An Efficiency Prediction Method for ATM Multiplexers

ffiffiflyffin ffiffi Yffi Pffiffi Renu Anthony ffiffiffiooperffiffird YMffi ffiu
ffi ffiell ffioffiffiunicffitions ffieseffirchffi ffied ffiffi
Phoneffi ffiffi ffiffiffi ffiffiffi ffiffiffiffiffi
Yffiolleffie of ffioffiputinffiffi ffieorffiiffi ffinstitut
Atlffintffiffi ffiA ffiffiffiffiffiffiffiffiffiffiffi Phonef
ffiffixffi ffiffi ffiffiffi ffiffiffi ffiffiffiffiffi effiffiilffi ffi

*A*b

work dimensioning methodology needed to achieve economic levels of network equipment usage while simultaneously maintaining satisfactory levels of network performance. The methodology described in this paper addresses a key element of this ATM network dimensioning challenge — the prediction of ATM multiplexing efficiencies achievable under realistic performance objectives while operating with traffic loads imposed by real traffic sources. This methodology is most directly applicable for dimensioning ATM networks and network equipment used in support of Permanent Virtual Connection applications. When used in combination with additional methods for treating the blocking experienced by ATM connections of diverse capacities, one may reasonably anticipate the extension of this methodology for dimensioning ATM networks supporting Switched Virtual Connection applications.

A variety of approaches have been advanced for addressing this ATM network dimensioning challenge, of which Bensaou (1990), Hui (1990) and Elwalid (1995) could be viewed as representative. In addition, significant evidence exists concerning the adverse performance impacts of improperly dimensioned statistical multiplexers operating on traffic that possesses the self-similar temporal characteristic (Leland (1994), Erramilli (1994)). It is therefore prudent for any ATM dimensioning methodology to include verification of its multiplexer performance predictions by testing with real traffic loads that possess this self-similar property. Such verification is provided for the method presented here. We remark that an extension of this measurement-oriented verification process can establish a basis for using routine traffic measurements in support of ATM network dimensioning.

Our study develops in several phases a method that predicts maximum safe load levels for ATM multiplexers. This methodology is developed using several distinct types of traffic source models and a high capacity simulator that is capable of completing simulation runs on the order of 10^{10} cell transfers in less than one hour of real time (Nikolaidis (1993), Nikolaidis (1994), Fujimoto (1995)). The unique benefits obtainable with this high capacity simulator become apparent when examining the clumped nature of cell losses.

In the first phase of this study, a traffic source model based upon a low order Markov Chain is used in a simulation-based process to develop a set of loading curves that are intended to provide conservative, or worst-case, predictions of multiplexer efficiency. These loading curves are parameterized by certain dimensionless ratios relating to both multiplexer and traffic source characteristics. In the second study phase, predictions of multiplexing efficiency based on these loading curves and suitable parameter matching techniques are validated using results from simulations driven by samples of Local Area Network (LAN) traffic.

The remainder of this paper is organized as follows. First, our goals and approach are described and the rather general statistical multiplexing model used for this study is identified. The Markov Chain-based synthetic traffic source model used in the initial phase of this work is then cited, together with the preliminary simulation-based results used to select "worst case" parameters for this traffic source model. Next considered are simulation techniques based on this initial traffic source model and the generation of loading curves. We then describe the simulation technique used in the second phase of this study, which uses a traffic source model based upon Ethernet LAN traffic samples. Comparison between the multiplexing efficiencies predicted from the loading curves and the multiplexing efficiencies achievable with the LAN traffic samples demonstrates the conservative character of this methodology.

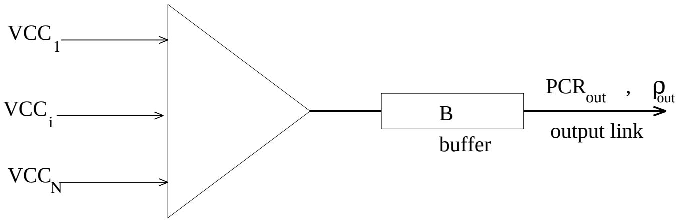

$PCR(VCC_i)$ = Peak Cell Rate of VCC_i, $i = 1, 2, \cdots, N$

PCR_{out} = Peak Cell Rate supportable on output link

$$G = \frac{\sum_{i=1}^{N} PCR(VCC_i)}{PCR_{out}}$$

approximated by an equivalent amount of output buffering. It is generally desirable to achieve large values for N and for two related measures of multiplexer efficiency, which are the statistical multiplexing gain and the output link utilization. It is taken as mandatory that the CLR for this system remain less than its objective value, so that all such efficiency measures are maximized subject to this constraint.

Three dependent variables are useful for characterizing the efficiency of an ATM multiplexer. If we let N be the number of inputs, or traffic sources, driving the multiplexer, these dependent variables are:

$$\begin{array}{rl} N^* = & \text{maximum value of } N \text{ that meets the CLR performance objective for a} \\ & \text{system configuration} \\ G \;= & \text{multiplexing gain, which is relatable to } N^* \text{ as } G = \frac{N^* \cdot PCR(VCC)}{PCR_{out}} = \frac{N^*}{K} \\ \rho_{out} = & \text{output link utilization, which is relatable to } G \text{ as } \rho_{out} = \frac{G \cdot SCR(VCC)}{PCR(VCC)} = \frac{G}{b} \end{array} \qquad (1)$$

where it is convenient to assume that $PCR(VCC_i) = PCR(VCC)$ and $SCR(VCC_i) = SCR(VCC)$ for $i = 1, \ldots, N$, and where K and b are dimensionless ratios whose definitions follow from the relations provided in (1).

The first type of traffic source model used here is based upon a low order Markov Chain that includes sufficient detail to permit independent adjustment of the first order and second order burst length statistics for traffic sources having ON/OFF emission characteristics. A four-state Markov Chain is selected for this purpose. Most of the parameters associated with this source model and the associated multiplexer are next set to "worst-case" values, which are intended to yield the most pessimistic simulation-based results possible for the largest achievable N^*, G and ρ_{out} when the multiplexing system is operating in compliance with its CLR performance objective. Then the remaining parameters of this worst-case configuration (most of which are expressed as dimensionless ratios) are varied to produce a set of loading curves. It is hypothesized that, when the statistics of real traffic sources are appropriately matched with the parameters that index these loading curves, safe predictions can be made concerning the largest values of N^*, G and ρ_{out} that will be achievable.

The second type of traffic source model used in this study incorporates time-based records, or traces, of Ethernet LAN traffic (Leland (1994)). We select this type of traffic for testing the loading curve predictions because such traffic has been shown to possess the self-similar temporal characteristic that makes its statistical multiplexing difficult (Leland (1994), Erramilli (1994)).

Simulation of high speed ATM multiplexers typically requires long execution times, owing to the large number of cell transfer events to be simulated in support of a practical CLR performance objective. While conventional sequential simulators are capable of simulating on the order of 10^7 cell transfers in a reasonable amount of processing time, our results indicate that significantly higher simulation capacity is desired to study cell losses commensurate with a CLR objective of 10^{-7} and to provide adequate consideration for the demonstrably nonindependent nature of cell losses.

Our research makes use of appropriately structured parallel simulation techniques that extend the methods of Fujimoto (1995) and allow about 10^{10} cell transfers to be simulated in about one hour of wall clock time. Our implementation is based on a 32 processor Kendall Square Research shared-memory multiprocessor. This simulator allows the ATM multiplexer's input links to be driven by a variety of source models. The two source models used in the current study are a four-state Markov Chain model, and a LAN trace-driven model.

3 ROLE OF A FOUR-STATE MARKOV CHAIN SOURCE MODEL

A four-state Markov Chain model having an ON/OFF operating mode is selected as our initial traffic source model because it provides several desirable features. First, it permits the independent adjustment of both first order and second order burst-oriented statistics by suitable treatment of the sojourn times in each of its four states. Its characteristics and application to this type of study are reasonably well understood (Eliazov (1990)). Furthermore, it may support analytically tractable analysis of the queuing aspects associated with the multiplexing system under consideration. Since our goal is to relate intermediate results in the form of loading curves obtained with this source model to the statistical multiplexing efficiencies achieved with real traffic sources (and for which analytically tractable models are not generally available), the focus of this work is oriented towards simulation-based analysis.

In general, each of the dependent variables N^*, G and ρ_{out} will be a function of a number of independent variables that characterize various aspects of the traffic sources and multiplexing system under consideration. The more significant of these independent variables are:

- B = size of output buffer, in cells
- K = capacity factor for a traffic source and this output link = $\frac{PCR_{out}}{PCR(VCC)}$
- CLR_{obj} = Cell Loss Ratio objective = 10^{-7} for this study
- MBS = Maximum Burst Size = bound on a traffic source's active period or burst length
- $c^2(A)$ = squared coefficient of variance of the active (ON) period of a traffic source
- $c^2(S)$ = squared coefficient of variance of the silent (OFF) period of a traffic source
- $m(A)$ = mean active (ON) period of a traffic source, in cell emission times referenced to $PCR(VCC)$
- $m(S)$ = mean silent (OFF) period of a traffic source, in cell emission times referenced to $PCR(VCC)$
- b = burstiness factor of a traffic source = $\frac{PCR(VCC)}{SCR(VCC)} = \frac{m(A)+m(S)}{m(A)}$

Observe that any two of the last three parameters listed are sufficient. We elect to use $m(A)$ and b.

Consider now the functional dependencies of N^*, G and ρ_{out}. The value of N^* can be expressed as

$$N^* = N^*(\ B,\ K,\ CLR_{obj},\ MBS,\ c^2(A),\ c^2(S),\ m(A),\ b\ ) \tag{2}$$

and analogous dependencies exist for G and ρ_{out}. We wish to reduce the number of independent variables in (2). This reduction can be accomplished by fixing certain parameters as system constants, and by setting other parameters to values that yield the lowest, or most pessimistic value, of N^*. The selection of model parameters to yield such a worst-case estimate is consistent with our objective of establishing a set of loading curves that can be used to provide conservative estimates of statistical multiplexing efficiency. It can be deduced from the relations provided in (1) that any selection of independent variable values to maximize N^* will also maximize G and ρ_{out}.

The independent variables B, CLR_{obj} and $m(A)$ can be fixed as constants. The value of B is set by the available multiplexing equipment. CLR_{obj} is a limiting value fixed by the operator of that multiplexing equipment, and is assumed here to be 10^{-7}. The value of $m(A)$ is set to approximate the mean Protocol Data Unit length associated with the type of traffic sources under consideration. Since these results are intended for validation against samples of Ethernet LAN traffic, $m(A)$ is taken here to be 28 cells[†].

The independent variables MBS, $c^2(A)$ and $c^2(S)$ are selected to yield worst case values of N^*, as follows.

A number of considerations, including a preliminary set of simulation runs, indicate that N^* generally increases as MBS is made smaller. Therefore setting MBS to infinity results in the worst-case treatment of this independent variable. Since the source model based on a four-state traffic Markov Chain yields unbounded burst lengths, such worst case treatment is provided.

Preliminary simulation runs, as well as prior investigation (Eliazov (1990)), indicate that setting $c^2(A) = c^2(S) \equiv c^2$ has little affect on the resulting value of N^*. Hence this is done here. It remains to select appropriate values for c^2.

From simulation runs using representative parameter settings, the variation of N^* with respect to c^2 was observed (the details of these sample runs are omitted here due to space limitations). From such experiments, it was observed that after an initial decrease in N^*, a point exists beyond which further increases in c^2 do not yield a significant further decrease in N^*. The value of c^2 at this point is selected for establishing simulation-based estimates of N^*, G and ρ_{out}. Once this is accomplished, the functional dependency of the dependent variables upon c^2 is eliminated.

4 SIMULATION ASPECTS AND LOADING CURVES

We now examine some further aspects of this simulation including characterization and treatment of cell loss events, and exhibit some representative loading curves for the initial phase of this methodology. Our simulator software is structured so that the simulated cell transfers (and cell loss events) associated with a single run are captured in a number of separately treatable "segments".

For one simulation run of about 10^{10} cell transfers, our simulator uses roughly 1,500 segments, each of which contains approximately 8×10^6 cell transfers. The observed CLR for one simulation run is calculated as the ratio of the number of cells lost to the total number of cells offered during that run.

On runs for which this observed CLR is in the neighborhood of 10^{-7} , it is found that all cell losses occur in a very few (typically less than 10) of these roughly 1,500 available segments. This is evidence of the tendency of cell losses to be clumped.

In our study, we make use of a "stopping rule" that permits the conclusion with reasonable confidence that $CLR_{obj} \equiv 10^{-7}$ is not exceeded on a particular simulation run, or on its indefinite extension. For a given set of independent variables, an efficient binary search procedure is used to select N for each simulation run, and to control convergence to N^* under the reasonable assumption that CLR is monotonically nondecreasing with N.

L

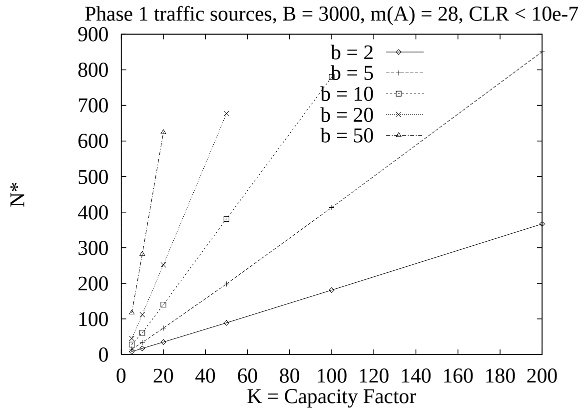

Figure 2:
Loading Curves for
Maximum Number of Traffic Sources

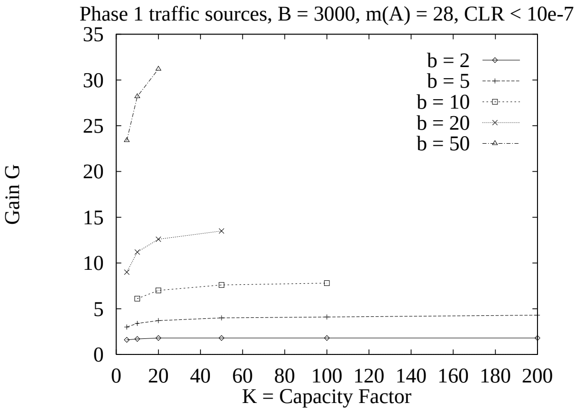

Figure 3:
Loading Curves for
Statistical Multiplexing Gain

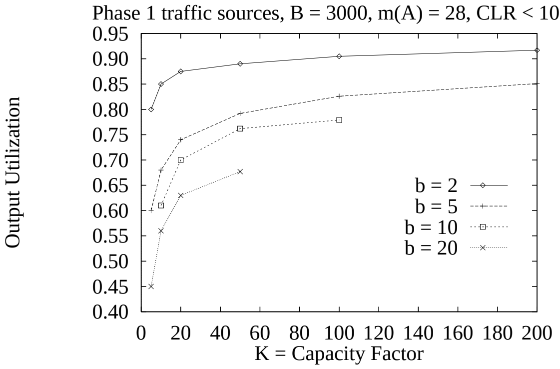

Figure 4:
Loading Curves for
Output Link Utilization

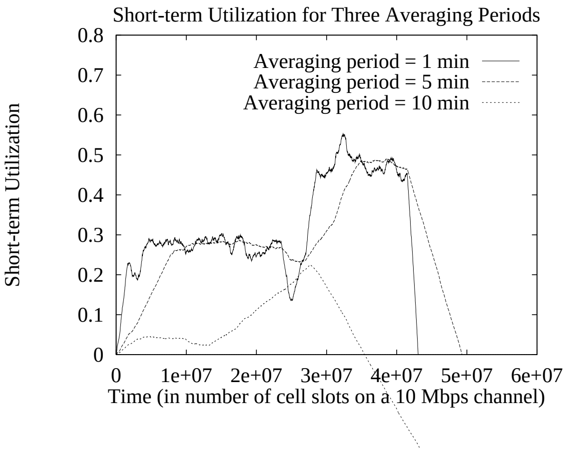

Figure 5:
Short-term Utilization of a Representative
Trace of LAN Traffic

of starting points is confined by our procedures to the initial portions of this trace, no source uses the trace more than twice when generating results reported here.

need to further pursue such investigations with considerably larger bases of real traffic data.

A number of statistics were obtained for the Ethernet trace under consideration after a preliminary conversion from packet-oriented data to cell-oriented data. This conversion is accomplished by treating each packet as a cell burst having suitable length and occurring at a $PCR(VCC)$ that is equivalent to 10 Mbits/second (for the Ethernet LAN). The average utilization, ρ_{source}^{long}, of this converted trace over its 30 minute duration is 0.33. The following measured values for this converted trace are noted:

- $m(A) = 24.8$ cells
- $m(S) = 49.7$ cells
- $c^2(A) = 4.63$
- $c^2(S) = 3.80$

Figure 5 displays three plots of the short term utilization for this converted trace that were obtained by averaging over sliding windows of 1 minute, 5 minutes and 10 minutes, respectively. A point on such a utilization plot is found through dividing the number of trace-originated cell slots occurring in a particular window (whose right end point is plotted on the abscissa) by the number of cell slots available in that window at the selected $PCR(VCC)$. The primary interest here concerns the central portions of these plots, with the end ramps being artifacts of data exhaustion in the sliding window averaging procedure. These plots of Figure 5 will be relevant when matching the multiplexing efficiencies achieved with trace-based sources to the efficiencies predicted by the loading curves given in Figures 2, 3 and 4.

D