A Sufficient Condition for Instability of Buffer Priority Policies in Re-Entrant Lines

Chuang Lin, Mingwei Xu, Dan C. Marinescu, Fengyuan Ren, and Zhiguang Shan

***Abstract*—We use a buffer-boundedness approach to study the stability of re-entrant lines with a buffer priority scheduling policy. Using Petri net models we establish a sufficient condition for instability of such systems having a positive feedback loop. An example of unstable systems is also given.**

***Index Terms*—Buffer boundedness, Petri net, priority scheduling, re-entrant lines, stability.**

I. Introduction

When designing and building a complex system it is necessary to select and enforce a scheduling policy for each shared resource. The selected scheduling policy helps the system achieve some objective function. For example, in the case of a manufacturing system, priority scheduling may allow us to respond more promptly to particular classes of customers.

Stability is a critical property of a scheduling policy. A queuing system is stable if the time in the system is bounded or, equivalently, if the number of customers in the system is bounded. The mean time in the system and the mean number of customers in the system are related by Little's law.

We assume all service times and arrival times are deterministic. In discrete-time dynamic systems, the notion "stability" commonly means "asymptotic stability," i.e., convergence of sample paths to a fixed and stable point. For Markovian systems, stability means the existence of a steady-state distribution, i.e., positive recurrence [1].

It is generally taken for granted that as long as the overall traffic intensity is less than unity [2] a network of queues is stable. However, it has been demonstrated that for a system with multiple classes of service and a deterministic scheduling policy based on priorities, instability may occur even for loads less than unity [3], [4].

Traditionally, system stability is studied based upon time boundedness of queuing models [1], i.e., the time in the system and the time spent by the customer in all states prior to its departure are bounded. In this note, we take a different approach. We study the buffer boundedness with the aid of Petri net [5], [6] models, and investigate the number of customers in the system based upon the markings of the Petri net model of the system. A scheduling policy is considered stable if the markings of all the places in the Petri net model of the system are bounded at all times [7].

Petri nets are good models for describing priority scheduling, as well as nondeterministic and asynchronous behavior. In addition, Petri net models can intuitively reveal a positive feedback structure, one of the major causes of instability. Here, we use Petri nets to model systems with buffer priority scheduling policies. We derive a sufficient condition for instability of systems containing a positive feedback loop (PFL).

This note is organized as follows. Re-entrant lines [4] and the conditions necessary for their stability are introduced in Section II. Section III outlines the properties for basic Petri net structures of systems with a PFL, and the relations of marking variation between two neighboring buffers. A sufficient condition for a system with a PFL to be unstable is presented in Section IV. At last, we give an example of unstable systems.

II. System Model

In this note, we study the stability of re-entrant networks of queues [2]. In our model, we assume that:

i) the routing is deterministic;
ii) the service time and the arrival time are deterministic;
iii) batch arrivals are allowed;
iv) static priority scheduling is supported;
v) scheduling is nonpreemptive.

The assumptions are as follows.

- There are S service stations.
- Each service station i consists of m_i identical servers (machines/processors/CPUs) that can run in parallel.
- There are K buffers.
- Customers in buffer k are served at service station $s(k)$, with service time t_K, and the service can be provided by one of the $m_{s(k)}$ machines located at the service station.
- Since routing is deterministic, a customer first arrives at buffer 1. After service is completed, the customer moves to buffer 2, and so on, until it finally reaches buffer K. After service is completed at this final buffer, the customer leaves the system.

For each $i = 1, \ldots, S$, the time units of work per machine at service station i, required by a customer is [3]

$$w_i = \sum_{\{k|s(k)=i\}} \frac{t_k}{m_i}. \quad (1)$$

To guarantee any practical form of stability for the system, the capacity constraint must be satisfied [3]

$$\rho = \max(\lambda w_i) < 1, \qquad \text{for } i = 1, \ldots, S \quad (2)$$

where ρ is the load of the system. Condition (2) is necessary but not sufficient for the stability of a system using priority scheduling [5].

Let $B_i = \{k|s(k) = i\}$ denote the set of buffers at service station i. The customers in different buffers at station i are served by the machines based on a scheduling policy. Several scheduling policies are discussed in [1], [3], and [8].

In a Petri net model of re-entrant lines, the buffers are represented by places, and the server processes are represented by transitions. In our model, every transition except for the first and the last connects two neighboring buffers, and acts as an output for one buffer and an input for the other. Customer arrivals to the system are represented by a transition with no input places and with the first buffer of the system as its output place. Customer departures from the system are represented by a transition with the last buffer as its input place and no output places. In such a Petri net model, the subnet obtained by deleting the first transition and the last transition is structurally bounded [6], i.e., under any initial marking, the submodel is always bounded.

Manuscript received June 28, 2001; revised March 27, 2002 and November 5, 2002. Recommended by Associate Editor R. S. Sreenivas. This work was supported by the National Natural Science Foundation of China under Grant 60173012 and Grant 90104002, the National Grand Fundamental Research 973 Program of China under Grant G1999032707, and the NSFC and RGC under Grant 60218003. The work of D. C. Marinescu was supported by the National Science Foundation under Grant MCB 9527131 and Grant DBI 9986316.

C. Lin, M. Xu, F. Ren, and Z. Shan are with the Department of Computer Science and Technology, Tsinghua University, Beijing 100084, China (chlin@tsinghua.edu.cn; xmw@csnet1.cs.tsinghua.edu.cn).

D. C. Marinescu is with the Computer Sciences Department, Purdue University, West Lafayette, IN 47907 USA (e-mail: dcm@cs.purdue.edu).

Digital Object Identifier 10.1109/TAC.2003.814268

The server process transitions are denoted by p_i, the corresponding firing rate by μ_i, and the delay by $t_i (t_i = 1/\mu_i)$. The transition corresponding to customer arrivals has a subscript *in*, its firing rate is λ, and the delay is $t_{in} (t_{in} = 1/\lambda)$. Each buffer i is represented by a place b_i. The customers are represented by tokens.

In a timed Petri net model without inhibitor and variable arcs, the throughput of the output transition will tend toward a limit, as the number of tokens in the input places increases, if the transition firing rate is independent of the marking of the places [7]. Indeed, the average token flow f_i through the transition p_i must be less than the average firing rate μ_i of the transition

$$f_i \leq \mu_i. \tag{3}$$

III. Stability of Basic Petri Net Models of Re-Entrant Lines With a PFL

Now, we discuss the characteristics of the basic Petri net structures encountered in modeling re-entrant systems with buffer priority scheduling. First, let us define the terminology useful in specifying re-entrant lines.

Note that the buffers are numbered according to the order in which they are visited by the customer. Without loss of generality, we model buffer b_i by a place labeled b_i in the Petri net model. We use the terminology buffer b_i and place b_i interchangeably in this note, when the two need not be distinguished. The server process that transfers customers (tokens) from buffer b_i to buffer b_{i+1} is modeled by a transition labeled p_i.

Definition 1: Buffer Sequence: A set of buffers $L = b_i, \ldots, b_n (1 \leq i < n)$ is called a buffer sequence if any pair of consecutive buffers b_k and b_{k+1} in the sequence are directly connected by a transition $p_k (i \leq k < n)$. ◇

In the Petri net model, inhibitor arcs are used to specify priority ordering between two buffers. An inhibitor arc drawn from place b_i to transition p_j, is called a *flow direction* inhibitor arc if $i < j$, else it is referred to as a feedback inhibitor arc. Flow direction inhibitor arcs are drawn in the direction of customer flow, while *feedback* arcs are in the opposite direction to customer flow.

Definition 2: Buffer Loop: If $L = b_i, \ldots, b_n (1 \leq i < n)$ is a buffer sequence and b_i and b_n are associated by a feedback inhibitor arc from b_n to transition p_i, L is called a buffer loop. When a buffer loop contains only a single feedback inhibitor arc, the buffer loop is called a simple buffer loop. ◇

Definition 3: PFL: Let $L = b_i, \ldots, b_n (1 \leq i < n)$ be a buffer loop. Furthermore, for every pair of neighboring buffers b_k and $b_{k+1} (i \leq k < n)$ either

1) there exists a flow direction inhibitor arc from b_k to transition p_{k+1}; or
2) the input arc that connects b_k to transition p_k and the output arc that connects p_k to b_{k+1} are variable arcs.

When there exists at least one flow direction inhibitor arc from buffer b_j in L to transition $p_{j+1} (i \leq j < n)$, then L is called a PFL. When $n - i = 2$, L is a special case of the PFL and is called a direct positive feedback loop (DPFL). ◇

To simplify the graphs, in the following models the places corresponding to the machines are removed unless explicitly specified. The marking of a buffer place b_i is a function of time t, denoted by $M(b_i, t)$.

In dynamic systems, stability is commonly used to mean "asymptotic stability," i.e., the convergence of a sample path to a fixed and stable point. In timed Petri net models, the samples are represented by the number of tokens in the places at different time points. Let us examine the relationship of changes in the number of tokens with time between two neighboring buffers, and the effect of the inhibitor arc.

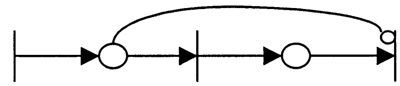

Fig. 1. Submodel with an inhibitor arc.

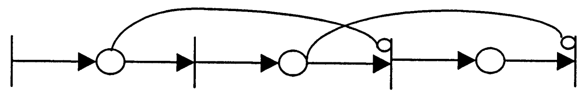

Fig. 2. Submodel with sequential inhibitor arcs.

In the following discussion, we assume that the parameters of the submodels satisfy (2).

Fig. 1 shows a submodel in which transition p_i has a higher priority than p_{i+1}, i.e., when b_i contains tokens, p_i is enabled but p_{i+1} is disabled until b_i is empty.

Let $M(b_i, t) > 0$, $M(b_{i+1}, t) = 0$ and let τ_i be a real number representing a time duration such that at time $t + \tau_i$ buffer b_i becomes empty, i.e., $M(b_i, t + \tau_i) = 0$. Then, τ_i satisfies

$$\tau_i = t_i \left(M(b_i, t) + \lfloor \lambda \tau_i \rfloor \right) \tag{4}$$

where $\lfloor x \rfloor$ denotes the largest integer that is smaller than or equal to x and t_i is the firing time of transition p_i. In (4), $\lfloor \lambda \tau_i \rfloor$ represents the number of tokens that enter b_i during τ_i. From (4), we have the following relationships:

$$\lambda \tau_i \geq \frac{\tau_i}{t_i} - M(b_i, t)$$
$$\lambda \tau_i < \frac{\tau_i}{t_i} - M(b_i, t) + 1$$

From these two relations and (4), we have

$$\frac{M(b_i, t) - 1}{\mu_i - \lambda} < \tau_i \leq \frac{M(b_i, t)}{\mu_i - \lambda}. \tag{5}$$

Therefore, from (4) at time $t + \tau_i$, we have

$$M(b_{i+1}, t + \tau_i) = M(b_i, t) + \lfloor \lambda \tau_i \rfloor \geq M(b_i, t). \tag{6}$$

Substituting (5) in (6), we get

$$\left\lfloor \lambda \left(\frac{M(b_i, t) - 1}{\mu_i - \lambda} \right) \right\rfloor + M(b_i, t) < M(b_{i+1}, t + \tau_i)$$
$$\leq M(b_i, t) + \left\lfloor \lambda \left(\frac{M(b_i, t)}{\mu_i - \lambda} \right) \right\rfloor. \tag{7}$$

Expression (4) should be revised when multiple inhibitor arcs are connected sequential. Fig. 2 shows an example where p_{i-1} has higher priority than p_i and, in turn, p_i has higher priority than p_{i+1}.

For the model in Fig. 2, let $M(b_i, t) > 0$, $M(b_{i+1}, t) = 0$, $M(b_{i-1}, t) = 0$ and let τ_i represent a time duration such that at time $t + \tau_i$ buffer b_i becomes empty, i.e., $M(b_i, t + \tau_i) = 0$. Then, τ_i satisfies

$$\tau_i = t_i \left(M(b_i, t) + \lfloor \lambda \tau_i \rfloor \right) + t_{i-1} \lfloor \lambda \tau_i \rfloor. \tag{8}$$

In (8), the last term represents the time when p_i is disabled because of the presence of tokens in b_{i-1}. This delay is equal to the time taken by p_{i-1} to move all the tokens from b_{i-1} to b_i. If there are multiple transitions $p_k (k < i)$ which have higher priority than p_i, i.e., there are multiple flow direction inhibitor arcs that end in p_i, the coefficient t_{i-1}

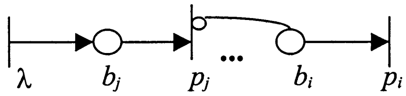

Fig. 3. Subnet model with a feedback inhibitor arc.

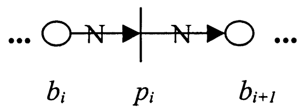

Fig. 4. Subnet structure with variable arcs.

of the second term should be replaced by the sum of the firing times of all these transitions $p_k (k < i)$. Equation (8) can be expressed as

$$\frac{t_i M(b_i,t) - (t_i + t_{i-1})}{1 - \lambda(t_i + t_{i-1})} < \tau_i \leq \frac{t_i M(b_i,t)}{1 - \lambda(t_i + t_{i-1})}. \tag{9}$$

When p_i is serving the last customer in buffer b_i, a customer arriving into b_{i-1} cannot contribute to the waiting time of p_i because the scheduling is nonpreemptive. Thus, the last term $t_{i-1}\lfloor \lambda\tau_i \rfloor$ becomes $t_{i-1}(\lfloor \lambda\tau_i \rfloor - 1)$, which is greater than or equal to zero. In this case, (9) becomes

$$\frac{t_i M(b_i,t) - (t_i + 2t_{i-1})}{1 - \lambda(t_i + t_{i-1})} < \tau_i \leq \frac{t_i M(b_i,t) - t_{i-1}}{1 - \lambda(t_i + t_{i-1})}. \tag{10}$$

Let us now consider a subnet with a feedback inhibitor arc, as shown in Fig. 3.

For this model, let $M(b_i,t) > 0$, $M(b_j,t) = 0$ and let τ_i represent a time duration such that at time $t + \tau_i$, buffer b_i becomes empty, i.e., $M(b_i, t + \tau_i) = 0$. Then, τ_i satisfies

$$\tau_i = t_i \left(M(b_i,t) + f(k) \right). \tag{11}$$

Where $f(k)$ denotes the number of customers which enter b_i during τ_i. Here, k depends on the size of τ_i, the sum of the delay times for transitions p_j to p_{i-1}, and the structure of the path (e.g., the existence of inhibitor arcs) from p_j to p_{i-1}. Therefore, at time $t + \tau_i$ we have

$$M(b_j, t + \tau_i) = \lfloor \lambda\tau_i \rfloor \geq \lfloor \lambda t_i M(b_i,t) \rfloor. \tag{12}$$

Fig. 4 shows a subnet where transition p_i connects two buffers b_i and b_{i+1} with a pair of variable arcs. This represents the case where the machines associated with server process p_i can work on a batch of customers at a time.

Let $M(b_i,t) > 0$, $M(b_{i+1},t) = 0$ and let τ_i represent a time duration such that at time t + τ_i buffer b_i becomes empty, i.e., $M(b_i, t + \tau_i) = 0$. We have

$$M(b_{i+1}, t + \tau_i) = M(b_i, t), \qquad \tau_i = t_i \tag{13}$$

Definition 4: Stability of Buffer Priority Scheduling Policies: Given a scheduling policy, if for any arrival rate λ for a system satisfying (2)

$$M(b_i,t) \leq N \qquad \forall\, b_i \text{ at any time instant } t \tag{14}$$

for some positive integer N (which could depend on the initial marking as well as λ), we say that the scheduling policy is stable. ⋄

When a system is unstable, there is at least one buffer b_i whose contents are unbounded in time, i.e., $M(b_i,t)$ may go to infinity. There are only two possible reasons for the unboundedness of b_i in such a system model. The first reason could be that the capacity of the output transition p_i of buffer b_i or the sum of the capacities of the set of consecutive transitions which have last-buffer–first-served priorities in the same service station are smaller than the system throughput; that is, (2) is not satisfied. However, this cannot be true, since we assumed that the system satisfies (2). The second reason could be that there is a positive feedback circulation for $M(b_i,t)$. In such a system model, only feedback inhibitor arcs can cause the positive feedback.

IV. Instability of Re-Entrant Lines With a PFL

We now derive the conditions for instability in re-entrant systems based on the properties of the basic Petri net structures given in Section III. If the system satisfies (2), and the net structure contains a PFL, we can determine whether the system is unstable using the following theorem.

Theorem: Sufficient Condition for Instability of Re-Entrant Systems With a PFL: Let a system $L = b_1, \ldots, b_K (K > 2)$ contain at least one PFL $L_1 = b_1, \ldots, b_n, \ldots, b_m (1 \leq n < m \leq K)$ in which there is a flow direction inhibitor arc connecting buffer b_n to transition p_{n+1}. The rate of customer arrivals to the system is λ and the system satisfies (2). If there exists an integer $N (N > 1)$ such that (15) holds, then the system is unstable

$$\left\lfloor \frac{\lambda}{\mu_m} \left(\left\lfloor \lambda \left(\frac{N-1}{\mu_n - \lambda} \right) \right\rfloor + N \right) \right\rfloor > N. \tag{15}$$

Proof: According to (7)–(13), for a PFL there exists a time sequence $\tau_1 < \tau_2 <, \cdots, < \tau_n <, \cdots, < \tau_{m-1}$ Such that $M(b_1,t) \leq M(b_2, t+\tau_1) \leq \cdots \leq M(b_{n+1}, t+\tau_n) \leq \cdots \leq M(b_m, t + \tau_{m-1})$ and

$$M(b_{n+1}, t+\tau_n) > \left(\left\lfloor \lambda \left(\frac{M(b_n, t+\tau_{n-1}) - 1}{\mu_n - \lambda} \right) \right\rfloor + M(b_n, t+\tau_{n-1}) \right).$$

Because $M(b_1,t) \leq M(b_n, t+\tau_{n-1})$ and $M(b_{n+1}, t+\tau_n) \leq M(b_m, t+\tau_{m-1})$, we have

$$M(b_{n+1}, t+\tau_n) > \left(\left\lfloor \lambda \left(\frac{M(b_1,t) - 1}{\mu_n - \lambda} \right) \right\rfloor + M(b_1,t) \right)$$

and

$$M(b_m, t+\tau_{m-1}) > \left(\left\lfloor \lambda \left(\frac{M(b_1,t) - 1}{\mu_n - \lambda} \right) \right\rfloor + M(b_1,t) \right).$$

Let $M(b_1,t) = N$, and τ_m represent a time duration such that $M(b_m, t+\tau_m) = 0$ at time instant $t + \tau_m (\tau_m > \tau_m - 1)$. Then, according to (12), we have

$$\begin{aligned} M(b_1, t+\tau_m) &\geq \left\lfloor \left(\left\lfloor \lambda \left(\frac{M(b_1,t) - 1}{\mu_n - \lambda} \right) \right\rfloor + M(b_1,t) \right) \frac{\lambda}{\mu_m} \right\rfloor \\ &= \left\lfloor \frac{\lambda}{\mu_m} \left(\left\lfloor \lambda \left(\frac{N-1}{\mu_n - \lambda} \right) \right\rfloor + N \right) \right\rfloor. \end{aligned}$$

The previous equation together with (15) implies that $M(b_1, t+\tau_m) > M(b_1,t)$. Thus the number of customers in the buffers increases as a function of time; that is, the system is unstable. ⋄

For (15) to hold, the initial marking $M(b_1,t)$ must be greater than 1. When the PFL includes a number of flow direction inhibitor arcs, the sufficient condition for instability can have a looser constraint. In the theorem, (15) is suitable for any PFL.

V. Example: An Unstable Client-Server System

In this example, we consider the client-server system shown in Fig. 5. Both the client and the server have two processes associated with them. We denote the processes at the client as p_1 and p_4, and those at the server as p_2 and p_3. Processes $p_1, \ldots, p_4$ are associated with buffers $b_1, \ldots, b_4$, respectively. Process $p_4 (p_2)$ has a higher priority (nonpreemptive) over $p_1 (p_3)$. First, we assume that transactions arrive at the

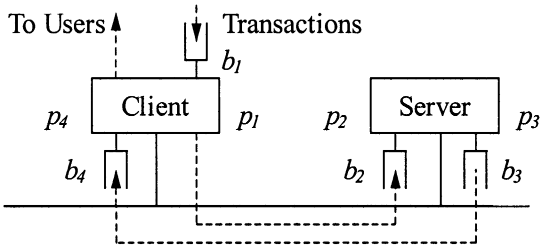

Fig. 5. Client-server system.

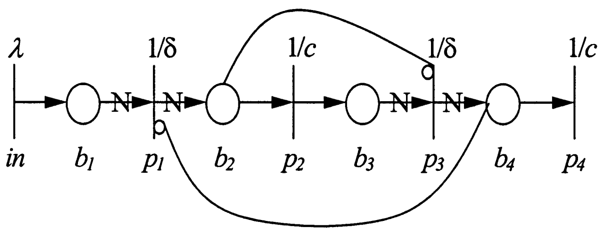

Fig. 6. Petri net model of the client-server system.

TABLE I
BUFFER MARKING FUNCTIONS

t	$M(b_1)$	$M(b_2)$	$M(b_3)$	$M(b_4)$
0^+	m	0	0	0
1/6	0	m	0	0
1/6+2m-2+2/3	0	0	3m-2	0
4/3+2m-2	0	1	0	3m-2
2/3+4m-2	2m-1	0	0	0

client in a deterministic fashion at the rate $\lambda = 1$. We further assume that each transmission of p_1 and p_3 takes the same CPU time δ, and the processing times of p_2 and p_4 for each transaction are equal and denoted by c. Finally, we assume that $c + \delta < 1$, i.e., (2) is satisfied.

A timed Petri net model of the client-server system is shown in Fig. 6. Transition *in* models the arrival of transactions to the client buffer b_1 at the rate λ. Transition p_i models process p_i for $i = 1, 2, 3, 4$. The inhibitor arc from b_4 to p_1 models the priority for process p_4 over p_1, and the inhibitor arc from b_2 to p_3 models the priority for p_2 over p_3. The cardinality of the variable arc associated with p_1 is equal to $M(b_1)$, and the cardinality of the variable arc associated with p_3 is equal to $M(b_3)$.

In the model shown in Fig. 6, a feedback loop is formed by the inhibitor arc from b_4 to p_1. This loop includes a pair of variable arcs between b_1 and b_2, a flow direction inhibitor arc connecting b_2 and b_3 and a pair of variable arcs between b_3 and b_4. So, the loop is a PFL. Let us set the initial marking, $M(b_1, 0^+) = m$ and $M(b_i, 0^+) = 0$ for i = 2,3,4, and set $\delta = 1/6$ and $c = 2/3$. Table I lists the buffer contents at different points. These results are obtained using (5), (7), and (11)–(13). At time $2/3 + 4\ m - 2$, the number of transactions waiting in b_1 increases to $2m - 1$. This shows that there is a positive feedback for the number of transactions in b_1 when $m > 1$. With increase in time this cycle repeats, and the number of transactions in b_1 grows indefinitely, i.e., the system is not stable. When $\lambda = 1$, $\mu_n = 3/2$ and $\mu_m = 3/2$, we can see from (15) that the sufficient condition for this system to be unstable is $m > 1$.

REFERENCES

[1] P. R. Kumar and S. P. Meyn, "Stability of queueing networks and scheduling policies," *IEEE Trans. Automat. Contr.*, vol. 40, pp. 251–260, Feb. 1995.
[2] L. Kleinrock, *Queueing Systems Volume I: Theory*. New York: Wiley, 1975.
[3] S. H. Lu and P. R. Kumar, "Distributed scheduling based on due dates and buffer priorities," *IEEE Trans. Automat. Contr.*, vol. 36, pp. 1406–1415, Dec. 1991.
[4] J. G. Dai and G. Weiss, "Stability and instability of fluid models for re-entrant lines," *Math. Oper. Res.*, vol. 21, pp. 115–134, 1996.
[5] T. Murata, "Petri nets: properties, analysis and applications," *Proc. IEEE*, vol. 77, pp. 541–580, Apr. 1989.
[6] C. Lin and D. C. Marinescu, "Stochastic high level Petri nets and applications," *IEEE Trans. Comput.*, vol. 37, pp. 815–825, July 1988.
[7] M. K. Molloy, "Fast bounds for stochastic Petri nets," in *Proc. Int. Workshop Timed Petri Nets*, Torino, Italy, July 1985, pp. 244–249.
[8] D. Gamarnik, "On deciding stability of scheduling policies in queueing systems," in *Proc, 11th Annu. ACM-SIAM Symp. Discrete Algorithms*, San Francisco, CA, 2000, pp. 467–476.

Stability of Data Networks Under an Optimization-Based Bandwidth Allocation

Heng-Qing Ye

***Abstract*—It is known that a data network may not be stable at the connection level under some unfair bandwidth allocation policies, even when the normal offered load condition is satisfied, i.e., the average traffic load at each link is less than its capacity. In this note, we show that, under the normal offered load condition, a data network is stable when the bandwidth of the network is allocated so as to maximize a class of general utility functions. Using the microscopic model proposed by Kelly for a transmission control protocol (TCP) congestion control algorithm, we argue that the bandwidth allocation in the network dominated by this algorithm can be modeled as our bandwidth allocation model, and hence that the network is stable under the normal offered load condition. This result may shed light on the stability issue of the Internet since the majority of its data traffic is dominated by the TCP.**

***Index Terms*—Bandwidth allocation, data network, Lyapunov function, stability, transmission control protocol (TCP).**

I. INTRODUCTION

There is no doubt that the Internet has been one of the most exciting and revolutionary technological developments in the past decade. The information flows along the Internet are still increasing dramatically, and the traffic control of the information flows has been an important issue in both the academics and the telecommunication industry. Currently, the majority of the Internet traffic is dominated by various versions of the transmission control protocol (TCP); see, for example, [7],

Manuscript received August 29, 2001; revised May 24, 2002 and October 31, 2002. Recommended by Associate Editor X. Zhou. This work was supported in part by a grant from the Academic Research Fund and the Center for E-Business of the National University Singapore. Part of this work was done while the author was visiting the Statistical Laboratory, Cambridge University, Cambridge, U.K.

The author is with the School of Business, National University of Singapore, 117591 Singapore (e-mail: bizyehq@nus.edu.sg).

Digital Object Identifier 10.1109/TAC.2003.814269