Performance and Traffic Measurement Scenarios for the AAI Network

Luiz A. DaSilva
Hongbo Zhu
Victor S. Frost
Joseph B. Evans
Douglas Niehaus
David W. Petr

Technical Report TISL-10980-3

Telecommunication and Information Science Lababoratory
Department of Electrical Engineering and Computer Science
The University of Kansas

March 13, 1995

Project Sponsor:
Sprint Corporation
Under prime contract to
Defense Advanced research Projects Agency
Contract number DABT 63-94-C-0068

Abstract

This technical report provides details of the measurement scenarios to be used in the performance and traffic measurements for the AAI network. For each of the experiments described in [1, 2], we list the topologies in which the experiment will take place, as well as its predicted duration and whether the experiment uses in-service or out-of-service techniques.

For each of the performance measurements, we discuss qualitative and/or quantitative aspecs of the expected results. Whenever possible, we determine upper bounds for the parameters, at which performance would be acceptable for all services provided by the AAI network.

1 Introduction

This report can be seen as a plan for the performance experiments to be conducted on the AAI network. As a result of these experiments, we expect to quantify the characteristics of the AAI as well as identify potential bottlenecks that may impact network performance. As a final result, in the future we will suggest admission, flow and congestion control techniques that may be used to improve this performance.

The scope of this report is to identify the scenarios (topologies, duration, etc.) in which the experiments will be performed, as well as to discuss the results we expect to observe.

The report starts by re-introducing all the experiment topologies (slightly modified from [2]). In the following sections, metrics for each network layer (Application, TCP, AAL5, ATM) are discussed. Finally, we present preliminary results for

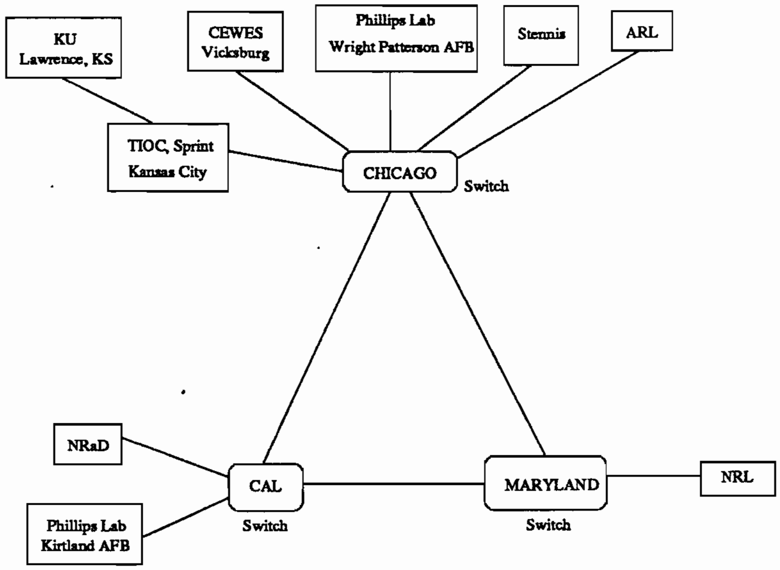

Figure 1: AAI Network Topology

the experiments which have already been performed and preview future research areas.

2 Description of Topologies

The AAI network topology can be described as in Figure 1. There are three Wide Area Network (WAN) switches connecting eight main sites.

The measurement sites include:

1. NRaD - San Diego
2. Phillips Lab - Kirtland AFB
3. Phillips Lab - Wright Patterson AFB
4. NRL - Stennis
5. CEWES - Vicksburg
6. NRL - Washington,D.C.

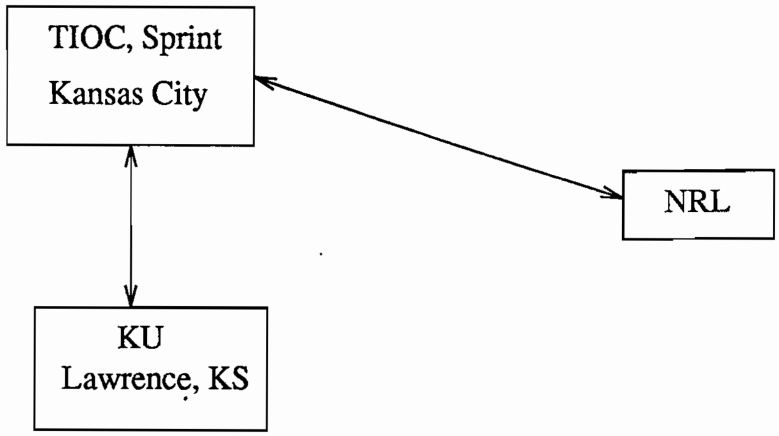

Figure 2: KU/TIOC/NRL Three-Site Configuration

7. ARL - Aberdeen, MD

8. TIOC, Sprint - Kansas City, KS

The University of Kansas site, shown in Figure 1, is the site that will originate the measurements. The scenarios in which these measurements will take place are described below.

2.1 KU/TIOC/NRL Three-Site Configuration

Traffic will be set up simultaneously among the KU, TIOC and NRL sites, in the configuration shown in Figure 2. As described in [1, 2], the NETSPEC application [3] will be used for testing this configuration. The effects of cell pacing and large TCP window sizes will be evaluated. Techniques to improve the throughput will be identified and recommended for use at other AAI sites.

2.2 Point-to-Point Full Duplex Configuration

There are 28 different possible pairs of point-to-point configurations in the AAI network. Performance testing will be done point-to-point separately in an underload state to determine the maximum attainable rates and other parameters.

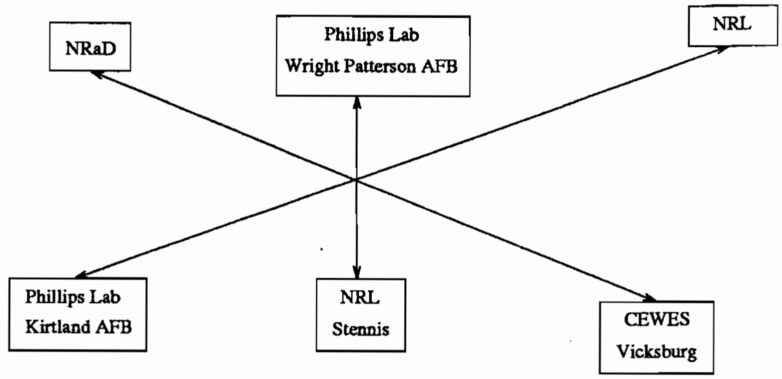

Figure 3: Three-Site Pair Full-Duplex Configuration

2.3 Three-Site Pair Full-Duplex Configuration

Traffic will be set up simultaneously between 3 pairs of measurement sites in the under-load state. There are a total of 28 such combinations; one of these combinations is shown in Figure 3.

2.4 Switch-by-Switch 2-way Full-Duplex Configuration

We will set up traffic between one given site and all other sites connected to the same switch, as shown in Figure 4. In this case, all traffic must go through only one switch. Therefore, there are five possible combinations.

In the current AAI topology (Figure 1), this scenario is only possible for the Chicago switch, since for the California switch it would be equivalent to the Point-to-Point Full Duplex Configuration, and there is only one main site directly connected to the Maryland switch.

2.5 Two-way Full-Duplex Configuration

Traffic will be set up between one site and every other site in the network in the under-load state. An example of this type of configuration is shown in Figure 5, where connections for all links are set up simultaneously.

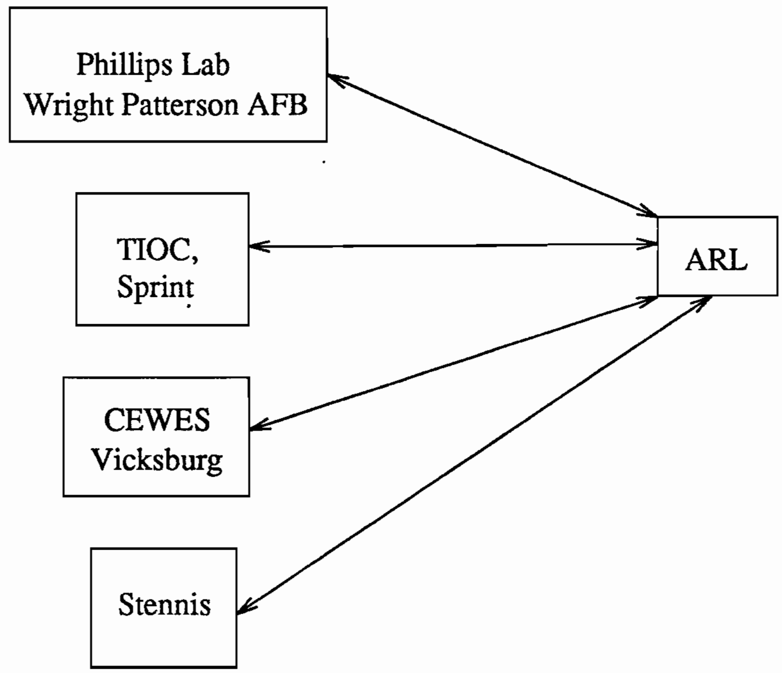

Figure 4: Switch-by-Switch Two-way Full Duplex Configuration

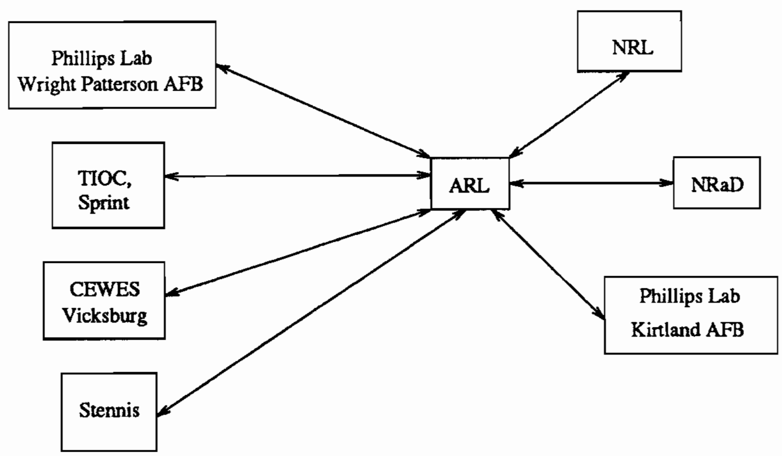

Figure 5: Two-way Full-Duplex Configuration

2.6 All-Site Full-Duplex Mesh Configuration

In this scenario, traffic will be set up simultaneously on all measurement sites. Each measurement site sends traffic to every other site, resulting in a mesh configuration. Notice this configuration is equivalent to performing all combinations of point-to-point testing simultaneously. Under this configuration, traffic will travel through all three switches.

Application Layer				
		Traffic Statistics	Performance Measurements	
		Frame Size	Frame Delay	Throughput
Topology	KU/TIOC/NRL	√	√	√
	Point-to Point		√	√
	Three-site Pair		√	√
	Switch-by-Switch		√	√
	2-Way		√	√
	All-site Mesh		√	√
Duration		5 min	5 min	5 min
In/Out of Serv		In	In	In

Table 1: Application Layer Scenarios

3 Application Layer Scenarios

Traffic statistics and performance measurements will be collected for the Application Layer. To the measurements described in [2], we have added the throughput of a given connection. Table 1 summarizes the scenarios in which the experiments will take place. These scenarios include the topology of the test network, the duration of the measurements, and whether they will be taken using in-service or out-of-service techniques.

The applications which will be used for these experiments include:

- Over-the-Shoulder (OTS) - remote screen viewing application developed at The University of Kansas as part of the MAGIC project;
- TerraVision
- EDC Image Visualization Application

As seen in Table 1, the experiment to determine the size of a frame in bytes will only be performed for one of the configurations, since the frame size is topology-independent.

On the other hand, the frame delay is expected to vary under different scenarios. As we increase the traffic competing for the same link, queueing effects will cause the frame delay to increase. In this fashion, supposing all other traffic in the network remains constant, we expect to measure the longest frame delay in the All-Site Mesh configuration and the shortest in the Point-to-Point configurations.

Since the AAI network must provide a wide variety of services, the frame delay must not exceed the highest acceptable delay for these services. The services which are most sensitive to end-to-end delay are telephony (transmission of speech) and distributed computing. We expect to verify acceptable delays are achievable for the services which we will be testing (high-speed data transmission, image applications, etc.).

We can obtain the throughput of a given connection directly from NETSPEC. We will transmit data using NETSPEC under two conditions: constant bit rate (CBR) and full speed. The throughput should remain constant while the network experiences acceptable loads and decrease rapidly at the onset of congestion. Between each of the WAN Switches and the main sites there exists a DS3 connection; the bandwidth of these connections is in the order of 34 Mbps. Therefore, we expect to measure throughput to be close to 34 Mbps as long as the network is not congested.

TCP Layer				
		Traffic Statistics	Performance Measurements	
		Packet Size	Packet Interarrival Time	Packet Error Ratio
Topology	KU/TIOC/NRL	√	√	√
	Point-to Point		√	√
	Three-site Pair		√	
	Switch-by-Switch		√	
	2-Way		√	
	All-site Mesh		√	√
Duration		5 min	5 min	30 min
In/Out of Serv		In	In	In

Table 2: TCP Layer Scenarios

4 TCP Layer Scenarios

The metrics that will be collected for the TCP Layer include Packet Error Ratio, Packet Interarrival Time, and Packet Size. Table 2 summarizes the scenarios in which these measurements will take place.

Again, the packet size is clearly topology-independent, so it will be determined using the KU/TIOC/NRL three-point configuration.

In the TCP layer, we define the Packet Error Ratio as:

$$PER = \frac{Number\ of\ Packets\ with\ Checksum\ Errors}{Number\ of\ Received\ Packets} \tag{1}$$

Considerations on lost and misinserted packets are left to the lower layers, as we will see in the following sections. This metric is not highly dependent on the network load; for this reason, we will run the experiment on only a few representative topologies.

On the other hand, the Packet Interarrival Time is highly dependent on the test topology. As the load on a particular segment of the network is increased, we expect to see a higher variance on the Interarrival Time. Among our chosen topologies, the highest variance should occur for the All-Site Mesh Configuration.

We should also bear in mind that the Packet Interarrival Time can be interpreted as an indication of delay jitter. If each packet is timestamped at the transmitter and we record the arrival times at the receiver, we can calculate the delay variance, independently of transmitter-receiver synchronization.

Let's suppose n packets are transmitted at times $\{t_{t1}, t_{t2}, \ldots, t_{tn}\}$ and received at times $\{t_{r1}, t_{r2}, \ldots, t_{rn}\}$. We can calculate the "relative delay" t_{di} as:

$$t_{di} = t_{ri} - t_{ti} \ , \quad 1 \leq i \leq n \tag{2}$$

Note that when the receiver and the transmitter are synchronized this relative delay is equal to the packet delay.

Then, the variance of the relative delay is a measure of the delay jitter. Notice that for this measurement of delay jitter, it is not necessary that the transmitter and receiver clocks be synchronized; therefore, no synchronization techniques, such as the Network Time Protocol (NTP), are assumed.

AAL 5 Layer					
		Performance Measurements			
		PDU Interarrival Times	PDU Error Ratio	BER	Throughput
Topology	KU/TIOC/NRL	√			√
	Point-to Point	√	√	√	√
	Three-site Pair				√
	Switch-by-Switch				√
	2-Way				√
	All-site Mesh	√			√
Duration		5 min	30 min	30 min	5 min
In/Out of Serv		In	In	Out	In

Table 3: AAL 5 Performance Testing Scenarios

5 AAL 5 Layer Measurement Scenarios

The measurement metrics for the AAL 5 include AAL 5 PDU interarrival times, AAL 5 PDU size, AAL 5 PDU error ratio, bit error ratio, and AAL 5 throughput. The PDU size measurement belongs to the traffic statistics collection category, while the rest of the metrics are for performance testing purposes.

Table 3 summarizes the scenarios for performance testing in the AAL 5.

Statistically, AAL 5 PDU size is topology-independent, so it will be measured using the KU/TIOC/NRL three-point configuration.

In the AAL 5, we define the PDU error ratio as:

$$PDU\ ErrorRatio = \frac{Number\ of\ PDU's\ with\ CRC\ Errors}{Total\ Number\ of\ Received\ PDU's} \tag{3}$$

It is expected that the PDU error ratio is not dependent on measurement topology. For this reason, we chose to test this metric under the point-to-point configuration only.

The bit error ratio (BER) is not dependent on test topology. Instead, it depends heavily on the quality of the hardware that the testing bit stream traverses. For this reason, we chose to test the BER metric under the point-to-point configuration only. Also notice that the BER measurement can only be tested by using an out-of-service technique. The reason is that the ATM protocol does not support bit error checking at the cell level. A bit stream with a known pattern will be sent from the transmitter, and the receiver will test for bit errors by comparing the received bit pattern with the known pattern.

Distinct services allow for different ranges of BER. The RACE (Research on Advanced Communication in Europe) consortium 1022 (Technology for ATM) considers bit error ratios in the order of 10^{-5} acceptable for remote process control and hifi sound, 10^{-6} for broadcast video, and 10^{-7} for data transmission and telephony [4].

AAL 5 PDU interarrival (inter-transmission) times can be an indicator of AAL 5 PDU delay jitter (variation). If we know the timestamps for both the PDU departure and arrival events for the same AAL 5 PDU source, the delay variation can be evaluated even though the network clocks may not be synchronized. The way of estimating the AAL 5 PDU delay jitter is similar to that of the TCP layer packet delay variation estimation in the previous section.

The PDU interarrival times are expected to be dependent on the test topology. As the load on a particular link of the network is increased, we expect to see a higher variance on the PDU delay. Among our chosen topologies, the highest variance should occur for the All-Site Mesh Configuration.

Throughput performance will be tested under all topological configurations by using the ATMSTAT program. We expect the throughput performance to be dependent on the measurement topology. The ideal throughput should be close to the capacity of the network links in use for any given topology.

ATM Layer							
		Performance Measurements					
		MCTD	2-Point CDV	CLR	CMR	CER	Throughput
Topology	KU/TIOC/NRL	√	√	√	√		√
	Point-to Point					√	√
	Three-site Pair	√	√	√	√		√
	Switch-by-Switch						√
	2-Way						√
	All-site Mesh	√	√	√	√		√
Duration		5 min	5 min	1 hr	1 hr	30 min	5 min
In/Out of Serv		In	In	In	In	Out	In

Table 4: ATM layer Performance Testing Scenarios

6 ATM Layer Measurement Scenarios

The measurement metrics for the ATM layer include mean cell transfer delay (MCTD), 2-point cell delay variation (2-point CDV), cell loss ratio (CLR), cell misinsertion rate (CMR), cell error ratio (CER), cell throughput, and cell interarrival times.

The cell interarrival times measurement belongs to the traffic statistics collection category, while the rest of the metrics are for performance testing purposes.

The cell interarrival times will be recorded under the mesh topology only. The duration of the measurement might not be very long due to the high speed link and the lack of data storage capability.

Table 4 summarizes the scenarios for performance testing in the ATM layer.

The cell error ratio is defined as:

$$CER = \frac{Number\ of\ Cells\ Received\ In\ Error}{Number\ of\ Received\ Cells} \quad (4)$$

The cell error ratio (CER) performance is not dependent on testing topology. Instead, it depends heavily on the quality of the hardware that the testing bit stream traverses. For this reason, we chose to test the CER metric under the point-to-point configuration only. Also notice that the CER can only be measured by using an out-of-service technique, because the ATM protocol does not support error checking at the cell level.

We expect the mean cell transfer delay, cell delay variation, cell loss ratio and the cell misinsertion rate to be highly dependent on measurement topology. Therefore we chose to test these metrics under several topologies.

The Cell Loss Ratio is defined as:

$$CLR = \frac{Number\ of\ Lost\ Cells}{Number\ of\ Transmitted\ Cells} = \frac{N_t - N_r}{N_t}\ ,\ if\ N_t \geq N_r \qquad (5)$$

where N_t is the number of information cells transmitted and N_r is the number of information cells received. The transmitter will insert OAM cells into a user information stream at suitable intervals. Each OAM cell will contain a count of the number of information cells transmitted since the last OAM cell. In this fashion, we are able to obtain the quantities N_r and N_t at the receiver. This method undercounts CLR if misinsertion occurs during the measurement period.

The Cell Misinsertion Rate is defined as:

$$CMR = \frac{Number\ of\ Misinserted\ Cells}{Time\ Interval} = \frac{N_r - N_t}{T_m}\ ,\ if\ N_r \geq N_t \qquad (6)$$

where Tm is the duration of the experiment. This method underestimates CMR if cell loss events occur during the measurement period.

Again, different services supported on an ATM network have different attributes. In Table 5 we list the maximum acceptable cell loss and misinsertion ratios for five types of services (from [4]).

Notice that the Cell Loss Ratio mentioned in Table 5 is equivalent to our CLR

Service	Cell Loss Ratio	Cell Misinsertion Ratio
Telephony	10^{-3}	10^{-3}
Data Transmission	10^{-6}	10^{-6}
Broadcast Video	10^{-8}	10^{-8}
Hifi Sound	10^{-7}	10^{-7}
Remote Process Control	10^{-3}	10^{-3}

Table 5: Service Attributes for an ATM Network

measurement. The Cell Misinsertion Rate in the Table would be defined as:

$$CIR = \frac{Number\ of\ Misinserted\ Cells}{Number\ of\ Transmitted\ Cells} = \frac{N_r - N_t}{N_t},\ if\ N_r \geq N_t \quad (7)$$

and therefore can also be easily determined from our experiments.

The cell throughput will also depend on the measurement topology. We plan to test this metric under all topologies. Under ideal conditions, the throughput under different topologies should be all close to the link capacity.

References

[1] Hongbo Zhu, Victor S. Frost, "Specification of Measurement Metrics for the AAI Network," *Technical Memorandum TISL-10980-1*, 1995.

[2] Hongbo Zhu, Victor S. Frost, Joseph B. Evans, Douglas Niehaus, and David Petr, "Performance and Traffic Measurement Techniques for the AAI Network," *Technical Memorandum TISL-10980-2*, 1995.

[3] Roel Jonkman, "NetSpec Specification," *TISL Document*, The University of Kansas, 1995.

[4] Martin de Prycker, *Asynchronous Transfer Mode - Solution for Broadband ISDN*, Ellis Horwood, 2nd edition, 1993.