Generalized Access Control Strategies for Integrated Services Token Passing Systems

Joseph W. M. Pang, Fouad A. Tobagi, *Fellow, IEEE,* and Stephen Boyd

***Abstract*—The demand for integrated services local area networks is increasing at a rapid pace with the advent of many new and exciting applications: office and factory automation, distributed computing, and multimedia communications. To support these new applications, it is imperative to integrate traffic with diverse statistical characteristics and differing delay requirements on the same network. An attractive approach for integrating traffic has been adopted in two token passing local area network standards, the IEEE 802.4 token bus standard [1] and FDDI [2]. The idea is to control the transmissions of each station based on a distributed timing algorithm, so as to achieve the following goals: (i) to limit the token cycles so that time-critical traffic can be accommodated, and (ii) to allocate pre-specified bandwidths to different stations when the network is overloaded. We have investigated the analysis and design of this protocol in [3]. In this paper, we generalize the transmission control algorithm used in [1]–[2]. The major advantages of the generalization over the original protocol are: (i) it provides a much expanded design space, (ii) it guarantees convergent behavior, and (iii) it gives meaningful insights into the dynamics of the basic control algorithm.**

I. Introduction

With many new emerging network applications, it becomes increasingly important to support diverse traffic types (*e.g.*, traffic with time constraints and/or bandwidth constraints) on the same network. Recently, two major token passing local area network standards, the IEEE 802.4 token bus standard [1] and FDDI [2], have adopted the same approach to integrate different traffic types. This approach is based on a dynamic, distributed transmission control algorithm using timers. In the following paragraphs, a generic description of the algorithm is given.

Under normal operations of a token network with no insertion or deletion of stations, N stations (indexed from 1 to N) access the channel in a cyclic order $(1, 2, \ldots, N, 1, 2, \ldots)$. The activity on the channel can be depicted in Fig. 1. $T_i^{(k)}$, known as the *service time*, denotes the transmission time of station i at the k^{th} token reception of that station. $C_i^{(k)}$, known as the *token cycle*, denotes the period between the k^{th} and $(k+1)^{\text{st}}$ receptions of the token at station i. $W_i^{(k)}$, known as the *walk time*, denotes the token passing overhead from station i to station $i+1$ (mod N if $i+1 > N$) in the k^{th} token cycle. Thus, by definition, $C_i^{(k)} = \sum_{j=i}^{N}(T_j^{(k)} + W_j^{(k)}) + \sum_{j=1}^{i-1}(T_j^{(k+1)} + W_j^{(k+1)})$. All summations are defined as zero if the lower index exceeds the upper index. We further define $Q_i^{(k)}$ to be the transmission quota computed by station i when the token arrives at that station in the k^{th} cycle. The transmission quota is honored by ensuring the service time in that cycle does not exceed the computed quota, *i.e.*, $T_i^{(k)} \le Q_i^{(k)}$.

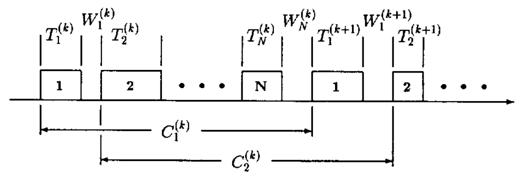

Fig. 1. Time diagram of a token passing channel.

The control algorithm used in [1]–[2] can be generically described by specifying how the quota $Q_i^{(k)}$ is computed. Stations are classified as either type I or type II. Station i, if it is of type I is assigned a parameter denoted by THT_i (Token Hold Time), else is assigned a parameter denoted by TRT_i (Target Token Rotation Time). The quota is computed by

$$Q_i^{(k)} = f_i(C_i^{(k-1)}) = \begin{cases} \text{THT}_i, & \text{if station } i \text{ is of type I,} \\ \max(\text{TRT}_i - C_i^{(k-1)}, 0), & \text{if station } i \text{ is of type II.} \end{cases} \tag{1.1}$$

The functions $f_i(\cdot)$ are called quota functions. The above control mechanism can be thought as imposing a quota that is either a constant or a ramp function of the previous token cycle, as illustrated in Fig. 2. In fact, the control algorithm is a distributed back off algorithm to limit the token cycle length. If a station experiences a long cycle, it will refrain from using the channel in the next cycle. Of course, stations back off according to the station type and the assigned control parameters. Type I stations do not back off at all, while type II stations back off according to their TRT's.

The rationale behind the above integrated access protocol is (i) to limit the token cycles so as to support stations with time-critical traffic (*e.g.*, voice, real-time control data), (ii) to allow other stations to access the channel efficiently while keeping the token cycles below a certain bound, and (iii) to allocate pre-specified bandwidths to the stations when the

Paper approved by J. W. Wong, the Editor for Wide Area Networks of the IEEE Communications Society. Manuscript received April 2, 1990; revised March 11, 1991 and April 6, 1992.

This work was supported by the National Aeronautics and Space Administration under Grant NAGW-419. This paper was presented at IEEE INFOCOM'89, Ottawa, Ont., Canada.

J. W. M. Pang is with Starlight Networks, Mountain View, CA 94043 USA.

F. A. Tobagi is with the Department of Electrical Engineering, Computer Systems Laboratory, Stanford University, Stanford, CA 94305 USA.

S. Boyd is with the Information Systems Laboratory, Stanford University, Stanford, CA 94305 USA.

IEEE Log Number 9401952

0090-6778/94$04.00

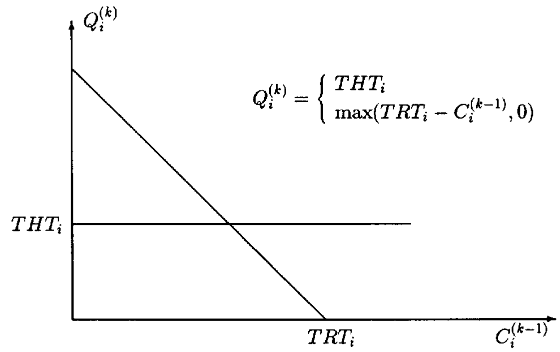

Fig. 2. Quota function for the standard protocol.

system is overloaded. In [3], we have studied the properties of the above protocol regarding the bandwidth allocation under overloading conditions. It was shown that, under overloading conditions, (i) the station transmission times will oscillate in one of many periodic patterns, (ii) the bandwidths allocated to the stations can be computed by the solution of a set of nonlinear equations, independent of the oscillatory pattern. A procedure for determining the bandwidth allocated to each station, as a function of the assigned parameters, was also given in [3].

By viewing the control algorithm as no more than specifying the quota functions f_i, such as that in (1.1), it is natural to inquire whether there are other quota functions that yield desirable behavior and feasible implementation. Without considering implementation complexity, we have explored different generalizations of the quota functions. The earliest studies were focused on simple modifications of (1.1) which resulted in functions like those in Fig. 3. It has become apparent later that the magnitude of the slope of the quota function must be kept below unity to ensure convergent behavior under heavy load. This observation has led to the conclusion that the slope is analogous to the feedback gain of the classical feedback control system. Equipped with this understanding, we were then able to generalize the quota functions further by allowing more general dependence on the past station transmission times rather than the token cycle while retaining the desirable convergent behavior. With further work, we discovered techniques and analytic results for computing token cycle bounds for the new class of quota functions. This paper is the result of many iterations of educated guess work followed by modeling and analysis.

The class of generalized quota functions presented in this paper is indeed very large. They all yield the same favorable properties of convergent behavior and tightly bounded cycles. However, not all generalized quota functions are easy to implement nor are all of them simple to engineer. The main goal of this paper is not to furnish complete implementation and engineering details of a new class of quota functions but to provide insights and understanding of the quota control mechanism with useful general analytic results. Discovery of different subclasses with different properties, say, hardware implementable or particularly fast convergence, represents future work beyond the scope of this paper. The very important issue of implementation of any specific types of generalized quota functions is also left for future studies. It is worth mentioning that although implementation feasibility is not explicitly considered in this paper, we do provide illustrations through some simple, practical and interesting examples of the generalized quota functions that are feasible for implementation.

In the next section, we shall introduce the generalization. In Section III, we investigate the stability properties of the generalization under overloading conditions. In Section IV, we provide an upper bound on the token cycles, together with a design example illustrating the optimization between delay constraint and efficiency. In section 5, we present an approximation for station throughputs under general load. Finally, we conclude the paper with a summary.

II. The Generalized Access Control Strategies

Under the *generalized access control strategy*, the quota functions are given by

$$Q_i^{(k)} = f_i(T_{i-1}^{(k)}, T_{i-2}^{(k)}, \ldots, T_1^{(k)}, T_N^{(k-1)}, \ldots, T_i^{(k-1)}) \quad (2.1)$$

where f_i is non-negative and there exists $1 > \delta > 0$ such that for all vectors $\vec{x}$, $\vec{y}$, there exists a *dependency vector* $\vec{a}$ where $1 - \delta \geq a_1 \geq a_2 \geq \cdots \geq a_N \geq 0$ such that

$$f_i(\vec{x}) - f_i(\vec{y}) = -\vec{a} \cdot (\vec{x} - \vec{y}) \quad (2.2)$$

The dependency vector $\vec{a}$, which may not be unique, depends on i, $\vec{x}$ and $\vec{y}$ but we have omitted the dependence in our notations for simplicity. From (2.2), we can infer that f_i is continuous. Furthermore, the differential dependence of f_i on the first component is stronger than that on the second and in turn stronger than that on the third and so on. This means that the dependence of $Q_i^{(k)}$ on the previous N service times is stronger if the service time is more recent. Notice that the standard control protocol using (1.1) with at least one type II queue does not belong to the class of quota schemes described by (2.2).

Let us consider two examples where the first is a special case of the second.

$$Q_i^{(k)} = \gamma_i \max(M_i - C_i^{(k-1)}, 0), \quad 1 > \gamma_i \geq 0, \quad (2.3)$$

$$Q_i^{(k)} = \gamma_i \min(U_i, \max(M_i - C_i^{(k-1)}, 0)), \quad 1 > \gamma_i \geq 0, M_i, U_i \geq 0 \quad (2.4)$$

(2.4) is depicted in Fig. 3. It is a simple matter to show that the above quota functions satisfy (2.2). Let us compare the above examples with the quota functions in (1.1). If $\gamma_i = 1$ (which is not allowed here), either $M_i = \infty$ or $U_i \geq M_i$ would reduce (2.4) to (1.1). U_i is analogous to THT_i and M_i to TRT_i. Thus, (2.4) is essentially the same as the standard protocol except for the multiplicative constant γ_i. This factor controls the slope of the ramp function and in turn dictates the convergence behavior of the system (see section 3.4). The factor also controls the ratio of the bandwidths allocated to the stations under heavy load (see section 4.2),

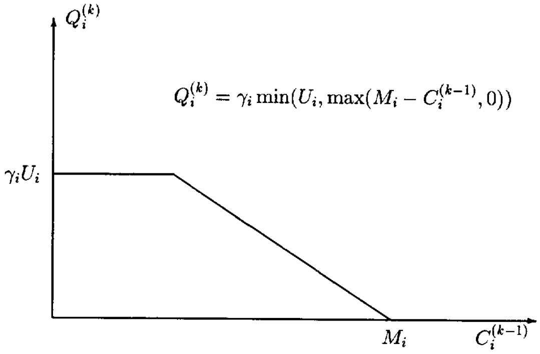

Fig. 3. An example of the generalized quota function.

which is a feature not provided by the standard protocol. Furthermore, it turns out that the ceiling $\gamma_i U_i$ yields a more favorable upper bound on token cycles but lowers the system efficiency (see section 4.2). In any case, both (2.3) and (2.4) are rather simple modifications of the standard protocol and hence require little computation overhead. Throughout this paper, illustrations of analytic results will be made with reference to examples (2.3) and (2.4). We have chosen these examples to be functions of the previous cycle rather than the individual station transmission times for practical reasons. In practice, it is easier to measure the cycle than the individual station transmission times. Furthermore, it turns out that the bandwidths allocated to the stations under heavy load are not sensitive to station re-ordering if the quota functions are dependent only on the previous token cycle. Both properties render quotas that are dependent only on the previous cycle desirable in practical design.

We conclude this section by presenting two more examples where the quota are dependent not only on the previous cycle but on individual service times. A more general case of (2.3) and (2.4) is $Q_i^{(k)} = \gamma_i \min(U_i, \max(M_i - \rho_{i,1} T_{i-1}^{(k)} - \rho_{i,2} T_{i-2}^{(k)} - \cdots - \rho_{i,N} T_i^{(k-1)}, 0))$ where $1 > \rho_{i,1} \geq \rho_{i,2} \geq \cdots \geq \rho_{i,N} \geq 0$, $1 > \gamma_i \geq 0$ and $U_i \geq 0$. A more interesting example is $Q_i^{(k)} = \gamma_i \min(U_i, \max(M_i - \alpha_i C_i^{(k-1)} - \beta_i C_i^{(k-2)}, 0))$ where $1 > \alpha_i \geq \beta_i \geq 0$, $1 > \gamma_i \geq 0$ and $U_i \geq 0$. The quota functions are actually functions of two previous cycles. These are valid quota functions if we consider a new system of $2N$ stations formed by replicating once the original system of N stations. Each station in the original system is visited twice in a cycle of the new system. With the same idea, one can extend the dependence of the quota to arbitrarily old service times. This type of quota functions constitutes a very interesting class of the generalized quota functions. However, due to space limitations, we shall not discuss specific properties of this class of quota functions in this paper.

III. Stability Under Overloading Conditions

In this section, we investigate the stability properties of the generalized access control strategies under the same assumptions adopted in [3]: (i) heavy load (*i.e.*, every station always has data buffered for transmission), (ii) constant token passing overheads (*i.e.*, $W_i^{(k)}$ is independent of k), and (iii) negligible overflow transmissions (*i.e.*, we ignore any transmissions beyond the computed quota). With these assumptions, we can write

$$C_i^{(k)} = W + \sum_{j=i}^{N} T_j^{(k)} + \sum_{j=1}^{i-1} T_j^{(k+1)} \tag{3.1}$$

$$T_i^{(k+1)} = f_i(T_{i-1}^{(k+1)}, \ldots, T_1^{(k+1)}, T_N^{(k)}, \ldots, T_i^{(k)}) \tag{3.2}$$

where $W = \sum_{i=1}^{N} W_i^{(k)}$ denotes the total token passing overhead in a cycle. Summations are defined as zero if the lower index exceeds the upper index. Equations (3.1) and (3.2) describe a dynamic system with system state $\vec{T}^{(k)} = (T_1^{(k)}, T_2^{(k)}, \ldots, T_N^{(k)})$. Our objective in this section is to explore the stability of this dynamic system. We shall first discuss the convergence of the token cycles. Next, we show that $\vec{T}^{(k)}$ converges as $k \to \infty$ for all initial conditions. Then we show that there is a unique equilibrium state for $\vec{T}^{(k)}$.

A. Convergence of Token Cycles

From (3.1), we have $C_{i+1}^{(k)} - C_i^{(k)} = T_i^{(k+1)} - T_i^{(k)}$. From (3.2) and the property of f_i, we have $T_i^{(k+1)} - T_i^{(k)} = -\sum_{j=1}^{i-1} a_j (T_{i-j}^{(k+1)} - T_{i-j}^{(k)}) - \sum_{j=i}^{N} a_j (T_{N+i-j}^{(k)} - T_{N+i-j}^{(k-1)})$. Thus, we have $C_{i+1}^{(k)} = \sum_{j=1}^{i} b_j C_{i-j+1}^{(k)} + \sum_{j=i+1}^{N+1} b_j C_{N+i-j+1}^{(k-1)}$ where $b_1 = 1 - a_1$, $b_j = a_{j-1} - a_j$ for $j = 2, 3, \ldots, N$ and $b_{N+1} = a_N$. Thus, the sequence $C_1^{(1)}, C_2^{(1)}, \ldots, C_N^{(1)}, C_1^{(2)}, C_2^{(2)}, \ldots$ is a moving average sequence with time-varying weights. Even with time-varying weights, the moving average sequence converges when the leading weight is always bounded away from zero. The proof is given in [7]. Furthermore, the proof shows that the difference between the sequence of cycles and its limiting value is bounded by a geometric sequence with parameter $(1 - \delta^N)^{\frac{1}{N}}$.

The sequence of cycles does not converge for the quota functions given by (1.1), because it is possible that $b_{N+1} = 1$ while all other weights are zero. In this case, the sequence oscillates with a period of $N + 1$, unless there are imbedded repetitive patterns, in which case the period will be a divisor of $N + 1$. It is now clear why there can be multiple periodic patterns for the protocol used in [1]–[2]; see section 3.4 for an illustration.

B. Convergence of Service Times

With geometric convergence in the cycles, the vector $\vec{T}^{(k)}$ also converges geometrically with the same parameter. This is because $T_i^{(l)} - T_i^{(k)} = \sum_{j=k}^{l-1} (C_{i+1}^{(j)} - C_i^{(j)})$ for $l > k$. Since the sequence of cycles converges with a geometric bound, $T_i^{(l)} - T_i^{(k)}$ must also be bounded by a geometric sequence with the same parameter. Thus, $\vec{T}^{(k)}$ is a Cauchy sequence and hence converges as $k \to \infty$ for all initial conditions. Furthermore, the convergence is geometric with parameter $(1 - \delta^N)^{\frac{1}{N}}$. If $\vec{T}^{(k)}$ converges, then there must exist at least one equilibrium state, that is, there must be at least one solution to the set of nonlinear equations $T_i^{(\mathrm{eqm})} = f_i(T_{i-1}^{(\mathrm{eqm})}, \ldots, T_1^{(\mathrm{eqm})}, T_N^{(\mathrm{eqm})}, \ldots, T_i^{(\mathrm{eqm})})$ for $i = 1, 2, \ldots, N$.

C. Uniqueness of Equilibrium State

Let $\vec{T}^* = (T_1^*, T_2^*, \ldots, T_N^*)$ be an equilibrium state and let us consider the difference between the sequence of consecutive service times $T_1^{(1)}, T_2^{(1)}, \ldots, T_N^{(1)}, T_1^{(2)}, T_2^{(2)}, \ldots$ and the equilibrium sequence $T_1^{(\mathrm{eqm})}, T_2^{(\mathrm{eqm})}, \ldots, T_N^{(\mathrm{eqm})}, T_1^{(\mathrm{eqm})}, T_2^{(\mathrm{eqm})}, \ldots$ Let $\Delta\vec{T}_{k\cdot N+i}$ be a column vector containing N consecutive difference terms described above starting at the service of queue i in the kth cycle. Using an argument similar to that in the last sub-section, we have $\Delta\vec{T}_{l+1} = \mathbf{A}_l \Delta\vec{T}_l$ where the matrix $\mathbf{A}_l$ is given by

$$\mathbf{A}_l = \begin{bmatrix} -a_{l,1} & -a_{l,2} & \cdots & -a_{l,N-1} & -a_{l,N} \\ 1 & 0 & \cdots & 0 & 0 \\ 0 & 1 & \ddots & \vdots & \vdots \\ \vdots & \ddots & \ddots & 0 & 0 \\ 0 & \cdots & 0 & 1 & 0 \end{bmatrix}$$

with $1 - \delta \geq a_{l,1} \geq a_{l,2} \geq \cdots \geq a_{l,N} \geq 0$. Again, the $a_{l,j}$'s depend on $\Delta\vec{T}_l$ but we omit the dependence in our notation.

We define the max-contiguous sum of a vector $\vec{u}$ by $N(\vec{u}) = \max_{1 \leq m \leq n \leq N} \left|\sum_{j=m}^{n} u_j\right|$. In [7], we show that $N(\cdot)$ is a vector norm and that $N(\Delta\vec{T}_l) \to 0$. Hence $\Delta\vec{T}_l \to 0$ as $l \to \infty$. Similar to the token cycles, the convergence is geometric with parameter $(1 - \delta^N)^{\frac{1}{N}}$. This result clearly shows the uniqueness of the equilibrium state, for if there were another equilibrium state $\vec{T}^{**}$, then we would have $N(\vec{T}^* - \vec{T}^{**}) = 0$, implying $\vec{T}^* = \vec{T}^{**}$.

D. Discussion

Through an example, we would like to illustrate the desirable properties of the generalized quota functions. Consider a system with quota functions given by (2.3) where $W = 0$, $N = 3$, $M_1 = 10$, $M_2 = 7$, $M_3 = 5$ and $\gamma_1 = \gamma_2 = \gamma_3 = \gamma$. We shall investigate the cases where $\gamma = 0.5, 1.0, 1.5$ for the two different initial states $(T_1^{(1)}, T_2^{(1)}, T_3^{(1)}) = (0,0,0), (4,1,0)$. Starting from these initial states, we generate two sets of trajectories for $(T_1^{(k)}, T_2^{(k)}, T_3^{(k)})$ for $\gamma = 0.5, 1.0, 1.5$. These trajectories are displayed graphically in Figs. 4 to 6 with the corresponding token cycle lengths.

It can be shown that there is a unique equilibrium state for each of the three cases $\gamma = 0.5$, 1.0, 1.5, and they are (2.8000, 1.3000, 0.3000), (4.3333, 1.3333, 0.0000), (5.4375, 0.9375, 0.0000) respectively. For $\gamma = 0.5$, both trajectories converge to the equilibrium state. For $\gamma = 1.0$, the trajectories exhibit different periodic patterns, but the bandwidth allocation for these patterns are the same as that of the equilibrium state. For $\gamma = 1.5$, we find a "limit cycle" in the first trajectory but there is no readily identifiable pattern in the other trajectory; moreover, the bandwidth allocations for both trajectories are different from that of the equilibrium state.

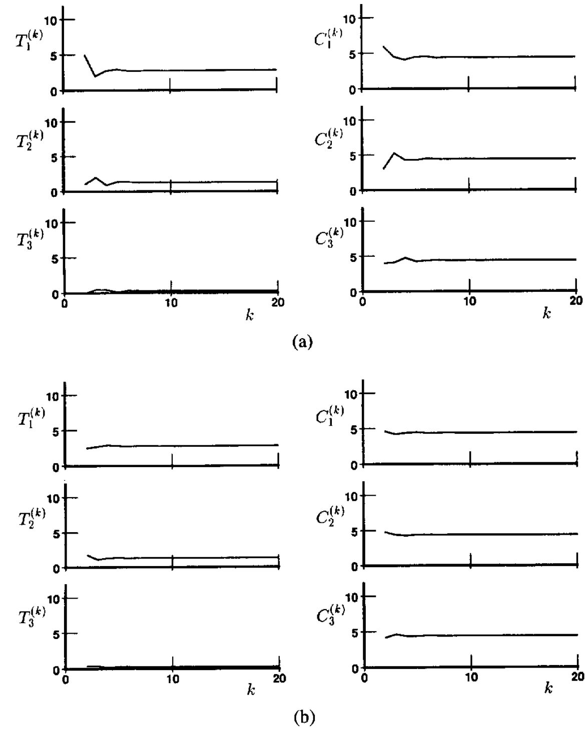

Fig. 4. Service times and cycle lengths for quota functions $0.5\max(M_i - C_i^{(k)}, 0)$ with $M_1 = 10$, $M_2 = 7$, $M_3 = 5$, and initial service vectors (a) (0, 0, 0), (b) (4, 1, 0).

The above comparison illustrates the undesirable properties of non-convergent systems. Indeed, the generalized access control strategies have been specially designed to avoid the unpredictable behavior of non-convergent systems observed at high load.

The monotone property of the dependency vector (*i.e.*, $1 - \delta \geq a_1 \geq a_2 \geq \cdots \geq a_N$) is crucial for the convergence behavior of the first system. As a result of this property, the token cycles constitute a moving average sequence with non-negative time-varying weights summing to unity. The convergence of this moving average sequence is clearly governed by those weights. The condition $\delta > 0$ forces the leading weight to be non-zero and this is sufficient to ensure convergence Fig. 4 shows that the convergence rate is much faster than the geometric bound with parameter $(1 - \delta^N)^{\frac{1}{N}}$ derived in this paper until near equilibrium. This is generally true because when there is a large variation in the token cycles, the true value of the leading weight in the moving average is much larger than δ, resulting in fast convergence in the token cycles. If $\delta = 0$, as in the second system, then the leading weight may be zero and the cycles may not converge. Nevertheless, the cycles oscillate within a "band" such that the slope of each quota function is either -1 or 0. One can check that both periodic patterns in Fig. 5 have token cycles oscillate between M_2 and M_3. If $\delta < 0$, as in the third system, then the weights can be negative; in this case, the token cycles and station transmission times can exhibit very strange behavior.

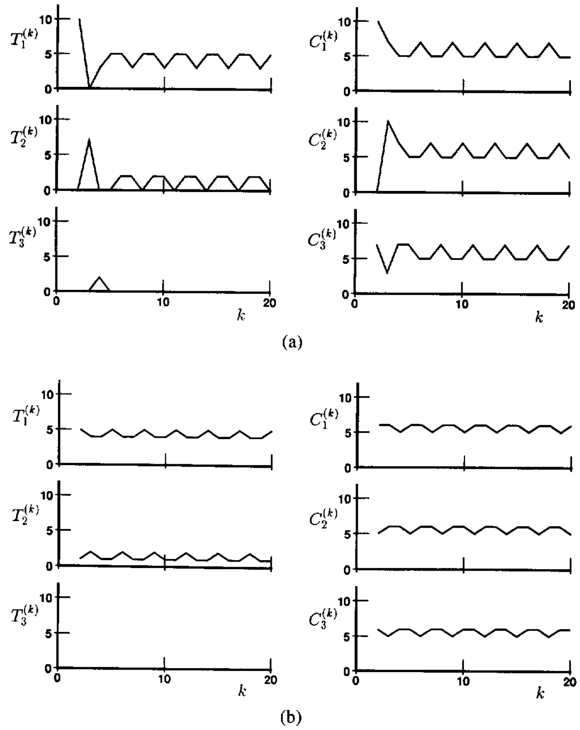

Fig. 5. Service times and cycle lengths for quota functions $1.0\max(M_i - C_i^{(k)}, 0)$ with $M_1 = 10$, $M_2 = 7$, $M_3 = 5$, and initial service vectors (a) (0, 0, 0), (b) (4, 1, 0).

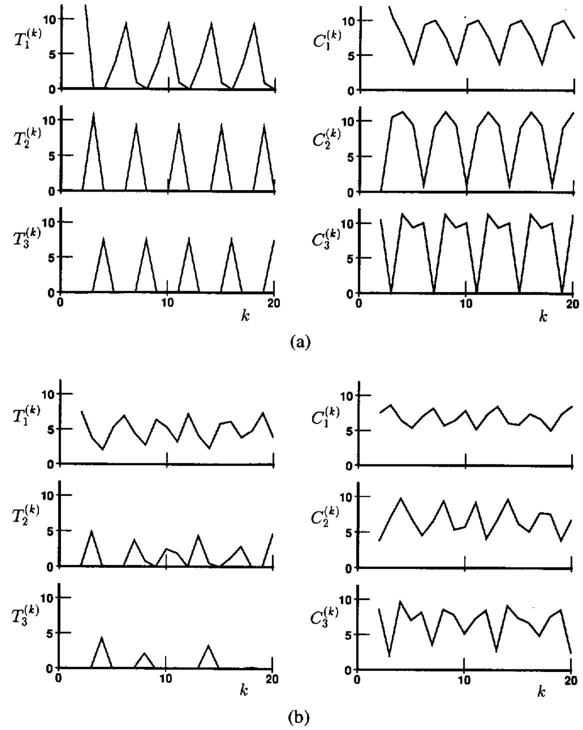

Fig. 6. Service times and cycle lengths for quota functions $1.5\max(M_i - C_i^{(k)}, 0)$ with $M_1 = 10$, $M_2 = 7$, $M_3 = 5$, and initial service vectors (a) (0, 0, 0), (b) (4, 1, 0).

In particular, we can see from Fig. 6 that the trajectories are strongly dependent on the initial conditions and all stations experience large fluctuations in their cycles and transmission times.

Under overloading conditions, the generalized control strategies provide TDM-like services with bandwidth allocation uniquely determined by the quota functions. Thus, we can specify the overload bandwidth allocation by appropriate assignment of the quota functions. Unfortunately, there are no known analytic techniques for solving the equilibrium state, except for special quota functions such as those given by (2.3) or (2.4). Nevertheless, it is worthwhile to explore a general property of the generalized control strategies.

Let $C^{(\text{eqm})} = W + \sum_{j=1}^{N} T_j^{(\text{eqm})}$ be the equilibrium cycle. The throughput of station i under heavy load is given by $\rho_i = T_i^{(\text{eqm})}/C^{(\text{eqm})}$, and the total system throughput under heavy load is given by $\rho = \sum_{j=1}^{N} \rho_j = 1 - W/C^{(\text{eqm})}$. It is noted that ρ_i is the minimum throughput of station i regardless of the load at other stations provided station i is heavily loaded. Furthermore, ρ is the maximum possible throughput of the system and it is often known as the *efficiency* of the system. A general result, which is very interesting and intuitive, is that the efficiency of a system increases as the quota functions increase, see [7] for the proof. This means that, from the efficiency point of view, the best system is the one with infinite quota functions (*i.e.*, exhaustive service). However, the token cycles will increase without bound as the load increases.

IV. Upper Bounds on Token Cycles Under Arbitrary Load

Stability is certainly not the only important issue. After all, the main purpose of access control is to avoid long token cycles that disrupt time-critical applications. One of the earliest approach to design token passing systems consistent with a constraint on the token cycle length was to enforce a fixed quota at each station such that the longest possible cycle is less than a specified maximum delay requirement. It was soon realized that this is far from being the best solution because (i) the delay constraint severely limits the total number of stations, and (ii) the system is under-utilized when few stations are active. A much better solution was adopted in [8] by dynamically limiting data transmissions that are not time-critical to relax the limitations on the number of stations imposed by the delay requirement. The idea behind the generalized control strategies is similar to that in [8] but with notable improvements. The objective of this section is to derive analytic results that facilitate practical design of time-constrained systems, that is, to select the parameters associated with the quota functions so as to guarantee all token cycles to be below a specified limit. To tackle the design problem, we first consider the problem of determining upper bounds on token cycles under arbitrary loading conditions, given the quota functions. We do retain the assumptions of constant token passing overheads and negligible overflow transmissions.

Before we embark on a detailed discussion, let us point out that a trivial upper bound on all token cycles is $W + \sum_{j=1}^{N} f_j(0,0,\ldots,0)$ since $T_j^{(k)}$ is bounded above by $f_j(0,0,\ldots,0)$ for $j = 1,2,\ldots,N$ and for all k. As we can see, this bound grows linearly in N, the total number of stations. If we were to use this bound as the design criterion, that is, to ensure that this bound is below the maximum delay requirement for real-time traffic, then we have to limit the total number of stations accordingly. It turns out that this is a very severe constraint on the total number of stations. It is true that for constant quota functions, this trivial bound is the tightest upper bound; but for many other quota functions, high correlations exist among the components of $\vec{T}^{(k)}$ and as a result, a much tighter bound, typically insensitive to the total number of stations, can be obtained.

Without loss of generality, we consider the problem of finding an upper bound for the token cycles experienced by station 1. Our approach is to consider first an upper bound on token cycles conditioned on the station transmission times in the previous cycle, and then the unconditional upper bound.

A. Conditional and Unconditional Upper Bounds

We first define, for an arbitrary vector $\vec{z} = (z_1, z_2, \ldots, z_N)$, $P_i(\vec{z}) = \sum_{j=1}^{i} z_j$ to be the ith partial sum. Given the k^{th} service vector $\vec{T}^{(k)}$, define the *overload vector* $\vec{U}^{(k+1)}$ and an *inflated overload* vector $\vec{V}^{(k+1)}$ by their respective components, $(U_1^{(k+1)}, U_2^{(k+1)}, \ldots, U_N^{(k+1)})$ and $(V_1^{(k+1)}, V_2^{(k+1)}, \ldots, V_N^{(k+1)})$, determined successively as follows:

$$U_j^{(k+1)} = f_j(U_{j-1}^{(k+1)}, \ldots, U_1^{(k+1)}, T_N^{(k)}, \ldots, T_j^{(k)}), \quad j = 1,2,\ldots,N \tag{4.1}$$

$$V_j^{(k+1)} \geq f_j(V_{j-1}^{(k+1)}, \ldots, V_1^{(k+1)}, T_N^{(k)}, \ldots, T_j^{(k)}), \quad j = 1,2,\ldots,N \tag{4.2}$$

The components of the vector $\vec{U}^{(k+1)}$ represent the transmission times in the $(k+1)^{\text{st}}$ cycle given that all stations use up their quotas. The components of the vector $\vec{V}^{(k+1)}$ are obtained in a similar way except the stations are successively allowed to overrun their quotas in any arbitrary fashion.

In [7], we show that $P_i(\vec{T}^{(k+1)}) \leq P_i(\vec{U}^{(k+1)}) \leq P_i(\vec{V}^{(k+1)})$ for $i = 1,2,\ldots,N$. In particular, we have $C_1^{(k+1)} = P_N(\vec{T}^{(k+1)}) \leq P_N(\vec{U}^{(k+1)}) \leq P_N(\vec{V}^{(k+1)})$. Furthermore, $P_i(\vec{U}^{(k+1)})$ is a decreasing function of each component of $\vec{T}^{(k)}$. This intriguing result states that, (i) a sudden overload of the system will yield the largest possible next cycle, (ii) sudden overloading of the empty system will yield the largest possible next cycle regardless of the service times in the current cycle, and (iii) a larger cycle would result if stations are allowed to overrun their quotas. The above results provide a straightforward technique for determining the tightest upper bound on token cycles as well as looser, but analytically simpler, bounds.

B. Discussion

To illustrate the use of the results in previous sub-sections, let us consider a system with quota functions given in (2.4) with $N = 3$, $W = 2$, $\gamma_1 = \gamma_2 = \gamma_3 = 0.8$, $M_1 = M_2 = M_3 = 22$, $U_1 = 15$ and $U_2 = U_3 = 10$. Starting from an empty system, we calculate the overload vector with respect to queue 1 and get (12, 6.4, 1.28). Summing the components of the overload vector and the total walk time, we find that the longest possible cycle experienced by queue 1 is 21.68. Similarly, the longest possible cycles for queues 2 and 3 are 21.2 and 21.52, respectively. Notice that the longest possible cycle experienced by different queues are not the same. By using the inflated overload vector $(\max_{1\leq i\leq 3}\{M_i\} - W, 0, 0) = (20, 0, 0)$, we show that 22 is a bound on all cycles experienced by all stations.

Using similar inflated overload vectors as above, it is not difficult to show that $\max(W, \max_i M_i)$ is a bound on all cycles, independent of the number of stations in the system. This is clearly a major advantage over the classical scheme using fixed quotas since the delay constraint does not pose any limit on the total number of stations. Of course, stations with time-critical data should not be delayed due to insufficient quota and hence they must have very large M_i's, possibly infinity. If $M_i = \infty$ for some i, then the maximum is clearly not a very good bound. By using a different inflated vector, it is possible to derive an alternative bound (see [7]) $\max(W, \max_{j\in R_1} M_j) + \sum_{j\in R_2} \gamma_j U_j$ where R_1, R_2 is an arbitrary partition of the set of stations. One possible choice of (R_1, R_2) is $R_2 = \{i \mid M_i = \infty\}$. In analogous to the standard protocol, R_1 is like the set of type I stations where the quota is fixed and R_2 is like the set of type II stations where the quota is subject to the token rotation time. Other choices may also be favorable, for example, including queues with small $\gamma_i U_i$ and large M_i in R_2. In general, we see that the delay constraint will primarily limit the number of time-critical stations rather than the total number of stations.

We have concentrated on the worst case cycle bound. In particular, we found that the worst situation occurs when an empty system is suddenly overloaded. It is unlikely that all stations will be overloaded at the same time. It is likely that, for bursty computer traffic, one or few stations become suddenly overloaded. To take advantage of this type of traffic characteristics, it may be reasonable to assume that at most a few of the stations can simultaneously become overloaded, and thus produce a smaller upper bound. In this kind of situations, the quota function given by (2.4) is more favorable over that given by (2.3) because if only one station, say i, becomes overloaded when the system is empty, then (2.4) yields $W + \gamma_i U_i$ for the next cycle but (2.3) yields the larger value of $\gamma_i(M_i - W)$ provided $M_i \geq W$. However, (2.3) does yield better efficiency.

In general, to satisfy a constraint on token cycle lengths and to maximize system efficiency are conflicting goals. We present a simple design example to illustrate this point. Consider a 10 Mbps token passing system with $W = 50$ μs. There are four groups of stations in the system, S_1 to S_4, each consisting of L stations, yielding a total of $4L$ stations. Stations in S_1 are 64kbps voice stations, generating voice packets of 6.4 μs long, while all other stations are data stations. Voice packets have a maximum delay constraint of 1ms. Among data stations, we would like to allocate bandwidths under overload in the ratio of 4:2:1 to stations in S_2, S_3, S_4 respectively, with a minimum of 100kbps to each station in S_2. The design problem is to maximize the total number of stations.

We shall use the basic time unit of 1 μs and omit all units from here on. As before, we consider quota functions given in example 2.2. For $i \in S_1$, we let $M_i = \infty$ and $\gamma_i U_i = 6.4$. For data stations, we let $U_i = M$, and $\gamma_i = 4l, 2l, l$ for $i \in S_2, S_3, S_4$, respectively. It is clear that the bandwidth constraints among data stations are satisfied by the chosen ratio in the γ_i's; of course, we must have $l < 0.25$ to ensure that the monotone property (2.2) is satisfied. The constraint on delay and bandwidth are $M + 6.4L \le 1000$, $\frac{4l(M - C^{(\text{eqm})})}{C^{(\text{eqm})}} \ge 10^{-3}$, where $C^{(\text{eqm})} - W = (6.4 + 7l(M - C^{(\text{eqm})}))L$. Maximizing L subject to the above constraints, we have $L^{(\max)} = 36$, $M = 769.6$ μs and l is quite arbitrary as long as $7lL^{(\max)} \gg 1$.

V. An Approximation for Station Throughputs

The central issues of this section are stability and station throughput under general load. Stability, in this section, refers to the situation where no stations have unbounded queues in the long run. We have derived useful overload throughput results. However, they are applicable only under restricted conditions. As an additional evaluation tool, we present a simple and yet very accurate approximation for station throughputs under general load. The approximation is based on the assumption that the variances of the station transmission times are small. We shall first develop the approximation and then explain its accuracy. Then we discuss an important application of the approximation. Finally, we present some simulation results for validating the approximation.

A. The Approximation

Consider a token system using the generalized quota functions, with infinite buffer at each station. Let r_i be the long run average input rate of work, defined as the product of the average packet arrival rate and the average packet transmission time, at station i. Let $\overline{C}_i \triangleq \lim_{L\to\infty} \frac{1}{L}\sum_{k=1}^{L} C_i^{(k)}$ be the long run average cycle of station i, provided the limit exists. $\overline{W}_i$, $\overline{T}_i$ and $\overline{Q}_i$ are defined analogously for the average token passing overhead, average station transmission time and average quota, respectively, for station i. Since $\overline{C}_i$ is independent of i, we use $\overline{C}$ for the average cycle of any station. The throughput of station i is defined by $\rho_i \triangleq \frac{\overline{T}_i}{\overline{C}}$ and the system throughput is defined by $\rho \triangleq \sum_{i=1}^{N} \rho_i$. If we further let $\overline{W} \triangleq \sum_{i=1}^{N} \overline{W}_i$, then it can be shown easily that $\overline{C} = \frac{\overline{W}}{1-\rho}$.

We make the assumption that a station in our system has either a finite or an infinite queue length. If a station has finite queue length, then $\rho_i = r_i$, else the station will use up its quota in each cycle and thus $\rho_i = \frac{\overline{Q}_i}{\overline{C}}$. Hence, $\rho_i = \min\left(r_i, \frac{\overline{Q}_i}{\overline{C}}\right)$. If the variances of the station service times are small, then $\overline{Q}_i$, the average of a quota function, can be accurately approximated by $f_i(\overline{T}_{i-1}, \ldots, \overline{T}_1, \overline{T}_N, \ldots, \overline{T}_i)$, the function of the averages. Indeed, if the variances of the station service times are zero, the approximation becomes exact. Using this approximation and the relations given in the last paragraph, we have

$$\rho_i' = \min\left(r_i, \frac{f_i(\rho_{i-1}'\overline{C}', \ldots, \rho_1'\overline{C}', \rho_N'\overline{C}', \ldots, \rho_i'\overline{C}')}{\overline{C}'}\right), \quad i = 1, 2, \ldots, N \tag{5.1}$$

where $\overline{C}' = \frac{\overline{W}}{1-\rho'}$ and $\rho' = \sum_{i=1}^{N} \rho_i'$. The quantities ρ_i', ρ' and $\overline{C}'$ denote the approximations for ρ_i, ρ and $\overline{C}$, respectively. Given $\{r_i\}$ and $\overline{W}$, (5.1) yields a unique solution for $\{\rho_i'\}$ and $\overline{C}'$. The proof can be found in [7], together with a numerical solution method. Furthermore, ρ' is a non-decreasing function of the input rates $\{r_i\}$, which shows that the approximation is consistent with the conflict-free nature of the token passing protocol.

The approximation is exact when the input rates are sufficiently large since the variances of the station transmission times are zero. In fact, the approximation is also exact for the cases where each r_i is either zero or very large. When all r_i's are small, the variances of the station transmission times are also small, and thus the approximation is very good. A more important observation is that when the input rate r_i is small, compared with the quantity $\frac{f_i(\rho_{i-1}'\overline{C}', \ldots, \rho_1'\overline{C}', \rho_N'\overline{C}', \ldots, \rho_i'\overline{C}')}{\overline{C}'}$, the right hand side of (5.1) will give the correct value, relatively independent of the error in estimating $\overline{Q}_i$ by $f_i(\overline{T}_{i-1}, \ldots, \overline{T}_1, \overline{T}_N, \ldots, \overline{T}_i)$. This explains why the approximation is still very accurate under many intermediate loads where the variances of service times may not be small. There is a further reason for the high accuracy of our approximation. Recall that the quota functions are decreasing with partial derivatives bounded between zero and minus one. Qualitatively, these quota functions tend to have slowly-varying derivatives, thus yielding a high accuracy in our approximation.

In passing, we note that the approximation is also applicable to the protocol used in [1]–[2]. When applied to the protocol used in [1]–[2], our approximation yields results equivalent to a recently proposed approximation [13]. However, the approximation in [13] is not applicable for the generalized access control strategies.

B. Applications of the Approximation

By solving (5.1) and comparing $\{r_i\}$ with $\{\rho_i'\}$, we can determine approximately whether a station is saturated, and if not, its throughput. If no stations are saturated, then there is a very useful quantity for measuring how close the system is to saturation. Let $\gamma_{\max}'$ be the largest value of γ such that $\{\gamma r_i\}$ yields a non-saturated system, determined by the approximation (5.1). Clearly, $\gamma_{\max}' > 1$ and the closer it is to unity, the closer is the system to saturation.

For a special class of generalized quota functions, there are additional implications of the approximation (5.1). A generalized quota function $f(x_1, x_2, \ldots, x_N)$ is convex if for any joint random variable $(X_1, X_2, \ldots, X_N)$ with positive components, $Ef(X_1, X_2, \ldots, X_N)$ is greater than or equal to $f(EX_1, EX_2, \ldots, EX_N)$ where the E operator refers to

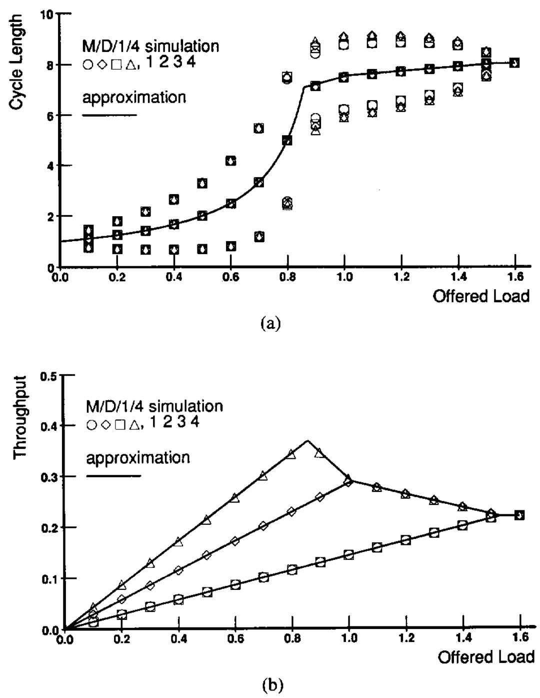

Fig. 7. (a) Cycle length versus offered load (M/D/1/4). (b) Station throughput versus offered load (M/D/1/4).

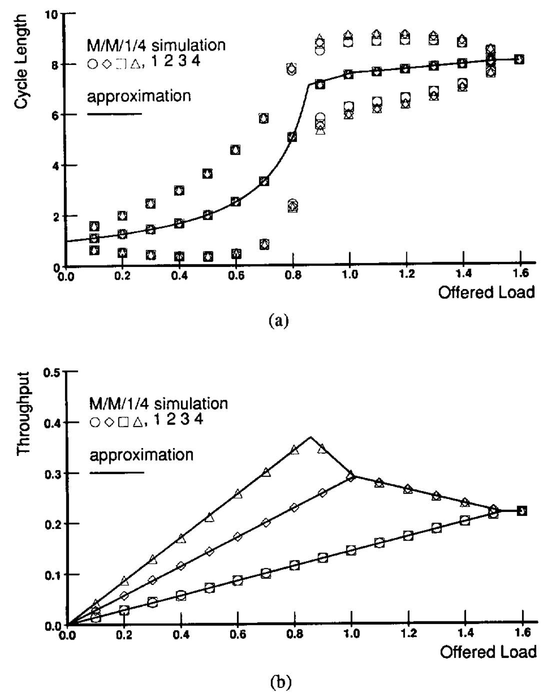

Fig. 8. (a) Cycle length versus offered load (M/M/1/4). (b) Station throughput versus offered load (M/M/1/4).

taking expectation. Two results, shown in [7], state that if the quota functions f_i are convex for all i, then (i) the approximate system throughput ρ' will never overestimate the true system throughput ρ, and (ii) a sufficient condition for stability is that

$$r_i < \frac{f_i(r_{i-1}\overline{C}'', \ldots, r_1\overline{C}'', r_N\overline{C}'', \ldots, r_i\overline{C}'')}{\overline{C}''}, \quad i = 1, 2, \ldots, N \tag{5.2}$$

where $\overline{C}'' = \frac{\overline{W}}{1-r}$ and $r = \sum_{i=1}^{N} r_i$. Thus, given the input rates, we can directly verify the stability of a system using convex quota functions, and also provide an upper bound on its maximum throughput.

C. Discussion

To validate our approximation, we present simulation results for two systems, A and B. Both systems have $\overline{W} = 1$ and $f_i = \gamma_i \min(U_i, \max(M_i - C_i^{(k-1)}, 0))$.

System A has four stations with $\gamma_i = 0.9$, $M_i = 10$ and $U_i = 5$ for $i = 1, 2, 3, 4$. The mean packet size is 1 and the mean packet arrival rates to the stations are in the ratio of 1:2:1:3 in their respective order. We simulate the system by varying the total arrival rates from zero to overload. To check the insensitivity of our approximation to high order moments of the input traffic, the simulations are performed for three cases:

(1) exponential inter-arrival times and fixed packet length,
(2) exponential inter-arrival times and exponential packet length,
(3) hyper-exponential (H_2) inter-arrival times and exponential packet length.

We obtain, from simulation, estimates for the mean and the standard deviation of the token cycles, and also for the station throughputs. We plot the mean, the mean plus one standard deviation, and the mean minus one standard deviation of the token cycles versus the total offered load in Figs. 7(a), 8(a), and 9(a) for cases (1), (2), and (3) respectively. On the same figures, the mean token cycles obtained from our approximation are also plotted. Similarly, the station throughputs from simulation are also plotted along with the approximation in Figs. 7(b), 8(b), and 9(b). From these plots, we can see that the station throughputs are indeed quite insensitive to the high order moments of the input traffic. Furthermore, the approximate results are virtually indistinguishable from the simulation results for a wide range of input load.

System B has eight stations with $\gamma_i = 0.9$ for all i, $M_i = 10$ and $U_i = 5$ for $i = 1, 2, 3, 4$, and $M_i = U_i = 8$ for $i = 5, 6, 7, 8$. The ratio of the packet arrival rates is 1:1:1:1:2:2:2:2 in their respective order. Packets have fixed length of 1 and the inter-arrival times are exponential. Simulations are run for total arrival rates from zero to overload. As before, the simulation results are plotted against the approximation in Fig. 10(a) and (b). From Fig. 10(a), we see that the approximation for average token cycle agrees very closely to the simulation results. When the total offered load is below 0.75, no stations are saturated, and the approximation for the individual station throughputs agree with simulation exactly. As the load increases beyond 0.75, stations 5 to 8 become saturated and the plots show

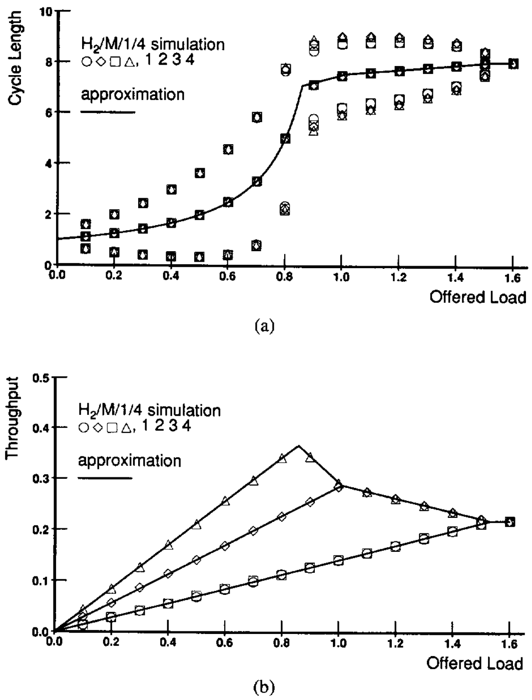

Fig. 9. (a) Cycle length versus offered load (H_2/M/1/4). (b) Station throughput versuss offered load (H_2/M/1/4).

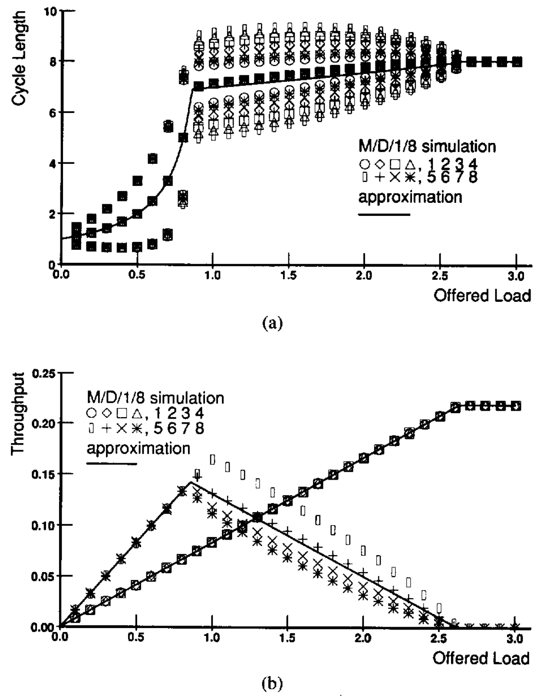

Fig. 10. (a) Cycle length versus offered load (M/D/1/8). (b) Station throughput versus offered load (M/D/1/8).

some discrepancies between the approximation and simulation. The approximation becomes exact again as the load further increases. The simulation reveals certain unfairness among stations 5 to 8. An intuitive reason is that station 5 has the earliest chance for picking up the "slack" offered by the first 4 unsaturated stations while station 8 has the least opportunity. Such spatial dependence of the station throughputs is not captured by the approximation. Nevertheless, the approximation still provides reasonably accurate results and more importantly, gives the correct trend at which the throughputs of the saturated stations decrease.

VI. Summary

In this paper, we have proposed and investigated a generalization of the integrated protocol adopted in [1]–[2]. The generalization is based on a distributed, dynamic control of the token cycles using generalized quota functions dependent on past station transmission times. We have identified the dynamics inherent in the control scheme as an important issue and shown, under overloading conditions that the proposed strategies have stable dynamics. We have given a comparison of several systems to illustrate the desirable stability properties of the generalization. By establishing results on the upper bounds, we have demonstrated the ability of the proposed strategies to control the token cycles. Examples have been given to illustrate the bounding techniques. The tradeoff between efficiency and delay constraint has also been demonstrated by an example. As an additional evaluation tool, we have proposed a simple and yet very good approximation for calculating station throughputs under arbitrary loading conditions. Simulations have been conducted to assess the accuracy of the approximation. The approximation allows us to explore the feasible region (*i.e.*, set of input traffic rates for which no stations are saturated), as well as the bandwidth allocation properties, for various control strategies.

In conclusion, the key advantages of the generalization are (i) to enrich the design space of access control strategies for token passing systems over those provided in the standards [1]–[2], so that we can design more efficient systems under more general constraints, (ii) to guarantee convergent behavior so that an accurate prediction of the station throughputs under overload can be made, and perhaps most importantly (iii) to provide an insightful understanding of the dynamic behavior inherent in the control scheme.

References

[1] The Institute of Electrical and Electronic Engineers, " Token-passing bus access method and physical layer specifications," American National Standard ANSI/IEEE std. 802.4-1985.

[2] *FDDI Token Ring Media Access Control*, draft proposed—ANSI Standard X3T9.5/83-16, rev. 10, Feb. 28, 1986.

[3] J. Pang and F. A. Tobagi, "Throughput analysis of a timer-controlled token-passing protocol under heavy load," *IEEE Trans. Commun.*, vol. COM-37, July 1989.

[4] J. Pang, F. A. Tobagi, and S. Boyd, "Generalized access control strategies for token passing systems," Comput. Syst. Lab., Stanford Univ., Stanford, CA, Tech. Rep. CSL-TR-88-371, Nov. 1988.

[5] J. Pang and F. A. Tobagi, "Generalized access control strategies for token passing systems," *IEEE INFOCOM'89*, Ottawa, Ont., Canada, Apr. 1989.

[6] J. Pang and F. A. Tobagi, "Throughput approximation for the generalized timer based protocol for token passing systems," Comput. Syst. Lab., Stanford Univ., Stanford CA, Tech. Rep. CSL-TR-89-375, Jan. 1989.

[7] J. Pang, "Access control strategies for token passing integrated services networks," Ph.D. dissertation, Dep. Elec. Eng., Stanford Univ., Stanford CA, 94305, Apr. 1989.

[8] F. A. Tobagi, F. Borgonovo, and L. Fratta, "EXPRESS-NET: A High Performance Integrated-Services Local Area Network," *IEEE J. Select. Areas Commun.*, vol. SAC-1, pp. 898–913, Nov. 1983.

[9] J. M. Ulm, "A timed token ring local area network and its performance characteristics," in *IEEE 7th Conf. Local Comput. Networks*, Minneapolis MN, Feb. 1982, pp. 229–236.

[10] K. Sevcik and M. J. Johnson, "Cycle time properties of the fddi token ring protocol," *IEEE Trans. Software Eng.*, vol. SE-13, pp. 376–385, Mar. 1987.

[11] M. J. Johnson, "Proof that timing requirements of the FDDI token ring protocol are satisfied," *IEEE Trans. Commun.*, vol. COM-35, pp. 620–625, June 1987.

[12] M. J. Johnson, "Fairness of channel access for non-time-criticial traffic using the FDDI token ring protocol," Res. Inst. for Advanced Comput. Sci., NASA Ames Res. Center, CA, Mar. 1986.

[13] D. Dykeman and W. Bux, "Analysis and tuning of the FDDI media access control protocol," *IEEE J. Select. Areas Commun.*, vol. SAC-6, pp. 997–1010, July 1988.

[14] A. P. Jayasumana, "Performance analysis of a token bus priority scheme," in *IEEE INFOCOM '87*, San Francisco CA, Mar. 1987.

[15] D. Janetzky and K. S. Watson, "Performance evaluation of the MAP token bus in real time applications," in *Advances in Local Area Networks*, K. Kummerle, J. O. Limb and F. A. Tobagi, Eds. New York: IEEE Press, 1987.

[16] C. A. Desoer and M. Vidyasagar, *Feedback Systems: Input-Output Properties.* New York: Academic, 1975.

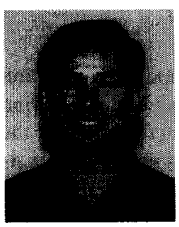

Joseph W. M. Pang received the B.A.Sc. and M.A.Sc. degrees from the University of British Columbia, Vancouver, B.C., Canada, in 1983 and 1985, respectively, and the Ph.D. degree from Stanford University, Stanford, CA, in 1989, all in electrical engineering.

From 1989 to 1992, he was with the Broadband Network Services group at Pacific Bell, San Ramon, CA. At Pacific Bell, he worked on advanced broadband network services, ATM, customer network management and metropolitan area networks. In 1992, he joined Starlight Networks, Mountain View, CA, doing research and development work on multimedia networking. His current research interests include system and protocol design for multimedia networking, LAN/WAN teleconferencing, high speed integrated services networks, hierarchical storage for digital video, and distance learning.

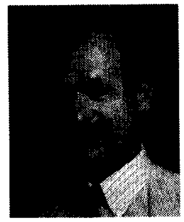

Fouad A. Tobagi (M'77–SM'83–F'85) received the engineering degree from Ecole Centrale des Arts et Manufactures, Paris, France, in 1970, and the M.S. and Ph.D. degrees in computer science from the University of California, Los Angeles, in 1971 and 1974, respectively.

From 1974 to 1978, he was a Research Staff Project Manager with the DARPA project at the Computer Science Department, University of California, Los Angeles, and engagedin modeling, analysis, and measurements of packet radio systems. In June 1978, he joined the faculty of the School of Engineering at Stanford University, Stanford, CA, where he is currently Professor of Electrical Engineering and Computer Science. He is also co-founder of Starlight Networks, a new venture concerned with multimedia networking, where he is serving as the Chief Technical Officer. His research interests include packet switching in ground radio and satellite networks, high-speed fiber optics local area networks, fast-packet switching, broadband integrated services digital networks, ATM, multimedia communications, modeling and performance eveluation, and VLSI implementation of network components.

Dr. Tobagi is the winner of the 1981 Leonard G. Abraham Prize Paper Award in the field of Communications Systems for his paper "Multiaccess Prtocols in Packet Communications Networks," and co-winner of the IEEE Communications Society 1984 Magazine Prize Paper Award for the paper "Packet Radio and Satellite Networks." He has served as Associate Editor for Computer Communications in the IEEE TRANSACTIONS ON COMMUNICATIONS for the period 1984–1986, Editor for Packet Radio and Satellite Networks in the *Journal of Telecommunications Networks* for the period 1981–1985, Co-Editor of the special issue on Local Area Networks of the IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS (November 1983), Co–Editor of the Special Issue on Packet Radio Networks of the PROCEEDINGS OF THE IEEE (January 1987), and Co-Editor of the Special Issue on Large Scale ATM Switching Systems for B-ISDN of the IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS (October 1991). He is also Co-Editor of *Advances in Local Area Networks*, a book in the series *Frontiers in Communications* (New York: IEEE Press). He is a member of the Association for Computing Machinery and has served as an ACM National Lecturer for the period 1982–1983.

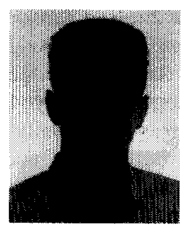

Stephen Boyd received the A. B. degree in mathematics, summa cum laude, from Harvard University in 1980, and the Ph.D. degree in EECS from the University of California, Berkeley, in 1985.

In 1985 he joined the Electrical Engineering Department at Stanford University, where he is now Associate Professor. In 1993 he became Director of the Information Systems Laboratory. His current research interests include numerical methods for convex optimization, especially interior-point methods; interactive engineering design; and computer-aided control system analysis and design.