SAND--96-8534C
CONF-960278--1

ATM_FORUM/96-0076

ATM FORUM: Traffic Management Subworking Group

TITLE: Effect of Self-Similar Traffic on the Performance and Buffer Requirements of ATM ABR Edge Devices

SOURCE:

Siamak Dastangoo, George J. Miller
sdas, gmiller@gateway.mitre.org
The MITRE Corporation,
Virginia

Helen Chen, Jim Brandt
hycsw, brandt@ca.sandia.gov
Sandia National Laboratories,
California

DATE: February 5, 1996 Los Angeles, CA

DISTRIBUTION TO: The ATM Forum Members

Abstract

Previous studies [1] [6] demonstrated that Ethernet local area network traffic is statistically self-similar and that the commonly used Poisson models are not able to capture the fractal characteristics of Ethernet traffic. This contribution uses simulated self-similar traffic traces from the MITRE Corporation [2] and Sandia's simulation software [3] to evaluate the ABR performance of an ATM backbone. The ATM backbone interconnects Ethernet LANs via edge devices such as routers and bridges. We evaluate the overall network performance in terms of throughput, response time, fairness, and buffer requirements. Because typical edge devices perform simple forwarding functions, their usual mechanism for signaling network congestion is packet dropping. Therefore, we believe that the proper provisioning of buffer resources in ATM edge devices is crucial to the overall network performance.

Acknowledgment

We would like to thank Chien Fang for the use of his ATM switch and host simulators. In addition, we express our appreciation to Jim Brandt and Rose Tsang for providing helpful suggestions to this contribution.

1. Introduction

This contribution evaluates the ABR performance using simulated self-similar traffic traces from the MITRE Corporation. We simulate ATM backbones that provide ABR services using a simple two-node fan-in/fan-out and a four-node parking-lot configurations. The self-similar model and its generation method are described in [2]. We based our evaluation on network performance on throughput, fairness, responsiveness, and buffer requirements.

2. Simulation

2.1 Network Topology and Evaluation Methodology

The self-similar traffic traces that we used are generated using a Hurst parameter of 0.75, and they have average burst rates of 10 Mbps (23,584 cells/s) and 76 Mbps (179,245 cells/s). The bursty characteristics of a 10-Mbps trace and a 76-Mbps trace are depicted in Figures 1 (a) and (b) respectively. Please note that the 70-Mbps model's burst-rate frequently exceed our simulation link rate of 155 Mbps (365,566 cells/s). For simplicity, we will reference these traffic traces as the 10-Mbps source and the 70-Mbps source throughout the remainder of this paper. In this study, we use a traffic source to represent the aggregate traffic from all Ethernet hosts having a common destination over the ATM backbone. Using the two-node configuration (Figure 2), we collected statistics for the baseline (single session), "two-session with staggered-start-time", and the "six-session-congestion" studies. In addition, a four-node parking-lot configuration (Figure 3) was used to conduct more studies on max-min fairness, network throughput, and buffer requirements. By varying the inter-switch distances, we performed our simulation study for both LAN and WAN scenarios. We simulated all LAN configurations using a 4-Km distance for all links. WAN scenarios were created by increasing the inter-switch link distances to 4000-Km for the two-node topology and 1333-Km for the three links in the four node parking-lot topology. All links are at OC3c rate.

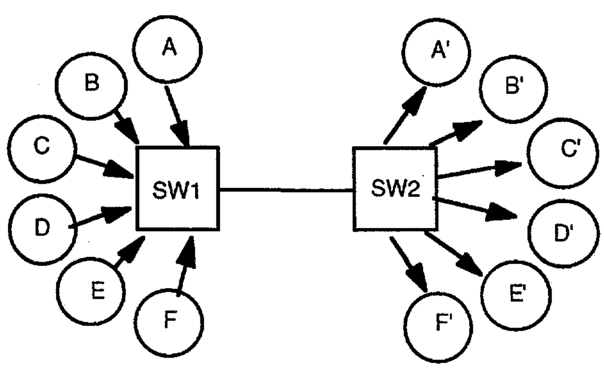

Figure 2. Two-Node Fan-in and Fan-out Configuration

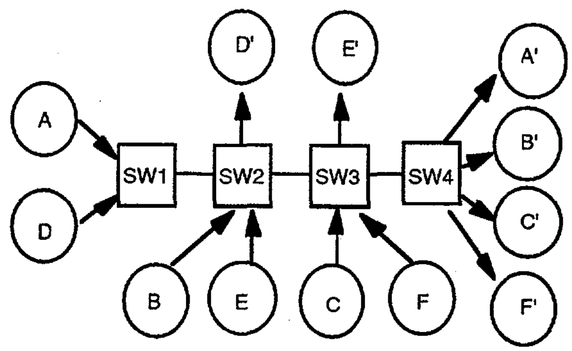

Figure 3. Parking-Lot Configuration

2.2 Simulation Parameters

The ABR EFCI switch that we simulated has unlimited output buffer but an EFCI threshold of 500 cells. This arrangement allows observation of EPRCA's rate control behavior as well as its switch buffer requirements under loss-less conditions. We also provide unlimited output buffer at the ATM edge devices, so that we can gain insight into their buffer requirements. The EPRCA parameters that we used are adopted from the study by C. Fang and A. Lin [4] and are listed in Table 1.

Table 1. EPRCA Simulation Parameters

PCR	AIR	ICR	MCR	RDF	Nrm	TOF	Trm	TCR	Mrm	Xrm	Xdf	TDFF
155	0.01892	10	0.5	512	32	2	100	10	2	32	0.5	0.5

2.3 Results and Analysis

2.3.1 Baseline Study

Results from the baseline studies of the 10-Mbps and 70-Mbps sources are plotted in Figures 4, 5, 6, and 7. Figures 4 (a) and (b) depict ACR statistic, the 70-Mbps source-burst-rate, and the ATM egress-rate for LAN and WAN respectively. The corresponding queue states of the LAN and WAN edge devices are plotted in Figures 5 (a) and (b). For both the LAN and WAN cases, the traffic source averaged around 179,245 cells/s or ~76 Mbps and burst rates frequently exceeded the OC3c link rates (Figure 4). This behavior resulted in some queue build-up at the edge device (Figure 5). The ACR curves in Figures 4 (a) and (b) show the slow response of the EPRCA - EFCI mechanism. The symptom is further amplified in the WAN scenario (Figure 4-b). The slow ramp-up of the ACR has resulted in the large initial queue buildup at the edge device (Figure 5). Also shown in Figure 4 are the LAN/WAN data egress rates. Even though the average egress rate approximates the average source rate, its peak burst rate was limited to the OC3c link speed.

Similar to Figure 4 and 5, and Figure 6 and 7 depict the source rate, ACR, ATM egress rate, and edge device queue states for the 10-Mbps baseline study. We observe the average source and egress rates at ~25,000 cells/s or 10 Mbps, the slow response of the

EPRCA-EFCI, and some queue buildup at the edge device. Of particular interest is the observed ACR behavior. As shown in Figure 6, the ACR in both LAN and WAN cases eventually grew to the link rate in the absence of congestion, even though their average source rates stayed at around 10 Mbps. If several such sessions, at steady state, increased their average burst rates simultaneously, it would be interesting to examine the effect on switch queue states

2.3.2 Two-sessions Study with Staggered-start

This experiment uses two 70-Mbps sources with their starting times staggered by 100 ms. Figure 8 plots the ACR curves of the two 70-Mbps sessions for both the LAN and WAN cases. Figure 9 depicts the egress rates of the same two sessions again for both LAN and WAN scenarios. As shown, fairness is achieved after reaching steady state. However, the results also demonstrate the slow response of the scheme. The time required to reach steady state is ~75 ms for the LAN configuration and ~225 ms for the WAN configuration. Figure 10 (a) and (b) show the output queue states of the source devices for the LAN and WAN cases. The queue sizes are much larger in the WAN case than that of the LAN case due to the larger feed back and slower ACR growth during startup. The switch queue states are depicted in Figures 11 (a) and (b) for the LAN and WAN switch respectively. As shown, they are well under control with the EPRCA parameters that we have simulated.

2.3.3 Six-session Congestion Study

This study uses the two-node configuration to connect six concurrent sessions (Figure 2). While sessions A, B, and C use the 10-Mbps sources, sessions D, E, and F use the 70-Mbps sources. We conducted this experiment to evaluate the fairness between sessions varying in bandwidth demands, 10 Mbps and 70 Mbps. Figures 12 and 13, depict ACRs and ATM egress rates of the LAN simulation while Figures 14 and 15 plot these for the WAN simulation. As Figures 12, 13, 14, and 15 show, the 10-Mbps sessions received their full bandwidth requirements and the 70-Mbps sessions were throttled to ~42 Mbps. For this configuration, our simulation demonstrated good fairness in sharing bandwidth among the competing sessions. However, the ACR plot of the 10 Mbps (Figure 12) shows that the ACR's of all three 10-Mbps sessions have grown to ~42 Mbps which equal the ACR's of the three 70-Mbps sources. This could result in sudden switch queue growth due to sudden increase in average burst rate of the 10-Mbps sources.

The output queue states of this experiment are presented in Figures 16 and 17 for LAN and WAN studies respectively. Because the three 10-Mbps sessions have received their demanded bandwidth, there are only small queue buildups due to occasional bursts that exceed their ACR. However, with the 70-Mbps sources, we observed linear queue growth in both the LAN and the WAN configuration. At 1500 ms into the simulation, the queue sizes had reached ~125,000 cells at each of the sources. It would be interesting to investigate the queue buildup for Poisson and continues sources with the same input rates as we believe that the burstiness of the traffic and its perturbation of the ACR might affect this. Again, the switch queue sizes (Figure 18) are well controlled with the EPRCA-EFCI implementation and this simple configuration

2.3.4 The MAX-MIN Fairness Study

We evaluated the EPRCA-EFCI scheme against the MAX-MIN fairness criteria using a simple four-node parking-lot configuration and six 70-Mbps sources (Figure 3). In this configuration, the theoretical fair share according to the MAX-MIN fairness guidelines are listed in Table 2. As described in section 2.1, this simulation is conducted for both LAN and WAN scenarios

Table 2. Theoretical MAX-MIN throughput in Mbps

A	B	C	D	E	F
38	38	38	115	75	38

The ACR and the ATM egress rate for this study are depicted in Figures 19 and 20. Results demonstrate that the EPRCA-EFCI scheme seemed to have achieved acceptable MAX-MIN fairness for this test. Similar to previous studies, the ACR curves show the slow response of the scheme and that the ACRs reflect source traffic's peak rather than average rate. And again there is a linear growth in output queue in the four sessions (A, B, C, and F) that bottlenecked at the last inter-switch link (Figure 21). Switch queue states in this experiment (Figure 22) demonstrate that the EPRCA-EFCI scheme was also able to keep the queue sizes within reasonable bounds using the parking-lot configuration.

3. Summary and Conclusion

This contribution evaluates the ABR performance using simulated self-similar traffic traces. We simulate ATM backbones that provide ABR services using the EPRCA-EFCI mechanism for flow control. We used a simple two-node fan-in/fan-out and a four-node parking-lot configuration to evaluate throughput, fairness, responsiveness, and buffer requirements.

Although we do not have hard numbers for fairness or link utilization, the plots are sufficiently clear to show that overall both fairness and link utilization were good in our test scenarios. In addition, the ABR flow control is able to protect the network during congestion. However, results demonstrate the slow response of the EPRCA-EFCI scheme which agrees with our previous study [5]. In addition, this study has revealed a concern due to the slow response of the ABR scheme. Using self-similar traffic sources, we observe that the steady state ACRs seem to reflect the peak cell rate during bursts, and is not adjusted downward fast enough to reflect the low average cell rate. If several such sessions increased their average rates simultaneously, it could cause the switch queue to overflow resulting in packet loss. This concern requires further study. We will investigate the possible interaction of the burstiness of the traffic and the ABR flow control, its perturbation to ACR, and the impact on buffer requirements. Furthermore, we also plan to examine the relative performance in the EPRCA-ER scheme and the EPRCA-EFCI mechanism with per VC queueing.

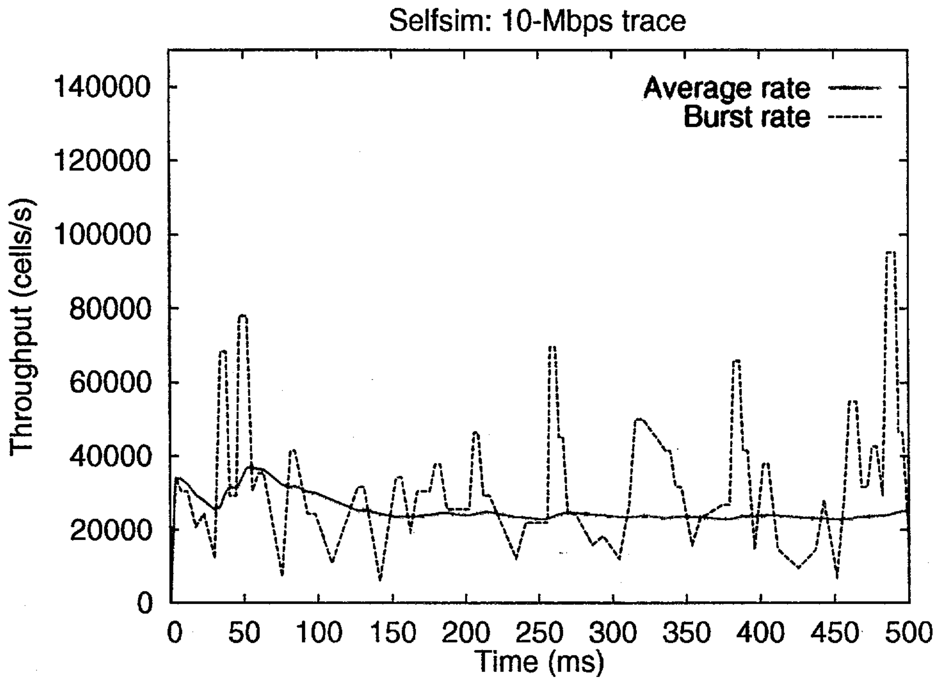

(a)

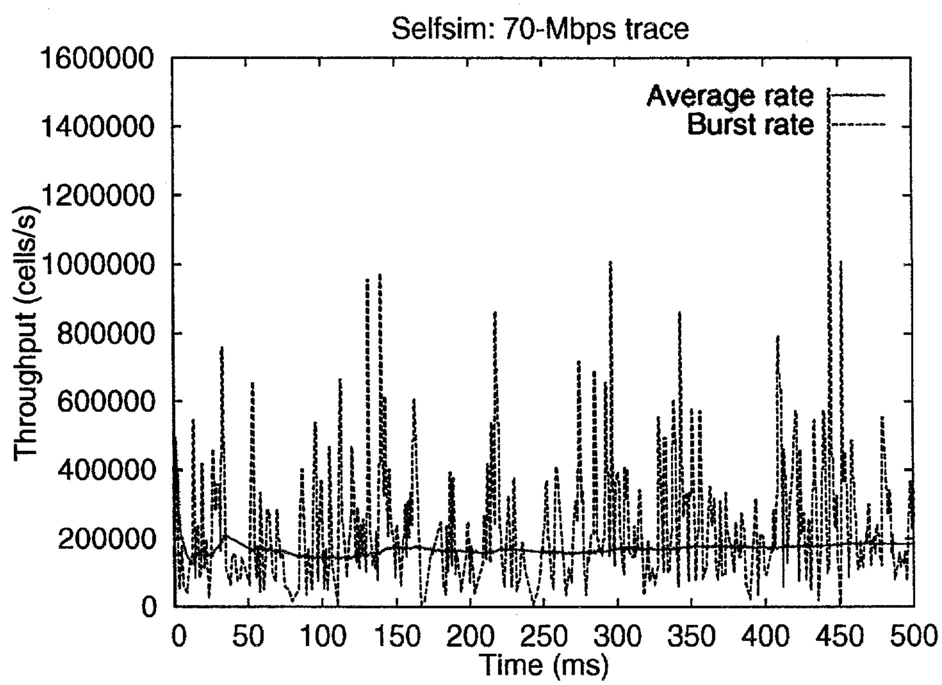

(b)

Figure 1. Self-similar traffic soutces with average burst rate of: (a) 10 Mbps and (b) 70 Mbps

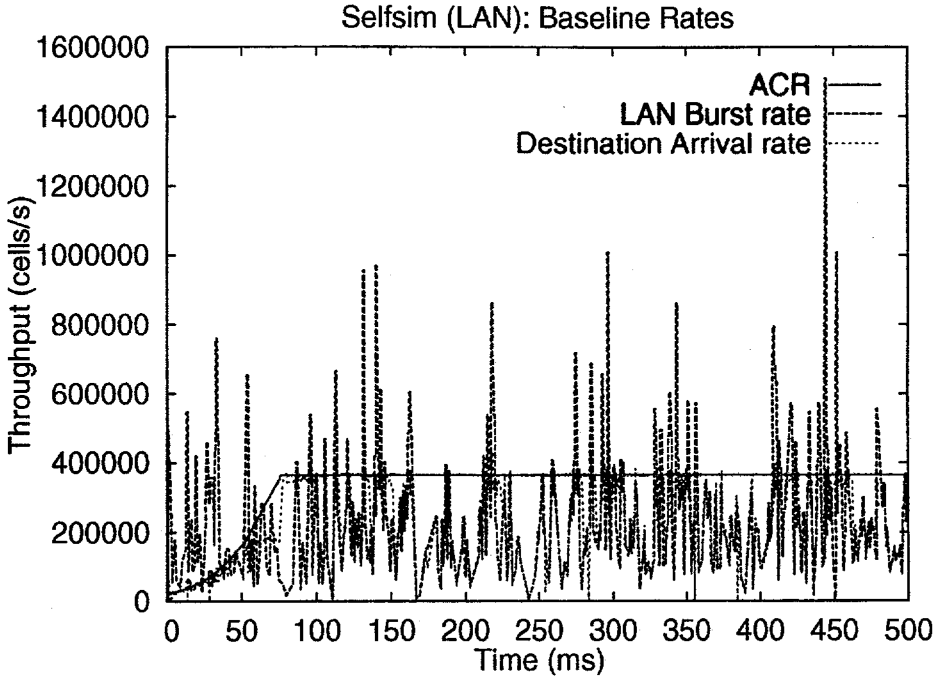

(a)

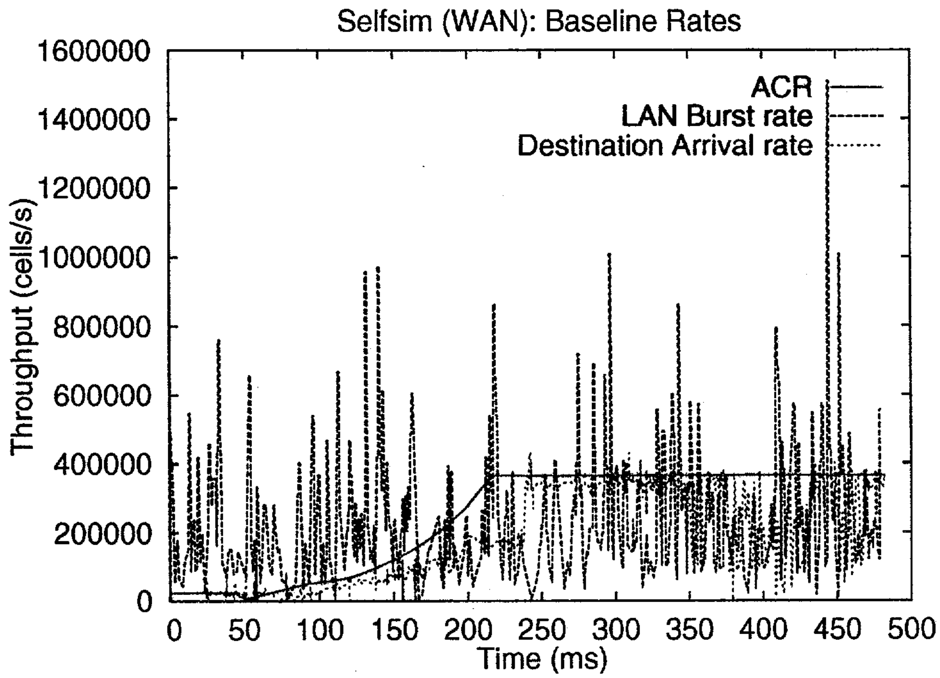

(b)

Figure 4. Baseline source rate, ACR, ATM egress rate for: (a) LAN and (b) WAN with 70-Mbps source

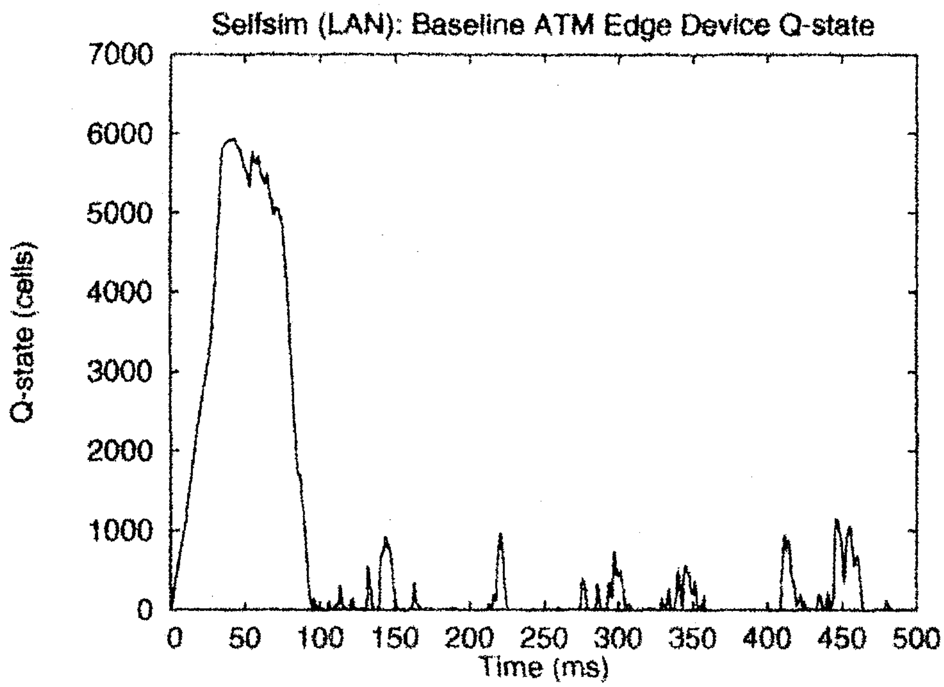

(a)

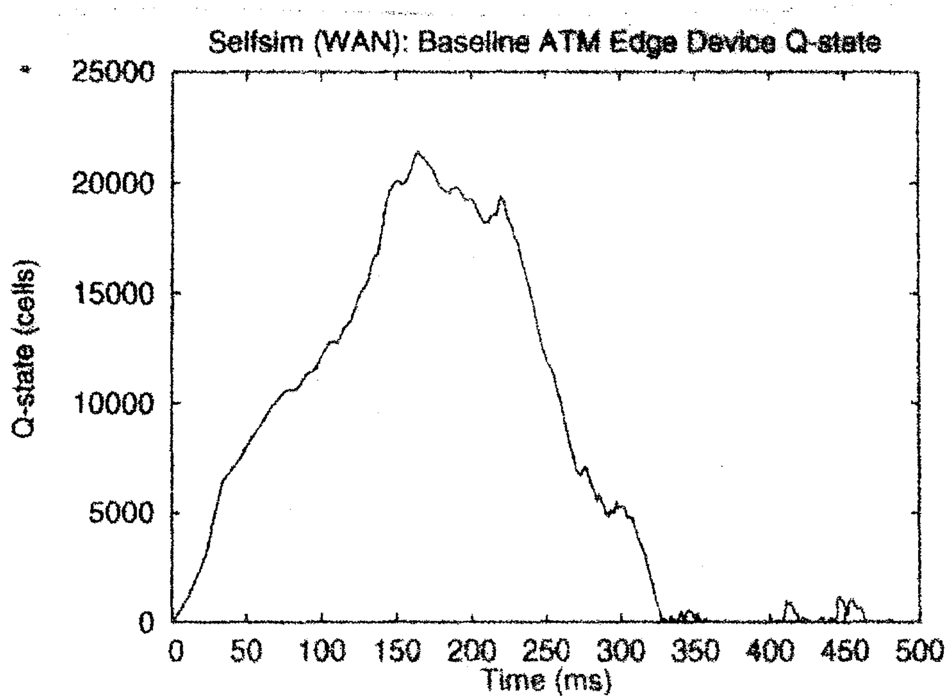

(b)

Figure 5. Baseline edge device queue state for: (a) LAN and (b) WAN with 70-Mbps source

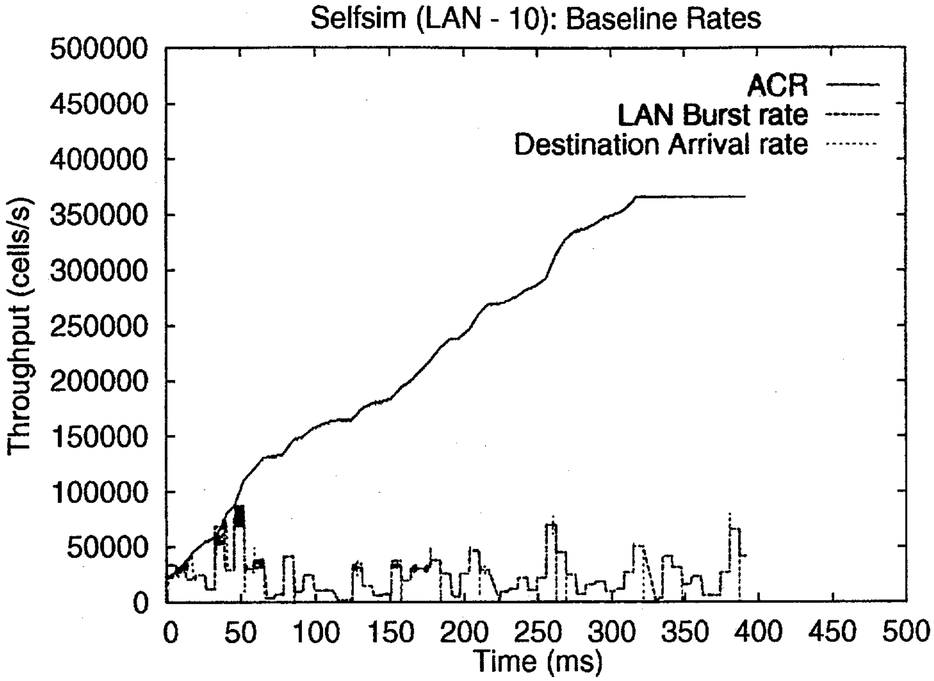

(a)

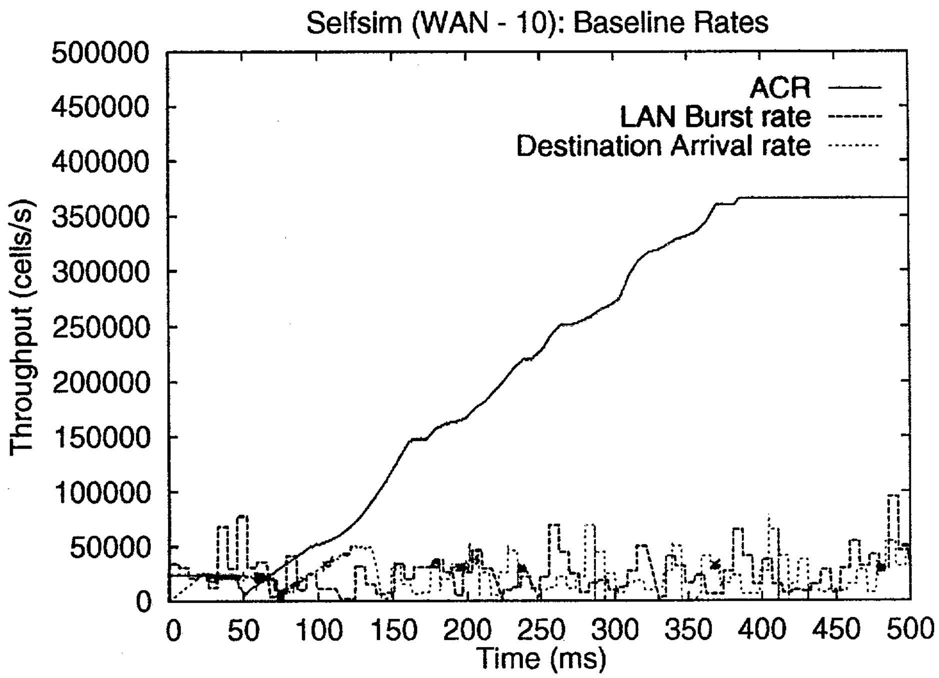

(b)

Figure 6. Baseline source rate, ACR, ATM egress rate for: (a) LAN and ***(b) WAN with 10-Mbps source***

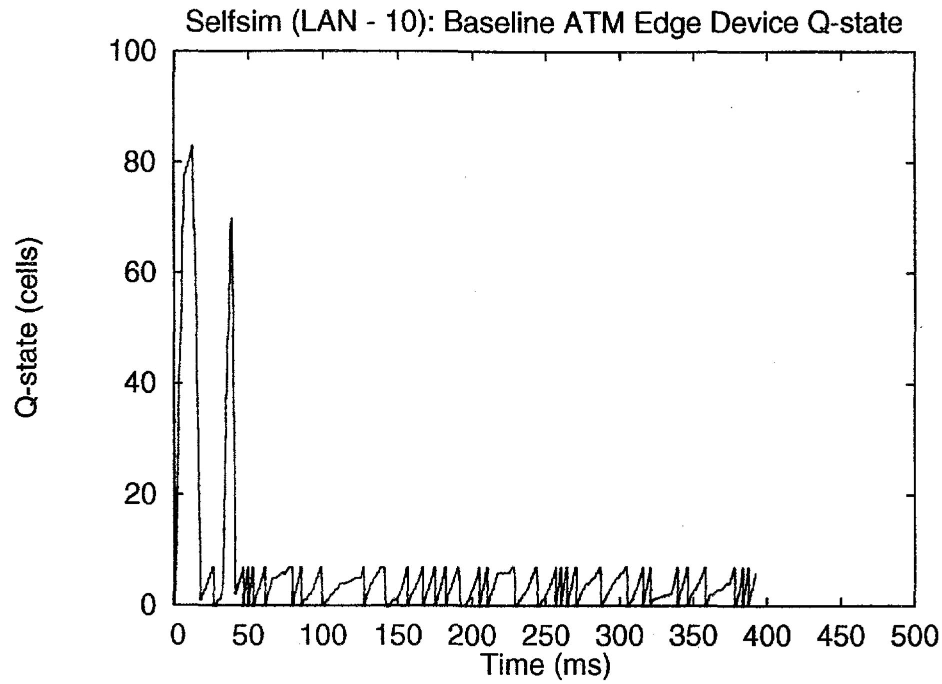

(a)

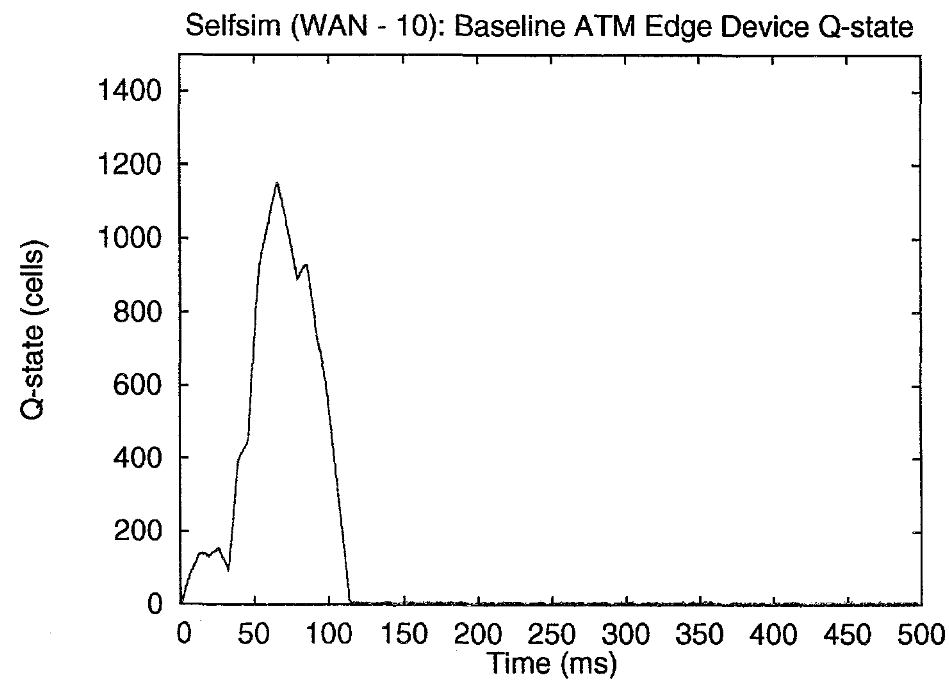

(b)

Figure 7. Baseline edge device queue state for: (a) LAN and (b) WAN with 10-Mbps source

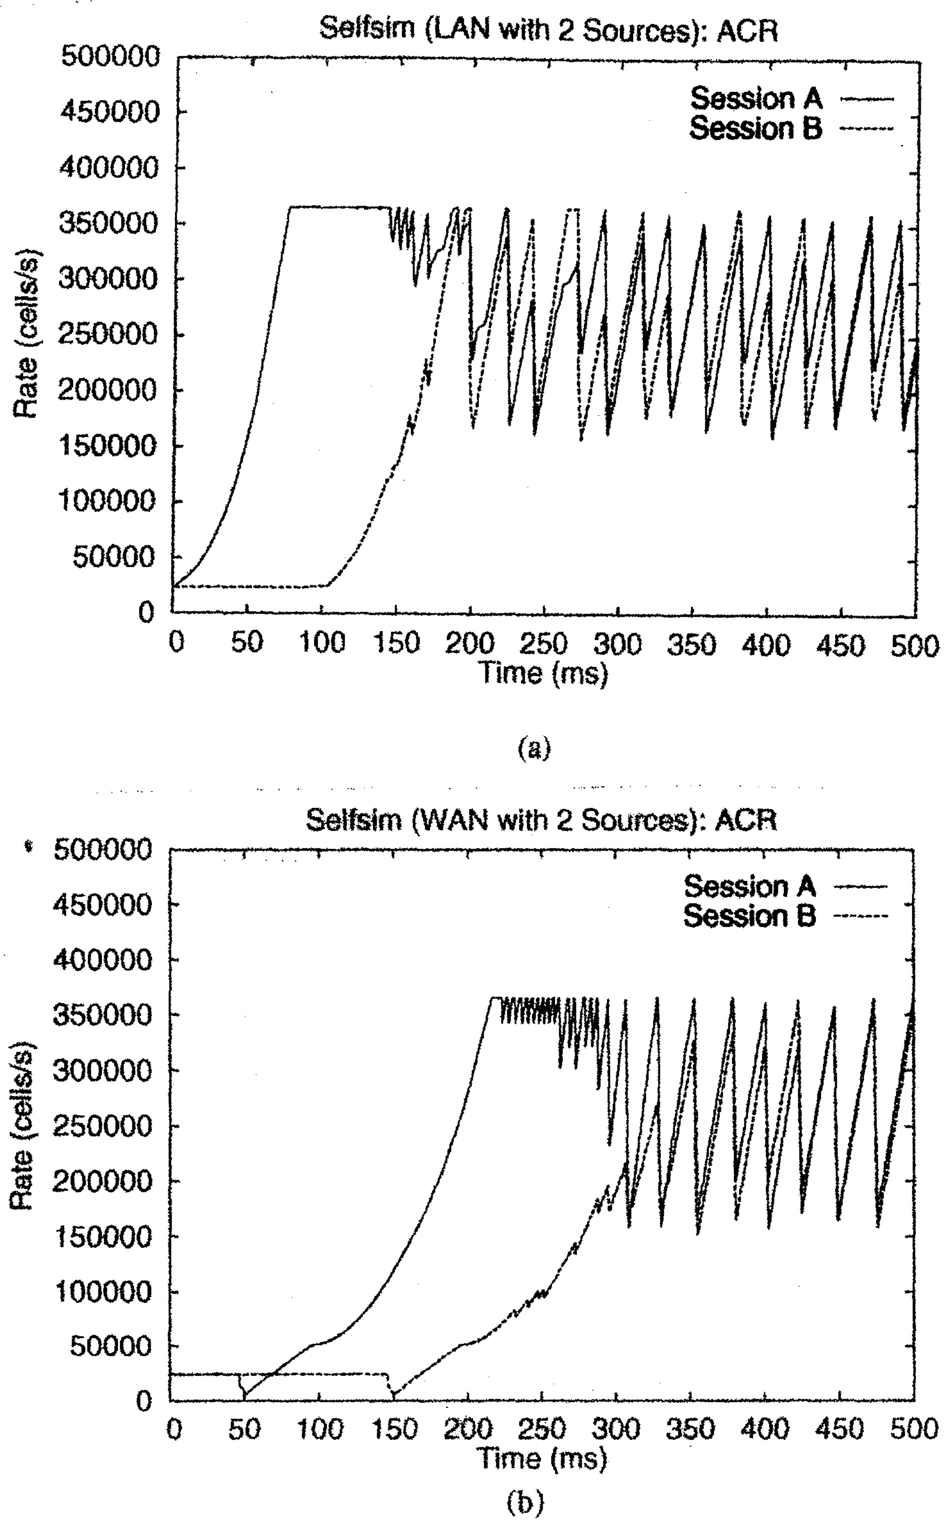

Figure 8. "Two 70-Mbps sources with staggered start" ACR: (a) LAN and (b) WAN

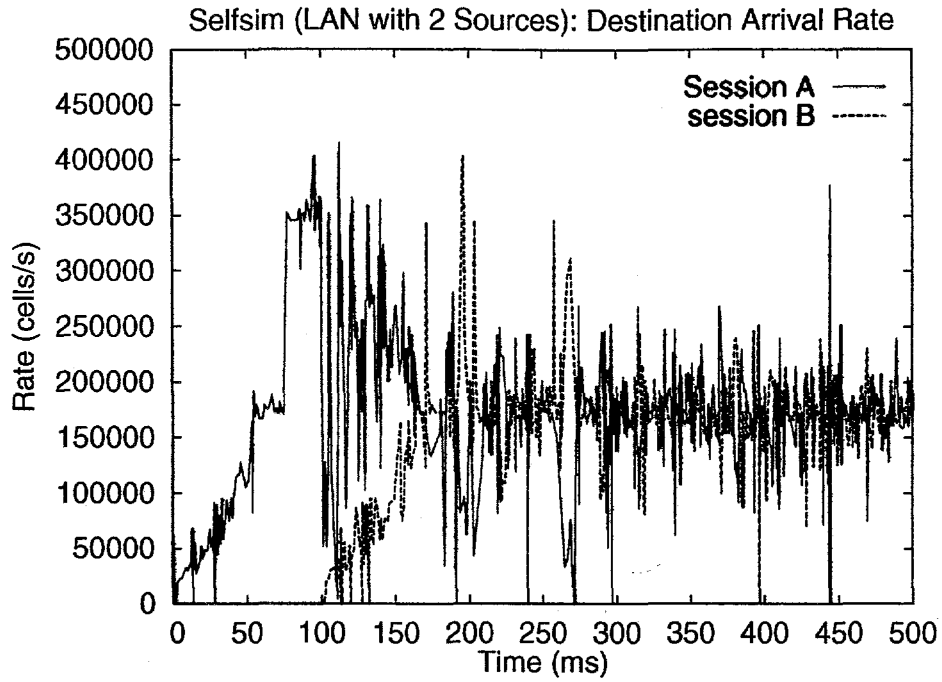

(a)

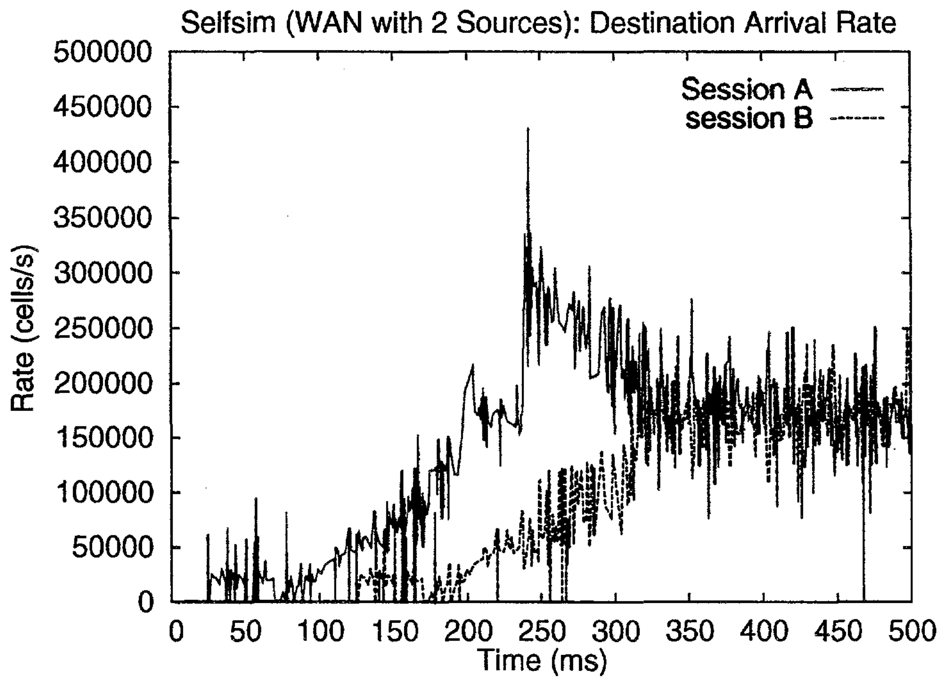

(b)

Figure 9 "Two 70-Mbps sources with staggered start" ATM egress rates: (a) LAN and (b) WAN

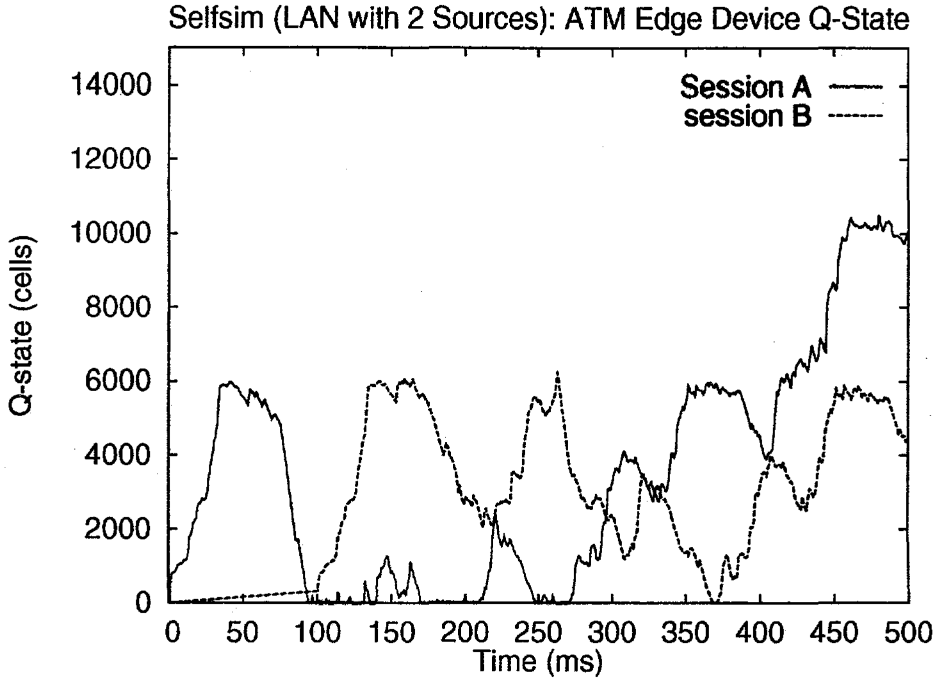

(a)

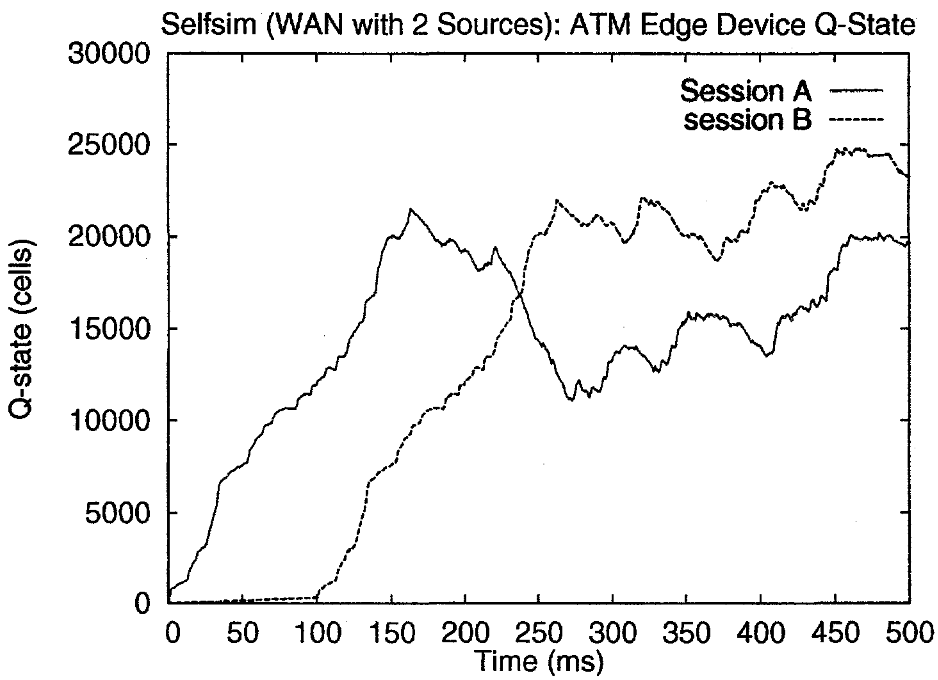

(b)

Figure 10 "Two 70-Mbps sources with staggered start" Edge device Q-state: (a) LAN and (b) WAN

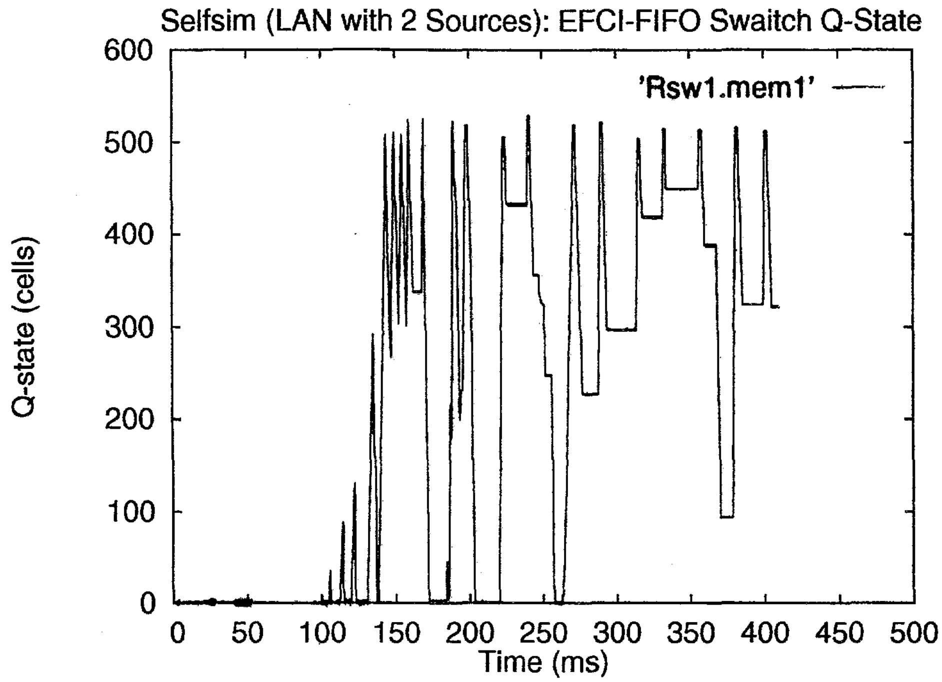

(a)

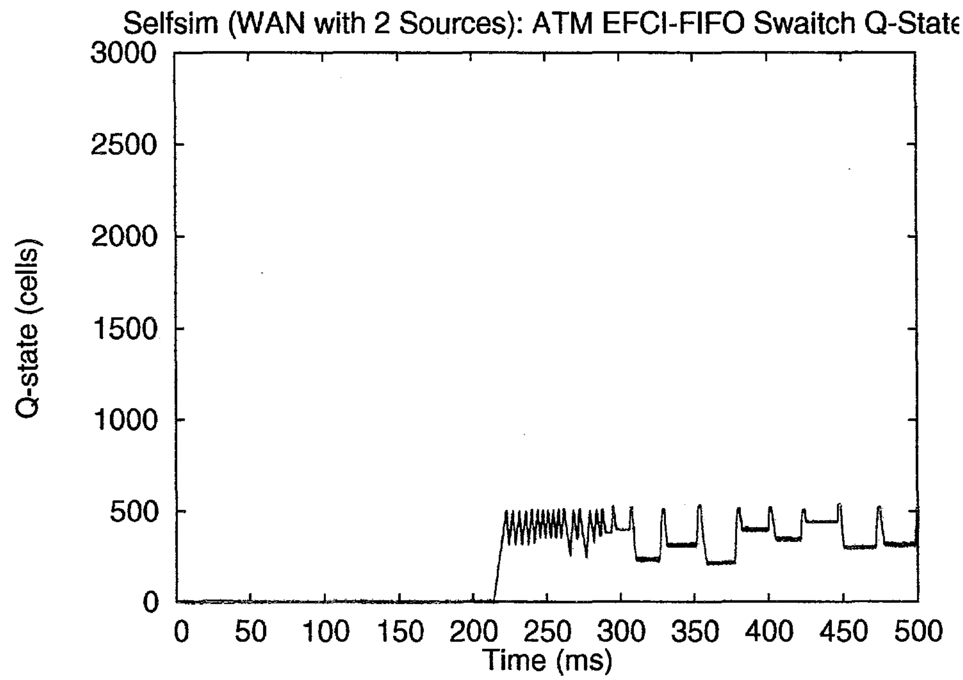

(b)

Figure 11 "two 70-Mbps sources with staggered start" Switch queue-state: (a) LAN and (b) WAN

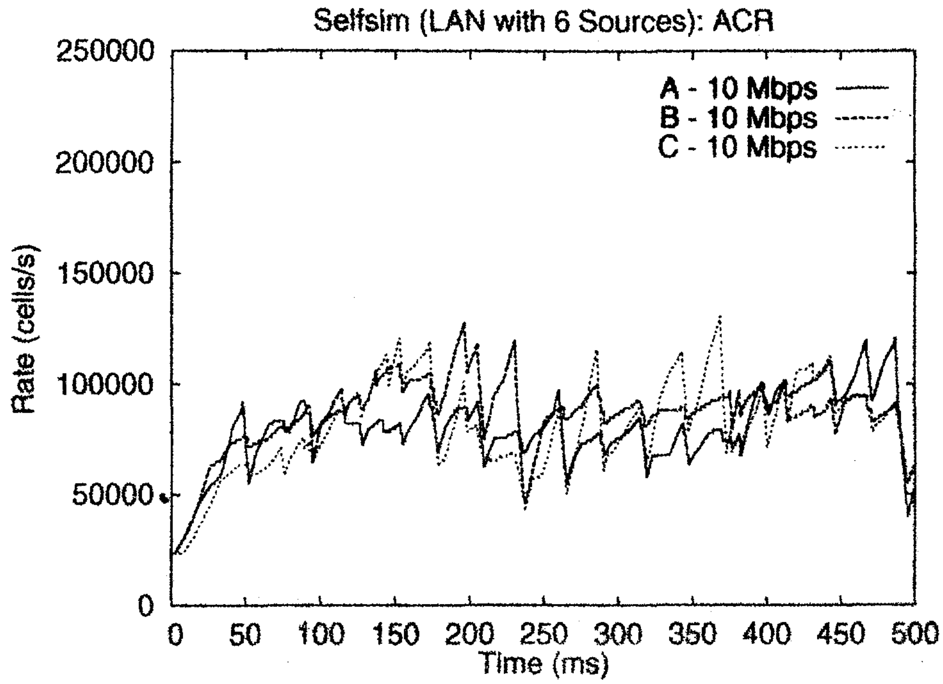

(a)

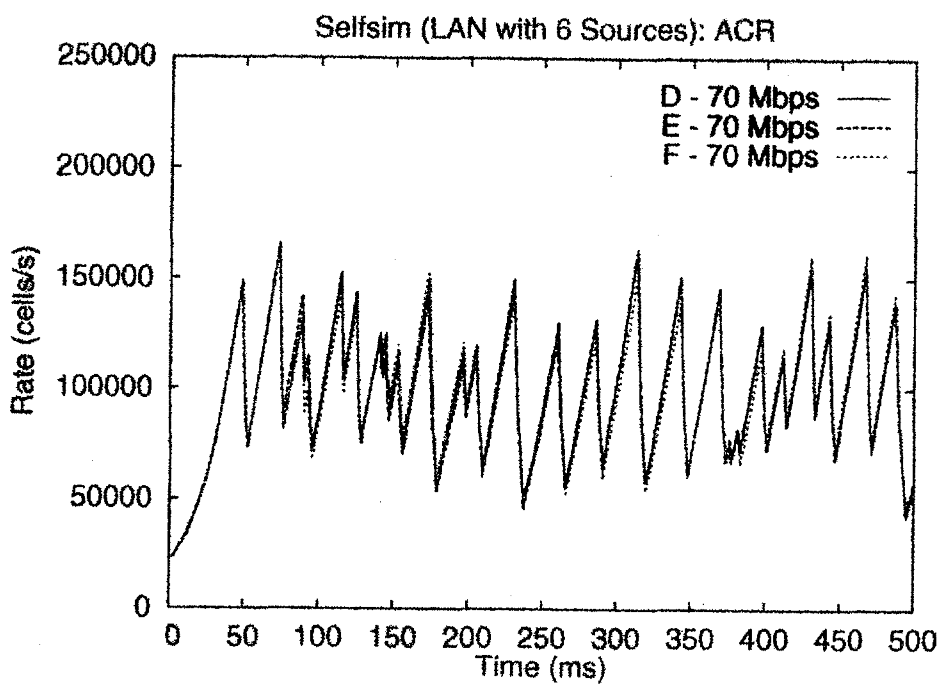

(b)

Figure 12 LAN "Six-session Congestion Study" ACR: (a) 10-Mbps sources and (b) 70-Mbps sources

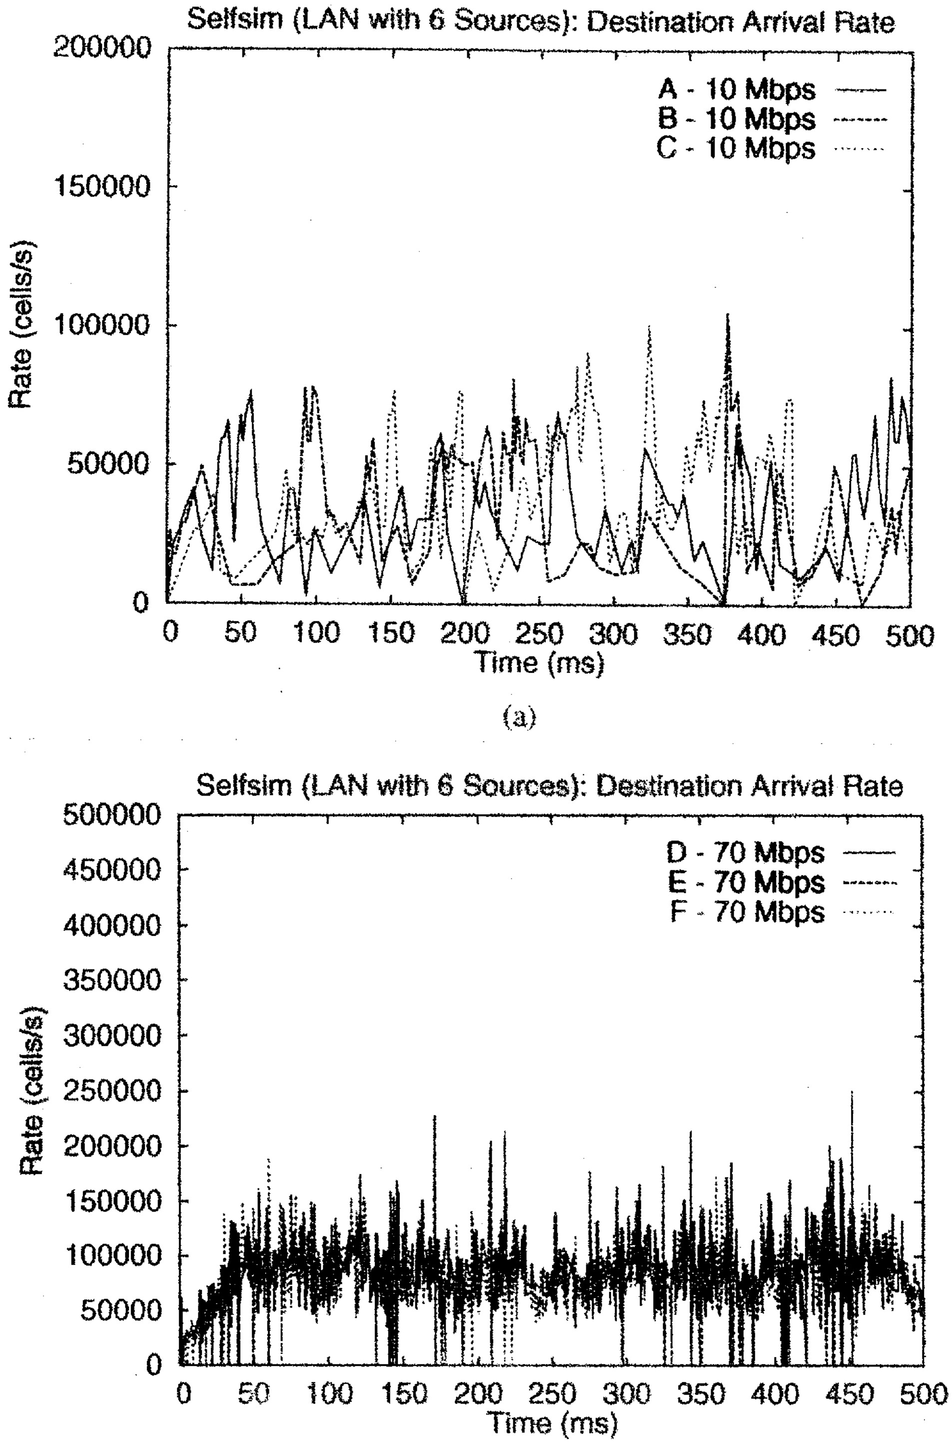

Figure 13 LAN "Six-session Congestion Study" ATM egress rates: (a) 10-Mbps sources and (b) 70-Mbps sources

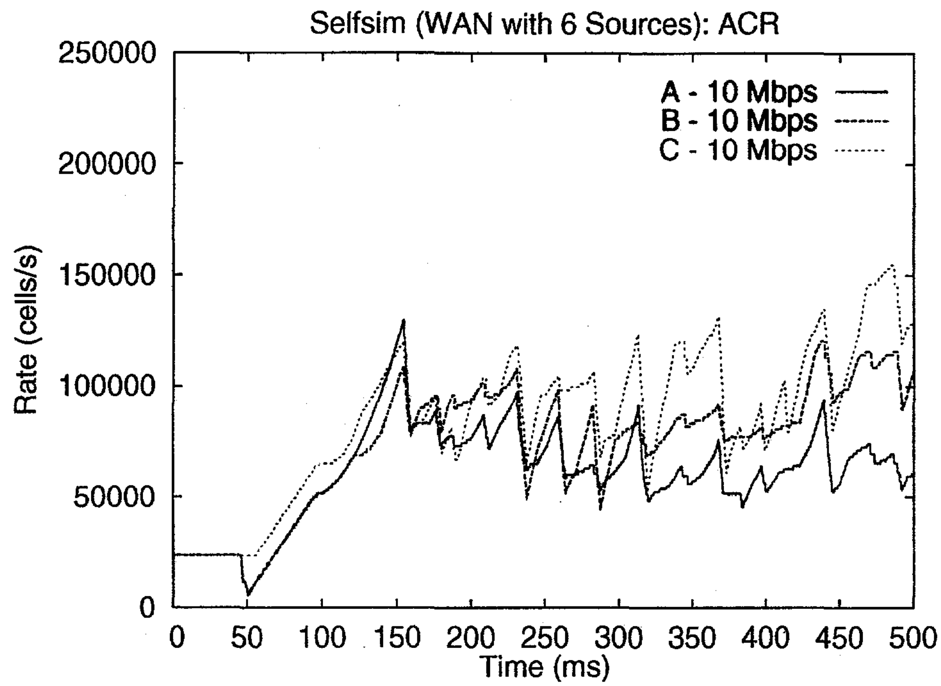

(a)

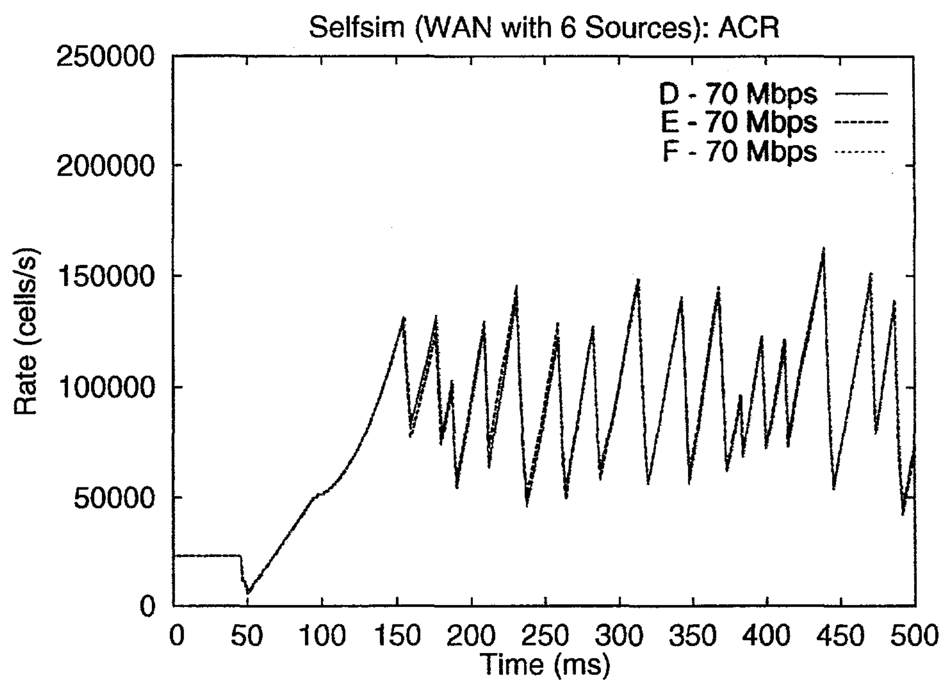

(b)

Figure 14 WAN "Six-session Congestion Study" ACR: (a) 10-Mbps sources and (b) 70-Mbps sources

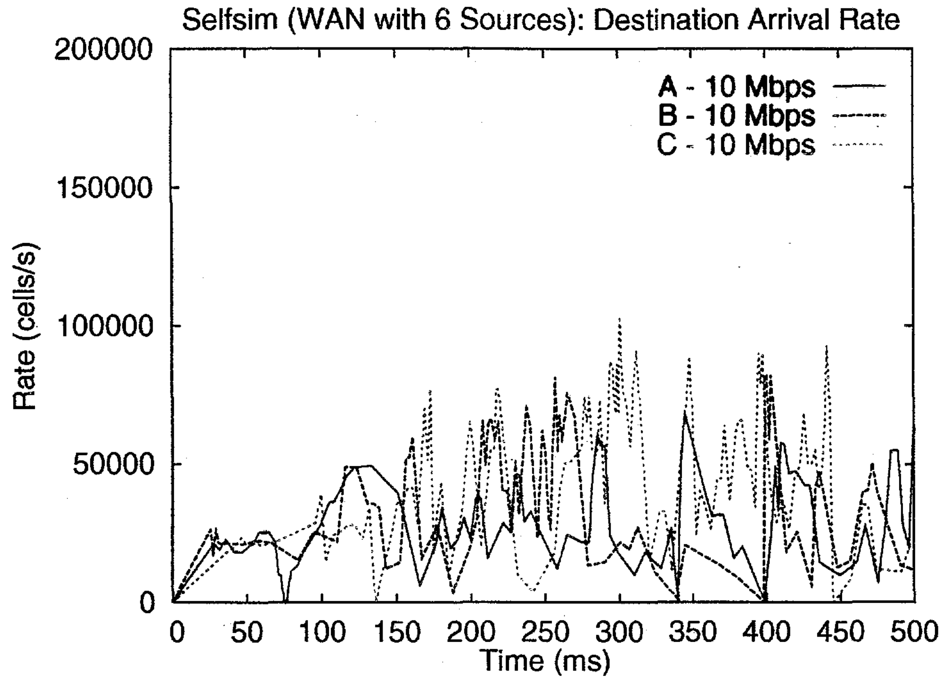

(a)

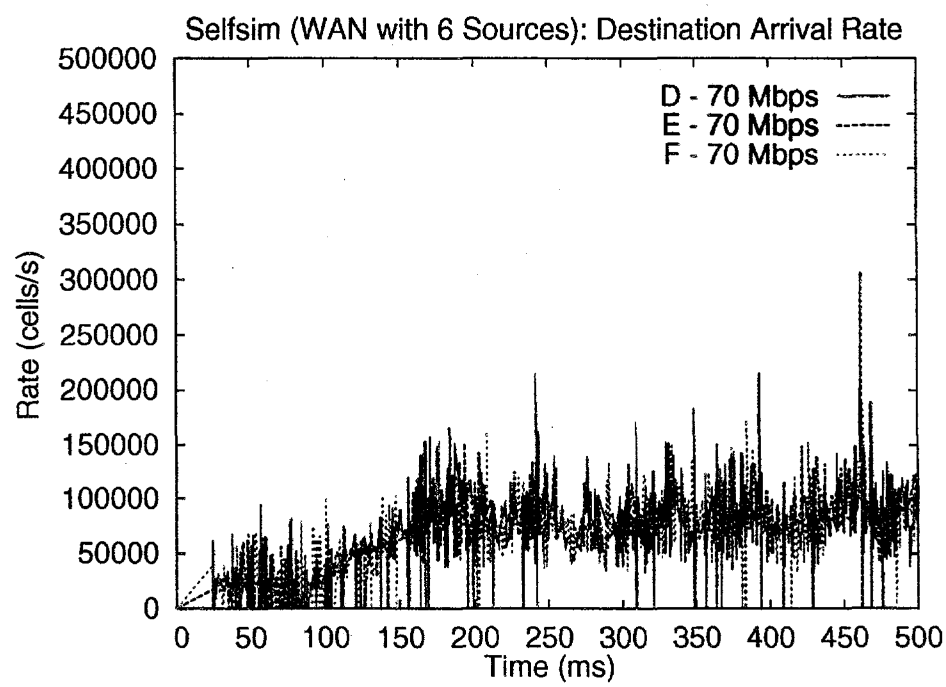

(b)

Figure 15 WAN "Six-session Congestion Study" ATM egress rates: (a) 10-Mbps sources and (b) 70-Mbps sources

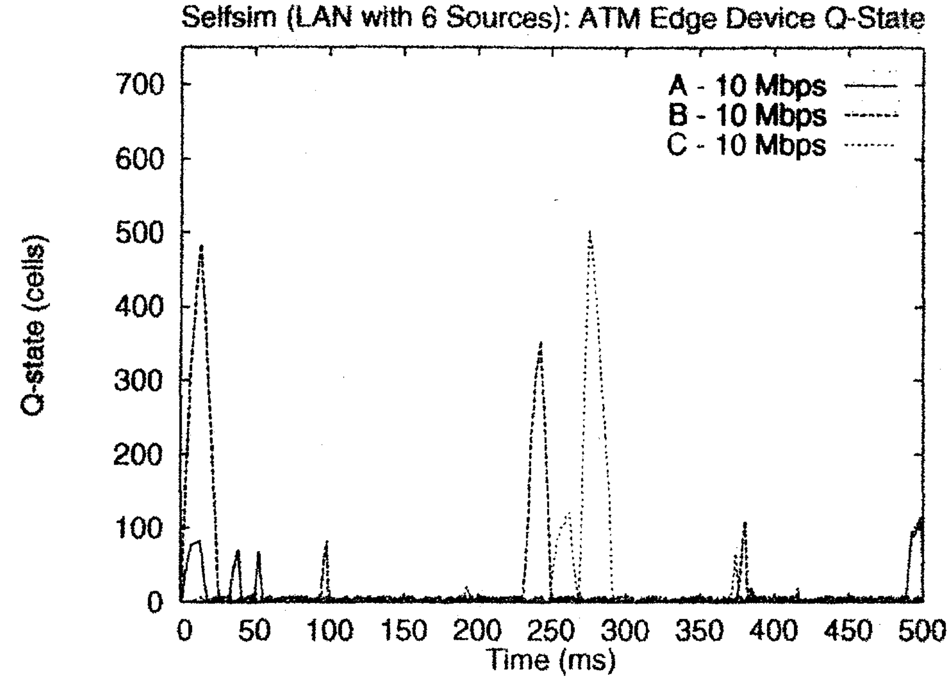

(a)

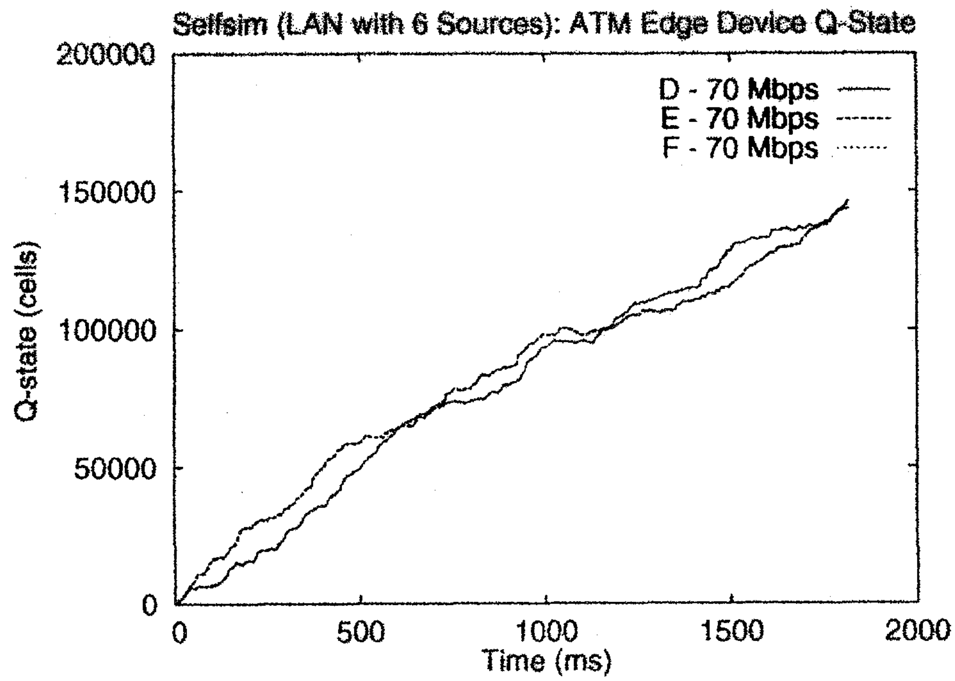

(b)

Figure 16 LAN "Six-session Congestion Study" edge device queue state: (a) 10-Mbps sources and (b) 70-Mbps sources

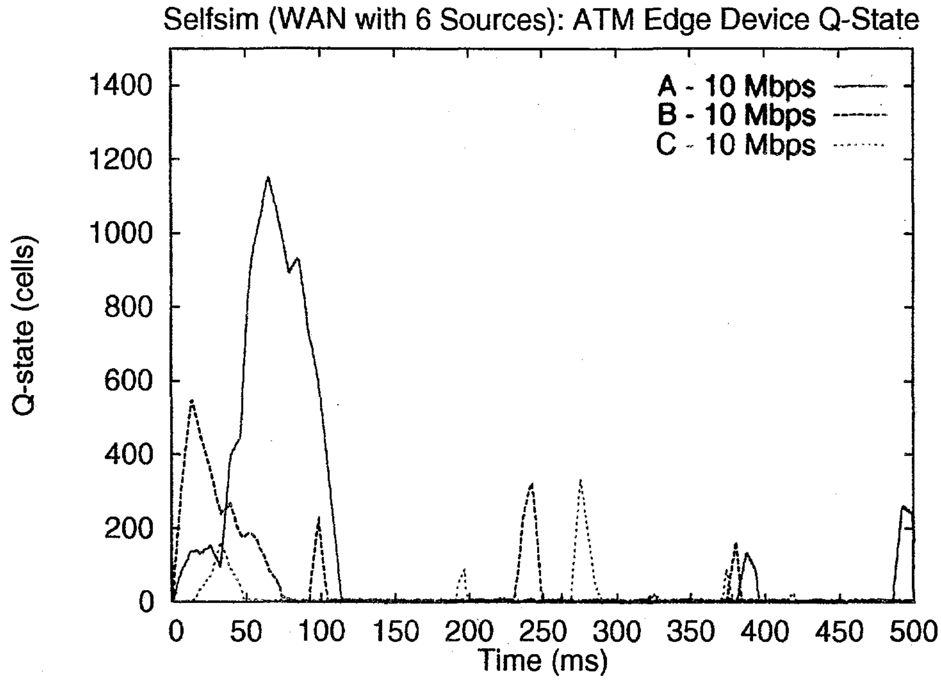

(a)

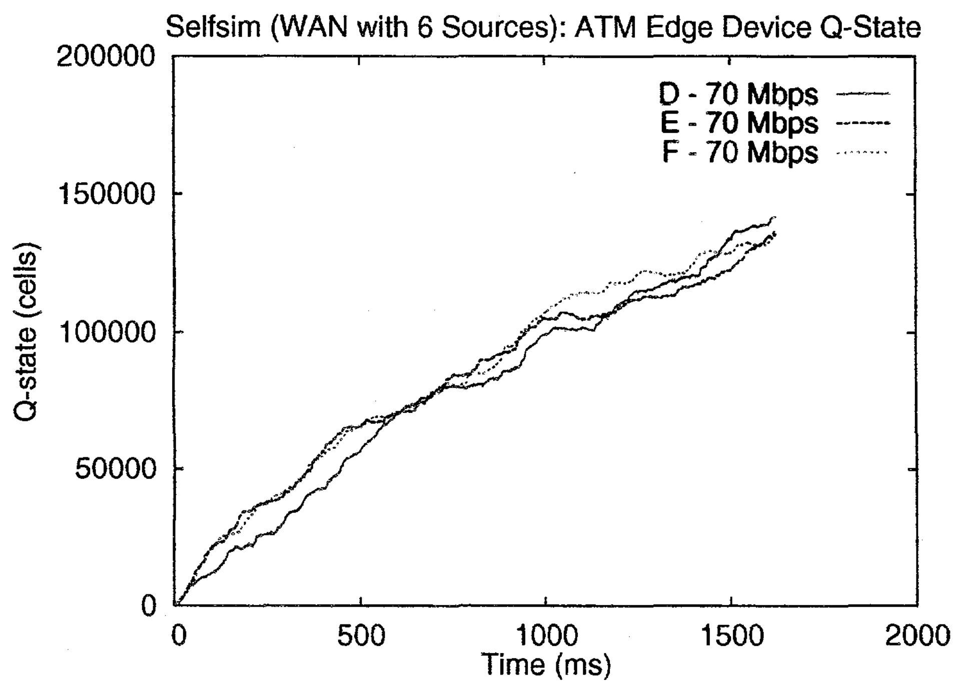

(b)

Figure 17 WAN "Six-session Congestion Stusy" edge device queue state: (a) 10-Mbps sources and (b) 70-Mbps

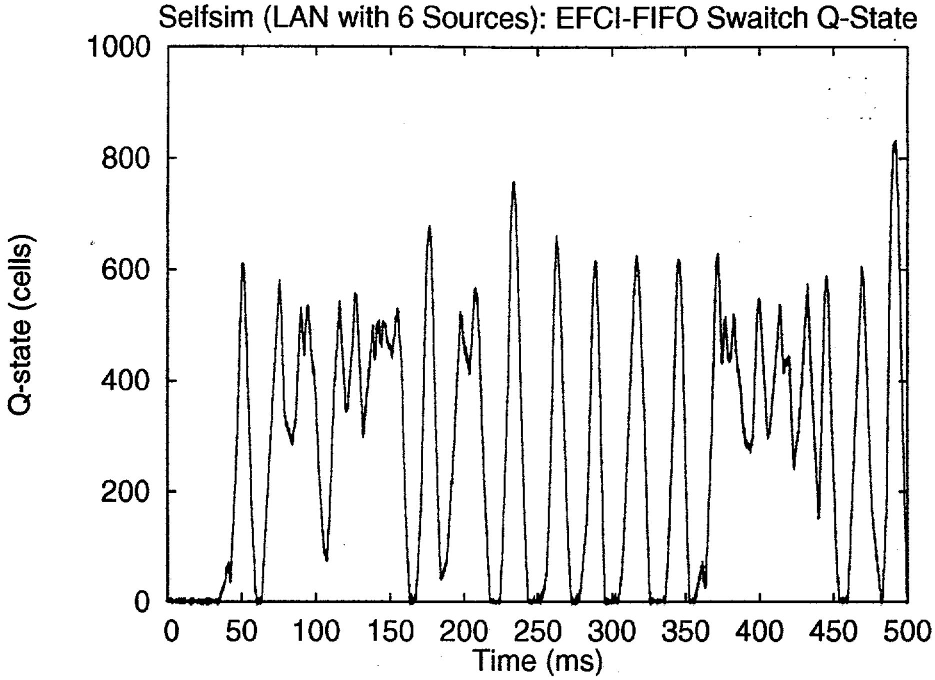

(a)

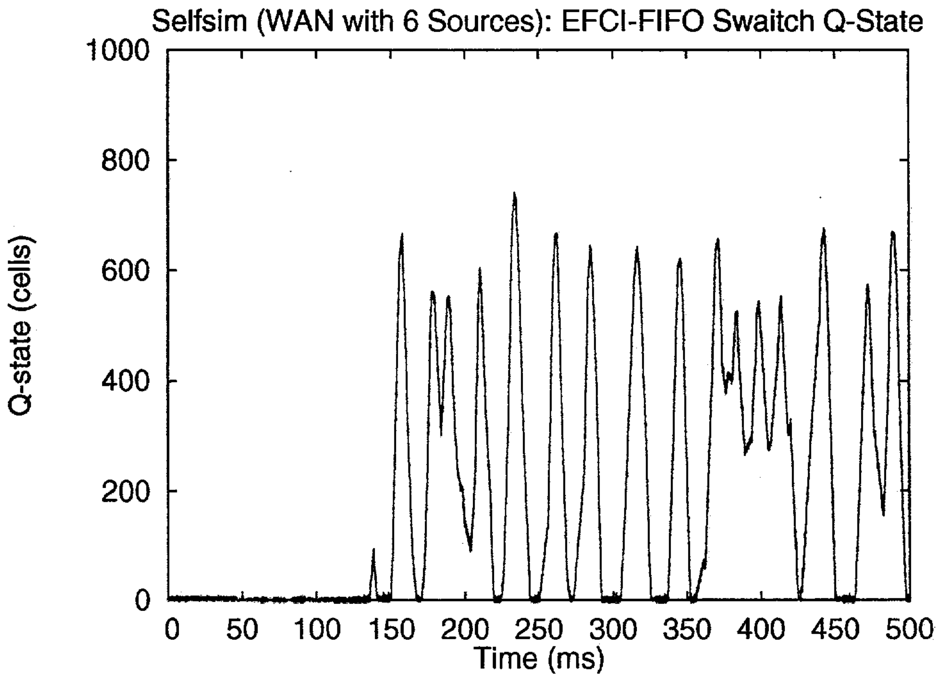

(b)

Figure 18 "Six-session Congestion Stusy" sitch queue state: (a) LAN and (b) WAN

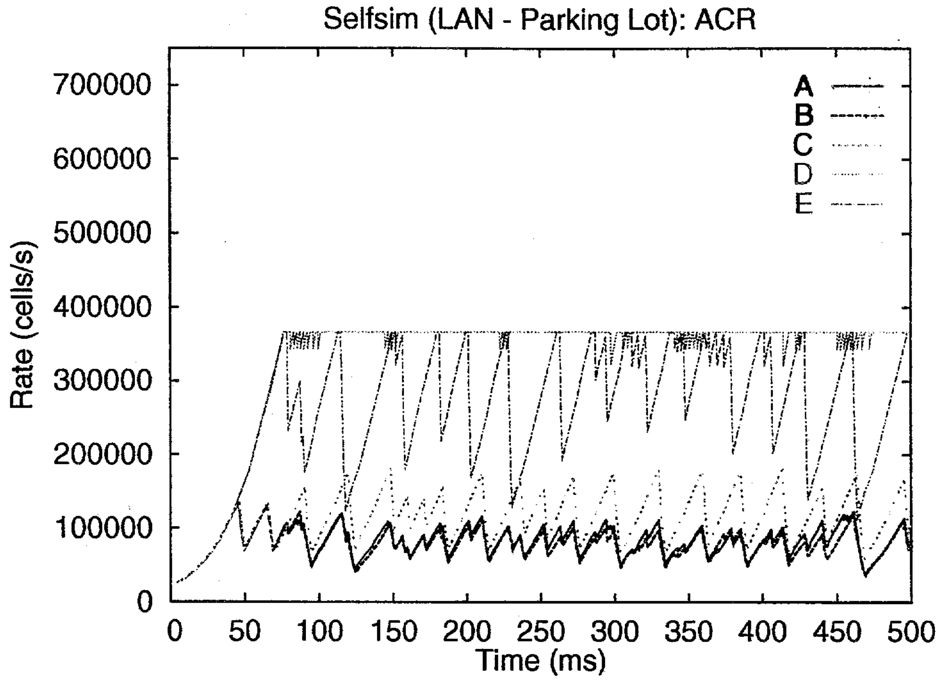

(a)

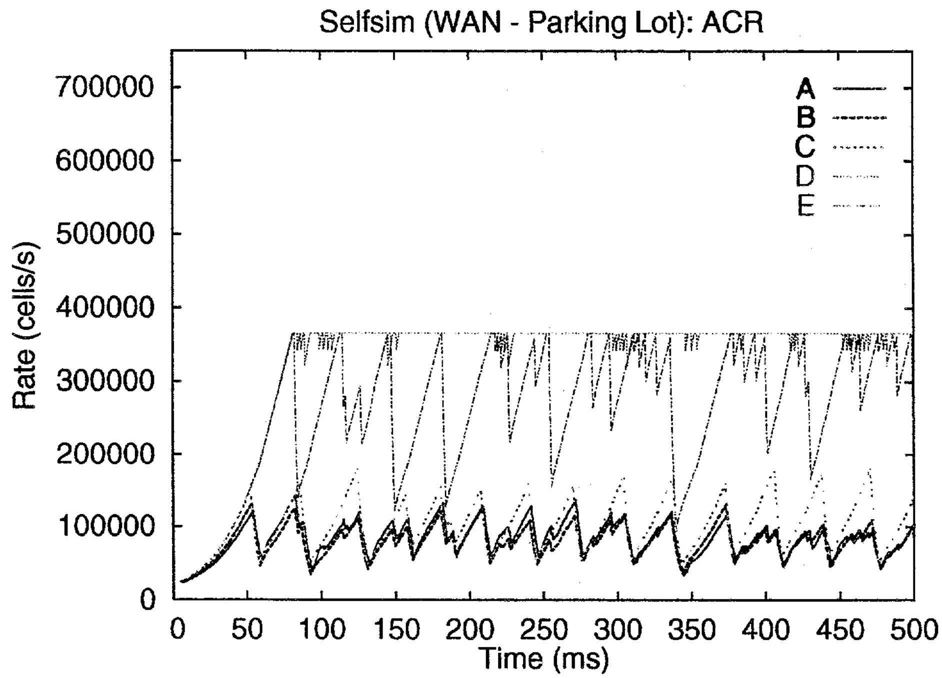

(b)

Figure 19 "MAX-MIN Fairness Study" ACR: (a) LAN and (b) WAN

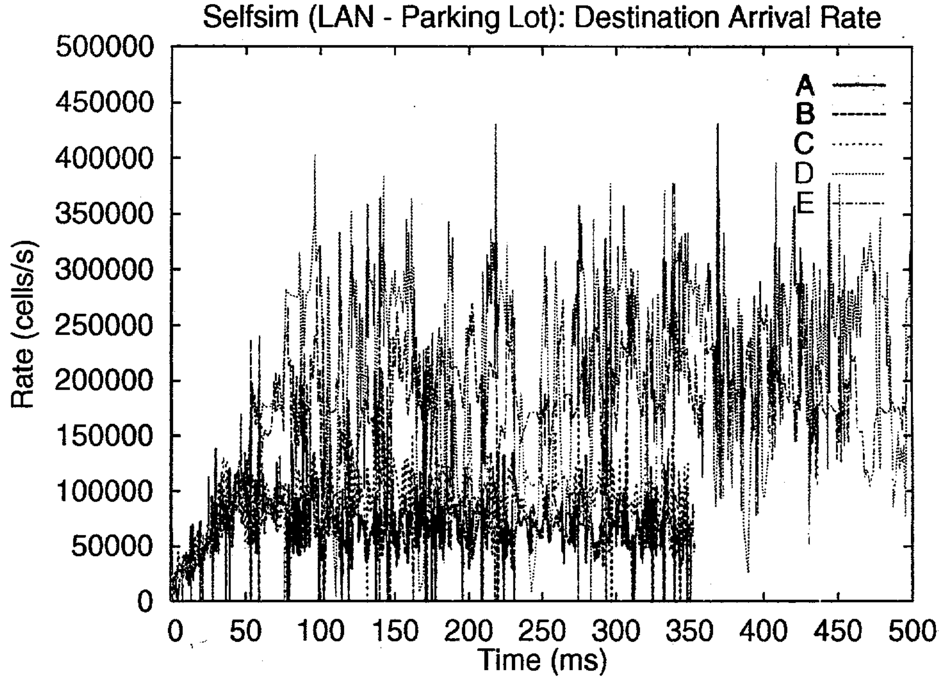

(a)

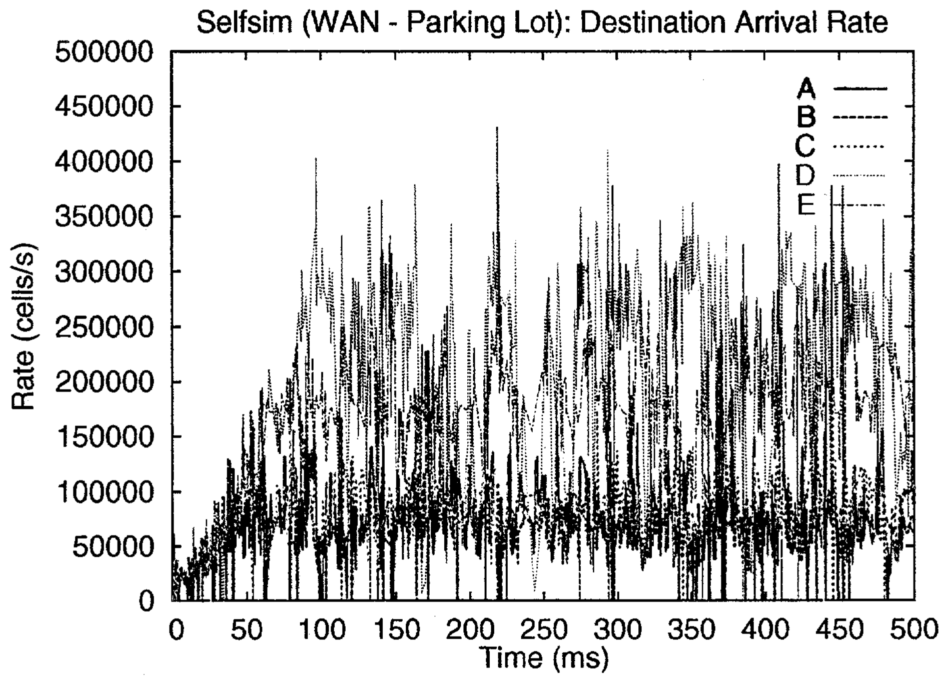

(b)

Figure 20 "The MAX-MIN Fairness Study" ATM egress rate: (a) LAN and (b) WAN

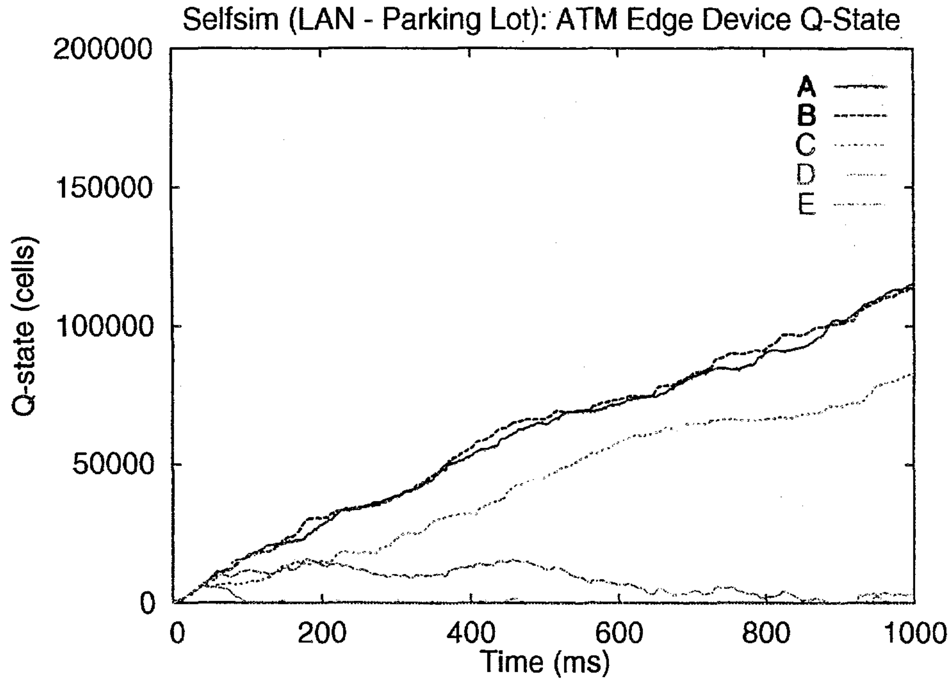

(a)

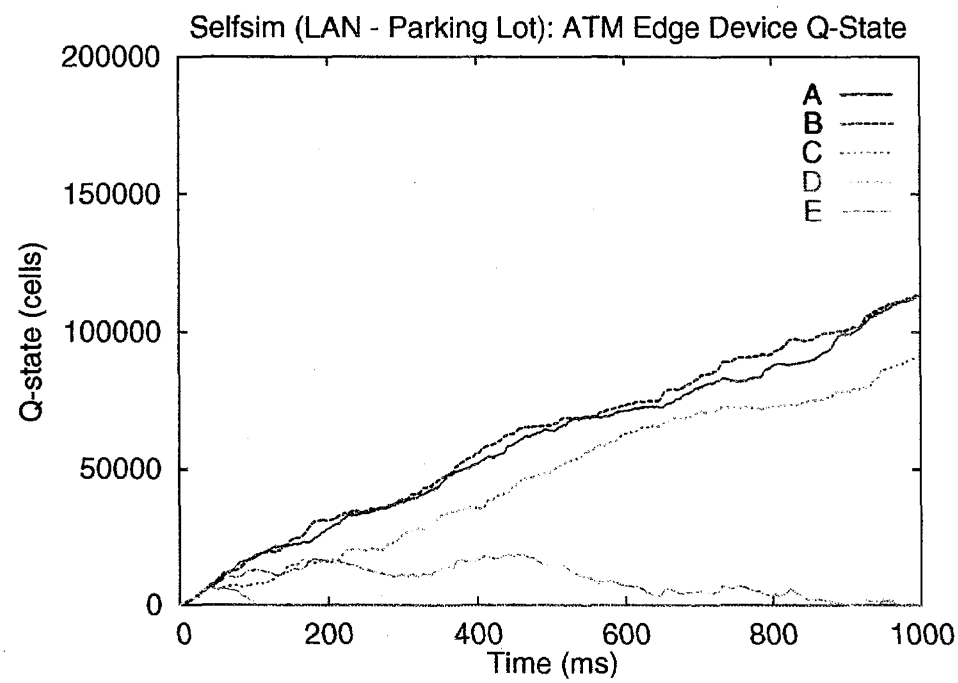

(b)

Figure 21 "The MAX-MIN Fairness Study" edge device queue state: (a) LAN and (b) WAN

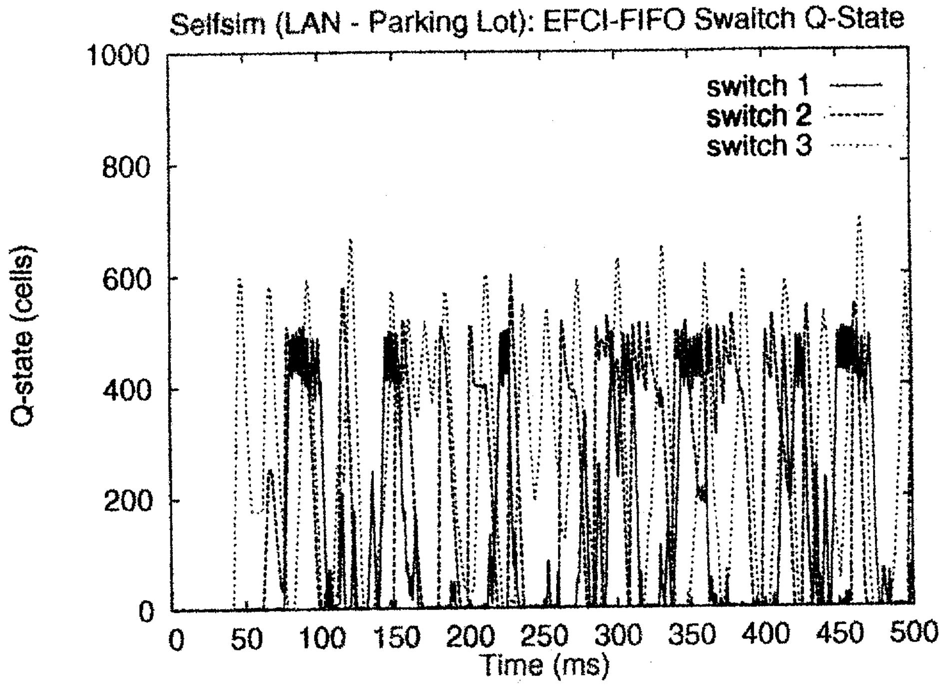

(a)

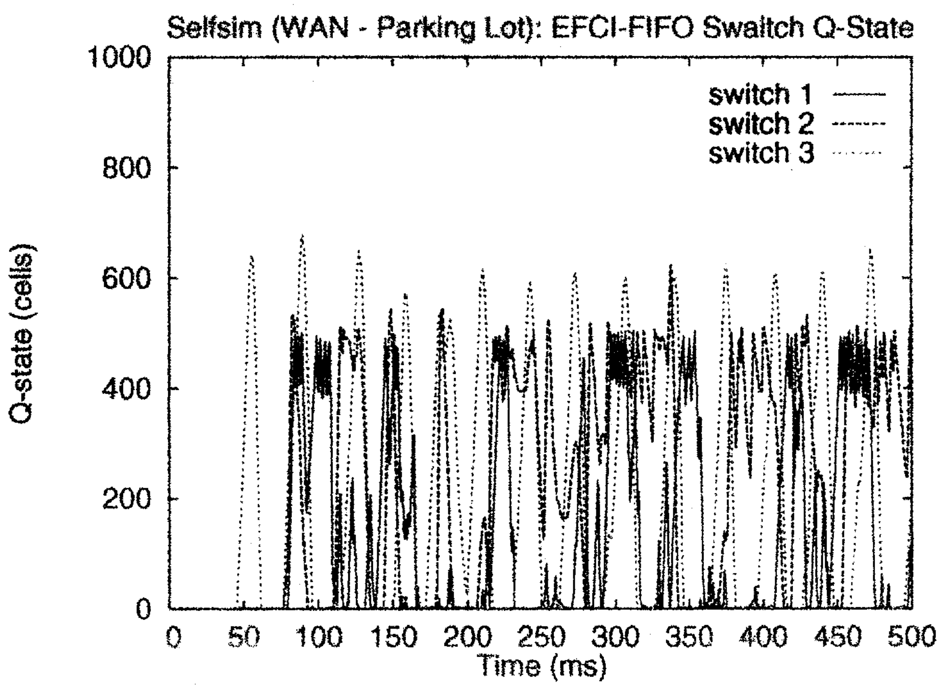

(b)

Figure 22 "The MAX-MIN Fairness Study" switch queue state: (a) LAN and (b) WAN

References

[1] W .E. Leland, M. S. Taqqu, W. Willinger, D. V. Wilson, On the Self-Similar Nature of Ethernet Traffic, Proc. SIGCOMM '93, September 1993

[2] S. Dastangoo, V. K. Samalam., *Simulation Results of Self-Similar Traffic over ATM ABR service*, AF/96-0026, February 5 1996

[3] C. Fang, H. Y. Chen, and J. A. Hutchins, *A Simulation Study of TCP Performance in ATM Networks*, Proc. IEEE GLOBECOM '94, pp. 1217-1223, November 1994.

[4] C. Fang, A. Lin, *On TCP Performance of UBR with EPD and UBR-EPD with a Fair Buffer Allocation Scheme*, AF/95-1328R, December 1995

[5] H. Y. Chen, Jim Brandt, *Performance Evaluation of Response Time in ATM LANs*, AF/95-1139, December 1995

[6] I. Norros, *On the Use of Fractional Brownian Motion in the Theory of Connectionless Networks*, IEEE Journal on Selected Areas in Communications, Vol. 13, No. 6, August 1995