2/28/1977
Lou Nelson
UCLA

THOUGHTS ON MULTI-NET EXPERIMENT CONTROL AND DATA COLLECTION FACILITIES

I. INTRODUCTION

This PSPWN was written to inform the group members about the current state of our thinking about measurement and control in the SIMP III and PRNET environments. Most of the issues presented here were identified at the small group meeting at BBN on January 20-21, 1977. Most of the opinions are mine and are not binding on anyone.

Before the end of the year we hope to have three networks which will be available for internetting experiments: the ARPANET; the Packet Radio Net (PRNET); and the SATNET. It is very desirable that an experimenter be able to run an experiment which utilizes the control and data collection tools for all three networks simultaneously from one control program and to have all the measurement data be collected at one site.

Currently we believe that a TENEX site on the ARPANET is the most convenient system from which to effect control of the experiments and the UCLA IBM 360 is the most convenient measurement data collection site for multi-net experiments. In the case of measurement experiments on the SATNET alone, any of the experimenters' hosts may serve as the measurement control and data collection site(s). The following discussion will first treat the requirements of the multi-net experiments.

II. MULTI-NET EXPERIMENT CONTROL

In all of the networks, experiment control will be exercised by sending parameter change messages of some form to the IMPs, SIMPs, and Station/Repeaters. In the ARPANET they are traditional parameter change messages with ARPANET leaders. In the SATNET they are traditional parameter change messages with internet headers. A traditional parameter change message contains parameter number and parameter value pairs as data. The parameter number is used as an index into the local parameter table for storing the new parameter value. For the PRNET we assume they will be special messages with internet headers.

Setting up and sending the parameter change messages is not difficult but is a tedious chore. It is very advantageous to have a program set up the messages for the experimenter by asking questions about the format of the experiment and then sending the appropriate parameter change messages to the subset(s) of the network(s) involved in the experiment. The format of the experiment may also be stored in a file for later minor modification and then be used as the source of the parameter change messages for a similar experiment. We have found that it is important that the control program run interactively as opposed to running in batch mode. It does not have to run with tight real time constraints but faster response to, say, a request to kill an experiment than five to ten minutes is required. For these reasons we believe a program or set of programs which runs on a TENEX system is desirable.

II.A MULTI-NET EXPERIMENT CONTROL EXAMPLE

In figure 1 (at the end of this note) a typical three network experiment is shown. The solid arrows indicate flow of control information from the TENEX to each of the three networks. The ARPANET parameter change messages are addressed to the parameter change fake host in the IMPs and are acknowledged via RFNMs.

The SATNET parameter change messages may be addressed to the fake host in the SIMPs or to a particular SATNET gateway. In the former case, the address of the fake host in the SIMP resides in the internet header in the packet and on top of that header is the ARPANET leader which contains the ARPANET address of the gateway into the SATNET that is to be used. The first acknowledgment for the message is the RFNM sent back by the gateway's ARPANET IMP. In addition the destination SIMP acknowledges by sending back to the TENEX all or part of the contents of its parameter table in an internet packet.

In the SATNET gateway case, the parameter change messages will be addressed to a measurement process in the gateway. That address resides in the internet header in the packet and, again, on top of that header is the ARPANET leader which contains the ARPANET address of the SATNET gateway. The first acknowledgment for the message is the RFNM sent back by the gateway's ARPANET IMP. The measurement process in the gateway then sends a special control message (whose contents could be the parameter change message itself) to its SIMP which broadcasts it to all the other SIMPs. A special parameter change message which requests that a SIMP send the contents of its parameter table back to the TENEX program is required to verify that all the parameter change messages arrive correctly.

The PRNET parameter change messages will be addressed to the measurement control "fake host" (?) in the gateway/station whose address will reside in the internet header. The gateway/station's address will reside in the ARPANET leader of the message.

For the PRNET the content of the special parameter change messages will be measurement control text lines which the

PRNET station software will translate into PRNET parameter change messages that will then be sent to the PRUs and TIUs using the SPP protocol. The first acknowledgment will be an ARPANET RFNM from the IMP connected to the gateway/station back to the TENEX. The final acknowledgment probably should be a message containing the gateway/station's PRNET parameter table that is sent in response to a special request or 'go' message from the TENEX control program.

An experiment commences when a 'go' parameter change message is sent from the TENEX control program to each network site involved in the experiment. This message triggers the message generators and the statistics fake hosts. The 'go' signal is also sent to he collection site(s). The experiment is stopped after a preset interval by a 'stop' message sent from the control program to each site to turn off the generators and statistics fake hosts. The collection site programs are turned off by a similar 'stop' message.

III. MULTI-NET EXPERIMENT DATA COLLECTION

Whereas experiment control is pretty much the same regardless of how many networks are involved, measurement data messages are of three diverse forms and come at different data rates.

III.A ARPANET MEASUREMENT DATA

We are very familiar with the ARPANET measurement statistics messages and know that they can be handled by the current collection software on the CCN IBM 360. In figure 1 the flow of the measurement packets is shown by the dashed arrows through the UCLA IMP into the IBM 360. No end to end protocol is used (no NCP in the IMPs).

III.B SATNET MEASUREMENT DATA

We don't yet know the exact form of the SATNET statistics messages in SIMP III, but we believe they are similar to the current SATNET Cummstats in size and the necessary data rates to get them to the collection site. They are raw internet packets and there are probably six to eight of them per Cummstat. They are generated by the statistics fake host in the SIMP and may be sent out in a group or one at a time. That is, a bit may be set in the header of each packet to request that the SIMP notify the fake host when the packet has been delivered to another host in the SATNET (the gateway in this case) if ordering of the packets is desired by the experimenter. The notification may be either explicit by an 'ACK' SIMP-host control message (which is triggered by the SATNET channel ack) or implicit by accepting the next packet from the fake host after several tries. No end to end protocol is used.

In figure 1 the flow of the packets is shown from the SATNET through the gateway into the ARPANET. They then go across the ARPANET to the UCLA IBM 360. Note that the packets are delivered to the UCLA IBM 360 in order of entry into the ARPANET if they are regular messages. So, ordering of the measurement packets may be preserved across both networks if desired.

III.C PRNET MEASUREMENT DATA

Measurement messages, both Cummstats, or their analogs, and pickup packets travel from the PRUs and TIUs to the gateway/station. There, they are put on disk AND are sent through the gateway into the ARPANET to the UCLA IBM 360. The packets may be spooled onto disk and then sent out over the ARPANET or may be put on disk and sent out over the ARPANET immediately. In either case Stan Lieberson at UCLA wants all the data stored on disk at the gateway/station regardless of their other destination in order to achieve reliable collection in case the link to the collection site or the collection host goes down. In figure 1 we show the measurement data flow again by the dashed arrows from the PRNET through the gateway/station into the ARPANET and thence to UCLA.

One issue is what if any end to end protocol should be used in transferring the data from the gateway/station to UCLA. The three alternatives are:

1. Use TCP 0. This has to exist in the gateway/station and the collection host.
2. Use nothing except the ARPA Subnet limitations on host pair bandwidth for ARPANET type 0 messages and rely on the ARPANET to deliver the packets in order. A maximum of 24KB may be "in the pipe" if full eight packet messages are used and 8KB if single packet messages are used. We should be able to achieve one second round trip times so data rates of 24KBPS and 8KBPS result.
3. Don't send the packets over the ARPANET while the experiment is in progress but wait until later and use the ELF system FTP or TCP0 with the PDP11 as an ordinary ARPANET host to send the data file to UCLA.

I prefer the second alternative of using no explicitly limiting protocol. Our experience has been that there are no difficulties in receiving the measurement data packets at a collection site during ARPANET and SATNET experiments. Probably only a much higher data rate requirement by the PRNET measurements could cause problems. If that is true then I favor the third alternative with perhaps the modification of using TCP 0 instead of FTP for transferring the data.

III.D COLLECTION PROGRAMS AT UCLA

The measurement data collection programs now have to accept and keep separate data packets from three networks, so three files are filled simultaneously. Separate packets of the multi-packet measurement messages are "reassembled" before being written out to a disk file by ordering the packets (if necessary) and stripping the internet headers from the second through n-th packets. The usual current timestamping of each message's arrival at the collection site is continued. All additional processing is delayed until the reduction programs are run after the experiment has been completed.

By looking at figure 1 again it is obvious that some signals have to be received by the collection programs to start and terminate data collection. The 'go' signal is a state record sent for each network from the TENEX control program which is to be written at the head of the measurement data file for each network. Any measurement data packets received before the 'go' message arrives are discarded. The 'stop' message may be either a special packet for each network from the TENEX control program or a 'trailer' network state record which is appended to the end of each file.

IV. SATNET EXPERIMENTS IN SIMP III ENVIRONMENT

The principles set forth above apply to the SATNET experiments with few modifications. They are typically single network experiments and the second network (the ARPANET) is only used for transporting measurement control messages to the SATNET and measurement data messages to UCLA or another collection site for reduction and analysis. A major issue not addressed above is the degree to which measurement data is segregated from other artificial traffic and kept out of the satellite channel during the experiment.

There are some obvious possibilities for handling the measurement data:

1. Send all the measurement data intermixed with other data on the channel to one gateway (ETAM or COMSAT) for transmission to the data collection site (UCLA or COMSAT). Note that all the data can be channeled to any gateway (e.g. NDRE or UCL also).

2. Send all the measurement data "out of band" on the satellite channel (similar to the way routing packets are sent in SIMP I) to one gateway as in 1.

3. Send measurement data over the ARPANET. Since each SIMP has a gateway host attached to it, each SIMP may route all its

measurement data to its gateway for transmission to the collection site. Currently it is planned that all the gateways except the COMSAT gateway are to be connected to the ARPANET so measurement data may go to the collection site without reentering the SATNET except for COMSAT.

One way to include the COMSAT SIMP's measurement data is to have the COMSAT host temporarily store the data until the 'stop' message is received from the control program which signals that the artificial traffic phase of the experiment is over. It then sends the measurement data to the primary collection site (UCLA, say). Reception of the 'stop' record at the end of the file (referred to in section III.D) signals the primary collection site that the experiment is over.

One open question is whether or not to use and end to end protocol for this data transfer such as TCP 0. I favor NOT using any end to end protocol.

If the primary collection site is the COMSAT host, another collection host such as UCLA can collect the measurement data from the other three SATNET SIMPs and reflect them to COMSAT in the last phase of the experiment.

We should decide soon which alternative should be used for collecting measurement data.