On measurement facilities in packet radio systems*

by FOUAD A. TOBAGI, STANLEY E. LIEBERSON and
LEONARD KLEINROCK
University of California,
Los Angeles, California

ABSTRACT

The growth of computer networks has proven both the need for and the success of resource sharing technology. A new resource sharing technique, utilizing broadcast channels, has been under development as a Packet Radio system and will shortly undergo testing. In this paper, we consider that Packet Radio system, and examine the measurement tasks necessary to support such important measurement goals as the validation of mathematical models, the evaluation of system protocols and the detection of design flaws. We describe the data necessary to measure the many aspects of network behavior, the tools needed to gather this data and the means of collecting it at a central location; all in a fashion consistent with the system protocols and hardware constraints, and with minimal impact on the system operation itself.

INTRODUCTION

This paper is primarily concerned with the unique measurement aspects of Packet Radio Systems as regards network evaluation, and considers the design of a set of measurement facilities, the development of data gathering techniques within the framework of the system design and the use of these measurements to evaluate the system performance and its operational algorithms.

The need for sharing of computer resources by organizing these resources into computer networks has been long recognized[1] and the feasibility of constructing such networks has been demonstrated by many successfully operating network systems. Perhaps the most prominent example is the ARPANET,[2] which utilizes the technique of packet-switching, appropriate for bursty computer network traffic, thus achieving better sharing of the communication resources.

The ARPANET emerged in 1969 as the first major packet-switching network experiment; since the essence of an experiment is measurement, and in line with Hamming's observation that "it is difficult to have a science without measurement", considerable care was taken from the beginning in the design and development effort to include the tools necessary and appropriate to satisfy the many measurement goals. As a result of well designed experiments on the ARPANET using these tools, valuable insight has been gained regarding the network usage and behavior.[3]

The Packet Radio System is another yet different example of a computer resource sharing network.[4] It is being developed by the Advanced Research Projects Agency in order to demonstrate the applicability of the packet radio concept in organizing computer resources into a computer communications network. It is this packet radio network which is of concern to us in this paper. The network is currently in its design phase,* and, as was the case with the ARPANET, care is being taken to include the ability to measure network behavior. UCLA is in charge of this measurement effort.

This concern for measurement is due to several factors. Firstly, these measurements provide a means to evaluate the performance of the operational protocols employed and the identification of their key parameters. Moreover, this realistic observation of the system behavior will assist in the validation and improvement of existing analytical models devised to study some of these operational schemes, such as the access modes and routing strategies.[5,6] Secondly, these measurements will allow for the detection of system inefficiencies and the identification of design flaws such as the inadvertent creation of a deadlock condition.[7] Thirdly, measurement facilities and data, when used to improve network design, are a valuable feedback process in which design deficiencies are detected and subsequently corrected. Wire networks differ from radio networks mainly in the omni-directional broadcast nature of the communication and consequently the protocols employed; therefore, it calls for new approaches in the design and implementation of the measurement facilities and their use.

---

* This work was supported by the Advanced Research Projects Agency of the Department of Defense (DAHC-15-73-C-0368).

* A preliminary demonstration of the system is under way. A prototype network is being set up in the Palo Alto, California, area.

In the following section, we present an overview of the packet radio system concepts and a brief description of the currently specified operational procedures. In a later section, we describe the network measurement facilities which consist of the measurement tools and the techniques for data collection. In the last section, we identify and discuss in some detail the desirable measurement functions to satisfy the need for validation and performance evaluation outlined above.

THE PACKET RADIO SYSTEM

Several papers have already appeared in the literature which describe the packet radio concept and discuss many of the issues involved in the system design.[1-6,8-10] In this section, we briefly describe these system components and operational procedures necessary to understand the measurement considerations presented below.

There are three basic functional components of a packet radio system:

(i) packet radio terminals—these are the sources and destinations of traffic on the packet radio network.

(ii) packet radio stations—these function as S/F switches for local traffic and as interfaces between the broadcast system and other computers or networks. Also, they perform directory, monitoring and control functions for the overall system, and they play a central role in that all traffic passes through the station, i.e., we have a centralized network.

(iii) packet radio repeaters—their function is to extend the effective range of terminals and stations by acting as Store-and-Forward relays.

The repeater, which has been developed by Collins Radio and is called a packet radio unit (PRU), consists of a radio transceiver and a microprocessor. The function of the PRU is to receive and transmit packets according to dynamic routing and control algorithms specified by the station. For simplicity and uniformity of design, the PRU is used as the front-end of terminal devices and of stations, interfacing them with the radio net. In Figure 1 we show an oversimplified picture of the PRU identifying its various sections: the radio transceiver, the store-and-forward software, the control process, and the measurement process.

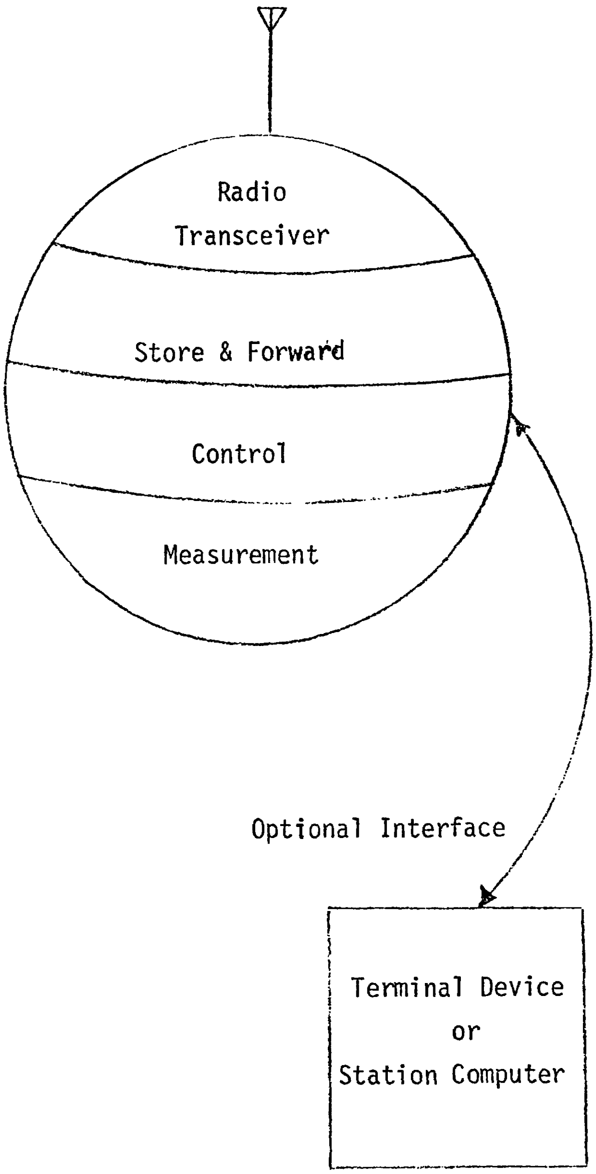

Figure 1—The packet radio unit

In this initial system, the terminals, stations and repeaters are linked together by a single broadcast channel using omni-directional antennas. The repeaters do not determine routes. All the routing computations are performed by the station. A hierarchical routing algorithm is used which makes the routing in the broadcast network resemble routing in a point-to-point network by forming a hierarchical tree structure. This structure is constructed by having the station assign to each repeater a label which defines its position in the tree. A packet is routed along the path determined by the tree, requiring the packet header to contain a string of appropriate repeater ID's, or labels. Thus, neighboring repeaters hearing the broadcasted packet but not on the determined path will reject the packet rather than relay it. However, this algorithm is flexible in that it allows the repeater to seek an alternate route for a packet when a path seems to be blocked. Moreover, the station with its monitoring procedures can dynamically restructure the tree by relabeling any of the repeaters in response to component failures or traffic congestion.

In order to achieve reliable packet transport, acknowledgment procedures are required. There are two types of acknowledgments; the end-to-end acknowledgments (FTE) between end devices, and hop-by-hop acknowledgments (HBH) between repeaters.[6] Except for the last hop on a packet's route, HBH acknowledgments are passive in that the relaying of a packet over a hop constitutes an acknowledgment of the transmission over the previous hop; this "echo acknowledgment" is due to the omni-directional broadcast property. At the last hop, an *active* HBH acknowledgment must be generated.

MEASUREMENT FACILITIES

Several factors exist in the packet radio system which do not allow for a simple transfer of ARPANET-like measurement facilities to a packet radio network. Although the latter utilizes the same technique of packet-switching, the packet radio concept is unique in the constraints it places on all system operations and the measurement effort in particular.

The radio broadcast nature of transmissions is such that the transmission of measurement data not only introduces overhead over its own path, but causes transmission interference at neighboring repeaters within hearing distances and creates additional overhead on those PRU's activities. Moreover, the desire to keep the components small and portable, as well as the limited speed of the IMP's CPU within the PRUs, place significant constraints on the measurement facilities and their usage. The available storage is extremely limited and the overhead placed on the PRU's CPU is of utmost importance in evaluating the feasibility of a measurement tool and of the collection of data in support of a measurement function. As the operational protocols of the net are different from wire nets, the measurement functions devised to support the evaluation of their performance are unique. Thus, the measurement effort consisted of identifying the measurement functions (as described in the following section) and devising the measurement facilities required to support those functions under the constraints that the system imposes. The development of the tools was an iterative design process seeking a balance between supporting the measurement functions and satisfying the system constraints, as well as making sure that the network communication protocols allow the implementation and proper functioning of those tools.

In this section, we describe the various types of statistics desired in the Packet Radio Net,* the traffic sources required in measurement experiments and the techniques available for measurement data collection. We shall postpone until the next section the detailed list of the quantities that will be measured by each of the types of statistics (tools).

Cumulative statistics (Cumstats)

As its name suggests these consist of data regarding a variety of events, accumulated over a given period of time, and provided in the form of sums, frequencies and histograms. We shall distinguish between those data collected at the PRUs (PRU based Cumstats) and those collected at the end devices (the end-to-end Cumstats). The PRU based Cumstats provide information about the *local* environment and behavior such as traffic load, channel access, routing performance, and repeater activity. Conversely, end-to-end statistics collected at network sources and sinks, that is stations and terminal devices, will reflect more global network behavior such as user delays and network throughput.

Trace statistics

The trace capability allows one to literally follow a packet through the network, and to trace the route which it takes and the delays which it encounters at each hop. In the ARPANET, selected IMPs gather data on packets to be traced (which may include any packet) and send this data to the collection point as a new packet. In the packet radio network, however, the collection of trace data at the repeaters is prohibited by the limited size of storage in the PRU. To overcome this problem, we have introduced a new type of packet called the Pickup Packet.* These packets are generated with an empty text field by traffic generators at end devices. As these packets flow normally in the network according to the transport protocols, selected repeaters will gather the trace statistics and will store them within the text field of the pickup packets themselves.

Snapshot statistics

Snapshots give an instantaneous peek at a PRU, showing its state at that moment with regard to buffer assignment and queue lengths. (In the ARPANET, which is a decentralized network in which each node contains routing algorithms and data, snapshots also include routing related information; in the Packet Radio Network, such information is available at the station). Changes to appropriate station tables will be time stamped and collected as the station's snapshot function.

Artificial traffic generators

Traffic sources

The creation of streams of packets between given points in the net, with given durations, intervals,

---

* These types of statistics, as well as traffic generators, which have been widely used in ARPANET measurement experiments, will differ significantly from those of the Packet Radio Network in regards to the specific quantities gathered and the means of collecting them at a central location.

* The notion of the pickup packet was first suggested by H. Opderbeck.

packet lengths, and packet types (Information and Pickup Packets) is clearly a requirement of any experimental system. While it might be desirable to provide each PRU with the capability of creating such traffic, this additional burden on the PRU software can be avoided if there exists a reasonable number of terminals with processors attached which, along with the station, will be programmed to provide the traffic-source functions indicated above.

Specifically, traffic-source features which the terminals (and Station) should provide are: (1) Information Packets—the user specifies the packet length, frequency, destination and duration of one or more streams of Information Packets. (The text content may be arbitrary.) (2) Pickup Packets—the user specifies the packet length, frequency, destination and duration of one or more streams of Pickup Packets.

In the initial system, there will be a limited number of system elements, making it desirable to simulate in a terminal a multi-terminal environment. That is, the traffic generated at a single terminal will emulate the traffic that would be generated by several separate sources. A great deal of complexity is introduced in the design of these devices because of the hardware and software capabilities required to support this function. Feasibility and techniques of achieving this is under investigation.

Station measurement process

Since the station is the central node and provides central control for the operation of the entire network it therefore plays a central role in the execution of measurement functions. It is through the station that the initiation and termination of measurement experiments is controlled. In particular, the station enables and disables the Cumstat and Pickup packets functions at the PRU's, and assigns to the various elements the intervals for Cumstat collections, and to the artificial traffic generator, their corresponding parameters. Moreover, it is to the station that all measurement data is ultimately destined; upon arrival at the station, the data is time-stamped and stored in a single measurement file for off-line reduction and analysis. In addition, all changes to the station's internal tables (routing, connectivity, PRU operational parameters, etc.) will be reflected by an entry into the measurement file, thus allowing the correlation of measurement results to the actual network configuration. A (measurement) process at the station will perform all of the above functions.

Measurement data collection

As mentioned earlier, pickup packets are generated at stations and terminals. Those packets generated at a terminal are destined to (and collected at) the station; those generated by the station will be returned by their destinations to the station as regular packets for collection into the measurement file.

Let us now discuss the techniques for centrally collecting cumulative statistics. The data, generated at the PRU's or terminal devices, must be transmitted to the station using the PR Net itself. One way of achieving this is to form at the PRU, at the end of each Cumstat interval, a measurement packet called the Cumstat packet, which is time-stamped and transmitted to the station. The second method consists of having the station send at regular intervals to appropriate PRU's an executable control packet* called an Examine packet which collects time stamped Cumstat data and which returns back to the station.

For purposes of analysis, it is desirable for the Cumstat data received at the station to correspond to equal length time intervals at the generating device. This can be achieved in the automatic method if reliable ETE transmission exists, i.e., if ETE acknowledgment capabilities are provided in all terminal devices and PRU's, preventing the loss of a Cumstat packet from a device on its way to the station. In the absence of the ETE capability in the PRU's, one may decrease the Cumstat intervals (thus increasing the frequency of transmitting Cumstat data), thereby decreasing the gaps between correctly received Cumstat packets. With the Examine method, variable length Cumstat intervals will occur since Examine packets, sent at regular intervals from the station, are subject to (i) the network random delays en route to the destination PRU, and (ii) the possibility of loss in either direction.

The choice of a collection method will have to take into consideration the overhead that it imposes on the PRUs and on the network.

MEASUREMENT FUNCTIONS

We have described in the previous section the measurement facilities that are desirable in a PRNET to support the measurement functions. In this section we shall identify and discuss these functions in some detail, determine the required data items and describe the role of these measurement facilities in supplying the data. These include: channel access, operational protocols, repeater performance, traffic characteristics, and the network's global performance.

Channel access

One of the main features that distinguish the Packet Radio Network from point-to-point networks is that

---

* An executable control packet is a packet that originates at the station and is destined to a PRU. It contains code to be executed by the destination PRU. In particular, the Examine packet contains the necessary code to load the contents of specified memory locations into the text of the packet for shipment back to the station.

devices transmit packets over a broadcast channel by using a random access scheme. These random access schemes are characterized by the sharing of a single channel in a multi-access fashion, thus allowing for packet interference to occur. Considerable progress has been made in analyzing these access modes, which include pure and slotted ALOHA and the more recently developed techniques of Carrier Sense Multiple Access (CSMA).[5,11-13] In the initial experimental system pure ALOHA, 1-persistent CSMA and non-persistent CSMA will be available. Our measurement aims are to validate the analytical models of the three access modes and to evaluate their performance in realistic environments.

In evaluating terminal access in a single hop system (a model commonly used in analysis), we consider an environment consisting of a single station and a population of terminals within range and in line-of-sight of the station. In order to determine the relationships between the network throughput (rate of successfully received packets at the station) and channel traffic (rate of packet transmissions over the channel), as well as the relationships between the network throughput and packet delays, the following quantities will be measured:

(a) the number of transmissions a packet incurs before success

(b) the one-hop packet delay: time elapsed since the packet is ready for transmission until it is acknowledged, i.e., until its acknowledgment packet is received from the station

(c) network throughput: average number of packets received at the station per unit time

Items (a) and (b) are obtained in the form of histograms by the Cumstat tools at the PRU and the end device respectively. Item (b) may also be obtained individually for each Pickup packet by having the originating device store in it its time of generation and its transmission times, and in the succeeding Pickup packet, store the time its acknowledgment arrived. Item (c) is obtained at the station from end-to-end cumulative statistics.

The task of measuring performance of terminal access techniques in multi-repeater environments differs from the previous one in that repeater-to-repeater traffic is present contending on the same channel. The environment consists of a number of repeaters and stations and a population of terminals, not necessarily all within range and in line-of-sight. The same quantities as listed above, measured over the terminal-to-repeater hop, will be collected using the same tools.

Operational protocols

Acknowledgment protocols

Echo acknowledgment suffers from packet interference. The delay until the echo acknowledgment is received at the transmitter is random. Thus, the packet may incur some additional transmissions beyond the first successful one creating additional overhead on the channel and in the PRUs. This number of additional transmissions is a measure of the inefficiency of echo acknowledgments; so too will be the number of packets discarded at the transmitter because of lack of reception of the echo acknowledgment. That is, the transmitter reached the maximum number of retransmissions of a packet before the echo acknowledgment was received; although the packet may have been successful, the transmitter declares itself unsuccessful in establishing communication!

Thus, we shall measure the efficiency of the Echo Acknowledgment protocol by measuring the number of additional transmissions beyond success incurred by a packet. To compute this number, a PRU must have two pieces of information; it must know how many times the packet has been transmitted, and it must also know which of those retransmissions was the one that reached the next repeater successfully. This information will be contained in two fields in each packet header, which we refer to here as fields A and B. Field B is used by the PRU to store the current transmission number of the packet. When the packet is successfully heard by the intended receiver, the contents of field B are saved by being stored into field A; when the Echo acknowledgment is successfully heard by the sending PRU, field A of the echo acknowledgment is compared with the current number of transmissions of the packet, i.e., the contents of field B in the sender's copy of the packet. If these two numbers differ, then the magnitude of that difference represents the number of times that the packet was retransmitted after it had already been successfully received at the next hop. This data is collected as part of the cumulative statistics of the sending PRU.

Routing protocols

Earlier we introduced the hierarchical routing scheme in use, which is based on a tree structure with the station as its root. The initial tree structure is created dynamically by the Initialization Procedure in which the station uses PRU connectivity information to create a tree that minimizes the number of hops between each repeater and the station. Thus the routing strategy initially performs shortest path (minimum hop) routing from repeaters to station and from station to repeaters. However, when the first choice shortest path cannot be used, the packet departs from this path and uses a shortest path from its new location. This will occur when a repeater has transmitted a packet over a hop the maximum number of times allowed without receiving an HBH acknowledgment; the repeater then alters the packet's header (to what is called the "ALL" label) so that any repeater within hearing distance and able to relay the packet in its intended direction will do so. This packet is then said to be alternately routed. It is retransmitted with its "ALL" header until either an HBH acknowledgment is received or the maximum number of retransmissions is once again reached, at which time the packet is discarded.

The analysis of a routing algorithm, particularly in a broadcast, and thus mobile, network, is a complex task, in that routing is topology- and load-dependent, and involves, with varying degrees of subtlety, all of the system's protocols. Thus, routing considerations are really a synthesis of most elements of the system design, and as such, the measurement of the algorithm involves at times the study of the interaction of the many system protocols.

Given the patterns of input load on the network, the *distribution of traffic flow* in the net is an indication of the behavior and efficiency of the routing and initialization algorithms. One may detect the concentration of traffic on specific routes creating congestion while alternate routes are not assigned; thus smaller delay routes may have been ignored in favor of the shorter routes provided by the initialization procedure.

To obtain the distribution of traffic flow, the following quantities are to be measured.

(a) the total number of packets received and transmitted at each repeater (obtained in the PRU Cumstats)

(b) the fraction of time the transceiver is busy (obtained by snapshot statistics, or in the PRU Cumstat by regular sampling of the transceiver's state)

Also, the point-to-point nature of this routing algorithm, restricting a packet at a given hop to a single repeater as its immediate destination, does not take advantage of the broadcast nature of the channel, in which several neighbors may actually hear the transmission and be capable of relaying the packet. Thus the following quantity is relevant:

(c) the number of packets correctly received and discarded because they are destined to other components in the net (obtained in the PRU Cumstat).

Moreover, to measure the potential of each neighboring repeater (say, repeater "n") as an immediate destination, it is essential to know the probability of success P(n) repeater n has to correctly receive a broadcast packet. This we do by maintaining in each PRU a table counting the number of successfully received packets from each immediate neighbor. The ratio of the number of packets correctly received from a given neighbor, to the number of packets transmitted by that neighbor, is a measure of P(n).

Another important feature of a routing algorithm is its adaptability to network changes: input traffic load, connectivity and component failure and repair. In evaluating the dynamics of such an algorithm, three factors must be examined: the time required to detect the network change, the time required to respond, and the quality of the response. The data items at each PRU necessary for these studies, which include some of those mentioned earlier, are:

(a) tables counting the number of packets correctly received from immediate neighbors

(b) number of packets alternately routed

(c) number of packets discarded, suggesting route congestion or component failure

(d) percent of time repeater is busy transmitting and receiving

These are obtained as cumulative statistics in the PRU.

In addition, the Pickup packet is a valuable tool in routing studies in that it contains the actual and complete route taken by the packet (pinpointing alternate routing), as well as time stamps to compute the queueing and transmission delays incurred at each repeater.

Repeater's performance

The evaluation of the performance of a repeater is most important in the analysis of network behavior; it allows us to break down key network measures (such as packet delay and throughput) into their elementary components and to examine the effects on these measures of the repeater activity and design (including buffer management, queueing discipline, and packet processing priorities).

The quantities relevant to packet delays are:

(a) The processing time of a packet flowing through a repeater; this is counted in Pickup Packets as the time lapse between the packet's arrival and the time it is placed on the transmission queue. This processing includes various checks such as checksum, packet type, routing labels, etc.

(b) the packet queueing delay at a repeater; this is also counted in Pickup Packets as the time elapsed from when the packet is placed on the transmission queue until it is considered for transmission (i.e., until it is at the head of the line, in a first-come-first-served discipline).

(c) the packet's service time; this is also counted in Pickup packets as the time elapsed from when the packet is at the head of the queue until its echo-acknowledgement is correctly received. Note that the actual service time (time until the packet is correctly received at the next repeater) is smaller than the one measured here due to the echo acknowledgment protocol used in this system. Note also that the service times of consecutive packets are correlated.

The quantities related to a repeater's communications activity are:

(d) percent of time the PRU transceiver is busy transmitting and receiving; this can be obtained in the PRU

Cumstat by regular sampling of the transceiver's state.
(e) the total number of transmitted packets at each repeater relative to the number of successfully transmitted packets. The latter number is obtained for each neighboring repeater by examining its table count which gives the number of packets correctly received from immediate neighbors.
(f) the percent of traffic received with checksum error (obtained in the PRU Cumstats).
(g) the percent of traffic received correctly but not intended for this repeater (obtained in the PRU Cumstat).

The quantities relevant to buffer management and occupancy are:

(h) the percent of time packet buffers are in a given state (free, queued for packet transmission, reserved for packet receive). This can be obtained in the PRU Cumstat by a regular sampling of the buffer states.
(i) the frequency of buffer overflow as a function of the load, and this is obtained also in the Cumstats by counting the number of packets discarded due to lack of buffer space.

Traffic characteristics

In determining the traffic characteristics, one should distinguish between *external traffic* (the input traffic generated by network users and traffic sources) and *internal traffic* (traffic relayed and generated at repeaters). The measurement functions determining the external traffic characteristics are not necessary when the entire traffic is artificially generated. They include:

(a) the geographical distribution of the input load (obtained in the end device Cumstats)
(b) characteristics of the terminal input processes (obtained in the form of histograms of packet intergeneration time from the end device Cumstats)
(c) the amount of traffic generated at repeaters for special purposes such as: control, measurement, etc. (i.e., overhead traffic) (obtained in the PRU Cumstats)

The characterization of internal traffic is crucial in the creation and validation of assumptions made in repeater models aimed at an analytic prediction of the performance of multi-repeater packet radio networks. To characterize this internal traffic, we may measure the following quantities at each repeater:

(a) interarrival time (defined as the time between the arrivals of two successive packets that have been correctly received and are destined to that repeater).
(b) interdeparture time (defined as the time elapsed between the acknowledgment of two consecutively transmitted packets).

Histograms of these quantities can be created from the information contained in the Pickup Packets.

Network's global performance

The ultimate goal of all system considerations is to create a network of high capacity providing minimal user (end-to-end) delay. We examine the success in achieving this goal by measuring the end-to-end delay and the network throughput (counted as the number of packets received at their respective destinations), under various patterns of input load, as well as the frequency of lost and duplicated packets.

It is important to note that these quantities are fundamentally affected by all the operational protocols. They allow us to obtain the main performance curves of throughput and delay.

The role of measurements in flow control

The station has the responsibility for centralized control over the entire network. To carry out this responsibility, the station requires various indications of network activity and performance. Some of this information will be acquired from incoming traffic, but much of this information must be specifically obtained by having monitoring procedures collect, from the various devices, a subset of the measurement items that have been seen presented throughout the paper.

CONCLUSION

In this paper, we have presented some of the results of our activities in the measurement aspect of the ARPA Packet Radio Project. We described the Packet Radio Network measurement facilities, consisting of the measurement tools and the techniques for data collection. We also identified and discussed the measurement functions required to gain insight into the behavior of this broadcast network. In so doing, we determined the data items required to support these functions and the means for their collection. This information is summarized in Table I.

In the design of these measurement facilities, a constant concern is to keep the overhead they create at the components and on the broadcast channel at a low level. An important activity will be to evaluate the cost of each element of the facilities in the prototype network, and to assess their impact on the network operation so as to design and conduct experiments in a manner that will minimize the bias introduced.

ACKNOWLEDGMENTS

We would like to acknowledge the Packet Radio Working Group and particularly Ralph Jones, David Retz, Ron Kunzelman and Don Nielson from Stanford Research Institute and Dick Sunlin from Collins Radio Corporation for the fruitful discussions that have taken place with them.

TABLE I—Summary of Measurement Items

Pickup Packets (at each PRU, the following data items are collected in the Pickup packet):

- time of arrival of the packet at the PRU
- time the Pickup packet was first placed on the transmit queue
- time of each transmission
- time HBH ack arrived (stored in next Pickup packet)
- the current PRU ID

PRU based Cumulative Statistics

- # of packets received in error
- # of packets received but not intended for this PRU
- histogram of # of transmissions per successful packet
- # of unsuccessful packets (dropped because of lack of ack)
- # of packets discarded because of lack of buffer space
- # of alternately routed ("ALL") packets received
- table counting number of correctly received packets from immediate neighbors
- # of transmissions beyond success
- # of packets incurring transmissions beyond success
- table sampling frequency of buffer states (and transceiver states)

End-Device Cumulative Statistics

- histogram of round-trip times
- # of packets transmitted
- # of duplicate packets detected
- # of packets discarded by the sender because of lack of ETE ack
- histogram of # of transmissions per successful (ETE) packet
- histogram of packet intergeneration time

Note: certain Cumstat items will distinguish between inbound (to the station) and outbound (from the station) traffic.

REFERENCES

1. Roberts, L. G. and B. D. Wessler, "Computer Network Development to Achieve Resource Sharing," Spring Joint Computer Conference, *AFIPS Conference Proceedings*, 1970, Vol. 36, pp. 543-549.
2. Kleinrock, L., *Queueing Systems, Vol. II, Computer Applications*, Wiley Interscience, 1976.
3. Kleinrock, L. and W. Naylor, "On Measured Behavior of the ARPA Network," in Nat. Comput. Conf., *AFIPS Conf. Proc.*, Vol. 43. Montvale, N.J., AFIPS Press, 1975, pp. 767-780.
4. Kahn, R. E., "The Organization of Computer Resources into a Packet Radio Network," in *Nat. Computer Conference, AFIPS Conference Proceedings*, Vol. 44. Montvale, N.J., AFIPS Press, 1975, pp. 177-186.
5. Kleinrock, L. and F. Tobagi, "Random Access Techniques for Data Transmission over Packet-Switched Radio Channels," in *Nat. Comput. Conf., AFIPS Conf. Proc.*, Vol. 44. Montvale, N.J., AFIPS Press, 1975, pp. 187-201.
6. Frank, H., I. Gitman and R. Van Slyke, "Packet Radio System—Network Considerations," *in Nat. Comput. Conf., AFIPS Conf. Proc.*, Vol. 44. Montvale, N.J., AFIPS Press, 1975, pp. 217-231.
7. Kleinrock, L. and H. Opderbeck, "Throughput in the ARPANET—Protocols and Measurement," *Fourth Data Communications Symposium*, IEEE Catalog Number 75 CH 1001-7 DATA, Quebec, October 1975.
8. Fralick, S. and J. Garrett, "Technological Considerations for Packet Radio Networks," in *Nat. Comput. Conf., AFIPS Conf. Proc.*, Vol. 44. Montvale, N.J., AFIPS Press, 1975, pp. 233-243.
9. Burchfiel, J., R. Tomlinson and M. Beeler, "Functions and Structure of a Packet Radio Station," in *Nat. Comput. Conf., AFIPS Conf. Proc.*, Vol. 44. Montvale, N.J., AFIPS Press, 1975, pp. 245-251.
10. Fralick, S., D. Brandin, F. Kuo and C. Harrison, "Digital Terminals for Packet Broadcasting," in *Nat. Comput. Conf., AFIPS Conf. Proc.*, Vol. 44. Montvale, N.J., AFIPS Press, 1975, pp. 253-261.
11. Tobagi, F., "Random Access Techniques for Data Transmission over Packet Switched Radio Networks," Ph.D. dissertation, Comput. Sci. Dep., School of Eng. and Appl. Sci., Univ. of California, Los Angeles, rep. UCLA-ENG 7499, Dec. 1974.
12. Kleinrock, L. and F. Tobagi, "Packet Switching in Radio Channels: Part I—Carrier Sense Multiple Access Modes and Their Throughput Delay Characteristics," *IEEE Trans. Commun.*, December 1975, pp. 1400-1416.
13. Tobagi, F. A. and L. Kleinrock, "Packet Switching in Radio Channels: Part II—The Hidden Terminal Problem in Carrier Sense Multiple Access and the Busy Tone Solution," *IEEE Trans. Comun.*, December 1975, pp. 1417-1433.