Proceedings of the 29th Conference
on Decision and Control
Honolulu, Hawaii • December 1990

A TRAFFIC/PERFORMANCE ANALYSIS OF THE BANDWIDTH MANAGEMENT THROUGHPUT-BURSTINESS FILTER

A. E. Eckberg
Room 3J-611
AT&T Bell Laboratories
Holmdel, NJ 07733 USA

D. M. Lucantoni
Room 3K-601
AT&T Bell Laboratories
Holmdel, NJ 07733 USA

Abstract

As an initial step in assessing the performance of an overall B-ISDN/ATM congestion control framework, the *Throughput-Burstiness Filter* (TBF) of the *Bandwidth Management* network congestion control capability is modeled and analyzed. The TBF is part of a distributed core control mechanism providing protection to, and fair allocation of, shared network resources. Through this model and analysis, the combinations of traffic throughput and burstiness that the TBF allows can be quantified.

1. Introduction and Summary

This paper is motivated by control issues in Broadband ISDN (*B-ISDN*), where information transport will be effected via Asynchronous Transfer Mode (*ATM*, a high-speed, "packet-based" transport mechanism; see [1] and [2]). B-ISDN/ATM networks will provide information transport for a rich mixture of services and applications, associated with which will be a broad spectrum of traffic types and transport performance needs. In order to be competitive with specialized private network alternatives in meeting these needs, it is expected that B-ISDN/ATM networks will need to be engineered to take full advantage of available shared network resources. The resulting combination of high resource utilizations and diverse traffic characteristics will require mechanisms such as service class treatments and both reactive and preventive congestion controls. Thus, a well-thought-out and comprehensive set of controls is essential to the viability of B-ISDN/ATM.

An overall strategy for congestion control, together with complementary flow and error controls, continues to be an outstanding technical challenge to B-ISDN/ATM. In particular, control schemes that have proven to be effective in lower-speed packet networks may not scale well to the high speeds, and to large propagation delays (in comparison with "packet" transmission times), that will be characteristic of B-ISDN/ATM. An overall control strategy needs to incorporate both preventive and reactive components, to provide robustness and the flexibility to accommodate diverse traffic and application types, to work at broadband speeds over wide-area distances, and to be implementable within the constraints imposed by the agreed-to structures of B-ISDN/ATM, many of which are even now being standardized.

An Overall Congestion Control Framework

Recently, in [3], a possible overall B-ISDN/ATM congestion control framework was described, consisting of a system of synergistic network congestion controls and end-terminal controls. This congestion control framework includes the notion of three "control domains" into which the network congestion controls may be partitioned:

i. Call-level controls — For each call, or *virtual connection* (VC), there should be the notion of a "service contract," specifying a set of service parameters mutually agreed-to between the network and the terminal(s). The service parameters are set during the call establishment procedure, and allow agreement with respect to items such as:
   - limits on the traffic which the network is expected to transport, and more specifically, a mechanism by means of which *excessive traffic* entering the network from a terminal can be discriminated in real-time from *non-excessive traffic*
   - performance that needs to be delivered by the network, e.g., throughput requirements, and delay, jitter, and loss tolerance/intolerance
   - special priority treatments that may need to be given to the call

ii. Network-wide ATM-level controls — With service parameters having been established via call-level controls, there needs to be a set of mechanisms for transferring these service agreements into practice during the transport of ATM cells. (The ATM *cell* is the basic "packet" transport unit in ATM.) This is the function of the network-wide ATM-level controls, with capabilities such as:
   - a capability for selectively shedding load under congestion conditions, and thus providing network resiliency to traffic uncertainties, implemented via an indicator in the ATM cell header (the *cell loss priority indicator*) which, if set, signifies that the cell may be *selectively* discarded in any network element along the VC path if the cell encounters congestion in that element above a threshold. The cell loss priority indicator may be set

a. by the sending terminal to signify that the cell carries nonessential information, or
b. by the network for "excessive-traffic" cells identified at the terminal-network interface (here we may term the cell loss priority indicator a *traffic violation tag* to emphasize this specific use of the cell loss priority)

Note that the traffic limits with respect to which ATM cells might be identified, and traffic-violation-tagged, as "excessive-traffic" cells are part of the "service contract" agreed to at the call establishment.

— a capability for conveying forward, along the VC to the ATM destination terminal, information concerning congestion encountered along this VC; this would be implemented via a *forward congestion indicator*, borne within the header of ATM cells, and set by any congested network element through which that cell passes.

— a capability for backward notification of relevant congestion onset/abatement conditions across an ATM interface via network-element-originated ATM cells.

iii. Network-element internal controls — Finally, the network-wide ATM-level controls described above require specific actions within various of the network elements; these include: traffic-monitoring and traffic-violation-tagging to discriminate between excessive and non-excessive traffic (only at the network interface), selective cell discard, setting of the forward congestion indicator, and triggering the backward notification of congestion. In addition, each network element will need to perform service scheduling for bandwidth and buffers consistent with the service parameters of the calls supported by that element, and to prevent one service class or VC from "locking out" another from transmission. (This establishes a system of "relative priorities" between VCs and/or service classes.)

The BWM Core Network Congestion Control

To achieve the robustness and flexibility that is required of a B-ISDN/ATM overall congestion control, all of the capabilities in the above framework are considered critical. However, within this framework, there is a *core* network congestion control, consisting of the traffic-monitoring, traffic-violation-tagging, and selective-cell-discarding capabilities, which we term *Bandwidth Management* (BWM). This core network congestion control operates in real-time, on each individual VC cell flow, and continuously provides protection to, and a fair allocation of, the shared network resources.

This is a system of *distributed* controls, with two basic control components:

i. the *Throughput-Burstiness Filter* (TBF), consisting of the traffic-monitoring and traffic-violation-tagging capabilities, and residing at the terminal-network interface
ii. the *Selective Cell Discard* (SCD), consisting of the discard of cells at low loss priority (e.g., traffic-violation-tag set) when congestion is encountered, and residing in all network elements

See [4] and [5] for more details. It is convenient to consider the TBF, with input equal to the flow of cells on a VC from the terminal, as composed of two components: (i) TBF1, with output equal to the flow of non-violation-tagged cells on the VC; and (ii) TBF2, with output equal to the flow of violation-tagged cells on the VC. It is the TBF1 component that we will address with modeling and analysis.

The Leaky-Bucket: A Specific Traffic-Monitoring Scheme

Many traffic monitoring schemes have been suggested and studied, but one that exhibits both simplicity of implementation and monitoring performance is the "leaky-bucket" monitoring algorithm (see, e.g., [6]). There are numerous variants of this algorithm, but a typical one involves maintaining a counter with simple operations at successive cell arrivals on a VC as follows:

i. The counter value, X, is initialized to 0.
ii. At each new cell arrival, the counter value is first decremented as $X \leftarrow max(X-cT, 0)$, where c is a specified parameter, and T is the elapsed time since the last cell arrival (but, as indicated, X is not decremented below 0).
iii. The counter value X is then compared with a prescribed value M, and
    a. if $X \geq M$, the cell is violation-tagged and passed on;
    b. if $X < M$, the cell is passed untagged, and the counter is incremented as $X \leftarrow X+1$,

This algorithm allows a sustained, untagged, throughput rate of c cells per time unit, and a burstiness that is controllable (for a given value of c) through the selection of the value of M. Note that the operation of the leaky-bucket is identical to that of a finite-capacity single server queue with a deterministic service time.

2. Modeling the Throughput-Burstiness Filter

Overview of Modeling Approach

Our objective is an effective methodology for modeling and analyzing the TBF1 module described above. To this end, we first define a model that can capture fairly general traffic characteristics of broadband traffic sources, and that lends itself to a quantitative methodology for assessing the effect on non-violation-tagged traffic throughput and burstiness due the BWM leaky-bucket traffic monitoring. We choose to use a Markov-modulated Poisson process (MMPP)[7] model for the source traffic, and to use the *peakedness* framework (see [8]) to quantitatively assess the non-violation-tagged burstiness. Thus, arrivals originate according to an N-state *MMPP*, and are offered to a single server queue with finite capacity M and deterministic service time $d=c^{-1}$, which models the dynamics of the BWM leaky-bucket. Arrivals that are allowed to enter the single server queue (i.e., cells that would not be violation-tagged) are simultaneously put onto an infinite-server system where service times are exponentially distributed with rate parameter μ. The focus of the analysis for this model is a joint characterization of the queue contents, Q, and the number in the infinite-server system, X, together with the *MMPP* Markov chain state.

In particular, a major objective is to quantify $E[X(\infty)]$, and $Var[X(\infty)]$, where $X(\infty)$ denotes the number in the infinite-server system at an arbitrary epoch in time. Then the throughput of the non-violation-tagged traffic is just $\mu E[X(\infty)]$, and the peakedness value at μ of this traffic is defined as $z(\mu) = Var[X(\infty)]/E[X(\infty)]$ (see [8]).

The parameter μ enters because the entire *peakedness function*, $z(\cdot)$, provides the type of burstiness characterization that is most useful in assessing the impacts of burstiness on transport performance (delays, losses, etc.). As is both qualitatively and quantitatively shown in [9], a value of μ roughly equal to the reciprocal of the delay in a network queue makes $z(\mu)$ most informative as a measure of burstiness from which performance impacts can be approximately predicted. Thus, we are interested in calculating the values of $z(\mu)$ for a range of μ values.

The analysis of the model first concentrates on an embedded Markov chain, then expands the analysis to an arbitrary epoch in time. We begin by presenting some results related to the *MMPP* arrival process.

Consider an irreducible N-state continuous-time Markov process with infinitesimal generator R and let the diagonal matrix Λ have diagonal entries $\{\lambda_1, \ldots, \lambda_N\}$, where λ_j is the rate of the Poisson arrival process when the Markov process is in state j. Let N_t be the number of arrivals in $(0,t]$ and J_t the state of the Markov process at time t. Now let

$$P_{ij}(n,t) = P\{N_t = n,\ J_t = j \mid N_0 = 0,\ J_0=i\}$$

be the (i,j) entry of a matrix $P(n,t)$. It is well known (see e.g., [10]) that the matrix generating function $P^*(z,t) = \sum_{n=0}^{\infty} P(n,t)z^n$ is explicitly given by

$$P^*(z,t) = e^{(R+\Lambda(z-1))t}, \qquad |z| \le 1,\ t \ge 0. \tag{1}$$

The stationary vector, $\boldsymbol{\theta}$, of the Markov process underlying the arrival process, satisfies the equations

$$\boldsymbol{\theta} R = \mathbf{0}, \qquad \boldsymbol{\theta} \boldsymbol{e} = 1, \tag{2}$$

where $\boldsymbol{e}$ is a column vector of ones. Define the matrix $\Phi(\mu)$ as

$$\Phi(\mu) = \int_0^\infty e^{-\mu t} P(0,t)\Lambda dt = \int_0^\infty e^{-\mu t} e^{(R-\Lambda)t}\Lambda dt$$

$$= (\mu I + \Lambda - R)^{-1}\Lambda$$

and note that its (i,j) element, $\phi_{ij}(\mu)$, is

$$\phi_{ij}(\mu) = \int_0^\infty e^{-\mu t} f_{ij}(t) dt$$

where $f_{ij}(t)$ is the probability density function of the time till the first arrival in the *MMPP* after time 0 and its state at the arrival is j, given that the state at time 0 was i. The matrix $U = \Phi(0) = (\Lambda - R)^{-1}\Lambda$ is stochastic and keeps track of the phase change during the first interarrival time.

The System Definition

Consider the triple $\{X(t), Q(t), S(t)\}$, where $X(t)$, $Q(t)$, and $S(t)$ are respectively, the number in the infinite system, the number in the bucket, and the phase of the arrival process at time t. Define $\tau_0=0$, and let τ_n, $n \ge 1$, be the epoch of the n-th occurrence of either a service completion in the bucket or the end of an idle period of the leaky bucket. We also define the embedded process $(X_n, Q_n, S_n) = (X(\tau_n^+), Q(\tau_n^+), S(\tau_n^+))$.

We now define

$$\pi_n(z,i,j) = E[z^{X_n} \mathbf{1}_{\{Q_n=i,\ S_n=j\}}]$$
$$= P[Q_n=i,\ S_n=j] E[z^{X_n} \mid Q_n=i,\ S_n=j], \tag{3}$$

where $\mathbf{1}_{\{\cdot\}}$ is the indicator function, and also define the vectors

$$\boldsymbol{\pi}_n(z,i) = (\pi_n(z,i,1),\ \pi_n(z,i,2), \ldots, \pi_n(z,i,N)).$$

Our first effort is to find a method for determining $\boldsymbol{\pi}_n(z,i)$, and in particular,

$$\boldsymbol{\pi}(z,i) = \lim_{n\to\infty} \boldsymbol{\pi}_n(z,i).$$

The number of customers present in the infinite system at time τ_{n+1} consists of the sum of two independent random variables. The first is the number of customers who were present at time τ_n, (if any), who are still present at time

τ_{n+1}. We refer to these customers as *decay* customers. The second random variable is the number of new customers who arrive to the infinite system during $(\tau_n, \tau_{n+1}]$ and are still there at τ_{n+1}. These are referred to as *immigration* customers. We note that if there are i tokens in the bucket at time τ_n then there are a maximum of $M-i$ immigration customers allowed into the infinite system in $(\tau_n, \tau_{n+1}]$. With this in mind we define the following quantities.

$[B_i(n,k,t)]_{jl} = P\{n$ immigration customers in service in the infinite server system at time t and k total immigration arrivals in $(0,t]$ and the arrival process in phase l at time t | the infinite server system started empty at time 0, a maximum of $M-i$ arrivals allowed in $(0,\infty)$ and the arrival process in phase j at time $0\}$,

for $0 \le i \le M-1$, $0 \le k \le M-i$, $0 \le n \le k$, $t \ge 0$. Let the matrix $B_i(n,k,t)$ have as its (j,l) element $[B_i(n,k,t)]_{jl}$ and define the matrices of generating functions

$$B_i(z,k,t) = \sum_{n=0}^{\infty} B_i(n,k,t) z^n.$$

Then it is easy to verify that

$$B_i(1,k,t) = \begin{cases} P(k,t), & k<M-i \\ \sum_{j=M-i}^{\infty} P(j,t), & k=M-i. \end{cases} \qquad (4)$$

We now have the following theorem.

Theorem 1: The generating functions $\pi(z,k,l) = \lim_{n\to\infty} \pi_n(z,k,l)$ satisfy

$$\pi(z,k,l) = \sum_{i=1}^{M-1} \sum_{j=1}^{N} \pi(1+e^{-\mu d}(z-1),i,j)[B_i(z,k-i+1,d)]_{jl} + \delta_{k,1} \sum_{j=1}^{N} z\pi(1+(\phi_{jl}(\mu)/U_{jl})(z-1),0,j)U_{jl}, \quad (5)$$

where $\delta_{k,1}=1$ when $k=1$ and 0, otherwise.

Proof: Due to space limitations, all lengthy proofs are omitted. We refer to [11] for the details.

Remark: Equation (5) has a simple intuitive interpretation. The first term is the sum over various initial conditions of the product of two generating functions. The first generating function represents the number of customers present at the end of a service who were present following the previous service. The second represents the number of new customers present at the end of a service who were not present following the previous service. The product of these two generating functions is the generating function of the convolution of these two independent random variables. The second term in (5) represents the generating function of the number of customers present at the end of an idle period who were also present at the end of the previous busy period. Only one new customer will be present since the first one ends the idle period.

By successively differentiating with respect to z in (5), setting $z=1$, and using vector notation, we have

$$\dot{\pi}(1,k) = \sum_{i=1}^{M-1} \left[\pi(1,i)\dot{B}_i(1,k-i+1,d) + e^{-\mu d}\dot{\pi}(1,i)B_i(1,k-i+1)\right] + \delta_{k,1}\left[\pi(1,0)U + \dot{\pi}(1,0)\Phi(\mu)\right], \qquad (6)$$

and

$$\ddot{\pi}(1,k) = \sum_{i=1}^{M-1} \Big[\pi(1,i)\ddot{B}_i(1,k-i+1,d) + 2e^{-\mu d}\dot{\pi}(1,i)\dot{B}_i(1,k-i+1,d) + e^{-2\mu d}\ddot{\pi}(1,i)B_i(1,k-i+1,d)\Big] + \delta_{k,1}\left[2\dot{\pi}(1,0)\Phi(\mu) + \ddot{\pi}(1,0)(\Phi(\mu)\circ(\Phi(\mu)/U))\right], \qquad (7)$$

where $\circ$ denotes the Schur (or entry-wise) product of two matrices, / denotes the entry-wise division of two matrices, and we use the notation of $\dot{x}$ and $\ddot{x}$ to denote derivatives with respect to z. From (6) and (7), we obtain the first two moments of the number of busy servers in the infinite system *at the embedded epochs* $\{\tau_n\}$. To obtain the peakedness of the untagged traffic we need the corresponding moments at an arbitrary time. To this end, we consider the embedded semi-Markov process (*SMP*) governed by $\{X_n, Q_n, S_n, \tau_{n+1}-\tau_n\}$. The fundamental mean transition time of this *SMP*, σ, is given by

$$\sigma = d(1-\pi(1,0)e) + \pi(1,0)(\Lambda-R)^{-1}e. \qquad (8)$$

For $n \ge 0$, $0 \le i \le M$, $1 \le j \le N$, define $y(n,i,j)$ to be

$$\lim_{t\to\infty} P\{X(t)=n,\ Q(t)=i,\ S(t)=j \mid X(0)=n',\ Q(0)=i',\ S(0)=j'\},$$

and let the vector $y(z,i)$ have j-th component $y(z,i,j)$, where

$$y(z,i,j) = \sum_{n=0}^{\infty} y(n,i,j) z^n, \qquad |z| \le 1.$$

By standard arguments based on the Key Renewal Theorem for Markov renewal processes (see e.g.,[12]) we obtain the following.

Theorem 2:

$$y(z,0) = \frac{1}{\sigma} \int_0^{\infty} \pi(1+e^{-\mu t}(z-1),0) e^{(R-\Lambda)t} dt, \qquad (9)$$

and for $i=1,...,M$,

$$y(z,i) = \frac{1}{\sigma} \sum_{m=1}^{i} \int_0^d \pi(1+e^{-\mu t}(z-1),m) B_m(z,i-m,t) dt, \quad (10)$$

where $\pi(z,M) \equiv \mathbf{0}$.

Proof: See [11]

This leads to

$$y(1,0) = \frac{1}{\sigma} \pi(1,0)(\Lambda-R)^{-1}, \qquad (11)$$

and for $i=1,...,M$,

$$y(1,i) = \frac{1}{\sigma} \sum_{m=1}^{i} \pi(1,m) \int_0^d B_m(1,i-m,t) dt. \qquad (12)$$

Using (4), we see that (11) and (12) are the expressions for the joint probability of the number of customers in the queue and the phase of the arrival process at an arbitrary point in time in the finite $MMPP/D/1$ queue.

Peakedness of the Non-Violation-Tagged Traffic

The peakedness $z(\mu) = Var[X(\infty)]/E[X(\infty)]$ is obtained from the following theorem.

Theorem 3:

$$E[X(\infty)] = \frac{1}{\sigma} \dot{\pi}(1,0)(\mu I+\Lambda-R)^{-1} e + \frac{1}{\sigma\mu}(1-e^{-\mu d}) \sum_{m=1}^{M} \dot{\pi}(1,m) e$$
$$+ \frac{1}{\sigma} \sum_{m=1}^{M} \pi(1,m) \sum_{i=0}^{M-m} \int_0^d \dot{B}_m(1,i,t) dt e, \qquad (13)$$

and

$$E[X^2(\infty)] = \frac{1}{\sigma} \ddot{\pi}(1,0)(2\mu I+\Lambda-R)^{-1} e$$
$$+ \frac{1}{2\sigma\mu}(1-e^{-2\mu d}) \sum_{m=1}^{M} \ddot{\pi}(1,m) e \qquad (14)$$
$$+ \frac{2}{\sigma} \sum_{m=1}^{M} \dot{\pi}(1,m) \sum_{i=0}^{M-m} \int_0^d e^{-\mu t} \dot{B}_m(1,i,t) dt e,$$
$$+ \frac{1}{\sigma} \sum_{m=1}^{M} \pi(1,m) \sum_{i=0}^{M-m} \int_0^d \ddot{B}_m(1,i,t) dt e,$$

Proof: Differentiating with respect to z in (9) and (10) and setting $z=1$ yields

$$\dot{y}(1,0) = \frac{1}{\sigma} \dot{\pi}(1,0)(\mu I+\Lambda-R)^{-1},$$

and

$$\dot{y}(1,i) = \frac{1}{\sigma} \sum_{m=1}^{i} \left[\dot{\pi}(1,m) \int_0^d e^{-\mu t} B_m(1,i-m,t) dt + \pi(1,m) \int_0^d \dot{B}_m(1,i-m,t) dt \right],$$

for $1 \le i \le M$. A second differentiation with respect to z in (9) and (10) and setting $z=1$ leads to

$$\ddot{y}(1,0) = \frac{1}{\sigma} \ddot{\pi}(1,0)(2\mu I+\Lambda-R)^{-1}$$

and

$$\ddot{y}(1,i) = \frac{1}{\sigma} \sum_{m=1}^{i} \left\{ \ddot{\pi}(1,m) \int_0^d e^{-2\mu t} B_m(1,i-m,t) dt + 2\dot{\pi}(1,m) \int_0^d e^{-\mu t} \dot{B}_m(1,i-m,t) dt + \pi(1,m) \int_0^d \ddot{B}_m(1,i-m,t) dt \right\},$$

for $1 \le i \le M$. By definition, we have

$$E[X(\infty)] = \sum_{i=0}^{M} \dot{y}(1,i) e,$$

and

$$E[X^2(\infty)] = \sum_{i=0}^{M} \ddot{y}(1,i) e.$$

The theorem now follows by recognizing that

$$\sum_{i=0}^{M-m} B_m(1,i,t) = e^{Rt}, \qquad t \geq 0. \qquad (15)$$

Remark: Evaluation of the expressions in (6), (7), (13) and (14) require the computation of the moment matrices $\dot{B}_m(1,i,t)$, $\ddot{B}_m(1,i,t)$, and various integrals of these. Efficient procedures for carrying out these computations are derived in [11].

3. Conclusion

We have described an overall B-ISDN/ATM congestion control framework, with a *Bandwidth Management* network congestion control core, and have focused modeling and analysis on one component of this BWM core, namely the portion of the *Throughput-Burstiness Filter* that allows ATM cells on a VC to pass non-violation-tagged. This analysis is being coupled with analyses of other congestion control components to derived a means for overall control performance assessment.

REFERENCES

1. "Broadband Aspects of ISDN," *CCITT COM XVIII Draft Recommendation I.121*, June, 1988.
2. "Meeting Report of Sub-working Party 8/1—ATM," *CCITT COM XVIII-TD.14-E*, Geneva, June, 1989.
3. A. E. Eckberg, "A B-ISDN/ATM Congestion Control Framework for Study and Refinement," T1S1 Contribution T1S1.5/90-158, San Francisco, June 25-29, 1990.
4. A. E. Eckberg, D. T. Luan, and D. M. Lucantoni, "An Approach to Controlling Congestion in ATM Networks," to appear in *Int. J. of Digital and Analog Communication Systems*, 1990.
5. A. E. Eckberg, D. T. Luan, and D. M. Lucantoni, "Bandwidth Management: A Congestion Control Strategy for Broadband Packet Networks — Characterizing the Throughput-Burstiness Filter," Int'l. Teletraffic Congress Specialists Seminar, Adelaide, 1989.
6. J. S. Turner, "New Directions in Communications (or Which Way to the Information Age?)," *IEEE Communications Magazine*, October, 1986.
7. H. Heffes, and D. M. Lucantoni, "A Markov modulated characterization of packetized voice and data traffic and related statistical multiplexer performance", *IEEE J. on Selected Areas in Communications, Special Issue on Network Performance evaluation*, SAC-4, 6, 856-868,1986.
8. A. E. Eckberg, "Generalized Peakedness of Teletraffic Processes," *Proceedings of the 10th International Teletraffic Congress*, Montreal, 1983.
9. A. E. Eckberg, "Approximations for Bursty (and Smoothed) Arrival Queueing Delays Based on Generalized Peakedness," *Proceedings of the 11th International Teletraffic Congress*, Kyoto, 1985.
10. M. F. Neuts, "The versatile Markovian point process", *J. Appl. Prob.*, vol.16, pp. 222-61, Mar. 1980.
11. A. E. Eckberg, and D. M. Lucantoni, "Modeling and Analysis of the Bandwidth Management Throughput-Burstiness Filter", to be submitted for publication.
12. E. Çinlar, *Introduction to Stochastic Processes*, Prentice-Hall, 1975.