Performance Evaluation of ATM Shortcut Connections in Overlaid IP/ATM Networks *

Victor Firoiu — vfiroiu@cs.umass.edu
Jim Kurose — kurose@cs.umass.edu
Don Towsley — towsley@cs.umass.edu

Department of Computer Science
University of Massachusetts
Amherst, MA 01003-4610 USA

CMPSCI Technical Report TR 97-40

July 29, 1997

Abstract

In this paper we present methods to evaluate the benefit of using direct ATM connections (shortcuts) between IP nodes in IP over ATM networks, and we evaluate the benefit of ATM shortcuts for several networks. We model an IP/ATM network with and without ATM shortcuts as two loss networks. We propose a metric, the Network Load Ratio, for network performance comparison, that gives the ratio of the number of flows accepted by two networks at the same network blocking probability. We derive an estimator of this metric, the Asymptotic Load Ratio, that has low computational complexity. We use this estimator as a basis for a methodology for network performance comparison, and use it to evaluate the benefit of ATM shortcuts in several concrete scenarios.

Keywords: IP/ATM networks, ATM shortcuts, network performance comparison, loss networks, network blocking probability.

*This material is based upon work supported in part by the National Science Foundation under Grant NCR-95-08274 and CDA-95-02639. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

1 Introduction

The exponential growth of interconnected data networks such as the Internet is probably the most important development in the telecommunication sector in the current decade. The important improvement in capacity of telecommunication networks through fiber optic and ATM technology has been an essential contribution to Internet's transition to a mass scale world network. Today, at most of its structural levels (Backbone, WAN, regional Internet Service Providers), the Internet is carried over telecommunication networks that are increasingly deploying ATM technology (e.g., NSF's very High Speed Backbone Network Service, vBNS, [5]). In such cases, the IP networks are configured as virtual networks over the ATM network infrastructure. The result, overlaid IP/ATM networks, feature a virtual separation between the IP and ATM networks.

The recent development and deployment of network layer ATM protocols such as ATM signaling [12] and PNNI [2] makes possible an interaction between IP and ATM networks. Recently, research has put forward several proposals to improve the efficiency of IP/ATM networks through an interaction between IP and ATM network layers. One proposal [10] is to *cut through* the IP processing of IP packets and transmit the data of an IP flow through a separate ATM virtual circuit (VC). The VC connects two IP routers at the edge of the ATM network, and the route of the ATM VC follows the IP route. Another set of proposals goes one step further and lets the ATM VC, called *ATM shortcut*, connecting two IP routers, be routed by the ATM network. ATM shortcuts are proposed by NHRP [9] for unicast IP flows, MARS [1] for multicast IP flows, and RSVP/ATM [3] for IP flows with Quality of Service (QoS) requirements. IP cut-through has proved to be beneficial in experimental systems and commercial products by increasing the average throughput of IP routers embedded in ATM networks. Additional performance improvements are possible in the ATM shortcut operation mode since routing IP flows in the ATM network provides the opportunity for a more efficient network resource utilization than the IP cut through mode. However, there has not yet been any theoretical or practical confirmation of the benefit of ATM shortcuts. In this paper we address the following question:

Given an overlaid IP/ATM network, what is the benefit of ATM shortcuts?

This question is important in IP/ATM network design, where ATM shortcut benefits can be weighed against the simplicity of a separated IP/ATM model. The purpose of this paper is to develop a methodology for evaluating ATM shortcuts, given an IP/ATM network. We find the benefits of ATM shortcutting to vary considerably depending on the interplay between the IP and ATM network topologies.

To answer the above question, we first model an IP/ATM network, with and without ATM shortcuts, as two distinct network graphs. We then propose a metric to compare the performance of the two networks: the Network Load Ratio. For a given IP/ATM configuration, the Network Load Ratio represents the additional load the IP/ATM network can accept on average when using ATM shortcuts (compared with the "IP cut through" case), for the same network blocking probability. We develop a method to compute the Network Load Ratio for a given network blocking probability based on the Fixed Point Method [11]. We also develop the notion of Asymptotic Load Ratio which is an approximation for the Network Load Ratio when the networks operate in underload conditions. We find experimentally through simulation that this is a good approximation (less that 12% error) for the Network Load Ratio for a range of network blocking probability (less than 0.01), thus giving a general indication of the relative performance of the two networks. This second method is very attractive due to its low computational complexity. We use this method to evaluate the benefit of ATM shortcuts in several concrete scenarios.

The paper is organized as follows. In Section 2 we provide a short background on overlaid IP/ATM networks and motivate the problem of evaluating the benefits of ATM shortcut. In Section 3 we propose a metric for network performance comparison, the Network Load Ratio and its limiting value for underloaded conditions, the Asymptotic Load Ratio, and two methods for their approximate computation. In Section 4 we evaluate the accuracy with which the Asymptotic Load Ratio estimates the Network Load Ratio. In Section 5 we use the Asymptotic Load Ratio to evaluate the benefit of ATM shortcuts for several IP/ATM networks. Section 6 concludes the paper.

2 Background and Motivation

2.1 IP over ATM

IP and ATM networks are increasingly coexisting on the same medium as overlaid networks. It is unlikely that either IP or ATM will disappear in the foreseeable future; rather both architectures are expanding and currently provide different types of service. IP is the established protocol for data communication networks, such as the Internet. ATM is progressively deployed in telecommunication networks that carry mainly voice traffic (for example ATT and MCI have had ATM backbones for years). Data networks using the IP protocol are currently carried over the telecom networks, but the two networks are virtually separated (such as NSF's very High Speed Backbone Network Service, vBNS [5]). It is possible, however, for IP to become aware of the underlying ATM infrastructure and to be able to take advantage of the capabilities of the underlying ATM network. Such IP-ATM interaction is investigated in the next section.

2.2 Issues in IP-ATM interaction

We consider IP flows transmitted on IP/ATM overlaid networks. We define an IP flow to be (see e.g. [10]) a sequence of IP packets having the same (IP source address, IP source port, IP destination address, IP destination port), and being transmitted within a determined time interval. We will not discuss the issue of determining the bounds of this time interval; we assume that they are either specified by the user application generating the flow, or are implicitly determined by the network (for example, the flow starts when the first packet is transmitted and the flow ends after a given silence period). By this definition, every IP packet belongs to an IP flow (see [10] for a discussion of this issue).

Consider now an ATM network (Fig. 1) where ATM switches are connected through ATM links, and some of the switches are also IP routers. The overlaid IP network consists of these IP routers connected through IP "links" that are ATM VC connections. There are three modes of IP-ATM interaction proposed for such overlaid IP/ATM networks, each offering a different type of service to IP flows.

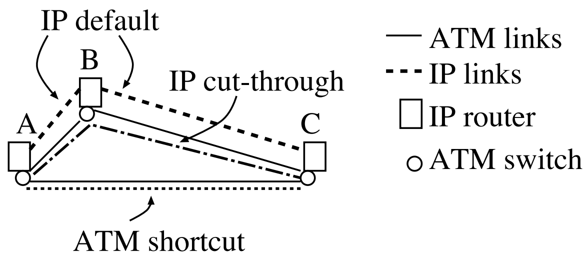

Figure 1: The three IP/ATM types of service for IP flows

In the simplest mode, the IP routers are connected through fixed ATM VCs that are IP links in the overlaid IP network. Data packets of an IP flow are routed and forwarded by the IP routers over these VCs. This type of service is called *IP default* (see Figure 1). For this service, no ATM connection setup is necessary. The overhead of IP packet processing is not significant for short lived flows due to the small number of packets in these flows.

For any IP flows, the overhead of IP packet fragmentation/reassembly and processing in each IP router constitutes a bottleneck in their transit. This bottleneck is a per-packet overhead, and its cumulative detrimental effect on the flow's throughput has a higher impact on long lived flows. [10] proposes to *cut through* the IP processing, once the IP route is known (see Figure 1). Separate VCs are created for a given IP flow between the adjacent IP routers on the flow's path and are spliced, thus removing the IP processing, and the ATM cells are forwarded on an end-to-end basis. The resulting improvement in data throughput has proven to be significant in practice for certain classes of IP flows. However, this IP/ATM operation mode does not take full advantage of the capabilities of the ATM network. The ATM VCs are constrained to follow a pre-established IP route, which might not efficiently use the resources of the ATM network.

Recently, several works [9, 1, 3] have proposed using the routing capability of ATM to benefit IP flows. An ATM VC established for an IP flow, called an *ATM shortcut*, is routed by the ATM network, and is not constrained by the default IP route (see Figure 1). NHRP and MARS allow IP applications to directly connect to an ATM network, and provide them the ATM address information required to establish ATM connections for unicast and multicast IP flows respectively. RSVP/ATM can be used to provide QoS guarantees to QoS IP flows that span IP/ATM networks. The QoS information included in RSVP messages is used to establish an ATM VC with required QoS guarantees. RSVP messages are also extended to convey the ATM addresses necessary to establish the VC.

So far, a significant amount of work has been done to specify the details of the above ATM shortcutting protocols, but the benefit of ATM shortcuts has not been quantitatively evaluated. No practical confirmation of the benefit exists either, as the protocols are still in the development phase. The main benefit of ATM shortcuts over IP cut-through comes from the better utilization of network resources when IP flows are routed in ATM, as can be seen in the following.

Consider a simple example in Fig 2. A flow with source A and destination C has route $A-B-C$ in IP and $A-C$ in ATM. If the flow requires 1 unit of bandwidth, then the network bandwidth consumption by the

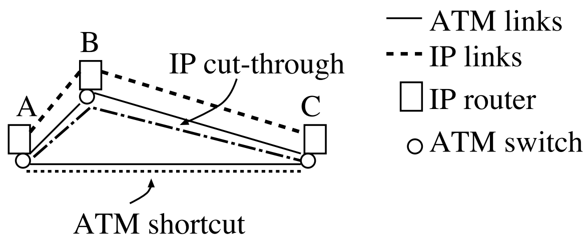

Figure 2: ATM shortcut versus IP cut through

given flow is 2 units when using IP routing and 1 unit when using ATM routing. It follows that ATM routing can yield a lower bandwidth consumption. This in turn results in a lower network blocking probability, or equivalently, more flows can be admitted for a given network blocking probability.

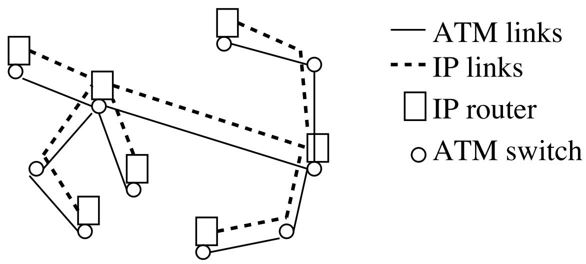

Figure 3: IP/ATM network with no ATM shortcut benefit

The benefit of ATM shortcuts depends on the topology of IP and ATM networks. Consider, for example, the network in Figure 3, where both IP and ATM networks have a tree structure. In this case, an IP flow between any pair of IP nodes has the same route in both IP and ATM network. Thus, for a given IP flow, the IP cut-through VC is identical to the corresponding ATM shortcut. We conclude that ATM shortcutting brings no benefit over IP cut-through in this network.

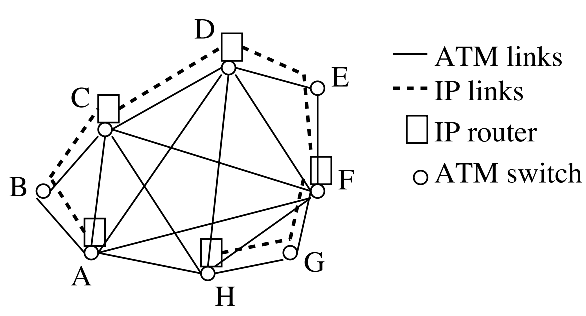

Figure 4: IP/ATM network with significant ATM shortcut benefit

At the other end of the spectrum, consider the example in Figure 4, where the IP network is a tandem (the IP nodes are arranged in series), and the ATM network is fully connected (forms a complete graph).

Given an IP flow between A and H, its IP route crosses seven ATM links, whereas its ATM route has only one link. In this case we expect a significant benefit of ATM shortcutting over IP cut-through.

Clearly, an evaluation of the benefit of ATM shortcuts is needed, based on the characteristics of the IP and ATM networks. Specifically, a systematic method is needed to identify the cases where ATM shortcuts are beneficial and to quantify such a benefit. The benefit of routing IP flows in an ATM network (ATM shortcuts) enables the network to admit more flows for the same network blocking probability. In the following, the evaluation of ATM shortcuts consists of comparing the flow loads offered to two networks having the same blocking probability. For the purpose of this comparison, the IP and ATM networks are considered separately, being characterized by their respective topologies and link capacities.

3 A Metric and Two Methods for Network Performance Comparison

In this section we propose a metric to compare the performance of two loss networks, and two methods for estimating it. The analysis in this section relies on the theory of loss networks; see [14, 6, 7, 11] and references therein.

We consider a network consisting of a set of nodes $\mathcal{V}$ and a set of links $\mathcal{L}$. A link $l \in \mathcal{L}$ has capacity C_l. The flows that arrive to the network are classified according to their source-destination pair of nodes, called *access nodes*, and to the type of flow (IP or ATM). All flows of class $k \in \mathcal{K}$ are routed on route r_k, where r_k is the sequence of links traversed by flows of class k. All flows have the same bandwidth requirement, irrespective of class. A flow of class k is admitted iff one unit of capacity can be reserved at each link $l \in r_k$. We define $\mathcal{R} = \{r_k | k \in \mathcal{K}\}$ to be the set of routes in the given network.

Flows arrive to the network according to a Poisson process with rate λ and have exponential holding times with rate μ. The intensity of offered load is $\rho = \lambda/\mu$. With probability ν_k a flow becomes a class k flow. The probability that a flow of class k is blocked (not admitted) is denoted by B_k, and the probability that a flow of any class is blocked at link j is denoted by L_j. The probability that a flow of any class is blocked, called the network blocking probability, is denoted by P, and is given by:

$$P = \sum_{k \in \mathcal{K}} \nu_k B_k = P(\rho) \tag{1}$$

where P depends on ρ through B_k. Let us denote

$$\rho_k = \nu_k \rho \quad \forall k \in \mathcal{K} \tag{2}$$

the intensity of offered load of class k. We also introduce the notations:

$$\rho_i = \sum_{k \in \mathcal{K}, i \in r_k} \rho_k \quad \forall i \in \mathcal{L} \tag{3}$$

and

$$\nu_i = \sum_{k \in \mathcal{K}, i \in r_k} \nu_k \quad \forall i \in \mathcal{L} \tag{4}$$

Let us now consider two networks N_1 and N_2 with link sets $\mathcal{L}_1$ and $\mathcal{L}_2$ and route sets $\mathcal{R}_1$ and $\mathcal{R}_2$, that have the same set of access nodes (i.e., the same set of flow end points) and the same offered load probabilities $(\nu_k)_{k\in\mathcal{K}}$. We define the following metric for comparing the performance of the two networks.

Definition 1 *The* <u>Network Load Ratio</u> *of networks N_1 and N_2 at blocking probability p, $R(N_1, N_2, p)$, is defined to be the solution of R such that*

$$P_1(R\rho) = P_2(\rho) = p \tag{5}$$

Network N_1 is said to perform R times better than N_2 at blocking probability p.

In other words, if N_1 and N_2 have a Network Load Ratio of R, N_1 can be offered R times more load that N_2 and present the same network blocking probability p. Observe that, in general, the Network Load Ratio R depends on the network blocking probability p: $R = R(p)$.

The Network Load Ratio R for a given network blocking probability p is given by:

$$R = \frac{P_1^{-1}(p)}{P_2^{-1}(p)} \tag{6}$$

In general, the network blocking probability cannot be computed exactly due to its computational complexity [8], but a good approximation is given by the Fixed Point method, also known as Reduced Load approximation [7]. Following the Fixed Point method, an approximation of the blocking probability B_k of flows in class k, B_k^*, is given by:

$$B_k^* = 1 - \prod_{j\in r_k}(1 - L_j^*) \tag{7}$$

where L_j^*, an approximation of the blocking probability L_j at link j, is the solution of the following system of equations:

$$L_j^* = Er(\sum_{k\in\mathcal{K}, j\in r_k} \rho_k \prod_{i\in r_k - \{j\}} (1 - L_i^*), C_j) \quad j \in \mathcal{L} \tag{8}$$

where Er is the Erlang loss formula:

$$Er(\rho, C) = \frac{\rho^C/C!}{\sum_{n=0}^{C} \rho^n/n!} \tag{9}$$

It follows that an approximation for the network blocking probability is:

$$P^* = \sum_{k\in\mathcal{K}} \nu_k B_k^* \tag{10}$$

The system of equations (8) can be solved by repeated substitutions, and is known to converge in most practical cases, giving a good approximation for B_k. The inverse of the network blocking probability function, $\rho = P^{-1}(p)$, can be computed using any numerical method for approximate root computation (e.g., the secant method) in conjunction with the Fixed Point method. The procedure based on equations (2), (7)-(10) allows us to compute the Network Load Ratio of two networks that operate at a given network blocking probability. It constitutes the first method that we propose for network performance comparison.

A problem with this method is its rather high computational complexity. The computation of Network Load Ratio based on equations (6-10) has complexity $O(L^2CKFS)$, where L is the number of links, C is the maximum of link capacities, K is the number of flow classes, F is the number of iterations in the Fixed Point method, and S is the number of iterations in the secant method. This complexity arises from the fact that the computation in (8) is done for all L links, with the product in (8) being $O(L)$, the sum being $O(K)$, and the computation of Erlang function being performed in $O(C)$. From our experiments in Section 4, for an approximation error of 0.01, typical values for F are between 3 and 100, and for S between 10 and 50.

Another problem with this network comparison method is that it gives the Network Load Ratio R of two networks for a given blocking probability p. The point value R might not give much insight into the relative behavior of the two networks for a range of network blocking probability values p.

After a series of empirical experiments with various networks of different topologies and loads, we have come to the conclusion that, in general, the Network Load Ratio R exhibits little variation in a range of values for p that are of interest to us. For example in Figure 5 we plot the network blocking probability of four networks with 17 access nodes: the complete graph, tandem, star and a model of NSF Backbone (see Figure 10) as of 1995. All networks have a uniform distribution of flow probabilities: for each pair of nodes there is one flow class, k, and $\nu_k = \nu \quad \forall k \in \mathcal{K}$. We observe that the Network Load Ratio is almost constant for blocking probabilities in the range $[10^{-5}, 0.01]$. In the following we provide a formal motivation for the above empirical observation and derive a second method for network comparison.

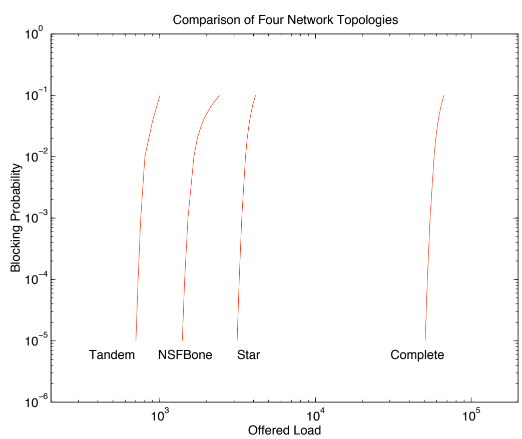

Figure 5: Comparison of four network topologies

Let us consider two networks N_1 and N_2 having all parameters defined above.

Definition 2 *The limit (if it exists) of the Network Load Ratio R for $\rho \to 0$,*

$$A = \lim_{\rho \to 0} R = \lim_{p \to 0} \frac{P_1^{-1}(p)}{P_2^{-1}(p)} \tag{11}$$

is called Asymptotic Load Ratio *of networks N_1 and N_2.*

Observe that A is a function of link capacities C_i. In Appendix A we show the following:

Proposition 1 *If all link capacities are the same in both networks, $C_i = C \quad \forall i \in \mathcal{L}_q, \quad q = 1, 2$, then*

1. *The Asymptotic Load Ratio is given by:*

$$A = (M_2/M_1)^{1+1/C} (\frac{n_2 + \sum_{i \in \mathcal{L}_2 - \mathcal{L}'_2} (\nu_{i2}/M_2)^{C+1}}{n_1 + \sum_{i \in \mathcal{L}_1 - \mathcal{L}'_1} (\nu_{i1}/M_1)^{C+1}})^{1/C} \tag{12}$$

where

$$M_q = \max_{i \in \mathcal{L}_q} \nu_{iq} \quad q = 1, 2 \tag{13}$$

$$\mathcal{L}'_q = \{i | i \in \mathcal{L}_q, \nu_{iq} = M_q\} \quad q = 1, 2 \tag{14}$$

$$n_q = |\mathcal{L}'_q| \quad q = 1, 2 \tag{15}$$

2. *The limit of the Asymptotic Load Ratio as $C \to \infty$ is:*

$$A^\infty = \lim_{C \to \infty} A = M_2/M_1 \tag{16}$$

Proposition 1 yields a simple approximation of the Asymptotic Load Ratio when link capacities are large ($C > 100$, as in most current networks).

For the general case, where the values of C_i are not restricted, we make the following conjecture:

Conjecture 1 *The Asymptotic Load Ratio is well approximated by:*

$$A^\infty = M_2/M_1 \tag{17}$$

where

$$M_q = \max_{i \in \mathcal{L}_q} \nu_i / C_i \quad q = 1, 2 \tag{18}$$

In Section 4 we present a set of experiments where we find the above expressions to be good approximations for the Asymptotic Load Ratio.

The computation of Asymptotic Load Ratio given in equations (16),(13), (17) and (18), constitutes our second proposed method for network comparison. Note that the complexity of the computation is $O(LK)$ where L is the number of links in the two networks and K is the number of flow classes. This makes the computation of Asymptotic Load Ratio much simpler than that of Network Load Ratio.

4 Accuracy of Asymptotic Load Ratio

In this section we evaluate the accuracy of Asymptotic Load Ratio. First, we verify that the Asymptotic Load Ratio, computed as in (16,13) and (17,18), is the limit of Network Load Ratio as the network blocking probability approaches zero. Second, we seek to determine the range of network blocking probability values for which the Asymptotic Load Ratio is a good approximation for the Network Load Ratio. Third, we measure and compare computation times for Network Load Ratio and Asymptotic Load Ratio.

We perform the evaluation through a set of simulation experiments, where the two methods are applied to pairs of randomly generated network topologies. We use the networks generated at Georgia Tech using the methods proposed in [13], here named RAND, and in [15], named Transit-Stub (TS).[1] The RAND networks have unstructured topologies, whereas the TS networks exhibit hierarchical structures and smaller diameters, which are claimed [15] to more accurately resemble real networks. All networks contain 100 nodes. All links in the RAND networks have OC3 (155 Mb/s) capacity, whereas only the backbone links in the TS networks have OC3 capacity and the rest, T3 (45 Mb/s) capacity. The bandwidth of a flow is 100 Kb/s. The shortest path routing policy is used to generate the routes.

In the first experiment, 10 networks of each network type were considered, and compared to each other using the two comparison methods. Specifically, the Network Load Ratio $R_{N_i,N_j}(p)$ is computed for each pair of networks N_i, N_j and for the following values of network blocking probability:

$$p \in \{10^{-5}, 10^{-4}, 10^{-3}, 10^{-2}, 0.02, 0.03, 0.04, 0.05, 0.07, 0.1\} \tag{19}$$

For each network pair, the Asymptotic Load Ratio A_{N_i,N_j} is also computed. We then compute the relative error of Asymptotic Load Ratio compared to Network Load Ratio for each network blocking probability value and each pair of networks:

$$e_{N_i,N_j}(p) = \frac{A_{N_i,N_j} - R_{N_i,N_j}(p)}{R_{N_i,N_j}(p)} \tag{20}$$

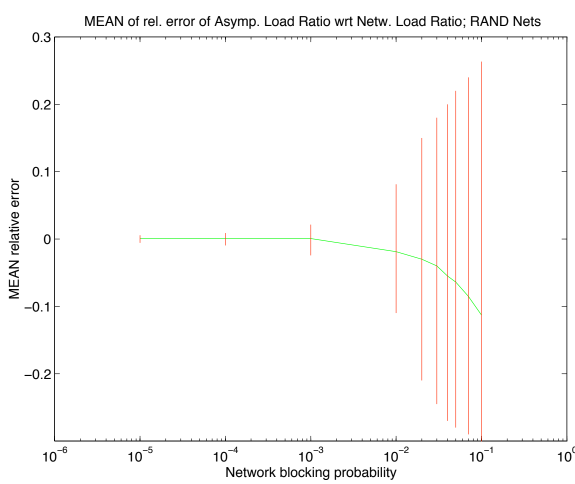

Figure 6: Mean of relative error of Asymptotic Load Ratio with respect to Network Load Ratio; RAND networks

In Figures 6 and 7 we plot, for each value of network blocking probability p, the average of the relative error values $e_{N_i,N_j}(p)$, over the set of all pairs (N_i, N_j) of networks. We also plot a vertical bar indicating

[1]The networks and the generating code can be found at http://www.cc.gatech.edu/fac/Ellen.Zegura/graphs.html

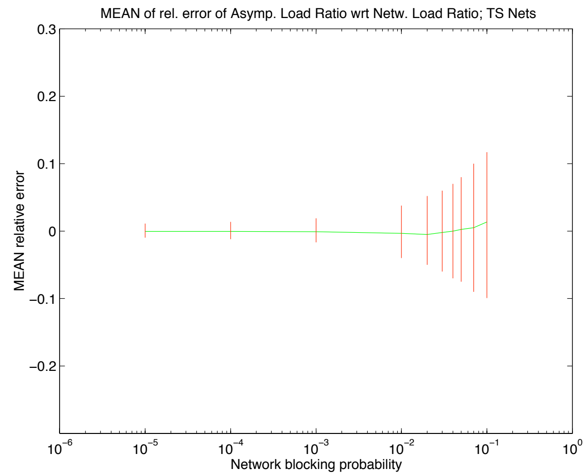

Figure 7: Mean of relative error of Asymptotic Load Ratio with respect to Network Load Ratio; TS networks

$\max i, j e_{N_i,N_j}(p)$ and $\min_{i,j} e_{N_i,N_j}(p)$. From the two graphs we observe that the relative error converges to zero as $p \to 0$, which confirms Proposition 1 and Conjecture 1 that the Asymptotic Load Ratio is the limit of Network Load Ratio.

Second, we note that the relative error is less than 0.12 for a network blocking probability less than 0.01, indicating that the Asymptotic Load Ratio is a good estimate for the Network Load Ratio in this range.

Last, we remark that the approximation is better for TS networks than for RAND, which suggests that the Asymptotic Load Ratio has more potential for "real" networks.

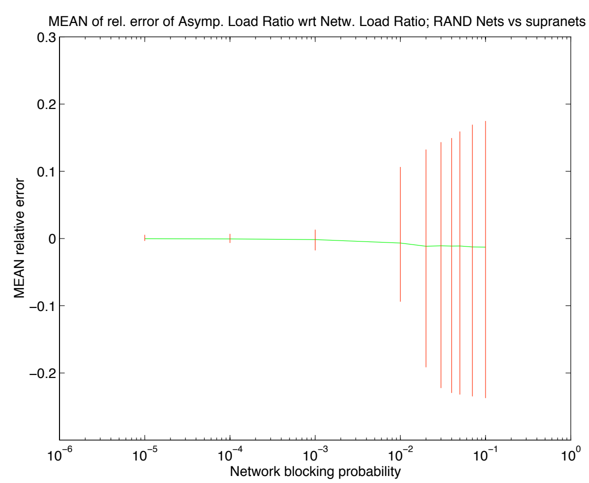

Figure 8: Mean of relative error of Asymptotic Load Ratio with respect to Network Load Ratio; RAND networks versus supranets

In the second experiment we treat each of the networks considered in the first experiment as a base network. Starting from each base network we create 10 "supranets" by adding a random number of links

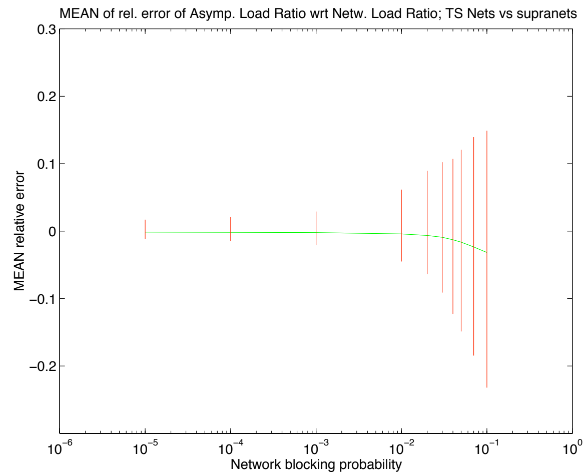

Figure 9: Mean of relative error of Asymptotic Load Ratio with respect to Network Load Ratio; TS networks versus supranets

(between 1 and 10). We then compute the Network Load Ratio and Asymptotic Load Ratio for every pair of base network and each of its supranets. The relative error statistics are processed as in the first experiment and displayed in Figures 8 and 9. We obtain similar results as before, the relative error converges to zero, and the Asymptotic Load Ratio is within 10% of Network Load Ratio for $p < 0.01$.

Following the above experiments, we conclude that the Asymptotic Load Ratio is a good approximation of the Network Load Ratio for $p < 0.01$, and thus is a good metric for network comparison.

For the above experiments we have also recorded the time required to compute the Network Load Ratio and Asymptotic Load Ratio for each pair of networks compared. The average computation time has been measured with the Unix code profiler *prof* on an 100MHz SGI workstation. Table 1 displays the average computation times for RAND and TS networks respectively. The difference in computation time in favor of Asymptotic Load Ratio confirms its simplicity predicted at the end of Section 3.

	RAND	TS
Network Load Ratio	103.927 s	129.676 s
Asymptotic Load Ratio	0.008 s	0.012 s

Table 1: Comparison of computation times for Network Load Ratio and Asymptotic Load Ratio

In conclusion, we consider that the comparison of network performance based on Asymptotic Load Ratio is preferable to the Network Load Ratio method for two reasons:

- The Asymptotic Load Ratio is empirically found to give a good approximation (within 0.1 relative error) for the relative performance of two networks (the Network Load Ratio) for a range of network blocking probability values ($p \in (0, 0.01)$).

- The Asymptotic Load Ratio is very simple computationally (four orders of magnitude) compared to the Network Load Ratio method.

5 Evaluating the Benefit of ATM Shortcuts

Informally, we have seen in Section 2.2 that the relative performance of two networks depends on their characteristics (topology, link capacities, routing). Thus, it is not possible to provide an answer to the question whether ATM shortcuts are beneficial in general. In this section we use the Asymptotic Load Ratio to evaluate the benefit of ATM shortcuts in the case of two concrete networks. Considering each network as an IP network, we compute the maximum benefit that it is possible to achieve from ATM shortcuts. We do this by allowing the ATM network to be fully connected. We also evaluate the benefit of ATM shortcuts for two concrete IP/ATM network configurations.

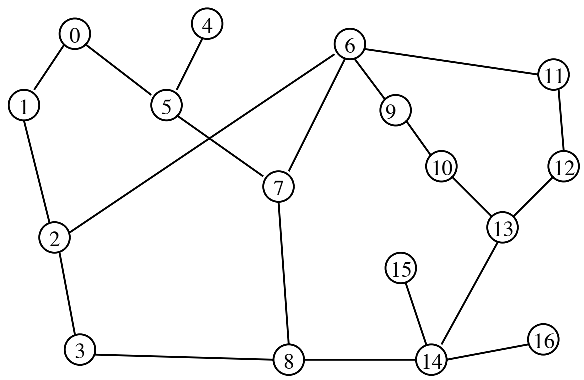

Figure 10: Topology of N17 (a model of NSF Backbone)

The first network, denoted by N17 is a model of the NSF Backbone as of 1995 (see Figure 10), having 17 nodes and links with OC3 (155 Mb/s) capacity. The flow load probabilities $(\nu_k)_{k\in\mathcal{K}}$ are the same for all pairs of nodes, all flows have bandwidth 100Kb/s, and the routing policy is shortest path (number of hops). The second network, denoted by N16, has 16 nodes, and is an abstraction of a commercial telecommunication network, where the topology, link capacities, link costs and flow load probabilities (only class 1) are given in [4]. The topology of N16 is shown in Figure 11.

In the first experiment, we treat N17 as an IP network that can be embedded in an ATM network (supragraph) having any number of nodes and links, with all link capacities being OC3. We ask the following question:

> "What is an upper bound on the performance improvement (Asymptotic Load Ratio) when the IP flows are routed in any such ATM network (supragraph) compared to the original network ?"

Given that we assume that both IP and ATM networks use fixed routing (i.e., not alternate routing), an upper bound is obtained from an ATM network that is a complete graph with 17 nodes, since link capacities are all limited to the same value (OC3). The computed value of Asymptotic Load Ratio is $A = 37$. This value

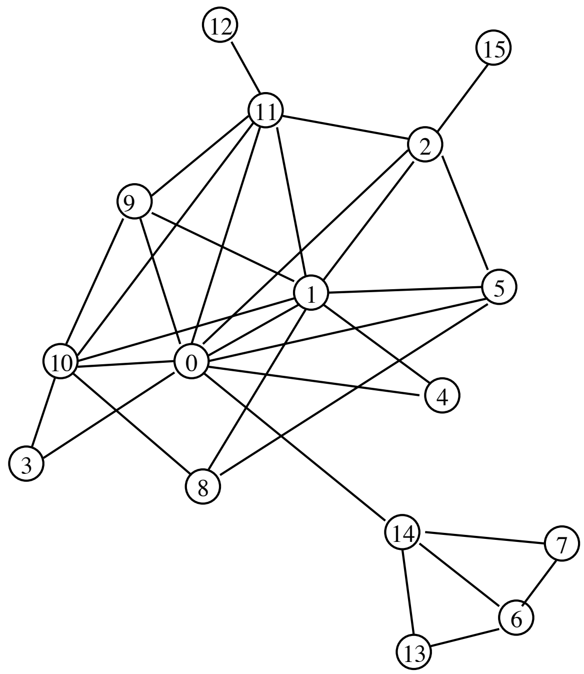

Figure 11: Topology of N16

indicates that N17 has a good potential for performance improvement when using ATM shortcuts. A graphic comparison between the performances of N17 and the complete graph with 17 nodes can be seen in Figure 5.

We perform the same comparison with N16, namely we compute the Asymptotic Load Ratio between the Complete graph with 16 nodes and all link capacities $C_l = 120$, and N16, while both having the same load probabilities, $(\nu_k)_{k\in\mathcal{K}}$ (given in Table 14 in [4]). The result, $A = 1.311$, shows a much smaller potential for performance improvement when using ATM shortcuts than N17. A qualitative explanation for this is that N16 is in some sense "more connected" than N17, the nodes have a larger degree, and the routes between two nodes are shorter; in a word N16 is closer to a Complete graph than N17.

In the second experiment, we look at the incremental gain as a few links and nodes are added to an IP network to form the embedding ATM network. We compute the maximum, minimum and average Asymptotic Load Ratio as 1,2,3,4 links and 0,1,2,3 nodes are added respectively, in all possible positions permitted by the base network topologies N17 and N16.[2] For N17 the capacity of added links is OC3, and for N16 is 120 units, which is the average of N16's link capacities. In Tables 2 and 3 we display the Asymptotic Load Ratio values for network N17 and N16 respectively.

First, by looking at the minimum values for Asymptotic Load Ratio in both tables, we can see that ATM shortcuts are not always beneficial. Since these values are smaller than 1, it follows that there are cases where ATM shortcuts are detrimental. This deterioration in network performance upon adding links and nodes occurs when the routing in the new topology concentrates a larger amount of traffic (ν_l) on some links l. Second, since the average values are all larger than 1, we might say that ATM shortcuts are beneficial "on

[2] We observe here that this exhaustive study was made possible by the low computational requirement of Asymptotic Load Ratio.

No. shortcuts	Maximum	Minimum	Average
1	1.541667	0.880952	1.135392
2	1.850000	0.740000	1.237412
3	2.055556	0.698113	1.325446
4	2.312500	0.685185	1.407431

Table 2: Incremental gain in N17

No. shortcuts	Maximum	Minimum	Average
1	1.143963	0.661319	1.056476
2	1.144125	0.317376	1.073013
3	1.144125	0.236685	1.064760
4	1.144125	0.234079	1.042132

Table 3: Incremental gain in N16

average". Third, we observe that the maximum increase in performance does not grow at the same rate as links and nodes are added. We expect that there is a threshold in the number of links added, above which little improvement is made, and that the threshold is well below the number of links of the complete graph. For example, N16 is likely to have reached this threshold at two links since the maximum value for A does not increase in Table 3 when adding more than two links, and because the upper bound on A's improvement is 1.311 (computed in the first experiment of this section) which is close to the maximum values in Table 3. On the other hand, for the N17 network, the maximum value of A does not stop increasing as more links are added. It concurs with the high upper bound of 37 computed in the previous experiment in predicting a significant ATM shortcut benefit for N17.

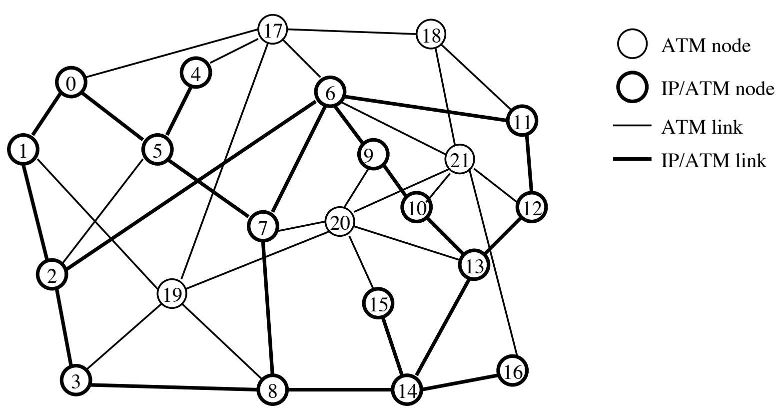

Figure 12: N17 embedded in an ATM network

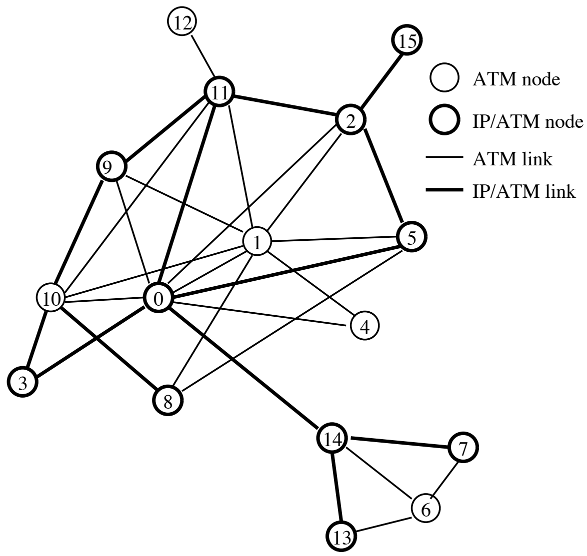

Figure 13: N16 embedding an IP network

In the last experiment, we first consider N17 as an IP network embedded in an ATM network given in Figure 12. We consider a set of loads for IP traffic with all load probabilities $(\nu_k)_{k\in\mathcal{K}}$ of the same value, having all IP nodes as access nodes. We compare the network performance when IP flows are routed in IP and ATM topologies respectively. We find that the Asymptotic Load Ratio is $A = 1.681$, which translates into more than 50% more IP and ATM flows being admitted when IP flows are routed in the ATM network.

Next, we consider N16 as an ATM network that embeds an IP network, as in Figure 13. The set of loads of IP traffic are a subset of the loads of class 1 given in [4] that have access nodes that are IP nodes (depicted with thicker lines in Figure 13). We compare again the network performance when the IP flows are routed in IP and ATM topologies respectively. We find that the Asymptotic Load Ratio is $A = 1.896$, confirming again that using ATM shortcuts in a more connected ATM network topology is likely to present benefits for IP flows.

6 Conclusion

In this paper we have considered the problem of evaluation of benefit of ATM shortcuts in IP/ATM networks, i.e., the benefit of ATM routing of IP flows in ATM networks. We proposed to measure the benefit of ATM shortcuts with the Network Load Ratio, that expresses the increase in the number of flows accepted by an IP/ATM network for the same network blocking probability. We developed a low complexity computation for Asymptotic Load Ratio, which estimates the Network Load Ratio relatively well (within 10% error) for underload network conditions (network blocking probability less that 0.01). We used the developed methodology to evaluate the benefit of ATM shortcuts for several concrete networks. The main result of this paper is a methodology for comparing network performance, which can be used to evaluate the benefit and tradeoff of ATM shortcuts, and in the more general context of network design.

We are currently working to extend the present work in several directions. First, we are attempting to prove the conjecture that extends the Asymptotic Load Ratio to networks with heterogeneous capacities. Second, the results and methods may also be extended to networks with alternate routing, flows with heterogeneous bandwidth requirements, and multicast flows.

A Derivation of Asymptotic Load Ratio

In this section we show the following:

Proposition 1 *If all link capacities are the same in both networks,* $C_i = C \quad \forall i \in \mathcal{L}_q, \quad q = 1, 2$*, then*

1. *The Asymptotic Load Ratio is given by:*

$$A = (M_2/M_1)^{1+1/C} \left(\frac{n_2 + \sum_{i \in \mathcal{L}_2 - \mathcal{L}'_2} (\nu_{i2}/M_2)^{C+1}}{n_1 + \sum_{i \in \mathcal{L}_1 - \mathcal{L}'_1} (\nu_{i1}/M_1)^{C+1}}\right)^{1/C} \tag{21}$$

where

$$M_q = \max_{i \in \mathcal{L}_q} \nu_i \quad q = 1, 2 \tag{22}$$

$$\mathcal{L}'_q = \{i | i \in \mathcal{L}_q, \nu_{iq} = M_q\} \quad q = 1, 2 \tag{23}$$

$$n_q = |\mathcal{L}'_q| \quad q = 1, 2 \tag{24}$$

2. *The limit of the Asymptotic Load Ratio as* $C \to \infty$ *is:*

$$\lim_{C \to \infty} A = M_2/M_1 \tag{25}$$

We start with the following statement given by Whitt in [14], Corollary 2.3, point (ii):

Proposition 2 *For any network,*

$$\lim_{\rho \to 0} L_i^*/Er(\rho_i, C_i) = 1 \quad \forall i \in \mathcal{L} \tag{26}$$

where

$$\rho_i = \sum_{k \in \mathcal{K}, i \in r_k} \rho_k \quad \forall i \in \mathcal{L} \tag{27}$$

The limit of class k flow blocking probability as $\rho \to 0$ is:

$$B_k^* = 1 - \prod_{i \in r_k} (1 - L_i^*) \to \sum_{i \in r_k} L_i^* \tag{28}$$

since $L_i^* \to 0$, and thus

$$B_k^* \to \sum_{i \in r_k} Er(\rho_i, C_i) \tag{29}$$

by Proposition 2. The limit of network blocking probability becomes:

$$\begin{aligned} P^* &= \sum_{k\in\mathcal{K}} \nu_k B_k^* \qquad \text{by (10)} && (30)\\ &\to \sum_{k\in\mathcal{K}} \nu_k \sum_{i\in r_k} Er(\rho_i, C_i) \qquad \text{by (29)} && (31)\\ &= \sum_{i\in\mathcal{L}} Er(\rho_i, C_i) \sum_{k\in\mathcal{K}, i\in r_k} \nu_k \qquad \text{by rearranging terms} && (32)\\ &= \sum_{i\in\mathcal{L}} \nu_i Er(\nu_i\rho, C_i) \qquad \text{by (4)} && (33) \end{aligned}$$

From the definition of Erlang formula (9) it is easy to see (for example [14], Corollary 2.3, point (i)) that, for $\rho \to 0$,

$$Er(\rho, C) \to \frac{\rho^C}{C!} \tag{34}$$

Then, for any two links $i, j \in \mathcal{L}$, we have:

$$\lim_{\rho\to 0} \frac{\nu_i Er(\nu_i\rho, C)}{\nu_j Er(\nu_j\rho, C)} = (\frac{\nu_i}{\nu_j})^{C+1} \tag{35}$$

Let M be the maximum link factor,

$$M = max_{i\in\mathcal{L}}\nu_i \tag{36}$$

let $\mathcal{L}'$ be the set of links with maximum link factor,

$$\mathcal{L}' = \{i | i \in \mathcal{L}, \nu_i = M\} \tag{37}$$

and let n be the number of links with maximum factor,

$$n = |\mathcal{L}'| \tag{38}$$

By factoring out the term with the maximum link factor in (33) and using (35), we have

$$P^*(\rho) \to M Er(M\rho, C)(n + \sum_{i\in\mathcal{L}-\mathcal{L}'} (\nu_i/M)^{C+1}) \triangleq \overline{P^*}(\rho) \tag{39}$$

Observe that $\nu_i/M < 1 \quad \forall i \in \mathcal{L} - \mathcal{L}'$.

Consider the networks N_1 and N_2 having all link capacities $C_{iq} = C \quad \forall q = 1, 2$. The limit of the Network Load Ratio R, $A = \lim_{\rho\to 0} R$, named Asymptotic Load Ratio, is given by the equation:

$$\overline{P_1^*}(A\rho) = \overline{P_2^*}(\rho) \tag{40}$$

equivalent to:

$$M_1 Er(M_1 A\rho, C)(n_1 + \sum_{i\in\mathcal{L}_1-\mathcal{L}_1'} (\nu_{i1}/M_1)^{C+1}) = M_2 Er(M_2\rho, C)(n_2 + \sum_{i\in\mathcal{L}_2-\mathcal{L}_2'} (\nu_{i2}/M_2)^{C+1}) \tag{41}$$

By replacing the Erlang formula with its limit from (34) we have:

$$M_1 \frac{(M_1 A\rho)^C}{C!}(n_1 + \sum_{i\in\mathcal{L}_1-\mathcal{L}_1'} (\nu_{i1}/M_1)^{C+1}) = M_2 \frac{(M_2\rho)^C}{C!}(n_2 + \sum_{i\in\mathcal{L}_2-\mathcal{L}_2'} (\nu_{i2}/M_2)^{C+1}) \tag{42}$$

It follows that

$$A = (M_2/M_1)^{1+1/C} \left(\frac{n_2 + \sum_{i \in \mathcal{L}_2 - \mathcal{L}_2'} (\nu_{i2}/M_2)^{C+1}}{n_1 + \sum_{i \in \mathcal{L}_1 - \mathcal{L}_1'} (\nu_{i1}/M_1)^{C+1}} \right)^{1/C} \tag{43}$$

We have

$$\lim_{C \to \infty} A = M_2/M_1 \tag{44}$$

since $\mu_{iq}/M_q < 1$ for $i \in \mathcal{L}_q - \mathcal{L}_q', q = 1, 2$.

References

[1] G. Armitage. Support for Multicast over UNI 3.0/3.1 based ATM Networks. Internet RFC 2022, November 1996.

[2] ATM Forum. ATM PNNI Draft Specification. ATM Forum 94-0471 R10, July 1995.

[3] A. Birman, V. Firoiu, R. Guérin, and D. Kandlur. Support for RSVP-based Services over ATM Networks. In *IEEE Global Internet*, 1996.

[4] A. Greenberg and R. Srikant. Computational Techniques for Accurate Performance Evaluation of Multirate, Multihop Communication Networks. *IEEE/ACM Transactions on Networking*, 1997.

[5] J. Jamison and R. Wilder. vBNS: The Internet Fast Lane for Research and Education. *IEEE Communications Magazine*, 35(1), Jan 1997.

[6] F. Kelly. Blocking Probabilities in Large Circuit-Switched Networks. *Adv. Appl. Prob.*, 18, 1986.

[7] F. Kelly. Loss Networks. *The Annals of Applied Probability*, 1(3), 1991.

[8] G. Louth. *Stochastic Networks: Complexity, Dependence and Routing*. PhD thesis, University of Cambridge, 1990.

[9] J. Luciani, D. Katz, D. Piscitello, and B. Cole. NBMA Next Hop Resolution Protocol (NHRP). Internet draft-ietf-rolc-nhrp-11.txt, 1997.

[10] P. Newman, T. Lyon, and G. Minshall. Flow Labelled IP: a Connectionless Approach to ATM. In *IEEE INFOCOM*, 1996.

[11] K. Ross. *Multiservice Loss Models for Broadband Telecommunication Networks*. Springer, 1995.

[12] P. Samudra. Draft of UNI Signaling 4.0. ATM Forum 94-1018 R4, July 1995.

[13] B. M. Waxman. Routing of Multipoint Connections. *IEEE Journal on Selected Areas in Communications*, 6(9):1617–1622, December 1988.

[14] W. Whitt. Blocking when Service is Required from Several Facilities Simultaneously. *AT&T Technical Journal*, 64(8), October 1985.

[15] E. Zegura, K. Calvert, and S. Bhattacharjee. How to Model an Internetwork. In *IEEE Infocom*, 1996.