#175

PACKET SATELLITE MULTIPLE ACCESS: MODELS AND MEASUREMENTS

by

Mario Gerla
Lou Nelson
Leonard Kleinrock
Computer Science Department
University of California at Los Angeles
Los Angeles, California
90024

Summary

This paper reports on a series of channel access measurements recently carried out on SATNET, an experimental packet satellite network, to evaluate and compare the performance of F-TDMA, R-TDMA and S-ALOHA channel access protocols. After a brief description of the protocols, the paper introduces objectives, methods and tools of the measurement activity; presents a selection of the most representative measurements; and derives general performance trends and comparisons among the three protocols.

1. Introduction

Satellite communications are becoming an increasingly attractive alternative to terrestrial communications both for the cost-effectiveness of satellite tariffs and services and the particular advantages offered by the multiple access and broadcast environment. Among the latter we mention: multiple destination broadcasting and conferencing capability; dynamic sharing of a common, centralized channel resource among a large population of users; fast deployment and easy relocation for systems equipped with small ground stations, etc.

The satellite broadcast channel is particularly cost-effective in data communications using the packet technology. One of the problems of extending the packet technology to satellite communications, however, is the dynamic sharing of the channel on a packet by packet basis. Various dynamic allocation algorithms have been proposed in recent years, starting with the slotted ALOHA contention scheme, developed primarily for use with a large number of earth stations having relatively low duty cycles[2,3] and evolving later to reservation schemes, where channel reservations can be made in a separate portion of the channel using pre-assigned slots[5] or contention slots[4]. Reservation schemes allow greater efficiencies than the simple S-ALOHA scheme when higher duty cycle stations are also present in the system.

To demonstrate the feasibility and cost-effectiveness of multiple access, packet satellite schemes, and, more generally, to investigate the issues of integration of satellite and terrestrial packet networks, an experimental effort was initiated in mid 1975. This effort lead to the development of SATNET, an experimental satellite network consisting of 4 ground stations (two in the Washington D.C. area, at Etam and Clarksburg; one in Goonhilly, England; and one in Tanum, Sweden) interconnected by a simplex, 64 kbps Intelsat IV-A SPADE channel. The ground station sites are equipped with satellite message processors, called Satellite IMPs (SIMPs)[6] which are an extension of the ARPANET IMP and implement channel access and network access protocols.

This research was supported by the Defense Advanced Research Project Agency of the Department of Defense under contract DAHC-15-73-C-0368

Gateway computers, implemented with PDP-11 hardware, connect SATNET to ARPANET to permit internetwork communications.

The participants in the SATNET experiment are Bolt, Beranek and Newman, Linkabit Corp., Comsat Corp., and U.C.L.A. in the U.S.; the University College of London in England; and the Norwegian Defense Research Establishment in Norway. The project is sponsored by the Defense Advanced Research Projects Agency, the Defense Communications Agency, the British Post Office and the Norwegian Telecommunications Administration.

One of the goals of the experiment is to test the feasibility and the efficiency of different channel access schemes. To this end, three among the most representative and better documented access schemes, namely, Fixed Time Division Multiple Access (F-TDMA), Reservation Time Division Multiple Access (R-TDMA)[5], and S-ALOHA[2], were implemented in SATNET, to gain experience with the implementation, operation and performance evaluation of packet satellite networks.

Other important goals of the SATNET experiment are: the efficient integration of speech and data traffic; the development of voice conferencing protocols appropriate for the satellite broadcasting environment; and the development of reliable Host-SATNET protocols which protect the network from congestion and permit its interconnection with other networks. To meet these goals, a more sophisticated channel protocol, the Contention-Priority Oriented Demand Assignment (C-PODA) protocol, has been developed to efficiently handle a varied mix of traffic requirements (i.e. interactive, batch, and digitized voice) with diversified delay and priority constraints.[7] Tests of C-PODA have just begun.

This paper reports on recent F-TDMA, S-ALOHA and R-TDMA experiments involving only two stations - Etam and Goonhilly. First, the protocols are reviewed, identifying critical parameters and important performance issues. Then, the measurement objectives, methods and tools are described. Finally, the experimental results are presented and are compared with analytical and simulation results to validate the models as well as the SIMP software implementation. The performance of different protocols is compared, and general performance trends are investigated.

2. Channel Protocols

2.1 The Channel

The protocols considered in this paper are all based on a slotted channel structure. Namely, the channel is subdivided into uniform time slots of 30 msec duration, i.e. sufficient to accommodate the maximum size packet. Slots are grouped into frames of 32 slots each. At the beginning of a frame, a slot is reserved for each station to broadcast a "routing" packet. The routing packet is used to broadcast routing information, maintain slot synchronization and return channel acknowledgements.

In a 2-station system the routing frame consits therefore of 2 routing slots and 30 information slots.

The maximum data field length in a packet is 1008 bits. Overhead is 264 bits (header + trailing checksum). Preamble and trailing carrier are 245 bits. Total packet length therefore, is 1517 bits, i.e. approximately 27 msec @ 56 Kbps. (Note: the SPADE channel initial capacity of 64 Kbps becomes effectively 56 Kbps after excluding SPADE overhead.) Considering that slot size is 30 msec, this leaves a guardband of 3 msec between consecutive packets. A large buffer pool (67 buffers) was dedicated to the satellite output channel in each SIMP to insure full channel utilization in spite of the long propagation delay (.25 sec, round trip).

SATNET communications are error protected by an explicit acknowledgment scheme which requires the retransmission of a packet if the acknowledgment is not received from the destination before an appropriate timeout. Channel acknowledgments are returned by the destination via routing packets or via regular packets (piggybacked acknowledgments). Sequencing and duplicate packet detection, however, are not provided by the channel protocol, and are the responsibility of higher level protocols. A SIMP also derives an implicit (or echo-) acknowledgment by monitoring its own transmissions after a round-trip time. If no echo-acknowledgment is received, the SIMP assumes that the transmission was corrupted by a collision or by uplink noise, and reschedules the packet for retransmission. This approach may, of course, introduce duplicates in the presence of local downlink noise.

2.2 Fixed-TDMA (F-TDMA) Protocol[10]

F-TDMA is the simplest and most robust protocol that can be implemented in a multiple access satellite system. Slot are equally subdivided among stations. The assignment is permanent, with no provision for dynamic reallocation of unused slots. In the two-station configurtion the slots are assigned as follows:

...R,R,E,G,E,G,... E,G,R,R...

where E and G denote slots owned by Etam or Goonhilly respectively, and R is a routing slot.

The permanent slot assignment creates N independent subchannels (where N = number of stations) of fixed bandwidth. This makes the analysis of the F-TDMA scheme easier than for dynamic allocation protocols, but not always trivial because of the interdependence between subchannels due to piggybacked ACK's. The main motivations for F-TDMA experiments are:

(1) The calibration of measurement tools (by comparing predicated and measured performance).

(2) The need for a well understood performance reference to which more complex schemes may be compared.

2.3 S-ALOHA Protocol[10]

In the S-ALOHA protocol each station maintains two output queues: The new queue (for new packets); and the retransmit queue (for packets that need to be retransmitted because of a previous conflict). At the beginning of a slot the station will transmit a packet from the retransmit queue with probability P_R (retransmit gate). If the retransmit queue is empty, the station will transmit a packet from the new queue with probability P_N (new gate). Furthermore, a packet arriving at an empty station (i.e. all queues = 0) is transmitted with probability $P = 1$. If two or more stations transmit in the same slot, their packets will overlap, creating a conflict and destroying the information content. The senders detect the conflict by monitoring the channel after a round trip time and promptly return a copy of the collided packet to the retransmit queue.

The version of S-ALOHA protocol discussed here does not provide for channel stability controls. A more recent S-ALOHA implementation, including dynamic control of the gates P_N and P_R as a function of channel load, is now being tested.[8]

Issues of interest in the S-ALOHA experiments are:

(1) The effect of load imbalance on throughput and delay performance.

(2) The impact of gate values P_N and P_R on performance.

(3) The possibility of "capture" situations in the case $P_N >> P_R$.

2.4 Reservation-TDMA (R-TDMA) Protocol[10]

The R-TDMA protocol establishes a permanent association (ownership) between slots and stations similar to F-TDMA. Unlike F-TDMA, however, the slots not claimed by the original owner may be reassigned on a Round-Robin basis to the stations that have traffic to send.

Each routing frame (32 slots) is subdivided into a certain number of reservation subframes, each subframe consisting of N reservation slots (where N is the number of stations), and a number of data slots. Reservation slots are smaller than regular slots, and up to three reservations can be packed into a regular slot. Therefore, only one regular slot is used for reservations in a two-station experiment.

Each station declares its backlog (i.e. packets awaiting transmission) using its reservation slot. Reservations are monitored by all stations and synchronized reservation tables are maintained in all stations showing the outstanding transmission requirements. The reservation table is used by the channel scheduler (a distributed algorithm that runs synchronously in all SIMPs) to assign future slots to users.

If at the beginning of a reservation subframe a SIMP does not hear all the reservations correctly (because of channel noise, for example) it declares itself out of synchronization and it switches from R-TDMA to F-TDMA for the duration of the following reservation subframe. Thus, the presence of noise in R-TDMA may cause not only the loss of some packets, but also the use of a less efficient channel assignment, namely F-TDMA.

Important issues to be investigated with the R-TDMA experiments are:

(1) The demand assignment properties in the presence of unbalanced traffic.

(2) The performance degradation caused by channel noise.

3. Performance evaluation: Methods and Techniques.

3.1 Objectives of the Measurement Activity

In the development and demonstration of SATNET protocols, measurements play an essential role in the following areas:

(1) Software Certification. This consists of verifying that the software is implemented according to specs. Although the implementer generally performs a systematic checkout of each module before release, a global verification is often necessary and is performed by comparing measurements with analytical and simulation predictions.

(2) Study of System Behavior. Via carefully planned experiments the performance of different access schemes is evaluated as a function of traffic and system parameters. A partial understanding of this behavior is possible also via simulation and, in some fortunate cases, via theory. However, the simplifying assumptions used in analytical and simulation models often cover or change some of the actual trends. A close interaction between models and measurements is therefore required, where models help in identifying the issues and measurements validate the modeling predictions.

(3) Optimization of System Parameters. The parameters of some of the more sophisticated protocols require careful tuning in order to achieve optimum channel efficiency, while satisfying the performance constraints imposed by the user. Analytical and simulation models may predict the range of the optimal parameters, and the impact of parameter changes on performance. Due to the limitations of the models, however, the final optimization of the parameters can be performed only on the actual system.

3.2 Measurement Planning

At the beginning of the project a general Satellite Measurement Plan was prepared specifying the performance criteria relevant to satellite access evaluation, and presenting a general methodology for the design of satellite experiments. The plan proposed the following performance measures for the evaluation and comparison of different access protocols:

(1) Throughput

(2) Delay

(3) System Stability

(4) Time to Recover after Loss of Synchronization

(5) Degree of Fairness

(6) Degree of Flexibility

Throughput and delay values can be measured directly, while the other variables are generally derived by combining elementary measures obtained with specially designed experiments.

The delay reported in our experimental results is the "one way" delay measured from the time the packet is accepted in the source SIMP buffer pool, to the time the packet is received without errors at the destination SIMP.

The throughput from A to B is defined as the fraction S of the available channel bandwidth (i.e. the bandwidth that remains after removing the routing slots) utilized by the traffic from A to B, and is calculated by dividing the successful packets from A to B by the elapsed slots, excluding routing slots. Thus, the maximum throughput achievable on the channel is S = 1 pkt/slot.

In addition to the general plan, detailed plans were prepared for each access scheme. Before preparing the plan for a specific scheme, analysis and simulation techniques were used to identify the most relevant properties and performance measures of the scheme, to determine the experimental range of some of the system parameters (e.g., R-TDMA reservation subframe size), and to determine the range of variation of some of the system variables (e.g., delay, queue length, etc.) Typically, before finalizing the measurement plan, a systematic series of simulation experiments were performed. Some of these experiments were then performed on the real channel, to allow for comparison of simulation and measurement data.

3.3 Measurement Tools

The tools used in the SATNET experiments may be divided into two categories: SIMP tools and Network Measurement Center (NMC) Tools.

The SIMP tools consist of the artificial traffic source and the cumulative statistics package. The artificial source generates single or multipacket messages at fixed or geometrically distributed intervals. (In the experiments reported here, however, only single packet messages with geometric arrival distribution were used.) Recently, the generator was augmented with a "cyclic generator" feature allowing the experimenter to change generation rate at prescheduled times during the experiment. This permits investigation of the impact of pulses and other time varying patterns on performance and stability of the various access schemes.

The SIMP accumulated statistics package is essentially an extension of the IMP package[9]. It accumulates average throughput and delay, queue length histograms and other relevant parameters during every collection period (specified by the experimenter and typically ranging from 3 to 15 seconds) and reports the data back to the Measurement Center in a statistics message. As a difference with respect to ARPANET statistics, the SIMP allows direct measurement of one-way delays (due to global SIMP slot synchronization) while ARPANET IMP permits only round trip delay measurements (since ARPANET IMP clocks are not synchronized).

The U.C.L.A. Measurement Center tools consist of a set of three programs implemented on the U.C.L.A. - CCN IBM 360/91: the control program; the collection program; and the data reduction program. The control program allows the experimenter to set up the experiment by specifying the various protocol and traffic parameters; to send these parameters via ARPANET to the SIMPs in SATNET; and to start and stop the experiment when desired. The collection program is responsible for collecting and storing on disk the accumulated statistics messages received from SATNET. The reduction program takes the accumulated statistics and derives global performance measures, averages, correlations, etc., prepares plot files, and tabulates the results in an easy to read format.

4. Performance Results

In this section we present a selection of the measurement results obtained during the channel protocol experiments. For each protocol we first develop, whenever possible an approximate analytical model for throughput and delay at various traffic loads. We then run the experiment on the simulator. Finally we perform the experiment on SATNET and compare measurement results with theory and simulation.

4.1 F-TDMA Protocol

We first develop an analytical model for the F-TDMA scheme considering two traffic cases:

(A) Light and medium load: the offered load is less than channel capacity so that buffer overflow probability is negligible.

(B) Heavy load: the load exceeds channel capacity and drives the channel to saturation.

A. Light and Medium Load

For this case the delay T between origin SIMP and destination SIMP is given by:

$$T = T_p + T_T + T_Q$$

Where: T_p = propagation delay = .25 sec

T_T = transmission delay = $.03 \times \frac{32}{30}$ = .032 sec

T_Q = queueing + latency delay

Using elementary queueing theory we obtain[1]:

$$T_Q = T_T \left(\frac{R}{1-2R} + .5 \ \frac{1-2R}{(1-R)^2} \right) \qquad (1)$$

Where: R = packet generation rate (pkts/slot)

B. Heavy Load

For this case we use Little's result[1] to obtain the following value of T (in seconds):

$$T = .032 \ \frac{\overline{q}}{S} - T_{ACK} \qquad (2)$$

Where: S = one way throughput of the station under consideration (pkts/slot)

$\overline{q}$ = average queue length

T_{ACK} = average time between the arrival of a packet at the destination and the return of the ACK to the origin.

The values of $\overline{q}$ must be calculated as a function of traffic pattern and input rates. Typical values range between 52 and 67. The value of T_{ACK} varies between .25 and .75 sec, depending on the traffic in the opposite direction. Equation (2) is very general (as general as Little's result!) and applies to any channel protocol scheme operating at saturation. In particular, it applies to the S-ALOHA and R-TDMA protocols. This general relationship is very useful in the experimentation of complex channel protocols because it provides a simple check even when more sophisticated analytical or simulation models are not available.

To verify the theory and calibrate the measurement software, a series of two way, balanced traffic experiments were carried out between Etam and Goonhilly. Input rates were the same at Etam and Goonhilly, and were ranging from R = .1 to R = 1 (pkt/slot). The duration of each experiment was 10 minutes, with statistics collection every 13 seconds. The delay versus throughput measurement results for Goonhilly are plotted in Fig.1 and are compared with theory and simulation. The nearly perfect agreement between analysis and measurements confirms the correctness of both theory and software implementation. The slight difference between simulated and measured throughput is due to a difference in generator implementations. Simulated delays, however, lie

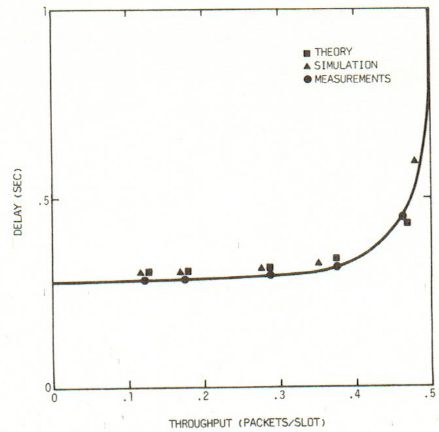

Fig.1 - F-TDMA. Two Way Traffic. Goonhilly Delay vs Throughput

on the measurement curve.

4.2 S-ALOHA Protocol

Some analytical solutions were investigated for a few simple cases to provide a term of comparison for simulation and measurement results. The following cases were analyzed:

(A) Two way traffic, light and medium load case (no buffer overflow), with symmetric ALOHA gates $P_N = P_R = P$.

(B) Two way traffic, heavy load case (saturated system), $P_N = P_R = P$.

A. Light and Medium Load Case.

The delay T (in seconds) is given by:

$$T = .25 + T_T \left\{ 1-\rho + \rho/P + \rho/P(1-\rho)\ (1-\rho_R) \right\} + \frac{G-S}{S}\ (.25 + T_T/P) \qquad (3)$$

Where:

T_T = Transmission delay = .032 sec

$S = G(1-G')$

$S' = G'\ (1-G)$

$\rho = \frac{G}{P}$

$\rho_R = \frac{G-S}{P}$

And:

S = total successful transmissions/slot from the station under consideration = R

G = total transmissions and retransmission per slot from the station under consideration

S'= total transmissions/slot from the opposite station = R'

G'= total transmissions and retransmissions per slot from the opposite station

T is readily evaluated after calculating G and G' from S and S', which are initially given.

B. Heavy Load Case, $P_N = P_R = P$.

The delay T (in seconds) is given by:

$$T = .032 \ \frac{\bar{q}}{S} - T_{ACK} \qquad (4)$$

Where:

$\bar{q}$ = average queue length

T_{ACK} = time between the arrival of a packet at the destination and the return of the acknowledgment to the origin.

S = throughput (packets/slot) = P(1-P)

Analytical results were derived only for the symmetric gate case $P_N = P_R = P$, since in the more general case the analysis is considerably more complex, involving the solution of an 8 x 8 Markov state diagram. For the heavy load case, however, we can still evaluate the delay T for non-symmetric gates ($P_N \neq P_R$), provided that the value of S is obtained via some other method (e.g. simulation).

A first series of experiments was aimed at verifying the basic S-ALOHA theory. The most representative case is the balanced traffic case with symmetric gates $P_N = P_R = .5$, and equal generation rate R at Goonhilly and Etam, for R ranging from .1 to 1. The delay versus throughput measurement results at Goonhilly are reported in Fig. 2, and are compared with theory and simulation, showing an excellent agreement. The theoretical asymptote at S = .25 was verified by the heavy load measurement data point (S = .247, T = 8.23) not shown in Fig.2 to avoid loss of detail at light load.

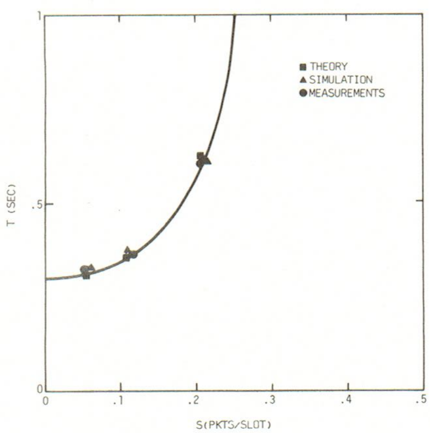

Fig.2 - S-ALOHA - Two Way Traffic Goonhilly Delay vs Throughput

For comparison with sysmmetric gate results, the next series of experiments investigates throughput and delay performance in a balanced traffic situation when non-symmetric ALOHA gates are used (i.e. $P_N \neq P_R$). The measurement results in Table I obtained for heavy input rate R = 1 at both sites show no performance improvement with respect to the symmetric case $P_N = P_R = .5$ (recall that throughput for each SIMP was S=.25 in the latter case). On the contrary, a severe performance degradation is suffered if $P_R \ll P_N$ (e.g. $P_N = .5$ and $P_R = .01$). In this case, a "capture" situation occurs, in which one station holds the channel and transmits at high rate P_N, while the other station is idling in the retransmit state characterized by very low retransmit rate P_R. Stations take turns in capturing the channel. The time history in both simulation and measurements shows that one SIMP may capture the channel for as long as 2 minutes (with some empty channel intermissions), while the other SIMP is quiescent. The variance of the capture periods is very high as compared with the average, suggesting that a typical 10 minute measurement experiment may not provide sufficient throughput and delay statistics. In fact, the "capture" experiment ($P_N = .5$, $P_R = .01$) reported in Table 1 shows a strong throughput and delay imbalance between Etam and Goonhilly, which would probably disappear by extending the duration of the experiment.

Gate Values		Throughput (pkt/slot)		Delay (sec)	
P_N	P_R	Goon-hilly	Etam	Goon-hilly	Etam
.5	.01	.061	.139	37.4	14.6
.5	.05	.146	.141	14.1	14.5
1.	.5	.218	.222	9.56	9.13

Table 1 - S-ALOHA Protocol. Capture behavior with non-symmetric gates.

Having concluded that non-symmetric gates are not attractive for balanced traffic, we then proceeded to study unbalanced traffic patterns suspecting that gate asymmetry may be advantageous in this case. To this end, we performed a series of experiments with input rate at Goonhilly $R_G = 1.$, and rate at Etam R_E varying from .1 to .5, for $P_N = P_R = .5$; and then repeated the experiments for $P_N = 1$, $P_R = .5$. Throughput measurements reported in Fig. 3 show that the non-symmetric gate selection is superior for $R_E \leq .15$ pkt/slot, leading to a 100% Goonhilly throughput improvement for $R_E = 0$. For $R_E > .25$, on the other hand, the non-symmetric choice introduces a 10% degradation.

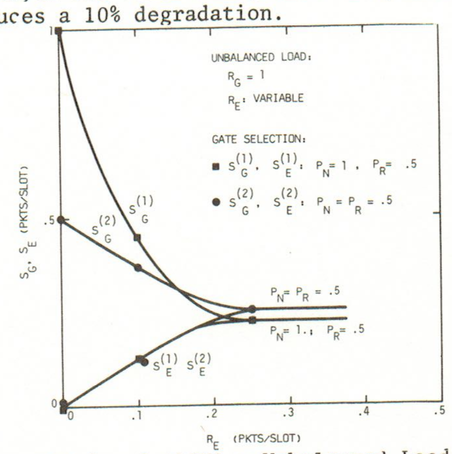

Fig.3 - S-ALOHA - Unbalanced Load. Performance sensitivity to Gate Selection.

The dependence of optimal gate selection on traffic pattern clearly indicates that system efficiency could be improved by dynamically adjusting the gates with a traffic sensitive control scheme. To this end, a closed loop control scheme[8] was developed for the S-ALOHA protocol, and is being investigated at the time of this writing.

4.2 R-TDMA Protocol

The analysis of R-TDMA is considerably more complex than that of previous protocols, and a complete theoretical development is beyond the scope of our work. Here we offer a few approximate results for specific load situations to permit the validation of selected measurement data points.

Specifically, we study three different load situations:

A. light load (with light opposite traffic)

B. light load (with heavy opposite traffic)

C. heavy load

Case A. Light load (with light opposite traffic)

We recall that a station may use its own slots without prior reservation, provided that such slots were not assigned to other backlogged stations[10]. Since we assume light load, the backlog is negligible, and stations transmit packets without need for reservations. The expression of the delay T is given therefore by:

$$T = 1.5\, T_S + T_P \qquad (5)$$

Where: T_P = propagation delay = .25 sec

T_S = effective slot time = $\frac{32}{32\text{-RR-RS}} \times .03$ (sec)

RR = number of routing slots in a routing frame = 2

RS = number of slots used for reservations in a routing frame

Case B. Light load (heavy opposite traffic)

In this case, the opposite station has heavy backlog, thus requiring reservations at all times. The delay expression becomes:

$$T = T_R + T_Q + T_T + T_P \qquad (6)$$

Where:

T_R = reservation delay = $\frac{.03 \times RF}{2} + .27$ sec

T_Q = queueing delay $\simeq \max\left(\frac{R \times RF - 1}{2}, 0\right) \times 2\, T_T$

T_T = transmission delay = .032 sec

T_P = propagation delay = .25 sec

And:

R = packet input rate (pkts/slot)

RF = reservation frame (slots)

Case C. Heavy load

In heavy load the following delay expression holds:

$$T = .032\, \frac{\bar{q}}{S} - T_{ACK} \qquad (7)$$

Where:

$\bar{q}$ = average queue length (packets)

S = one way throughput of the station under consideration (pkts/slot)

T_{ACK} = average time between the arrival of a packet at the destination and the return of the ACK to the origin (sec).

The first experiment was intended to calibrate the R-TDMA measurement and consisted of balanced traffic at Goonhilly and Etam with rate R varying from .1 to 1. Reservation subframe was set to 10 slots.

The measurements taken at Goonhilly and reported in Fig.4 show an excellent agreement with modeling predictions.

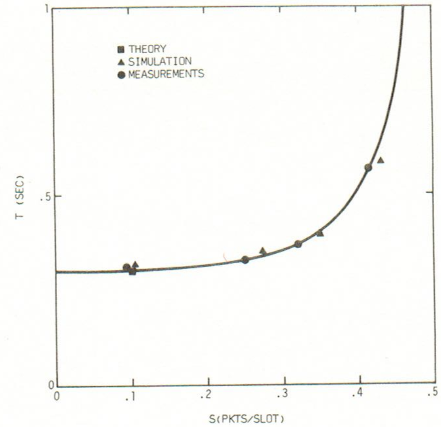

Fig.4 - R-TDMA- Two Way Balanced Traffic. Goonhilly Delay vs Throughput

By comparing Goonhilly and Etam performance during the balanced experiments we verified that the bandwidth is fairly apportioned between the two stations, and the delay performance is identical in both directions.

A more interesting demonstration of R-TDMA performance is offered by two sets of unbalanced traffic experiments, with Goonhilly input rate R_G fixed to the values .1 and 1 respectively, and Etam input rate R_E varying between .1 and 1. Reservation subframe is RF = 10 slots. Fig. 5 shows Etam delay versus throughput for the two experiments.

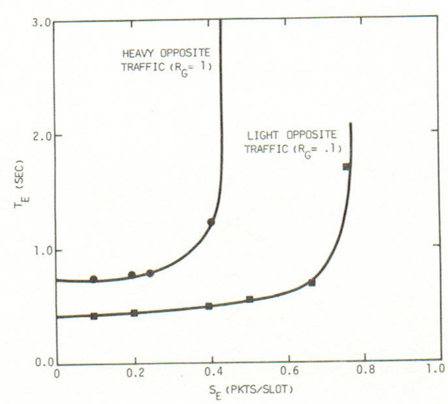

Fig.5 - R-TDMA - Two Way Traffic. ETAM Delay Vs. Throughput for Different Values of Opposite Traffic

The unbalanced traffic experiment emphasizes one of the most attractive properties of the R-TDMA scheme, namely the capacity to acquire unused slots originally owned by the other station. Considering the lower curve in Fig.5 we notice that the delays for $S_E \leq .4$ are comparable with the F-TDMA delays in Fig.1. This is because reservations are not needed at low traffic levels with negligible backlog. For increasing input rate R_E, Etam progressively acquires the portion of bandwidth not utilized by Goonhilly. Next, considering the upper curve in Fig.5 (i.e. R_G=1), we notice that the initial delay is higher because reservations are necessary for all packets. Etam must "force its way through" and compete for bandwidth with Goonhilly; but it will be granted its entire bandwidth request until the channel is equally apportioned between the two stations.

The final set of R-TDMA experiments involved the sensitivity of throughput and delay performance to channel noise. Downlink noise was simulated in a SIMP by discarding a fraction P of the received packets, where P is referred to as the "noise gate". Routing packets are assumed immune to noise, thus providing a safe return path for ACKs. Reservation noise gate is 1/3 P, since reservation packets are 1/3 of a slot.

The efficiency of the R-TDMA protocol is based on its ability to maintain synchronized reservation tables at the various SIMPs. However, channel noise may cause loss of synchronization and temporary reversion of channel scheduling to the F-TDMA mode until synchronization is regained. If only one of the two stations is out of sychronization, collisions may occur which further reduce performance. We may expect, therefore, that the F-TDMA protocol outperforms the R-TDMA protocol beyond some critical noise gate value.

Analytical models were developed to provide better understanding of the relationship between performance parameters and noise. For a station operating at saturation in a noisy channel, the delay T is independent of the particular protocol and may be expressed as follows:

$$T = .032\,\frac{\bar{q}}{S} - T_{ACK} - \frac{P_M}{1-P_M}\left(T_P + \frac{P_0}{1-P_0}\,(T_P + T_0(1-P_M))\right) \qquad (8)$$

Where:

$\bar{q}$ = average number of buffers in use

S = effective throughput

T_{ACK} = average time to return an ACK

P_M = transmitting station local noise gate ("my" noise)

P_0 = receiving station local noise gate ("other" noise)

T_P = propagation delay = .25 sec.

T_0 = ACK time-out = 2 sec.

Similar expressions may also be derived for lightly loaded stations[11].

The noise experiments are performed in an unbalanced traffic environment (R_G=1; R_E=.1) to better study R-TDMA performance degradation due to synchronization loss and collisions, since both such effects are emphasized by load imbalance. The reservation subframe is chosen to be RF=30.

R-TDMA measurement results as a function of noise level for Goonhilly (the heavy station) are reported in Fig.6. Also reported are the F-TDMA results obtained in identical traffic and noise conditions. Measured delays were also compared with delays calculated according to equation (8), to find that the discrepancy was rather small (less than 10%), thus confirming the validity of the model and the accuracy of the measurements.

Inspection of Fig.6 shows that the performance degradation of both R-TDMA and F-TDMA scheme is almost negligible for noise gate values up to P=.03. Beyond this threshold, the performance of R-TDMA degrades more rapidly than that of F-TDMA (as expected). However, for a noise level P=.3 R-TDMA is still superior to F-TDMA, confirming the robustness of R-TDMA to noise.

4.4 Comparing Different Protocols

The performance comparison of protocols is strongly

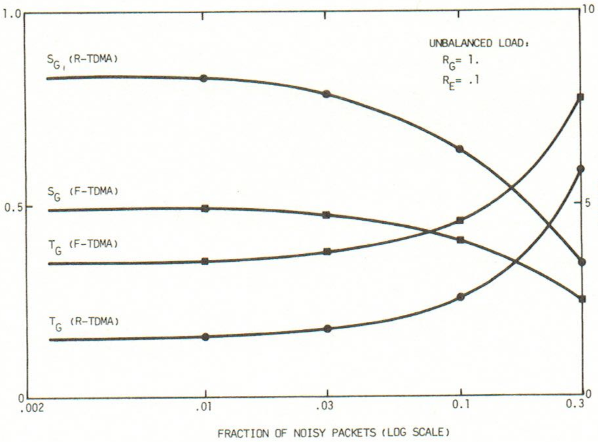

Fig.6 F-TDMA and R-TDMA performance in a noisy environment with unbalanced load.

dependent on system parameters such as the number of stations, traffic pattern, channel noise level, etc. Here we limit ourselves to the two station case, and consider a family of traffic patterns in which one station is heavily loaded (R_G=1), while the other station has a traffic volume R_E variable from 0 to .5.

Fig.7 shows Goonhilly and Etam throughput measurements for F-TDMA, R-TDMA and S-ALOHA (P_N=1; P_R=.5) protocols. Similarly, Fig.8 shows delay measurements for the three protocols. On the curves in Fig.7 and 8, the labels A, B and C refer to F-TDMA, R-TDMA and S-ALOHA respectively, while the subscripts 1 and 2 refer to Goonhilly and Etam respectively.

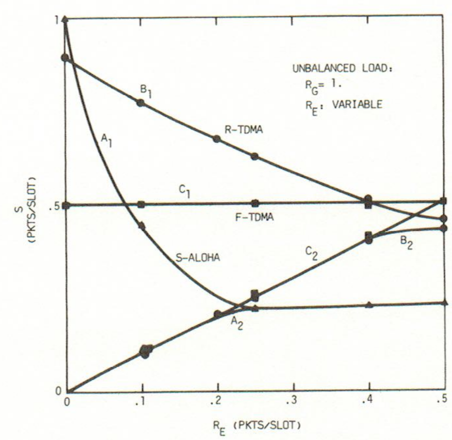

Fig.7 Unbalanced Load. Throughput Performance of Different Protocols.

From Fig.7 and 8 we notice that R-TDMA yields higher throughput and lower delay for a wide range of traffic patterns (i.e. $.01<R_E<.4$); and it is second to S-ALOHA for $R_E<.01$ and to F-TDMA for $R_E>.4$ (in both cases, however, by a minimal amount) mainly because of the overhead introduced by the reservation slots.

5. Conclusions and Directions of Future Experiments

The main objectives of the measurement activity, namely software checkout and study of system behaviour, were successfully met during the first phase of the SATNET experiment involving the testing of conventional channel access protocols (i.e. F-TDMA, S-ALOHA and R-TDMA). Measurements were debugged and validated with the assistance of analytical and simulation models,

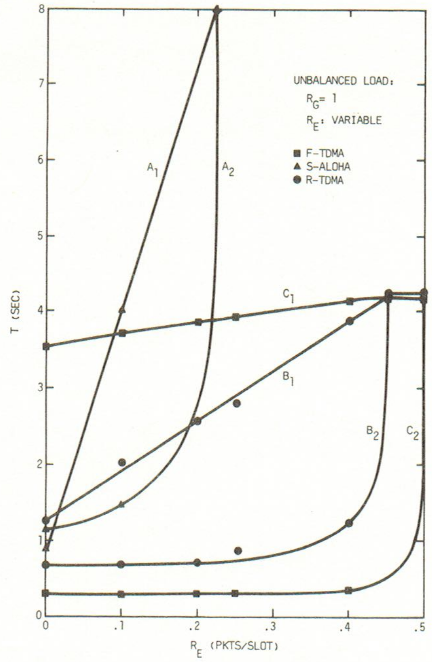

Fig.8 Unbalanced Load. Delay Performance of Different Protocols.

demonstrating once more the importance of modeling in the development and testing of experimental systems.

Important results of this first phase of the SATNET experiment are: (1) the superiority of the R-TDMA protocol over the other protocols for a very wide range of traffic patterns; (2) the extreme robustness of the R-TDMA reservation scheme to downlink random noise; and (3) the "capture" effect present in the S-ALOHA protocol for certain gate value selections (namely $P_N \gg P_R$).

Future channel access experiments will include the implementation and testing of more sophisticated protocols such as the stability controlled S-ALOHA scheme[7] and the C-PODA scheme[8]. During this phase, system parameter optimization will become an important aspect of the experiments since the future schemes contain a number of critical parameters which can only be properly tuned by experimentation.

Acknowledgement

The authors wish to express their gratitude to Elina Tseng, Warren Usui and David Clemans for their assistance in developing the simulation program, the network measurement tools and carrying out the experiments.

References

1. Kleinrock, L., "Queueing Systems, Vol. I: Theory," John Wiley & Sons, 1975.

2. N. Abramson, "Packet Switching with Satellites," National Computer Conference Proceedings, June 1973, Vol. 42.

3. L. Kleinrock and S.S. Lam, "Packet-Switching in a Slotted Satellite Channel," National Computer Conference Proceedings, June 1973, Vol. 42.

4. L.G. Roberts, "Dynamic Allocation of Satellite Capacity Through Packet Reservation, "National Computer Conference Proceedings, June 1973, Vol.42.

5. R. Binder, "A Dynamic Packet Switching System for Satellite Broadcast Channels", Proceedings of the International Conference on Communications, June 16-18, 1975, San Francisco, CA.

6. S.C. Butterfield, R.D. Rettberg, and D.C. Walden, "The Satellite IMP for the ARPA Network,"Proceedings of the Seventh Annual Hawaii International Conference on System Sciences, Honolulu, Hawaii, January 1974, Computer Nets Supplement, pp. 70-73.

7. Jacobs, I. et al, "C-PODA-A Demand Assignment Protocol for SATNET", Fifth Data Communications Symposium Proceedings, Snowbird, Utah, Sept. 77.

8. Gerla, M. and L. Kleinrock, "Closed Loop Stability Controls for S-ALOHA Satellite Communications", Fifth Data Communications Symposium Proceedings, Snowbird, Utah, Sept. 77.

9. Kleinrock L., and W. Naylor, "On Measured Behaviour of the ARPA Network", National Computer Conference Proceedings, Chicago, May 74.

10. Binder, R. "Current SIMP Protocols", Packet Satellite Project Working Note No.11, Bolt Beranek and Newman, April 76.

11. Gerla, M. et al. "F-TDMA, R-TDMA and S-ALOHA Measurements in SATNET", Packet Satellite Project Working Note No. 72, University of California, Los Angeles, May 77.