{\partial \lambda_1}[\lambda_1 P_1^d(\lambda_1)] = 1,$$

which implies that

$$\lambda_1^c = [1 - P_1^d(\lambda_1)] \cdot \left(\frac{\partial P_1^d}{\partial \lambda_1} \right)^{-1}$$

with the right hand side evaluated at $\lambda = \lambda_{crit}$. If this system is directly observed, an estimate of $P_1^d(\lambda_1)$ can be obtained under any admission policy yielding a flow λ_1. Suppose that, at the same time, an estimate of the sensitivity of the delinquency probability with respect to the flow, $(\partial P_1^d/\partial \lambda_1)$, can also be obtained. Then, one would immediately be able to determine whether the load threshold point shown in Fig. 3 has been reached or not, and hence determine the optimal flow allocation. Furthermore, in an adaptive sense, flows can be adjusted on-line until the optimal allocation is found. In many cases, this gradient estimation is accomplished without knowledge of the actual load λ or the link capacity μ_1 [10]. The problem of determining stochastic gradient estimates along a single observation interval is addressed through Perturbation Analysis [7] and the Likelihood Ratio technique [8]. However, results currently available apply primarily to performance measures such as throughput and mean message delay. The case of $(\partial P_1^d/\partial \lambda_1)$ is considered in [9] and is the subject of ongoing research.

4. EXAMPLES

In what follows, we present three examples for the case of M/M/1 link models. That is, we assume that messages arrive into the system according to a Poisson arrival process with parameter λ and are assigned to links through a probabilistic rejection and routing mechanism. In addition, we model all links as exponential servers with parameters μ_i, $i=1,\ldots,n$ and a FCFS queueing discipline.

4.1 Constant Deadline on Total Delay

Let the deadline always take on the value $D = \tau$. In this case $P_i^d = P[T>\tau]$, and

$$f_i(\lambda_i) = \lambda_i \, e^{-(\mu_i-\lambda_i)\tau}$$
$$h_i(\lambda_i) = (1+\lambda_i\tau) \, e^{-(\mu_i-\lambda_i)\tau}$$

for all $i=1,\ldots,n$, and $f_0(\lambda_0) = \lambda_0$ as before. Even though we are unable to obtain an explicit expression for $h_i^{-1}(\cdot)$, and hence for $g_i(\cdot)$, it is easy to show that $h_i^{-1}(\alpha)$ is the root of the equation

$$(1+\lambda_i\tau) \, e^{-(\mu_i-\lambda_i)\tau} - \alpha = 0$$

that lies in $[0,\mu_i)$. This root exists and is unique provided

$$e^{-\mu_i\tau} \leq \alpha < 1+\mu_i\tau\,. \tag{10}$$

In such a case, $g_i(\alpha) = h_i^{-1}(\alpha)$. Otherwise, if $\alpha < e^{-\mu_i\tau}$, then $g_i(\alpha) = 0$, and the case where $\alpha > 1 + \mu_i\tau$ does not occur because, as we have already shown, $\alpha \leq 1$.

It is easily shown that **A1-A4** are satisfied. Furthermore, the optimal solution results in rejections if and only if $\lambda > \Sigma_i h_i^{-1}(1)$, where $h_i^{-1}(\alpha)$ is the root of the equation

$$(1+\lambda_i\tau) \, e^{-(\mu_i-\lambda_i)\tau} - 1 = 0$$

that lies in $[0,\mu_i)$, which we know to exist and be unique since (10) holds for $\alpha = 1$.

4.2 Constant Deadline On Queueing Time

We assume that deadlines are deterministic and equal to τ, and let $P_i^d = P[Q>\tau]$ where Q is the queueing time of a random message. In this case,

$$f_i(\lambda_i) = \frac{\lambda_i^2}{\mu_i} \, e^{-(\mu_i-\lambda_i)\tau}$$

$$h_i(\lambda_i) = \frac{\lambda_i}{\mu_i} \, (2+\lambda_i\tau) \, e^{-(\mu_i-\lambda_i)\tau}$$

for $i = 1,\ldots,n$. Assumptions **A1-A4** are easily shown to hold for these functions. We are unable to obtain an explicit expression for $h_i^{-1}(\lambda_i)$, and hence for $g_i(\lambda_i)$. However, we can easily show that $h_i^{-1}(\alpha)$ is the root of the equation

$$\frac{\lambda_i}{\mu_i}(2+\lambda_i\tau)\, e^{-(\mu_i-\lambda_i)\tau} - \alpha = 0$$

that lies in $[0,\mu_i)$. This root exists and is unique provided $0 \le \alpha < 2 + \lambda_i\tau$, in which case we have $g_i(\alpha) = h_i^{-1}(\alpha)$. But since $0 < \alpha \le 1$, we can conclude that α is in the range, hence the root does exist uniquely for all choices of system parameters, and $\lambda_i^* > 0$, $i = 1,\ldots,n$.

4.3 Exponential Deadlines on Total Delay

We assume that D is an exponential random variable with mean $1/\gamma$, and, once again, P_i^d = P[T>D]. Then

$$f_i(\lambda_i) = \frac{\lambda_i\gamma}{\gamma+\mu_i-\lambda_i}$$

$$h_i(\lambda_i) = \frac{\gamma+\mu_i}{(\gamma+\mu_i-\lambda_i)^2}$$

for $i = 1,\ldots,n$. It is easy to show that P_i^d is such that **A1-A4** are satisfied. Furthermore, we can write the following expression for $g_i(\alpha)$,

$$g_i(\alpha) = \begin{cases} \gamma+\mu_i - \sqrt{\dfrac{\gamma+\mu_i}{\alpha}}, & \alpha \ge \dfrac{1}{\gamma+\mu_i}, \\ 0, & \alpha < \dfrac{1}{\gamma+\mu_i}. \end{cases}$$

If the optimal solution has $A = \{m+1,\ldots,n\}$, and $\lambda_0^* = 0$, then

$$\lambda_i^* = \gamma + \mu_i - (\gamma+\mu_i)^{1/2}\left(\sum_{i=m+1}^{n} \frac{\gamma+\mu_i}{(\gamma+\mu_i-\lambda_i)^2}\right)^{-1/2}$$

for $i = m+1,\ldots,n$, and 0 otherwise. If $\lambda_0^* > 0$, then

$$\lambda_i^* = \gamma + \mu_i - (\gamma+\mu_i)^{1/2}.$$

The optimal solution results in rejections if and only if

$$\lambda > \sum_{i=1}^{n}\left[\gamma + \mu_i - \sqrt{\gamma+\mu_i}\right].$$

For the special case of a single link with capacity μ, the optimal flow control policy, $(\lambda_0^*, \lambda_1^*)$ is

$$\lambda_0^* = \begin{cases} \lambda - (\gamma + \mu - \sqrt{\gamma+\mu}), & \lambda \geq \gamma + \mu - \sqrt{\gamma+\mu}\ , \\ 0 & \text{otherwise}\ , \end{cases}$$

$$\lambda_1^* = \begin{cases} \gamma + \mu - \sqrt{\gamma+\mu}, & \lambda \geq \gamma + \mu - \sqrt{\gamma+\mu}\ , \\ \lambda & \text{otherwise}\ . \end{cases}$$

If we do not apply an admission policy (i.e. never reject), then the optimal loss rate is

$$L^* = \begin{cases} \dfrac{\lambda\gamma}{\gamma+\mu-\lambda}, & \lambda \leq \mu\ , \\ \lambda & \text{otherwise}\ . \end{cases}$$

but allowing rejection gives

$$L^* = \begin{cases} \dfrac{\lambda\gamma}{\gamma+\mu-\lambda}, & \lambda \leq \gamma + \mu - \sqrt{\gamma+\mu}\ , \\ \lambda - \left[\dfrac{(\gamma+\mu-\sqrt{\gamma+\mu})\cdot(\sqrt{\gamma+\mu}-\gamma)}{\sqrt{\gamma+\mu}}\right], & \text{otherwise}\ . \end{cases}$$

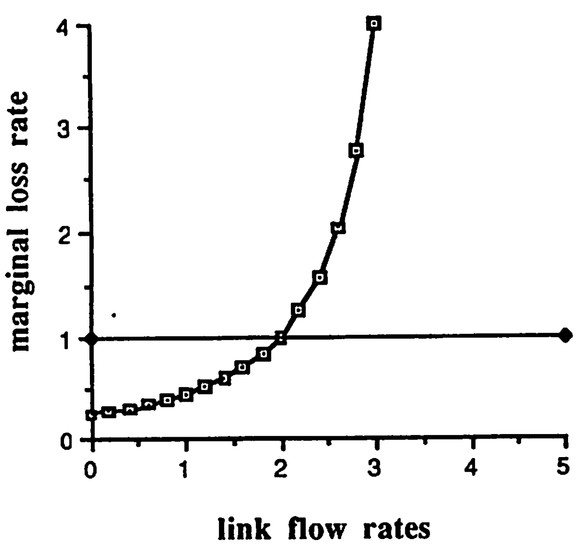

Figure 4a: *Marginal Loss Rates*

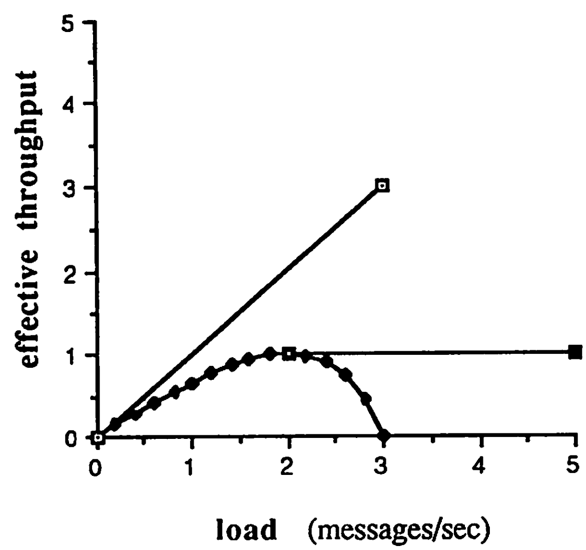

Figure 4b: *Effective Throughput*

In Figures 4a and 4b, we have plotted the marginal loss rates, $h_0(\lambda_0) = 1$ and $h_1(\lambda_1)$, and the effective throughput, $\lambda \cdot (1-P_L)$, for the case where $\mu = 3$, $\gamma = 1$. Note that Fig. 4b is similar to Fig. 3, and clearly shows the existence of a critical load threshold for this model.

5. CONCLUSIONS AND FUTURE WORK

We have shown that when the performance metric of interest is the probability of losing messages which exceed a given deadline, it is advantageous to apply flow control that deliberately rejects part of the incoming traffic. Thus, some messages which might be lost anyway are prevented from accessing the system, a fact which contributes to increasing the fraction of the accepted messages making their deadlines. The optimization problem we have studied applies to networks supporting real-time traffic, a situation which is expected to be commonly encountered. The solution to this optimization problem and its structural properties were analyzed in section 2. A specific criterion for determining load conditions under which it is indeed optimal to apply flow control was also provided (equation (9)). Approaches towards developing explicit algorithms to implement flow control and routing for the case of n parallel links were presented in section 3.

The analysis presented in this paper appears to be the point of departure for a variety of interesting extensions. First, as mentioned in the Introduction, we have considered the case where flow control is applied in the form of probabilistic message rejections at a source node. However, our analysis can easily be applied to the case where all messages must be assigned to a link. Second, we can introduce weights for the components $f_i(\lambda_i)$ of the performance metric so as to penalize rejection more than delinquency in some desirable fashion. This setup would model systems where messages are of some use when delinquent. Third, it is possible that our explicit results for the M/M/1 link models may be extended to more general situations, so as to obtain analytical solutions using the algorithm in section 3.1.

As discussed in section 3.2, of great interest is the wide applicability of stochastic gradient estimation and on-line optimization techniques that take advantage of observed system data with little or no information concerning parameters or distributions involved in an analytical model. Motivated by the results in [10],[11], we believe that such techniques will indeed be useful in this framework, in order to develop adaptive schemes for flow control and routing in networks with real-time traffic.

An obvious extension of great interest is the generalization of our results to more complex networks. Consider, for instance, a virtual circuit model, consisting of several links in tandem. Since a link is often shared by many virtual circuits, flow control at a source node presents the opportunity to improve performance for all real-time traffic sharing that link. Moreover, it is possible that rejection of messages can be implemented not only at the source node of a virtual circuit, but also at intermediate nodes. This appears to be a fruitful area for further research. In addition, the presence of finite buffers provides yet another opportunity to benefit from rejections by simultaneously reducing blocking probabilities.

Another challenging problem concerns the extension of our flow allocation problem for real-time traffic to the case of dynamic policies, i.e. policies taking into account some system state information (e.g. instantaneous queue lengths in making link assignments). Some of the techniques used in [9] seem promising for this situation.

REFERENCES

[1] J. F. Kurose and R. Chipalkatti, "Load Sharing in Soft Real-Time Distributed Computer Systems," *IEEE Trans. on Computers*, vol. C-36, 8, pp. 993-1000, 1987.

[2] P. Nain and K. W. Ross, "Optimal Priority Assignment with Hard Constraint," *IEEE Trans. on Automat. Contr.*, vol. AC-31, 10, pp. 883-888, 1986.

[3] A. Tantawi and D. Towsley, "Optimal Static Load Balancing in Distributed Computer Systems," *J. Ass. Comput. Mach.*, vol. 32, pp. 445-465, 1985.

[4] D. Eager, E. Lazowska, and J. Zahorjan, "Dynamic Sharing in Homogeneous Distributed Systems," *IEEE Trans. on Software Engin.*, vol. SE-12, pp. 662-675, 1986.

[5] B. Gavish and P. Schweitzer, "The Markovian Queue with Bounded Waiting Time," *Management Sci.*, vol. 23, pp. 1349-1357, 1977.

[6] C. G. Cassandras and S. G. Strickland, "On-Line Sensitivity Analysis of Markov Chains," *IEEE Trans. on Automat. Contr.*, vol. AC-34, 1, pp. 76-86, 1989.

[7] Y. C. Ho, "Performance Evaluation and Perturbation Analysis of Discrete Event Dynamic Systems," *IEEE Trans. on Automat. Contr.*, vol. AC-32, 7, pp. 563-572, 1987.

[8] M. I. Reiman and A. Weiss, "Sensitivity Analysis via Likelihood Ratios," *Proc. 1986*

Winter Simulation Conf., pp. 285-289, 1986.

[9] C. G. Cassandras and J.-I. Lee, "Applications of Perturbation Analysis to Optimal Resource Sharing in Discrete Event Systems," *Proc. 1988 American Control Conf.*, pp. 450-455, 1988.

[10] C. G. Cassandras, M.V. Abidi, and D. Towsley, "Distributed Routing with On-Line Marginal Delay Estimation," to appear in *IEEE Trans. on Commun.*, 1989.

[11] C. G. Cassandras, S. G. Strickland, and J.-I. Lee, "Discrete Event Systems with Real-Time Constraints: Modeling and Sensitivity Analysis," *Proc. 27th Conf. Decision and Control*, pp. 220-225, 1988.