Simulation of ATM Congestion Control Algorithms

Jesse A. Alama (alamaja@mrs.umn.edu)

A.A. Lopez (alopez@mrs.umn.edu)

University of Minnesota, Morris
Morris, MN 56267

Introduction

The increased reliance upon computer networking as a modern communication media has placed much pressure on scientists to develop robust, powerful, and flexible networking models. In this paper, we describe Asynchronous Transfer Mode (ATM) networks which were introduced in the last decade and currently carry a significant proportion of the global network traffic. We discuss the advancements ATM has made in the last few years. Finally, we observe that work remains to be done in the field of congestion control algorithms and describe our experience with a simulator that makes the study of these algorithms manageable.

What are ATM networks? and Why is this technology important?

ATM is a robust networking protocol that allows sharing of high speed links and switches to improve network reliability and throughput. ATM networks break messages up into 53 byte *cells* for ease of handling and retransmission of damaged packets. These networks can easily and flexibly set up services between two of its customers as compared with traditional networks which require dedicated, permanent services between customers.

ATM technology is important to the information technology world because it: i)provides guaranteed Quality of Service services, ii)uses telephone technology that we are familiar with, iii)provides for more robust transmission of information, iv)permits us to combine different types of traffic over high capacity links, v) makes better use of high capacity links, and vi)provides for more reliable and cost effective transmission of information.

One of the characteristics that separates ATM from most other networks is called quality of service (QoS). An application may request certain critical parameters for its session, such as available bandwidth; number of cells that arrive out of order; and maximum delay between two endpoints. Quality of Service is an important characteristic because as data network traffic grows, a higher percentage of it is time-sensitive. Time-sensitive traffic includes all forms of voice and video services including newer services such as Voice over IP (VoIP) and video conferencing. Voice traffic is particularly sensitive to delays, as the human ear has problems understanding spoken language if portions of messages are lost, distorted or delayed. QoS specifications will allows us to guarantee that a minimal throughput will be available for a given user and/or application.

ATM services are similar to telephone services in many ways. Both services set up virtual circuit (VC) between two locations (for example when you make a call between the USA and Japan) and request a guaranteed capacity (e.g. being able to clearly hear the other party on a telephone conversation). The services are also similar in carrying a combination of different traffic types (e.g. voice, financial data and video information). Because of these similarities we have seen a significant increase in the adoption of ATM services by major telecommunication companies such as AT&T, MCI and Sprint.

The robustness of ATM services is achieved by several factors. The primary factor is that the quality of new transmission media has improved by several orders of magnitude over the past three decades. A secondary factor is the simplicity of ATM services that breaks up messages into 53 byte cells that can be easily retransmitted. And finally, the fact that the ATM protocols have been carefully developed under the watchful eye of the ATM Forum [1] with input from many vendors and researchers.

ATM services are designed for the sharing of high speed links. Within the ATM model [2] different types of services are supported (Adaptation Layers, AAL 1-5) each with characteristics suited for a different type of user service. These individual services are designed to accommodate: Constant Bit Rate (CBR) traffic (such as that required when providing a circuit emulation service); Variable Bit Rate (VBR) traffic (as usually required by voice and video applications whose signals can be compressed) and Unspecified Bit Rate (UBR) traffic, which represents a 'best-effort' type of service suitable for file transfer and e-mail services. Some authors would further subdivide these services into categories such as real-time and non-real-time VBR and add an Available Bit Rate (ABR) service.

The utilization of a high capacity link can be improved by combining several of the above services over one link. In this scheme, if a VBR customer is not using part of her 'requested' bandwidth, the extra bandwidth can be temporarily assigned to ABR or UBR users that have a backlog of information to transmit. Whenever the VBR customer has need for all of her requested bandwidth, the ABR user will be throttled back immediately and the VBR customer will proceed unimpeded. To do this, the switch uses information contained in the headers of the many cells it receives per second.

In many cases the ATM traffic will be able to be combined over one link. Link sharing generally results in more economical transmission of information. Furthermore, as many of these high speed links are implemented using fast double-ring SONET/SDH technologies, the transmission will be more robust. SONET/SDH technology can send information in two directions in a ring. If any active component of the system senses that there has been a disturbance in its upstream flow, the direction of the information flow will be automatically and instantaneously reversed. SONET/SDH makes ATM services more economical and robust.

Some current uses of ATM technology .

Early proponents of ATM envisioned it as the omnipresent network technology. Even the desktop would be served directly by ATM. However, over the past four years the network market has split between local access to desktops and backbone/wide area networks. The last 100 meters of the network continues to be dominated by Ethernet networks which were prevalent throughout in the early 1990's. In fact, Ethernet has been enhanced by the addition of 100BT networks that increased the throughput ten-fold at a relatively modest 40% increase in cost of adapters and hubs. Because of the ease and economy with which existing 10 BT Ethernet networks can be converted to 100BT, ATM technology is no longer seen as a necessity in serving the desktop.

However, even in such an environment there are exceptions. Georgia Public Broadcasting (GPB) is using IP over ATM in its local area network [4]. Their work is made feasible by the recent recommendation from the ATM Forum for a Real Time over ATM (RMOA) specification and by GPB's desire for a protocol that will support video on demand and voice over their OC-3 (155 Mbps) environment.

On the other hand, backbone and wide area networks are starting to be dominated by ATM technology. For instance, the backbone in most large university campuses is made up of a series of ATM switches interconnected by high speed links. Likewise, ten of the eleven largest national Internet Service Providers have ATM switches at their core, according to an International Data Corp. study [4]. In addition, ATM services are starting to become widely available as public carriers start to offer switched virtual connections over ATM. In 1998, more than $1.7 billion was spent on ATM switching equipment and that amount is expected to double by the year 2000. Thus it appears that ATM networks will be with us for a long time to come.

Motivation for Study

The ATM networking model can be easily applied to simple ATM networks. In more complex network configurations, additional factors need to be taken into consideration. Consider a network with a CBR client and an ABR client. If cells from the CBR client are always present in the ATM switch's queues, and if CBR traffic always takes priority over ABR traffic, will the ABR cells ever go through? It is the goal of network designers to avoid *starvation* that might occur in our hypothetical network.

Or consider a network with dozens of applications. May the ATM switch allocate more resources that it has at any given time? If too conservative of approach is taken, resources may be wasted. If the strategy is too aggressive, the switch may allocate far more resources than it could ever handle, thus reducing the quality of service it provides.

Choosing a middle line between efficiency and guaranteed quality of service is the motivation behind developing *congestion control algorithms*. Congestion control algorithms are procedures for determining when a certain cell ought to have the resources of the network, and when it should not.

Our work concentrates on studying congestion control algorithms. To study congestion control algorithms, we use an ATM *simulator*.

The NIST ATM Simulator

One of the most visible signs of progress in ATM networking is the increased availability of ATM simulation tools. The National Institute for Standards and Technology [3] has developed a free software simulator for ATM networks. In this section, we describe our experience and work with this simulator.

The NIST ATM simulator is a graphical network simulation tool. One visually constructs network configurations (see Figure 1) and sets certain parameters, such as link speeds, traffic types, distance between nodes, and many others. When the simulator is run, the operation of an ATM network is simulated using the parameters entered earlier.

The NIST ATM simulator contains several built-in congestion control algorithms. The simulator also allows for other congestion control algorithms to be designed or integrated, thus allowing the researcher an opportunity to test their own congestion control algorithms.

Results

We designed a network to compare the performance of two congestion control algorithms for ABR traffic. In an ideal situation, ABR traffic will take up the remaining resources of the network after CBR and VBR applications have been satisfied. However, achieving ideal ABR behavior is difficult even for the best congestion control algorithms.

Much of the problem deals with the finite nature of ATM switches. An ATM switch can only process so many cells per second. As cells stream into the switch, they are stored in a queue. The switch determines which cells to handle from its many queues by adhering to a congestion control algorithm.

In Figure 2, we have provided a graph detailing ideal ABR behavior. One of the algorithms supported by the NIST simulator is Explicit Forward Congestion Indication (EFCI), which is intolerant of ABR traffic occupying space in the switch’s queues. Once the queue fills up, the switch denies the ABR application's requests for more bandwidth. In fact, the switch begins to cut back the amount of resources for the application. While the switch cuts the amount of resources for the ABR traffic, it must still process the CBR and VBR traffic. In situations of extreme amounts of ABR traffic, this algorithm can cause monotonic decreases in the resources given to the ABR application (see Figure 3).

Similar results may be seen in another algorithm, called NIST ER (Explicit Rate). This algorithm does not cut back the ABR resources in so extreme a manner as the ECFI algorithm (see Figure 4).

We were able to resolve the ABR reduced throughput when using the NIST algorithm by setting one of its variables to a higher value. After adjusting the variable, ABR traffic exhibited expected behavior. Throughout these simulations the remaining network resources were *fully allocated*.

```
(1)  if receive ABR cell
(2)    if queue-length < output-queue-length
(3)      add cell to output queue
(4)      count = count + 1
(5)      if (count % N) = 0
(6)        if output-queue-size > (old-output-queue-size + tau)
(7)           congested = TRUE
```

In the above algorithm, the variable *tau* controls congestion tolerance. It was initially set to 0. Another parameter, *count*, controls whether the check for congestion occurs. Since tau was set to such a low value, the conditional on line (6) was always true, thus setting the congestion flag. With this flag set, the switch proceeded to cut the resources allocated to ABR, since it is the lowest priority traffic type. Although this section of code is repeated every time ABR cells are received, but since the buffers were always full, the ABR traffic had its resources continually reduced.

The function which governs just how much an ATM switch can decrease the ABR traffic is given by the expression *allow cell rate = allow cell rate (*1 – *rate decrease factor)* [5]. This function decreases monotonically since the *rate decrease factor* is always positive. Hence the monotonic decrease seen in Figures 3 and 4.

Importance of the simulator

Research of ATM networks can be conducted using a physical network or through simulations. Physical ATM networks continue to be prohibitively expensive. Typical configurations still cost tens of thousands of dollars. As such, the only practical way of studying ATM networks is through simulation.
Yet commercial ATM simulators are very expensive. The ATM community is fortunate to have the NIST ATM simulator for free study.

For small colleges, the simulator is an ideal tool. Its availability and cost have proven to be the best way for us to continue our research. The simulator could also be used in computer network classes to expose the students to the intricacies of network protocols.

The NIST ATM simulator has helped us study congestion control algorithms in great detail. The graphical user interface and the logging tools have given feedback that a mathematical analysis might have left out.

Conclusion

We have given a brief introduction to ATM networks. We have discussed what ATM is, why it is important, where it appears, and how cost effective it is. We then discussed how the ATM field has progressed. The NIST ATM simulator was introduced and shown to be useful in studying congestion control algorithms, but also for the study of networking generally. We have argued that the simulator may be a valuable tool for undergraduate computer science education.

Bibliography

1 – ATM Forum, www.atmforum.com. ATM Forum, Worldwide Headquarters, Mountain View, CA 94040-1313

2 – High Performance Communication Networks, J. Walrand, P. Varaiya, Morgan Kaufman, 1996

3 – ATM/HFC Network Simulator, Operation and Programming Guide, Ver 4.0, N. Golmie, F. Mouveaux, et.al., U.S. Department of Commerce, Technology Administration, National Institute of Standards and Technology, Gaithersburg, MD 20899, December, 1998

4 - ATM Angles for IP Edge, D. Pappalardo, Network World, 16, 7, 1999, pp 1 and 61.

5 – Source Behavior for ATM ABR Traffic Management: An Explanation, R. Jain, et. al., IEEE Communication Magazine, 34, 11, pp 50-57, 1996.

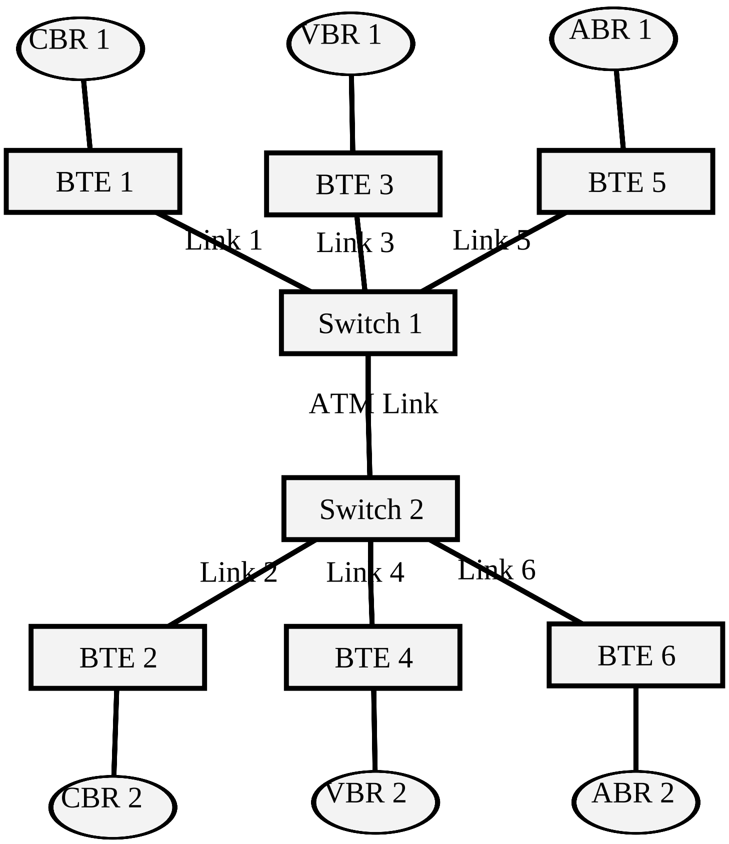

Figure 1: A sample simulator configuration

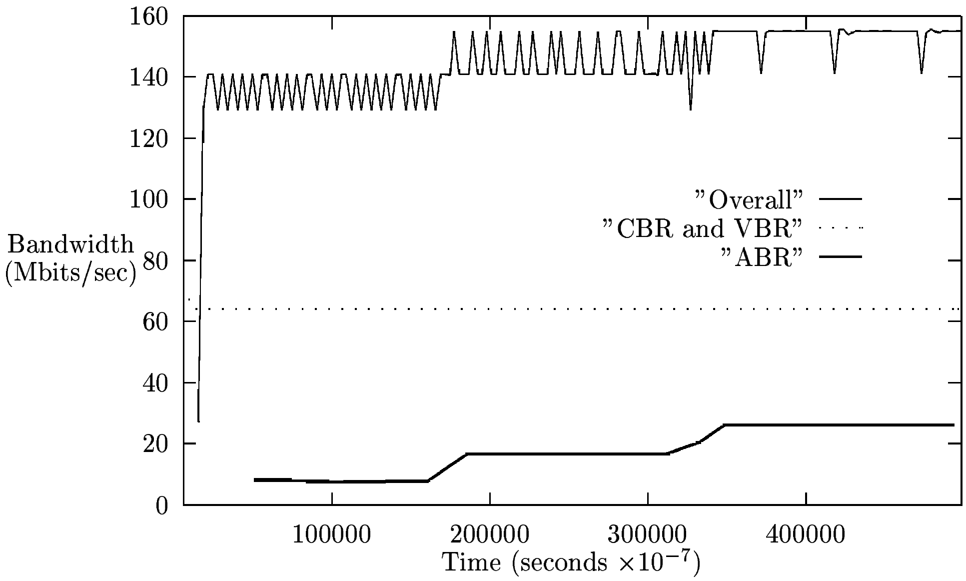

Figure 2: The ABR traffic increases until it saturates the network.

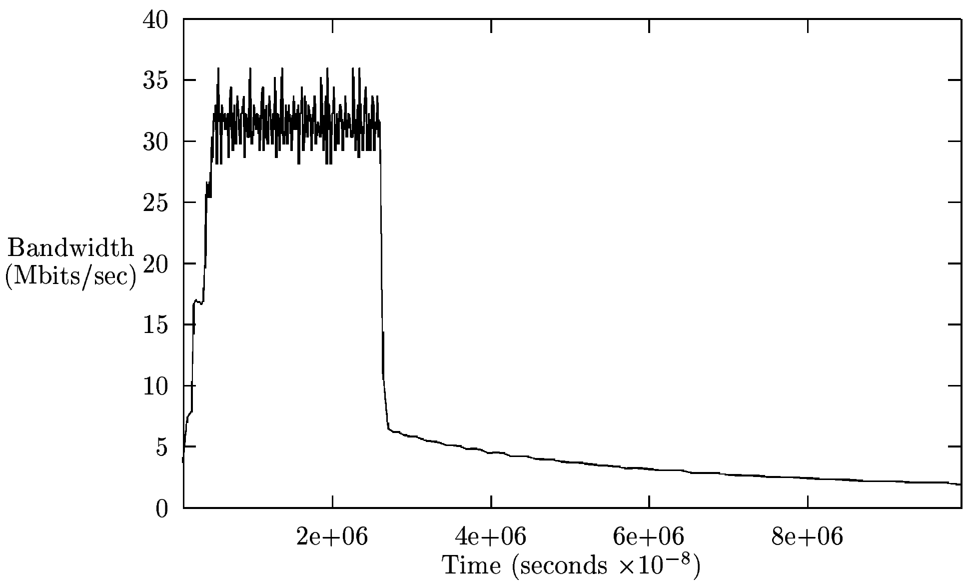

Figure 3: ABR traffic decreases monotonically

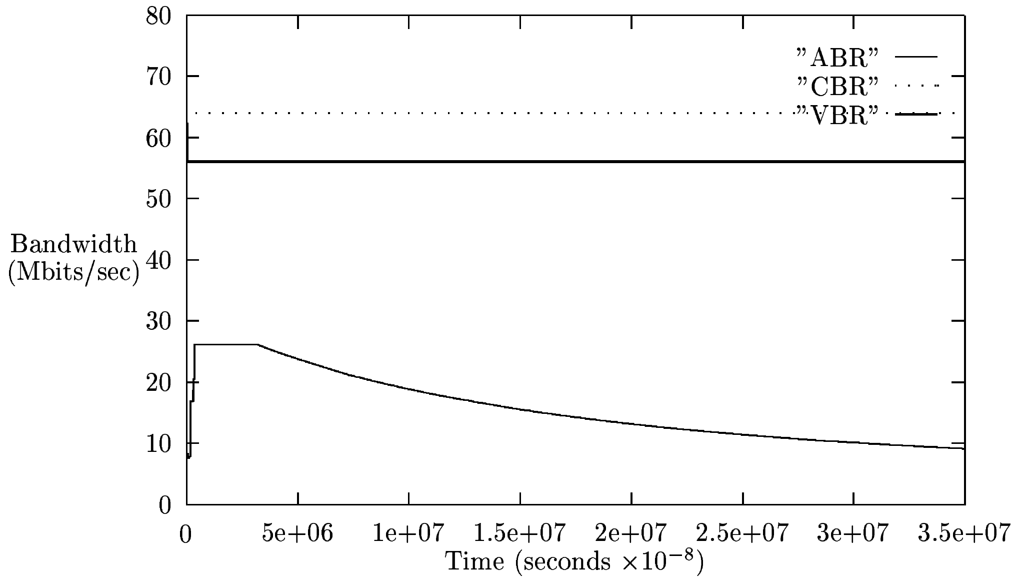

Figure 4: The NIST ER algorithm. ABR traffic still decreases monotonically.