



AD-A242 750



(2)

OFFICE OF NAVAL RESEARCH

FINAL REPORT

1 January, 1990 through 30 September, 1991

for

GRANT N00014-90-J-1289

Management, Control and Performance of Broadband
Integrated Networks

Principal Investigators:
Prof. Mischa Schwartz
Prof. Thomas E. Stern
Prof. Aurel Lazar

Department of Electrical Engineering
Columbia University
New York, NY 10027

Reproduction in whole, or in part, is permitted for any purpose of the United States
Government.

This document has been approved for public release and sale; its distribution is unlimited.

DEFENSE TECHNICAL INFORMATION CENTER



9115918

ONR Final Report

This report consists of a research summary (Part I) followed by an index of publications (Part II). The research summary covers the following areas of activity: Network Management and Control, and Asynchronous Network Performance Evaluation and Design. References cited in the research summary are keyed to the index of publications.

I. Research Summary

I.1 Network Management and Control

This area of research covered two topics, Access Control in Broadband Integrated Networks, with particular focus on video traffic streams; and fault management, extending previous work in this area.

Access control has become a contentious issue in the standards arena. It is clear that the introduction of high-speed networks makes possible the integration of services currently carried over separate networks. The vision of the next decade and beyond is that of providing users one access port, running at high bit rates, that allows the integration of a variety of traffic types. These include video and high-resolution images, in addition to the more common data and voice. Each of these services has a different quality of service requirement, in addition to widely differing bit rate (bandwidth) requirements.

There are two parts to the access control problem: First is admission control. A decision must be made whether or not to admit a given traffic stream - can the network guarantee to provide the desired quality of service without impacting adversely traffic streams already accessing the network? The second part is commonly called the "policing mechanism": once a decision has been made to allow various traffic types to access the network, how does the network control access to ensure any user does not violate its condition(s) of admission.

A variety of controls have been proposed in the standards bodies; there is currently no agreement on which to adopt. The basic problem is that there is a lack of understanding of the traffic characteristics of potentially new users of broadband networks. (Most important is the class of video traffic.) We therefore began a detailed study at the onset of this grant of appropriate means of characterizing video traffic, and, given such characterization, techniques for controlling the traffic. Our results thus far are summarized in Sec. I.1.1 following. Later in this report we summarize other work carried out that compares a number of traffic scheduling algorithms at an access port to ensure the quality of service of each traffic class is maintained.

In Sec. I.1.2 following, we summarize our work in fault management. This too has become a critical issue in managing both current networks and future high-speed networks. As more heterogeneous networks are interconnected, and the operation of the networks becomes more complex, means must be found for rapidly and automatically detecting, identifying, and correcting faults. Current fault management systems are ad hoc in nature, each keyed to a particular subsystem, system, or network. Vendors differ in their approaches to tackling this problem. We therefore began a study some years back of unifying methods of identifying faults. We chose an abstract finite state machine (FSM) model of a system under surveillance for a fault occurring, because it is commonly used in representing communication protocols. This model thus allows many existing systems to be so represented.

Accession For	
NTIS	UNCLASS
DTIC	TAB
Unrestricted	
Justification	
By	
Distribution/	
Availability Codes	
Avail and/or	
at	
A-1	

I.1.1. Video Characterization and Control

As noted above, the control of broadband integrated service networks is a serious problem. Asynchronous Transfer mode (ATM) has been selected as the international standard for carrying disparate traffic in these networks. In ATM, traffic accessing a network at an access port is packetized into "cells" 53 bytes long (48 bytes of data plus 5 bytes of header information) and then asynchronously transmitted into the network.

An ATM buffer at the access port serves to multiplex the various traffic types and enqueue them before transmitting them, cell by cell, into the network. The buffer may drop cells if full; delay is incurred as well. Both delay and probability of loss must be managed so that a given traffic stream, once admitted to the network, receives its specified quality of service without impacting adversely other users of the network.

We focused in our work on this grant on video characterization and control because so little is known of video in the ATM environment. For this purpose we developed a histogram model of variable-rate video traffic. We found experimentally using a number of compressed video sequences, that the sequence of bits, frame by frame, (i.e., the video bit arrival rate) could be represented by an 8-bin histogram. Playing out the bits in a given video frame over the entire 30-msec frame enabled us to characterize the access buffer problem, for a given bit arrival rate, as an M/D/N queueing problem, with N the size of the access buffer in cells. For the actual video sequence, approximated by an 8-bin arrival rate histogram, we assumed that as the arrival rate changed from frame to frame, the buffer system would quickly reach steady state compared to the duration of the frame. This is because the number of cells transmitted during a frame is typically very large. This enabled us to assume that the buffer occupancy distribution could be found by first solving the M/D/N problem conditioned on a given arrival rate and then unconditioning by weighting the resultant solutions by the appropriate histogram probability. Comparison of analytical with simulation results validated these assumptions.

This histogram model and M/D/N approximation based on it thus enabled us to accurately capture the queueing behavior due to the video traffic, allowing us to predict results such as buffer occupancy distributions and cell loss rates. Given the bit rate histograms of several individual video sources, a histogram for the aggregate arrival process can be found and system behavior predicted, quantifying the gain from statistical multiplexing.

Using this video histogram model, we developed a new call admission control scheme and policing algorithm, and showed that they can be used to prevent congestion in the network as well as protect calls that adhere to parameters given at call setup. Specifically a video user desiring access to the network provides an 8-number "worst case" histogram. This "worst case" histogram is defined as the histogram of the portion of the sequence which causes the worst system performance. For example, if system performance is given by loss probability, the worst case histogram is the one measured over the period of the call that provides the highest cell loss rate. The histogram selected depends on a specified window of video frames over which measured. Call admission is based on a convolution of this histogram with those of calls already in progress.

The policing mechanism proposed keeps a running histogram of the same window of frames, sliding the window and updating the histogram each frame time. This measured histogram is used to calculate an update system performance (call loss rate). If the system performance agrees with that stipulated at call admission, the actual histogram is acceptable.

If the calculated call loss rate is higher than that predicted at call admission, the policing mechanism either drops some calls or marks them as being in violation of the contract.

The performance of this admission control and policing mechanism was evaluated by simulation for seven multiplexed video sources. It was found to work very effectively as the number of noncompliant sources increased: it consistently tagged those violating their contract, allowing the others through. The only problem was that if tagging alone were used, a percentage of the complaint calls was lost due to buffer overflow. Dropping tagged calls rectified this situation.

A paper describing this work has been accepted for presentation at the forthcoming INFOCOM'91 [1].

1.1.2 Fault Management

As noted above, we have been using finite state machine (FSM) models to develop a unified approach to fault management in complex, interconnected networks. In earlier work we used the concept of a reduced-state FSM as an observer, capable of detecting specified types of faults in minimum time. In the work carried out on this grant we focused on a very simple group of FSM observers (of generally two states each) capable of detecting almost all faults in the system under observation. (The exceptions generally are deadlock or livelock situations). We developed a number of alternative algorithms for decomposing any given FSM into a specified number of equivalent independent two-state machines, running in parallel, each of which is capable of detecting a subset of the possible faults in the original FSM model. The total number of distinct states in the two-state machines must be greater than or equal to the number in the original FSM.

Transitions give rise to signals transmitted to the parallel set of observers. A faulty transition changes the signal sequence, and is guaranteed to be caught by at least one of the observers, except, as noted, for the cases of livelock or deadlock. The use of independent parallel observers implies that the fault detection/identification system itself degrades gracefully in the event one of the observers itself fails.

A potential problem with the two-state machines operating in parallel is that of maintaining synchronism between them. A procedure was thus developed enabling operation of the two-state machines without requiring knowledge of the current state of the original FSM model under observation. This decomposition technique was applied successfully to a number of communication examples. These involved the IEEE 802. logical link control used in LANs, and the OSI class 4 transport protocol. Details of this procedure appear in a paper accepted for presentation at INFOCOM'91, as well as for publication in the IEEE Transactions on Communications.[2].

1.2 Asynchronous Network Performance Evaluation and Design

1.2.1 Performance Evaluation

The work on asynchronous network performance evaluation has focused on integrated networks (for example, ATM) carrying bursty, heterogenous traffic from large numbers of sources. Three aspects of the problem have been considered: (a) Performance analysis and control using Markov-modulated rate process (MMRP) models, (b) Analysis and design of fast packet switches, and (c) Performance of full motion coded video sources in asynchronous networks.

a) In the work on MMRP models we have concentrated on a new approach to modeling using "separable" models. These are Markov-modulated rate models which are appropriate for accurately representing large classes of networking and statistical multiplexing problems, and which at the same time are computationally and analytically tractable. The fundamental property of separable models is that in systems involving the interaction of a large number of independent subsystems (e.g. many multiplexed sources) the computational complexity of the analysis problem is determined essentially by the size of the subsystems rather than the global system. This circumvents the "state explosion" problem common to most other modeling and analysis approaches. The work has covered a broad range of topics, including the fundamental separability properties in fluid models and their application to channel sharing systems [5], and the extension of the fluid model work to point processes [4,15]. This work is now being applied to flow control problems in asynchronous networks using some generalizations of the "leaky bucket" algorithm.

b) The work on fast packet switches is now complete. Its objective was to provide a methodology for determining performance of generic fast packet switches under complex loading conditions (including correlated sources and unbalanced traffic), and realistic operating constraints (including priority queueing disciplines and finite buffer capacities). Published results emanating from this work included an accurate means of analyzing output queueing in fast packet switches [9], optimal buffer allocation for switches with input and output queueing [11], and effects on non-uniform traffic on throughput [11, 14, 18].

c) The work on full motion video deals with how the video codec interacts with the characteristics of the asynchronous network that carries the coded traffic. This is the logical extension of earlier work we did on packet voice [3,6,8]. We are interested in how loss and delay in the asynchronous network will affect the quality of the received information, and how special coding procedures can be used to optimize performance in the face of these network impairments. Using the MMRP approach described above we have succeeded in constructing some fairly good models for two kinds of coded video sources: sub-band coded and hybrid motion compensation DCT coded video. The accuracy of these models has been validated through simulation based on empirical data from actual video scenes. Work currently underway involves using these models to predict the degradation effects due to limited channel capacity, queueing delays, and packet loss in ATM systems serving many multiplexed video sources.

1.2.2 Lightwave Architectures Supporting Broadband Asynchronous Networks

The broadband asynchronous networks of the future will ride on an underlying physical layer made up of optical fibers; i.e. a lightwave network. It is the opinion of most people (including ourselves) that high performance lightwave network architectures supporting broadband services must be radically different in structure from the networks currently deployed; that is, we cannot achieve high performance simply by replacing copper by fibers, as is currently being done. The work we have recently undertaken deals with a new type of architecture: the Linear Lightwave Network (LLN) [10,13] which is capable of providing enormous capacity to large groups of users with widely differing requirements, and spread over large geographical areas. These networks perform only linear operations on the optical signals, including controllable power combining, dividing and linear amplification. Being controllable, LLN's can respond dynamically to changing demand and to failures. By creating an asynchronous or packet switched "logical network" overlay on the LLN, the power and flexibility of the LLN architecture can be exploited to create a flexible, adaptable and reliable physical foundation for the logical network it supports. We have done considerable research on many issues in LLN's including routing [16,17,21], topological design [19,20], multicasting [23] and cross-talk caused by imperfect devices

[22]. Under other sponsorship, an experimental effort is underway to test out the concepts in a small scale prototype network.

II. Index of Publications

- [1] P. Skelly, S. Dixit, M. Schwartz, A Histogram-Based Model for Video Traffic Behavior in an ATM Network Node with an Application to Congestion Control, to be presented, INFOCOM'92 Florence, Italy, May 1992.
- [2] C. Wang and M. Schwartz, Fault Detection with Multiple Observers, to be presented, INFOCOM'91, Florence, Italy, May 1992; accepted for publication, IEEE Trans. on Communications.
- [3] S. Ganguly and T.E. Stern, Performance Evaluation of a Packet Voice System, IEEE Trans. on Comm., Jan. 1990.
- [4] A. Elwalid, D. Mitra, and T.E. Stern, A Theory of Statistical Multiplexing of Markovian Sources: Spectral Expansions and Algorithms, Proc. First International Workshop on the Numerical Solution of Markov Chains, N.C. State University, Raleigh, NC, Jan. 1990.
- [5] T.E. Stern and A. Elwalid, Analysis of Separable Markov-Modulated Rate Processes, Advances in Applied Probability, March 1991.
- [6] N. Yin and S.Q. Li, and T.E. Stern, Congestion Control for Packet Voice by Selective Packet Discarding, IEEE Trans. on Comm., May 1990.
- [7] T.E. Stern, A Linear Lightwave MAN Architecture, Proceedings of the NATO Advanced Research Workshop on Architecture and Performance Issues of High-Capacity Local and Metropolitan Area Networks, Sophia-Antipolis, France, June 1990. G. Pujolle (ed.), Springer, Berlin.
- [8] T. Brunner and T.E. Stern, Performance Analysis of a Packet Voice Messaging System, Proc. INFOCOM '90 San Francisco, June 1990.
- [9] J.S.-C. Chen, R. Guerin and T.E. Stern, Markov-Modulated Flow Model for the Output Queueing of a Packet Switch, Proc. IEEE INFOCOM '90, San Francisco, June 1990.
- [10] T.E. Stern, Linear Lightwave Networks: How Far Can They Go?, GLOBECOM '90, SAN DIEGO, DEC. 1990.
- [11] J.S.-C. Chen and T.E. Stern, Optimal Buffer Allocation for Packet Switches with Input and Output Queueing, GLOBECOM '90, San Diego, Dec. 1990.
- [12] J.S.-C. Chen and T.E. Stern, Throughput Analysis Optimal Buffer Allocation and Traffic Imbalance Study of a Generic Non-blocking Packet Switch, IEEE JSAC Special Issue on Teletraffic Analysis of Communication Systems. April 1991.
- [13] T.E. Stern, Linear Lightwave Networks, submitted to IEEE Trans. on Communications.

- [14] J.S.-C. Chen and T.E. Stern, Throughput Reduction due to Non-uniform Traffic in a Packet Switch with Input and Output Queueing, ICC'91.
- [15] D. Mitra, A.I. Elwalid and T.E. Stern, Statistical Multiplexing of Markov Modulated Sources: Theory and Computational Algorithms, Presented at the 13th International Teletraffic Conference, Copenhagen, June 1991.
- [16] Krishna Bala and T.E. Stern and Kavita Bala, Multiple Channel Routing in a Linear Lightwave Network, OFC '91, San Diego, CA, Feb. 1991.
- [17] Krishna Bala, T.E. Stern and Kavita Bala, Algorithms for Routing in a Linear Lightwave Network, IEEE INFOCOM '91, Miami, FL, April 1991.
- [18] J.S.-C. Chen and T.E. Stern, Performance of Generic Non-blocking Packet Switch, Proc. Fourth International Conference on Data Communication Systems and Their Performance, Barcelona, June 1990. Reprinted in Data Communication Systems and Their Performance, G. Pujolle and R. Puigjaner (eds.) North-Holland, Amsterdam, 1991.
- [19] K. Bala and T.E. Stern, Topologies for Linear Lightwave Networks, SPIE OE/Fibers '91, Boston, Sept. 1991.
- [20] T.E. Stern, S. Jiang and K. Bala, Graph-Theoretic Problems in Optical Network Control and Design, submitted for presentation at the 2nd ORSA Telecommunication Conference, Boca Raton, FL, Mar. 1992.
- [21] K. Bala, T.E. Stern and Kavita Bala, A Minimum Interference Routing Algorithm For A Linear Lightwave Network, IEEE GLOBECOM' 91, Phoenix, AZ, Dec. 1991.
- [22] Jacob Sharony, Kwok W. Cheung and T.E. Stern, Wavelength Dilated Switches (WDS) - A New Class of Suppressed Crosstalk, Dynamic Wavelength-Routing Crossconnects, OFC'92, San Jose, CA, 1992.
- [23] K. Bala, Konstantinos Petropoulos and T.E. Stern Multicasting in Linear Lightwave Networks, submitted IEEE INFOCOM'92, 1992.