be known, then only a total of ten parameters need to be estimated as part of the adaptation process. Of course, the necessity of averaging on-off gas-jet action presents well-known problems of its own (e.g., [3]).

In addition, the lightness requirements in space components may also present difficulties linked to the presence of low-frequency structural modes (see, e.g., [12]). In particular, while the previous discussion can be extended easily to control the rigid dynamics of manipulators mounted on the spacecraft, by applying the results of Slotine and Li [18] and this paper using the "virtual manipulator" formalism of Vafa and Dubowsky [22] or the approach of Alexander and Cannon [1], practical implementation will require flexibility issues to be explicitly addressed. In many space robotics applications, however, distributed flexibility effects can be adequately modeled using simple lumped approximations, which can in turn be easily handled using, e.g., singularly perturbed versions [11] of the original rigid-model adaptive control results.

Acknowledgment

The authors are grateful to R. Sanner for stimulating discussions and suggestions and for performing the extensive simulations summarized in Section IV. They would also like to thank A. Von Flotow for helpful clarifying discussions.

References

[1] H. L. Alexander and R. H. Cannon, "Experiments on the control of a satellite manipulator," Stanford Univ., Stanford, CA, 1987.
[2] H. Asada and J. J. E. Slotine, *Robot Analysis and Control.* New York: Wiley, 1986.
[3] E. V. Bergmann, B. K. Walker, and D. R. Levy, "Mass property estimation for control of asymmetrical satellites," *AIAA J. Guidance, Contr., Dynam.*, vol. 10, no. 5, 1987.
[4] A. E. Bryson, "Control of spacecraft and aircraft," Dep. Aeronaut. Astronaut. Rep., Stanford Univ., Stanford, CA, 1985.
[5] C. I. Byrnes and A. Isidori, "On the attitude stabilization of rigid spacecraft," preprint, 1986.
[6] P. Crouch, "Spacecraft attitude control and stabilization: Application of geometric control theory to rigid body models," *IEEE Trans. Automat. Contr.*, vol. AC-29, 1984.
[7] L. A. D'Amario and G. S. Stubbs, "A new single rotation axis autopilot for rapid spacecraft attitude maneuvers," *AIAA J. Guidance, Contr., Dynam.*, vol. 2, 1979.
[8] T. A. W. Dwyer, "Exact nonlinear control of large angle rotational maneuvers," *IEEE Trans. Automat. Contr.*, vol. AC-29, 1984.
[9] T. A. W. Dwyer and A. L. Batten, "Exact spacecraft detumbling and reorientation maneuvers with gimbaled thrusters and reaction wheels," *AASJ Astronaut. Sci.*, vol. 3, 1985.
[10] T. A. W. Dwyer, M. S. Fadali, and N. Chen, "Single step optimization of feedback-decoupled spacecraft attitude maneuvers," presented at the IEEE 24th Conf. Decision Contr., Fort Lauderdale, FL, 1985.
[11] F. Ghorbel, J. Y. Hung, and M. W. Spong, "Adaptive control of flexible joint manipulators," presented at the IEEE Int. Conf. Robot. Automation, Scottsdale, AZ, 1989.
[12] J. L. Junkins and J. D. Turner, *Optimal Spacecraft Rotational Maneuvers.* New York: Elsevier, 1986.
[13] S. Monaco and S. Stornelli, "A nonlinear attitude control law for a satellite with flexible appendages," presented at the IEEE 24th Conf. Decision Contr., Fort Lauderdale, FL, 1985.
[14] A. P. Morgan and K. S. Narendra, "On the uniform asymptotic stability of certain linear non-autonomous differential equations," *SIAM J. Contr. Optimiz.*, vol. 15, 1977.
[15] D. C. Redding and N. J. Adams, *AIAA J. Guidance, Contr., Dynam.*, vol. 10, no. 1, 1987.
[16] J. J. E. Slotine, "Putting physics in control," *IEEE Contr. Syst. Mag.*, vol. 8, 1988.
[17] J. J. E. Slotine and M. D. Di Benedetto, "Hamiltonian adaptive control of spacecraft," M.I.T.-N.S.L. Rep. 880603, 1988.
[18] J. J. E. Slotine and W. Li, "On the adaptive control of robot manipulators," *Int. J. Robot. Res.*, vol. 6, no. 3, 1987.
[19] ——, "Adaptive manipulator control: A case study," *IEEE Trans. Automat. Contr.*, vol. 33, Nov. 1988.
[20] ——, "Composite adaptive robot control," *Automatica*, 1989.
[21] M. Takegaki and S. Arimoto, "A new feedback method for dynamic control of manipulators," *ASME J. Dynam. Syst., Measurement, Contr.*, vol. 102, 1981.
[22] Z. Vafa and S. Dubowsky, "On the dynamics of the manipulators in space using the virtual manipulator approach," in *Proc. 1987 IEEE Int. Conf. Robot. Automation*, Raleigh, NC, 1987.
[23] B. Wie and P. M. Barba, "Quaternion feedback for spacecraft large angle maneuvers," *AIAA J. Guidance, Contr., Dynam.*, vol. 8, no. 3, 1985.

On Not Losing Track

ALAN WASHBURN

Abstract—**This paper uses tracking time as an explicit objective in designing tracking algorithms that function in clutter. Implications of using "time to the first mistake" as a measure of tracking time are explored. An optimal tracker is derived for a special circumstance where there is clutter, but no measurement error, and is compared to a tracker based on the Maximum Likelihood principle.**

I. Introduction

A tracker is given the position of a target at time 0, and then asked to track the target in the presence of false alarms (clutter) that occur randomly in space and time. Sooner or later, some seductive sequence of false alarms will lead the tracker astray, thus ending the tracking period. The phenomenon of "losing the bubble" is inevitable, but nonetheless hard to quantify. There may be *short* periods of *confusion* before track is lost *irrevocably*, and *accuracy* is involved. The proper definitions of the four italicized words are not obvious, so it should come as no surprise that there is no widely accepted scalar measure of effectiveness (MOE) for tracking in the presence of false alarms.

The lack of an MOE has not inhibited the development of tracking algorithms that function in the presence of false alarms and/or multiple targets. Bar-Shalom [1] gives a good survey, or see the more recent book by Bar-Shalom and Fortmann [2]. Most algorithms are founded on the Maximum Likelihood principle. Considerable testing has been done on real and simulated data, but no tracking algorithm has yet been shown to actually be optimal in the sense of maximizing some specific MOE when clutter is present. This state of affairs is not unreasonable, given the usual robustness of statistical procedures based on the Maximum Likelihood principle. Still, it should be of interest to compare the Maximum Likelihood procedure to one that is actually optimal in some well-defined sense, even if the difficulty of deriving optimal procedures forces the comparison to be done in a highly simplified setting. Making such a comparison is our main object here. In Section II, the setting is defined and an optimal tracker (MT) is derived. The MT is then compared to the corresponding Maximum Likelihood tracker SR1 in Section III.

II. The Memoryless Tracker

The memoryless tracker (MT) is presented with a sequence of measurements, one at a time, each of which is either a two-dimensional target contact or a clutter point. It can remember only the most recently accepted measurement, and must accept or reject each measurement as it occurs. Tracking begins at time 0 with MT assuming correctly that the target is at the origin, and ends at the time when MT makes its first mistake. Fig. 1 shows an illustration of the process where tracking ends at time 11, the time of the first target contact not included in the track. Inclusion of the clutter point at time 18 would also be a mistake, but not the first. The objective is to design a tracker that maximizes the average tracking time. The motion of the target is assumed to be diffusion without drift in two dimensions, with target contacts at times that constitute a Poisson process. Clutter is assumed to be a space-time Poisson process, the natural model for events happening with perfect randomness in space and time [3]. The important parameters and their physical units are

λ = Poisson rate of target contacts (time^{-1})
η = Poisson clutter rate ($\text{length}^{-2}\cdot\text{time}^{-1}$)
κ = diffusion constant ($\text{length}^{2}\cdot\text{time}^{-1}$).

Since there are three parameters and two physical dimensions, a single

Manuscript received August 8, 1988; revised June 26, 1989. Paper recommended by Past Associate Editor, B. Sridhar.
The author is with the Department of Operations Research, Naval Postgraduate School, Monterey, CA 93943.
IEEE Log Number 9035897.

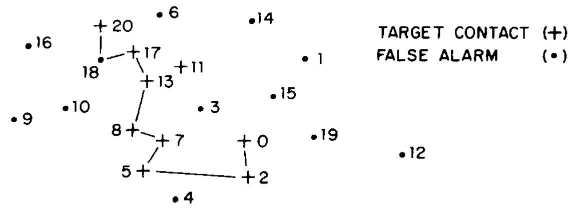

Fig. 1. Illustration of data association.

dimensionless parameter will suffice. It is convenient to let that parameter be

$$\alpha = \lambda/\sqrt{2\pi\eta\kappa} = \text{dimensionless contact rate.}$$

Given that tracking time ends if clutter is ever accepted, MT might as well make decisions conditioned on the assumption that all previously accepted measurements are target contacts. Since diffusion is a Markov process, with the current definition of tracking time MT actually has no *need* for memory because all accepted contacts other than the last are valueless for forecasting. Furthermore, the acceptance decision need depend only on the time t since the last acceptance and the distance of the latest measurement from the last acceptance. Let the latest measurement be accepted if and only if that distance is smaller than $r(t)$. The function $r(t)$ completely defines the tracker; it is this "window" function that must be chosen optimally to maximize the expected tracking time.

Beginning at time 0 or just after any accepted target contact, define the random variables

A = time to the next false alarm in the window

C = time to the next target contact

T = additional tracking time starting from the referenced time.

Then

$$T = \begin{cases} A & \text{if } A < C \\ C & \text{if } A > C \text{ and no acceptance at } C \\ C + T' & \text{if } A > C \text{ and acceptance at } C \end{cases} \tag{1}$$

where T' is independent of, but has the same distribution as, T.

Let $f(x)$ be the probability density function of A, let $p(t)$ be the probability of accepting a target contact at t, and let $\tau = E(T) = E(T')$ where $E(\cdot)$ is the expected value operator. Taking expectations on both sides of (1),

$$\tau = E\left(\int_0^C x f(x)\,dx + (C + \tau p(C)) \int_C^\infty f(x)\,dx\right). \tag{2}$$

Let $q(t) = \int_t^\infty f(x)\,dx$. Since $\int_0^C x f(x)\,dx + C \int_C^\infty f(x)\,dx = \int_0^C q(x)\,dx$ (integration by parts), (2) can be rearranged to obtain

$$\tau = \frac{E\left(\int_0^C q(x)\,dx\right)}{1 - E(p(C)q(C))}. \tag{3}$$

Since the rate at which false alarms occur in the window is $\pi\eta r^2(t)$, the average number of such occurrences up to time t is

$$Y(\lambda t) \equiv \pi\eta \int_0^t r^2(u)\,du. \tag{4}$$

Since the number of false alarms in the window is a Poisson process,

$$q(t) = \exp(-Y(\lambda t)). \tag{5}$$

Since the position of the target at time t relative to its position at time 0 is bivariate normal with mean zero and variance κt in each direction,

$$\begin{aligned} p(t) &= \int_0^{2\pi} \int_0^{r(t)} \frac{1}{2\pi\kappa t} \exp(-x^2/2\kappa t) x\,dx\,d\theta \\ &= 1 - \exp(-r^2(t)/(2\kappa t)) \end{aligned} \tag{6}$$

where the second equality follows easily after substituting $u = x^2/2\kappa t$.

Using (5), (6), and the fact that C is an exponential random variable with mean $1/\lambda$, (3) can be written

$$\lambda\tau = A/(1 - B) \tag{7}$$

where

$$\begin{aligned} A &= \lambda \int_0^\infty \lambda \exp(-\lambda t)\,dt \int_0^t \exp(-Y(\lambda v))\,dv \\ &= \int_0^\infty \exp(-x - Y(x))\,dx \end{aligned} \tag{8}$$

and

$$B = \int_0^\infty \exp(-x - Y(x))\{1 - \exp(-Z(x))\}\,dx \tag{9}$$

where

$$Z(x) \equiv \alpha^2 Y'(x)/x. \tag{10}$$

The second equality in (8) can be obtained through integration by parts, and $Z(\lambda t)$ is the same thing as the argument of the exponential in (6).

For any specific function $Y(\cdot)$, substitution into (8), (10), (9), and then (7) will produce a normalized tracking time $\lambda\tau$. Since choosing $Y(\cdot)$ is the same thing as choosing the window function, the problem now is to choose $Y(\cdot)$ to maximize $A/(1-B)$ where A and B are each of the form $\int_0^\infty F(Y(x), Y'(x), x)\,dx$, and $A > 0$ and $0 < B < 1$ for all $Y(\cdot)$. Suppose that the maximum value of $KA + B$ is 1 for some number K. Then it follows that the maximum value of $A/(1 - B)$ must be $1/K$ since $KA + B \leq 1$ if and only if $A/(1 - B) \leq 1/K$. Since maximization of expressions like $KA + B$ is a problem in optimal control theory, one method for maximizing $A/(1 - B)$ is to alternate the operations "maximize $KA+B$" and "let $K = (1-B)/A$" until convergence occurs. In practice, the maximization operation need be repeated only about three times [6]. For $\alpha = 4$, the MT curve in Fig. 2 shows the optimal $Y'(x)$ versus x or, equivalently, $\pi\eta r^2(t)$ versus λt (the SR1 curve will be discussed later). A complete graph of $Y'(x)$ for *MT* when $\alpha = 4$ would show it rising to a maximum of 1.125 at $x = 38$ and then decreasing to 0 at $x = 87$. If *MT* for some reason accepted no measurement over a time period of $87/\lambda$, then it would never accept another measurement. However, the probability of having no target contact over such a long period is only $\exp(-87)$, so the chances of "loss of track by stalling" are remote.

III. The Maximum Likelihood Principle

Let λ be the Poisson rate per unit time of target contacts and let η be the Poisson rate per unit time per unit measurement space of false alarms. Assume that the state of the target evolves according to a Gauss–Markov process and each measurement is a linear function of the state corrupted by noise, i.e., make the assumptions that permit Kalman filtering in the absence of false alarms. Let T be a subset of the measurements, $\bar{T}$ the complement of T, and let δ and Δ be infinitesimal units of time and measurement space, respectively, as in [4]. Then, using the fact that the innovations ν_j at each measurement are normal and independent [5], the infinitesimal probability of the measurements, given that T is the set of target contacts, is

$$P_T = \left[\prod_{j \in T} (\lambda\delta) \frac{\Delta}{s_j} \exp\left(-\frac{1}{2} \nu_j' S_j^{-1} \nu_j\right)\right] \left[\prod_{j \in \bar{T}} (\eta\delta\Delta)\right] \tag{11}$$

where S_j is the covariance of the innovation ν_j at measurement j under the assumption that T is the target's track, s_j is the square root of the determinant of $2\pi S_j$, and ν_j' is the transpose of ν_j. Letting $M = \prod_{j \in \bar{T} \cup T} (\eta\delta\Delta)$, this can also be written

$$M/P_T = \prod_{j \in T} y_j \exp\left(-\frac{1}{2} \nu_j' S_j^{-1} \nu_j\right) \tag{12}$$

where $y_j = \pi s_j/\lambda$.

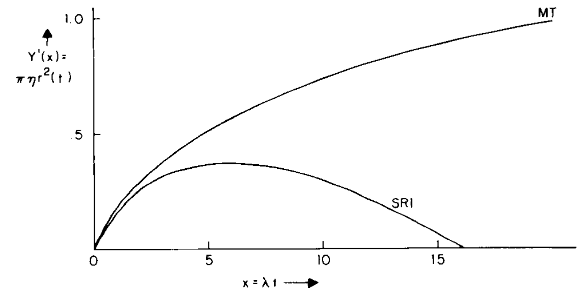

Fig. 2. Window functions for MT and SR1.

Finally, let $D_T = \ln(M/P_T)$ and $d_j = \ln y_j + \frac{1}{2}\nu_j' S_j^{-1}\nu_j$. Then

$$D_T = \sum_{j\in T} d_j. \tag{13}$$

The track with the smallest value of D_T is most likely in the sense that it maximizes P_T. Note that D_T can be negative, and also that $D_T = 0$ if T is empty.

Now consider the tracking situation of Section II, letting T consist of measurements $\cdots z_i, z_j, \cdots$ made at times $\cdots, t_i, t_j \cdots$, with $t_j > t_i$. Also, let z_0 be the initial target location, and let $t_0 = 0$. Then, for $j \neq 0$,

$$\nu_j = z_j - z_i,$$
$$S_j = \kappa(t_j - t_i)I, \tag{14}$$

and

$$d_j = \ln\left(\frac{2\pi\kappa(t_j - t_i)}{\lambda}\right) + \frac{|z_j - z_i|^2}{2\kappa(t_j - t_i)}$$

where I is a 2×2 identity matrix and $|\ |$ denotes Euclidean distance. d_j depends on only two of the measurements: the jth and the one previous to the jth in T. Therefore, without reference to T, every two measurements have a possibly negative "distance" from the earlier to the later, with the distance d_{ij} from i to j being the formula for d_j in (14). Furthermore, the most likely track is the shortest route leaving the origin that preserves the order of the measurement times. Imposing the constraint that each measurement must be irrevocably accepted or rejected as it occurs, one arrives at the SR1 tracker (SR1 for "shortest route with 1 remembered measurement"). If a new measurement is separated from the most recently accepted one by distance r and time t, then from (14), SR1 will accept it if and only if $d < 0$ where

$$d = \ln(2\pi\eta\kappa t/\lambda) + r^2/(2\kappa t). \tag{15}$$

Let

$$r^2(t) = 2\kappa t \ln(\lambda/2\pi\eta\kappa t)) = 2\kappa t \ln(\alpha^2/(\lambda t)). \tag{16}$$

Then $d < 0$ if and only if $r < r(t)$, so SR1 functions very much like MT, except that the window function $r(t)$ is different. Fig. 2 compares the two window functions. Both functions permit the possibility of "stalling," but SR1 stalls more quickly than MT, and in general is more reluctant to accept measurements. The normalized tracking time $\lambda\tau$ for SR1 is shown in Fig. 3 as the lowest of four curves. The normalized tracking time for MT is not shown because it is only slightly larger than the time for SR1. When $\alpha = 4$, for example, the normalized times for SR1 and MT are 5.42 and 5.46, respectively. This may seem remarkable in view of the great difference shown in Fig. 2, but the fact is that the window function is very seldom employed for the large arguments where MT and SR1 differ substantially. The SR1 tracker is nearly optimal in the sense of having a small average time to the first mistake, in spite of not having been designed with that specific objective in mind. Once again, the Maximum Likelihood principle has proved robust.

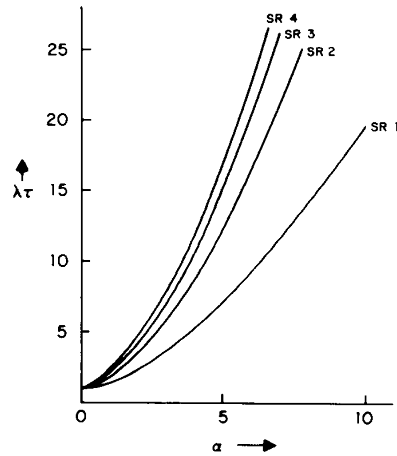

Fig. 3. Normalized tracking time ($\lambda\tau$) versus contact rate (α) for the SRN tracker.

IV. The Effects of Memory

It is easy enough to generalize SR1 to take advantage of tracker memory. Specifically, the SRN tracker remembers N candidate tracks between measurements. After each measurement, SRN constructs additional tracks by extension of each candidate to include the measurement, ranks the $2N$ tracks in order of distance from the origin, preferring new (extended) tracks to the old ones in case of ties, and remembers only the N shortest tracks. Actually, to make the required comparisons, SRN only needs to remember the distance of each track from the origin, plus the time and location of the terminus for each track. Thus, a total of four numbers are associated with each track in two dimensions.

Since SR1 is nearly optimal in the sense of maximizing the average time to the first mistake, it is tempting to study SRN with respect to the same MOE. However, "time to the first mistake" can be defined in more than one way when multiple tracks are remembered. Define the time to the first mistake in the strict sense (SN) to be the first time when the shortest track in storage is not the exact track of the target, or in the wide sense (WN) to be the first time when the exact track of the target is not stored. The strict definition might be thought natural if at some time a decision had to be taken that would succeed if and only if the tracker's best guess of the target track were actually correct. The wide definition would be more appropriate if the action (which might begin with the employment of other sensors to resolve the ambiguity) could be applied to all N stored tracks simultaneously. The choice of definition turns out to be crucial since the wide times benefit from memory, while the strict times are dominated by $S1$.

Theorem 1: In every instance, $SN \leq S1$.

Proof: Let $D(i)$ be the shortest distance from the origin to measurement i, with the distance between measurements being as defined in Section III. Also, let i be the terminus of the last correct track of SR1, and let j be the index of the measurement that causes SR1's first mistake. If SRN has not already made a mistake before t_j, then SRN must store the same track as SR1 just before the measurement at t_j. Suppose that SR1's mistake is to include a false alarm. Then $d_{ij} \leq 0$; consequently, $D(j) \leq D(i)$, and therefore the shortest track stored by SRN after processing the jth measurement cannot be the correct one that ends at i (recall that SRN prefers the new track in case of equality). If SR1's mistake is to not include a target contact, then $d_{ij} > 0$; consequently, $D(j) > D(i)$, and therefore the correct path ending at j will not be the shortest one stored by SRN after processing the jth measurement. Therefore, regardless of the mistake made by SR1, $SN \leq S1$. □

Since $SN \leq S1$ in every instance, it follows that $E(SN) \leq E(S1)$. If "mistake" is to be interpreted strictly, then the best of the SRN class is SR1. If $E(SN)$ is to be used as an MOE, then memory is useless.

Fig. 3 shows normalized wide-sense tracking time $\lambda E(WN)$ for $1 \leq N \leq 4$. The curves for $N = 2$, 3, and 4 were obtained by simulation, with each point being based on 10000 replications. The simulation curve for $N = 1$ agrees with the results of numerically integrating (8) and (9). It should be apparent that there is a large payoff for being able to remember a second track, with a reduced benefit for even more memory. With the wide-sense definition of tracking time, memory is beneficial.

V. Summary

There are two implications here for tracker design. Use of the Maximum Likelihood principle has been reinforced through the observation that SR1 is almost as good as MT in a problem where MT is optimal. On the other hand, the benefits of tracker memory seem hard to quantify without considering the ultimate application of the tracker.

References

[1] Y. Bar-Shalom, "Tracking methods in a multitarget environment," *IEEE Trans. Automat. Contr.*, vol. AC-23, pp. 618–626, 1978.
[2] Y. Bar-Shalom and T. Fortmann, *Tracking and Data Association*. New York: Academic, 1988, p. 172.
[3] W. Feller, *An Introduction to Probability Theory and Its Applications, Vol. I*, 2nd ed. New York: Wiley, 1957, pp. 146–154.
[4] R. W. Sittler, "An optimal data association problem in surveillance theory," *IEEE Trans. Mil. Electron.*, vol. MIL-8, pp. 125–139, 1964.
[5] P. Smith and G. Buechler, "A branching algorithm for discriminating and tracking multiple objects," *IEEE Trans. Automat. Contr.*, vol. AC-20, pp. 101–104, Feb. 1975.
[6] A. Washburn, "Analysis of the memoryless tracker," Naval Postgrad. School, Tech. Rep. NPS55-80-014, 1980.

Optimal Flow Control of Multiclass Queueing Networks with Partial Information

MAN-TUNG T. HSIAO and AUREL A. LAZAR

Abstract—**Structural results as well as explicit solutions for the optimal flow control problem of multiclass queueing networks with decentralized information are given. Two criteria are investigated: the network (respectively, user) optimization criterion maximizes the average network (user) throughput subject to an average network (user) time delay constraint. It is shown that these problems can be analyzed in terms of an equivalent network by using the generalized Norton's equivalent. The structure of the network (user) optimization problem is exploited to obtain further structural results, viz. a representation (separation) theorem. The optimal flow control under both criteria is solved using a linear programming formulation. The structure of the optimal control is shown to be of a window type in both cases. For load balanced networks, the optimal flow control is found explicitly in terms of the given system parameters.**

I. Introduction

Optimal flow control refers to the class of problems where a strategy which maximizes some optimization criterion is to be chosen in order to prevent degradation of services in computer communication networks. The choice of the optimization criterion depends on the particular application. In practice, the average throughput and the average time delay are two quantities of interest in considering the performance of a computer network and its underlying protocol. Relatively little work has been done in the area of decentralized control within the framework of general queueing networks and Markov processes [6], [7], [10], [17], [19]. While the delay shared information pattern models used in [10] and [19] are appropriate for multiple-access broadcast networks, it is more natural to consider only local information with no sharing in data communication and packet-switched networks [5], [6], [7], [15]. Beutler and Ross provided a general treatment of constrained semi-Markov decision processes [3], [18].

In this note, we consider the problem of preventing throughput degradation by controlling the input rate into the network. A multiclass queueing network model is employed for describing the packet-switched network. When a new user logs onto the network, a session is set up. The packets of this new user are distinguished from those of other existing source–destination pairs. The latter is modeled as an interfering traffic flow entering and leaving the network at arbitrary nodes. An optimal flow control strategy is to be derived for this new user based on the existing load in the system. The information available to the controller is local, i.e., the flow generated by interfering traffic is not directly observable. The objective is to maximize the average throughput subject to a constraint on the average time delay. Two measures are investigated. One is the average over all user packets in the network, and the other is the average over the user packets under flow control.

II. The Model

Consider a datagram or a virtual-circuit packet switching network (PSN). There are S switching nodes, each with negligible nodal processing delays and no nodal blocking (i.e., there are ample buffers available). The nodes are connected by M unidirectional links. The destination of a packet, upon completion of service at a station, is determined by a fixed probability distribution. It can be routed to another node within the network or it can leave the network entirely with certain probabilities. Two classes of packets are distinguished. One class belongs to a particular source–destination pair, for which the optimal flow control mechanism is to be found. It is assumed that there is a maximum of N_2 packets where N_2 is an arbitrarily large number.[1] The other class models the interfering traffic, and is composed of packets from all other source–destination pairs sharing the network. It is assumed that there is a maximum of N_1 interfering packets with i.i.d. exponential interarrival times. Packets are acknowledged individually by an end-to-end protocol. Acknowledgments may be piggybacked or standalone. It is assumed that negligible delay is incurred in returning an acknowledgment, partly because it is much shorter than data messages, and partly because it may have higher priority. (This assumption, however, can easily be relaxed and incorporated into our model.) The source under flow control feeds into a controller which determines the rate of allowing packets into the network based on the number of unacknowledged messages.

A. Queueing Model

The PSN described above is modeled using a queueing network. Each link is considered as a first-come first-served (FCFS) service station with exponential service rate μ^i, $1 \le i \le M$. The interfering traffic is modeled as a conditional Poissonian stream which enters the network at each of the nodes. Their routing through the network is probabilistic and may, in general, be different from the routing of the flow-controlled packets. The service times of all packets are assumed to be chosen independently at each stage (see, e.g., [11]). The queueing network model for our flow control problem is shown in Fig. 1.

This model consists of a network of M FCFS exponential stations with class independent service rates μ^i, $1 \le i \le M$. Two classes of packets enter the network. Class 1 (interfering traffic) arrives at a conditional Poisson rate δ. Class 2, modeling the source–destination traffic under flow control, enters the network at rate λ_k, depending on k, the number of class 2 packets already in the network. The routing is probabilistic. Class 1 packets enter the network through station i with probability $r^{1\cdot i}$. Upon completion of service at station i, they move to station j with probability r^{1ij}. These packets leave the network from station i with probability $r^{1i\cdot} = 1 - \sum_{j=1}^{M} r^{1ij}$. Similarly, the routing probabilities for class 2 packets are r^{2ij}. They enter station i from the bottom queue with probability $r^{2\cdot i}$ and join the bottom queue after completion of service at station i with probability $r^{2i\cdot} = 1 - \sum_{j=1}^{M} r^{2ij}$. The bottom queue models

[1] For sufficiently large N_2, the optimal control is independent of N_2.

Manuscript received July 15, 1986; revised July 15, 1987, October 31, 1988, and July 15, 1989. Paper recommended by Past Associate Editor, A. Ephremides. This work was supported in part by the National Science Foundation under Grant CDR-84-21420 and in part by the U.S. Army Research Office under Contract DAAG 29-83-K-0021.

M. T. Hsiao is with the School of Electrical Engineering, Purdue University, West Lafayette, IN 47907.

A. A. Lazar is with the Department of Electrical Engineering, Columbia University, New York, NY 10027.

IEEE Log Number 9035896.

0018-9286/90/0700-0855$01.00 © 1990 IEEE