Title
A DATA LINK BETWEEN SMALL COMPUTERS AND A CDC 6600

Authors
Andreae, Sypko W.
Lafore, Robert W.

Publication Date
1967-10-26

UCRL-17622

University of California

Ernest O. Lawrence Radiation Laboratory

A DATA LINK BETWEEN SMALL COMPUTERS AND A CDC 6600

Sypko W. Andreae and Robert W. Lafore Jr.

October 26, 1967

DISCLAIMER

Presented at the IEEE 14th Nuclear Science
Symposium, Oct. 31-Nov. 2, 1967, Los Angeles,
California

UCRL-17622
Preprint

UNIVERSITY OF CALIFORNIA

Lawrence Radiation Laboratory
Berkeley, California

AEC Contract No. W-7405-eng-48

A DATA LINK BETWEEN SMALL COMPUTERS AND A CDC 6600

Sypko W. Andreae and Robert W. Lafore Jr.

October 26, 1967

A DATA LINK BETWEEN SMALL COMPUTERS AND A CDC 6600*

Sypko W. Andreae and Robert W. Lafore Jr.

Lawrence Radiation Laboratory
University of California
Berkeley, California

At LRL small computers are used on-line with high-energy physics experiments. When the experimenter desires to estimate the relative value of the experimental data, the necessary analysis can only be performed on a large computer. Hence the need arose for a fast two-way Data Link between one or more small computers and a large computer.

The lack of interrupt facilities in the available large computer and the necessity of using twisted-pair telephone lines both posed unusual design problems.

The main components of the Data Link are interfaces for the two computers, a buffering and error-correcting system, and a transceiver system using 1-mile-long telephone lines. A unique demand and response system both maintains synchronization and supports a highly accurate error-correction system.

Control words communicate word count, transfer direction, etc., between the computer programs. Special high-reliability pulse-train signals transmit critical information, such as "error" and "end of record" between the interfaces.

Objectives of the Data-Link System

Introduction

At the Lawrence Radiation Laboratory in Berkeley, small computers, mostly PDP5's and 8's, are used on-line to gather data from high-energy physics experiments. The data is typically stored on tape by the small computer and later analyzed off-line by a CDC 6600. The small computer can provide the experimenter with some very simple checks of the incoming data--for example, whether the experimental equipment is working at all--but does not have the size or power for performing sufficient analysis to indicate whether the experiment is producing the desired data in terms of physics. Thus the experimenter must wait for the hand-carrying of the data tape from the experimental area to the 6600, and for the normal batch processing of the job, before he can make any necessary corrections in his experiment. This feedback loop might occupy hours or even days.

From this situation there arose the need for a Data Link between the small computers and a powerful machine capable of complete data analysis. The relatively small throughput which it was anticipated the link would carry precluded the use of a medium-size time-sharing computer, or a small computer to interface the link with the 6600. Instead, the link was to provide a direct connection, via telephone lines up to several miles long, between the small experimental computer and the CDC 6600. The link was to incorporate the following design objectives:

1. Modifications to the 6600 were to be minimized. Since the link would constitute a very small part of 6600 throughput, interference or serious degradation of the normal batch processing could not be tolerated. Also, modifications to the 6600's Chippewa Operating System were to be kept as simple as possible.
2. A reasonably high data rate was required, preferably exceeding that of the high speed tape units already used by the 6600.
3. Error-correction facilities were to be kept as much as possible within the data link itself, to avoid complicated software checking by both the small computer and the 6600. Also, since the data was to be carried by twisted-pair phone lines, the error-detection system needed to be quite powerful.
4. The link would need to be used only occasionally, no more often than every half hour or so, for the transmission of 20 to 30 000 words of data.
5. Rapid response by the 6600 to the link was unnecessary: a lapse of several minutes from the time the link requested service from the 6600 until the 6600 was able to respond was acceptable.
6. The link was to be kept as general-purpose as possible to enable it to be used with other types of computers if the need arose.

Experimental Environment

At this point it may be helpful, in order to clarify the intended function of the link, to review some typical experimental techniques employed by the physicist and see how the data generated is used.

Let us compare two much-used approaches to the recording of data from high-energy nuclear events, the bubble-chamber technique and the counting technique.

---

*This work was done under the auspices of the U. S. Atomic Energy Commission.

The bubble-chamber approach uses primarily a photographic process to record an event. The photographs are later analyzed by elaborate man-machine scanning systems which yield data in the appropriate digitized form. This data can be fed into a large computer, where the analysis is performed. The bubble-chamber film frames contain in general much more information than is abstracted from them during the first analysis. It is therefore possible to go over the same film frames again and again to analyze different phenomena from the same events.

By contrast, the counting approach uses a technique by which the experimental data is immediately translated into a digitized form. Many years ago counters were the main devices used to register the data. Today we see many other digital data-producing devices, for instance, spark chambers and photomultipliers. The success of a spark-chamber experiment depends largely on how correct the physicist was initially in assuming which phenomena were to be expected. His experiment is carefully aimed at one or perhaps a few phenomena, and if well aimed, his data will be rich in information on only those phenomena. Selecting phenomena from an event and digitizing them is here done on-line, in contrast to the approach taken with bubble-chamber film frames, where selecting and digitizing takes place during the off-line scanning process. Since the physicist involved in such counting physics experiments can never be completely sure about his assumptions, he needs to be able to adjust his experiment when he perceives that his aim was poor.

The limited capabilities of a small computer like the PDP-5 permit only the most rudimentary analysis of the data. However, simple checks can be performed to see if the experimental equipment works as expected and if the data is valid in a very general sense. The results of these simple checks can be displayed on an oscilloscope, and a program can be constructed to provide several displays which can be requested via teletype or the console of the small computer.

These displays may show the experimenter that his equipment is working correctly, but what it does not show him is the quality of the data in terms of physics: Does all this data bring him nearer to his experimental goal? It would be desirable to perform "analysis in depth" on-line.

The link system was therefore conceived to provide a reliable high-speed data-transfer medium using private telephone lines over distances of a few miles. The data is transmitted in records of a certain number of 12-bit words per record; each word is transmitted in parallel. In addition to data transmission a conversation is carried on between the input-output (I/O) programs of the small and the large computer to exchange relevant information on the data records before and after each data record is sent.

Aspects of the CDC 6600

Several aspects of the CDC 6600 system are considered in order to illustrate some of the problems encounterd in the design of the Data Link. The CDC 6600 system is constructed to protect the central processing unit (CPU) as much as possible from I/O interference. As Fig. 1 shows, the CPU can be thought of as being surrounded by a protective layer of a 131K 60-bit word core memory, around which ten peripheral processors (PP) are situated. The only way the CPU can communicate with the PP's is via the 131K core memory. Nearly all PP's are involved in I/O communications. Their tasks are assigned by a controlling I/O program on the basis of availability, which makes for a very efficient use of the PP's. Thus it is possible to keep all PP's busy for most of the time, alternating between a number of assignments that is far higher than the number of PP's. In this sense there exists a true time-sharing within the system. This same time-sharing approach is taken in many other areas of the CDC 6600, and given the correct software, this indeed results in a very fast computer.

Still this computer system lacks the essential (hardware) mechansims to make efficient time-sharing of the CDC 6600 possible between devices outside the CDC 6600. For example, there is no way in which the PP's can be interrupted by a signal from the outside world, in the way this is possible with interrupt hardware in many other machines. Of course, one can to a certain extent make up for that by software simulation of the desired mechanisms. In our situation this was virtually impossible. The systems software used (Chippewa) is made for batch processing of jobs and would need major alterations to implement such simulation. Another important possibility was to dedicate to the Data Link one PP which could perpetually test for conditions indicating whether the Data Link would be alive or not (simulation of an interrupt). However, dedication of a PP for one single I/O task obviously violates the PP time-sharing principle and therefore decreases the throughput of the system, especially when the job mix tends to be I/O-limited. It was therefore decided that a PP would not be dedicated but only temporarily assigned to the Data Link, just as with any other I/O equipment.

Use of the Data Link

The only way a conversation can start between programs on both sides of the Data Link is to give the CDC 6600 system the initiative. However, it is up to the experimenter or the program of the small computer to decide whether enough of the right data is gathered to be sent over to the analyzing Fortran program. To reconcile these conflicting requirements, a procedure is followed in which the human operator of the CDC 6600 is included in the interrupt system. In fact, the operator is the only part of the CDC 6600 system that can be interrupted from the outside world, using the current operating system. Near his console a light signal indicates that the small computer is making a service request. As soon as he can reserve a control point and enough space in core, he will load the Fortran Analysis program in central core memory, just as he would a batch-process job. However, this Fortran program is placed into the job stream from the side, avoiding the normal priority mechanisms, and slowing the normal job flow slightly. It should be noted that although operator intervention at this point will be time-consuming in terms of the computer time scale, it is still well within the acceptable overall response time of several minutes specified in the design objectives.

After the Fortran program is compiled it soon requires input or output. At this point control is transferred to a Chippewa subroutine which calls for a special I/O routine to be loaded into some PP. This I/O routine is made to deal with the Data Link for the duration of at least one physical record and its surrounding control communications. Depending on wait times and system needs, the originally assigned PP may quite possibly be withdrawn by the system to be used for some other assignment, and a different PP assigned back to the Data Link for the next physical record.

It is mentioned above that the large computer would preferably be the same computer already used for off-line processing of the experimental data. One reason for this preference is the existence of Fortran analysis programs and subroutines within the computer system, which it would be desirable to use unmodified for both on-line and off-line analysis.

Using the link system, these programs need not have their READ and WRITE statements changed at all. Instead, several control cards tell the Chippewa system during compilation to replace all READ TAPE statements with READ Data Link statements.

Higher-Level Program Interaction

The Data Link can be thought of as merely a vehicle for data and special instructions between the two programs. Since this vehicle is available, communications on a higher level are possible: Many kinds of interaction between the two programs can be invented to make the Data Link more versatile. For example, the small computer may send a record of data. This record, either by prearrangement, or by means of the accompanying command words, is to be interpreted by the Fortran program as a large set of instructions to itself. Branches can be modified, switches set, the different available analysis approaches chosen, and quantity, kind, and format of output results selected. Thus the experimenter has available to him, in a limited way, an extended on-line processing capability which includes some control over the analyzing process.

Design Aspects for the Link

The Data Link is designed as an asynchronous logic device. Basically it consists of a string of buffers, arranged more or less as 12 parallel shift registers. The transmitters, telephone lines and receivers are located halfway along this buffer string. When data transfer starts, the first buffer is loaded from one of the computers with a 12-bit word, and that buffer is then declared FULL. If the control logic senses that the next buffer in line is EMPTY, data is transferred from the first buffer into the next buffer. This process repeats itself for all the buffers in the string. When for some reason the last buffer is not emptied at the receiving end, it is possible to fill up all the buffers of the Data Link, at which point the transfer of data stops. As described here the control logic only needs to sense the condition of pairs of neighboring buffers in order to decide whether the transfer between the buffers should take place or not. A difficulty arises where two neighboring buffers are located on different sides of the telephone lines. It is desirable to avoid dividing the control logic for such a pair of buffers between the two parts of the Data Link on each side of the telephone line, because of the high number of control signals that would have to run back and forth between the two locations. In the Data Link the solution chosen is the use of a response signal. Each word transmitted over the telephone lines to the other location is received and echoed back (retransmitted) as a response signal. Echoing occurs only when a new word is welcome, that is, when the first receiver buffer is EMPTY. Thus, this response signal is used in the same way as are the FULL and EMPTY signals from buffers on the same side of the telephone lines.

In addition the echoed word is compared with the originally transmitted word. If the echoed word is in error, the same original word is transmitted again. This may be repeated many times until finally the echoed word returns successfully and the next word can be transmitted.

A more detailed description of this error-checking scheme is given later in this paper.

Speed of Data Transfer

From the previous paragraph it now becomes obvious that the travel time of the words along the telephone lines is one of the speed-limiting factors of the Data Link. Other factors are the channel performance of the two computers on each end of the Data Link, and the propagation, detection, and deskewing delays in the buffer logic of the Data Link itself. The delays in the Data Link are insignificant and will not be discussed here. Most modern small computers can transfer a record of data at the rate of one word per 2 μs. The PP of the CDC 6600 System is capable of transferring one 12-bit word per microsecond for the duration of one physical record.

If words are to be transmitted every 2 μs, the propagation delay of the telephone cables between the Data Link stations should not exceed 1 μs. Since this represents a cable length of only 600 ft, the Data Link system tends to be limited by the telephone line. The lines will be from 1/2 to 3 miles long, which corresponds with maximum transfer rates between one word per 8 μs and one word per 50 μs. From a system's point of view, these rates are quite acceptable, being of the same order as the data rates used in the existing magnetic tape units of the CDC 6600 system, as discussed above.

Organization of the Link

Link "Conversation"

Before the data itself can be transferred via the link, the I/O routines in the computers at each end of the telephone lines must exchange information regarding the format of the data to be transmitted. This is done by nondata words, which, borrowing from CDC terminology, are called "function words" if they go from the 6600 to the small computer, and "status words" if they are initiated by the small computer.

An initial status word from the small computer is used to signal the operator of a service request, via the console light, and a function word carries back his push-button response (see Fig. 2). Later, when the operator has loaded the Analysis program, the program itself initiates the sending of a function word to the small computer to inform it that the data may now be transmitted.

Since the Fortran Analysis program in the 6600 requests only logical records of a certain length, the PP I/O routine must obtain the physical record length from the small computer: This is sent as a single data word. To read in this word (as to read in any status or data words) the PP first sends a function word, in this case "GO-WORDCOUNT", whose only purpose is to cause the status word to be transmitted from the channel synchronizer (CS) to the PP (see Fig. 3).

Following the transfer of each physical record, the small computer informs the PP that either (a) more physical records must follow to complete the logical record, in which case an "end of physical record" status word is transmitted, or (b) the logical record has been completed, in which case an "end of logical record" is sent.

Data output (from the 6600) is analogous to data input except that the PP must inform the small computer at the beginning of each record whether the record will be composed of data or only an end of file.

Since the status and function words are to a large extent initiated and interpreted by software, the "elements of the conversation" described above can be altered to meet future changes in the 6600 operating system or even to provide the interface for entirely different computers at either end of the link.

Function and status words are distinguished from data by the presence of an extra bit which is transmitted as the 13th bit in parallel with each word. This 13th bit has the logical value "1" when a data word is sent, and "0" when a nondata word is sent. Since every word must contain at least one bit to be recognized by the receiver, the 13th bit also provides a means of recognizing an all-zero data word.

The Carrousel

The four registers connected directly to the long-line receivers and transmitters (see Fig. 3) are called the "carrousel" and constitute the heart of the error checking and correcting system.

During normal data transmission (for example. during input to the 6600), a word flows out of the small computer memory buffer into the B register (BR) in the device synchronizer (DS), then into the CR, and finally the DR, where it is both stored and transmitted to the CS. In the CS it is received in the CR, shifted to the DR, and retransmitted to the DS. Arriving in the CR of the DS, this echo is compared with the original word still in the DR. If the comparison succeeds, both CR and DR are cleared. During the round trip of this first word from DS to CS and back, a second word will have shifted from the small computer memory to the BR, where it awaits a successful comparison of the previous word. When both CR and DR are cleared, this word is free to shift into the CR, and the sequence is repeated again. A similar process takes place for data output from the 6600.

Since I/O operations between the computers and the link are concurrent with the transmission and echoeing of the previous word, the cycle times of the computers cease to effect the data rate of the system (assuming line length is the limiting factor). Also, if the computer on the receiving end is not able to accept the data fast enough, the process simply pauses until the received word is finally read in, thus preventing loss of data or the necessity for complex synchronization and timing devices.

Error Correction

If the comparison check on the echoed word should fail, only the CR is cleared, and a special signal called a "fuzzy signal" is sent to the receiving end to warn that the previously received word was in error. This clears the DR in the receiver and permits the same word to be retransmitted as before. The same word may be sent many times until a correct echo is finally received. However, if this echoeing process should continue too long, the link assumes that a serious fault has occurred and sends function and status words to this effect to the two I/O programs.

During the periods when the link is waiting for the computers or the 6600 operator to give it further instructions, it arranges that both the synchronizers will be in transmit mode. This ensures that any noise word received at either end will be compared with the zeros in the empty DR and erased, as with an ordinary error. (If the synchronizers were left in receive mode they would retransmit any noise word arriving in them and eventually fill up all registers with replicas of the noise word.)

The "fuzzy signal" referred to above for signalling errors is one of several signals that, because of their critical role in the operation of the link, must be as nearly as possible immune from noise. This is accomplished by sending a pulse train instead of a single pulse as with data or function/status words. Because these signals are very rarely sent in comparison with the number of data words, the additional time needed to provide a pulse train does not increase significantly the time necessary to transmit each record. Other fuzzy signals indicate the beginning and end of a physical record and provide a general reset.

State Diagrams

In the design of the data link considerable use was made of state diagrams to clarify the operation of the system (see Fig. 4). The position of various flip-flops (for example, those indicating whether a particular register is full or empty), constitute "conditions," a particular combination of which generates a "state." The state in turn alters the conditions by performing various operations, e.g., shifting a word from one register to another. These altered conditions then generate a new state, and so forth.

Since each synchronizer contains more than a dozen flip-flops, plus many external signals which also provide conditions, it becomes extremely difficult to keep in mind the exact sequences and timing. Therefore, to completely specify the logical operation of the machine, both for the initial design and for future debugging, a computer program was written to simulate the operation of the logic. This proved invaluable in tracking down some of the more obscure design and wiring errors.

Debug Boards

Debugging of the link required the development of another feature of the system: debug interrupt boards. In normal operation the transitions from one state to another are made very quickly and are frequently nonrepetitive: e.g., a single status word is sent, then the link waits for some response from the outside world. This situation would create extreme difficulties in debugging unless some method of slowing down the transition from state to state could be found. (If this were a clocked system, this effect could of course be achieved by merely slowing down the clock.) The interrupt boards fulfill this function by inserting a switch and a light into the output line from each state. Thus, with the switch open, the conditions for a particular state cause the corresponding light to appear; but the effects of that state will not be enacted until the switch is closed, whereupon the light for the next state in the sequence should appear. Since lights indicating the position of the major flip-flops and the contents of each register appear on the front of the synchronizer control panel, the debug interrupt boards provide a means of seeing the status of all active elements in slow motion. Used in conjunction with the computer printout of the sequence of conditions and states, this provides an easy and effective debugging method.

Conclusions

The data-link system provides a reliable and satisfactory solution in its present context to the problem of interfacing small data-gathering computers to a large computer for the analysis of a comparatively small quantity of data at infrequent intervals. Should this type of service become more popular in the future, the link can be readily modified to accept changes in the 6600 operating system (such as using a dedicated PP for all link data) or interfacing to a time-sharing computer devoted to link data handling.

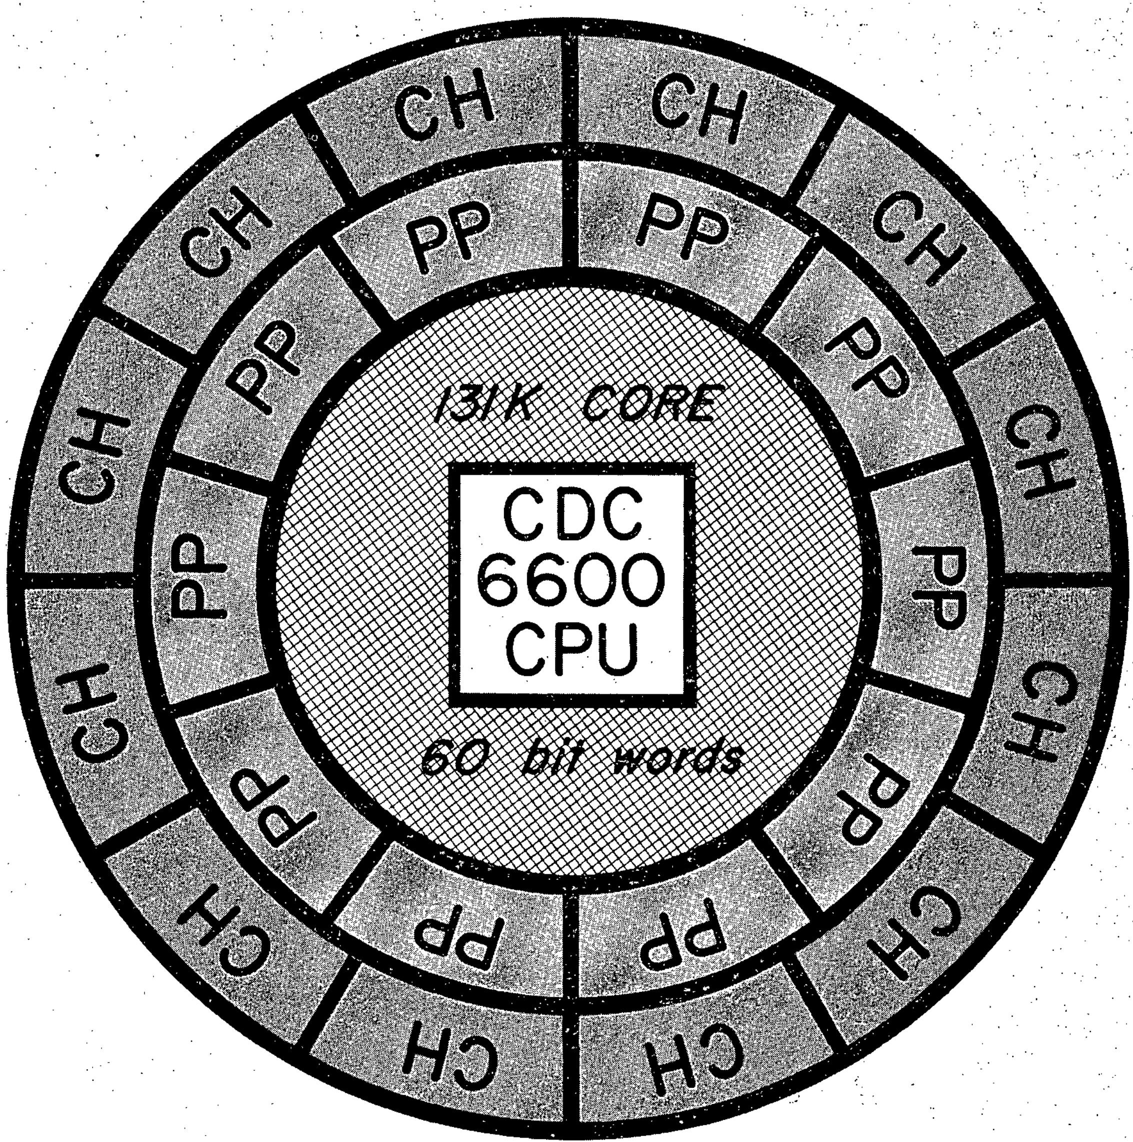

CBB 6710-6236

Fig. 1. CDC 6600.

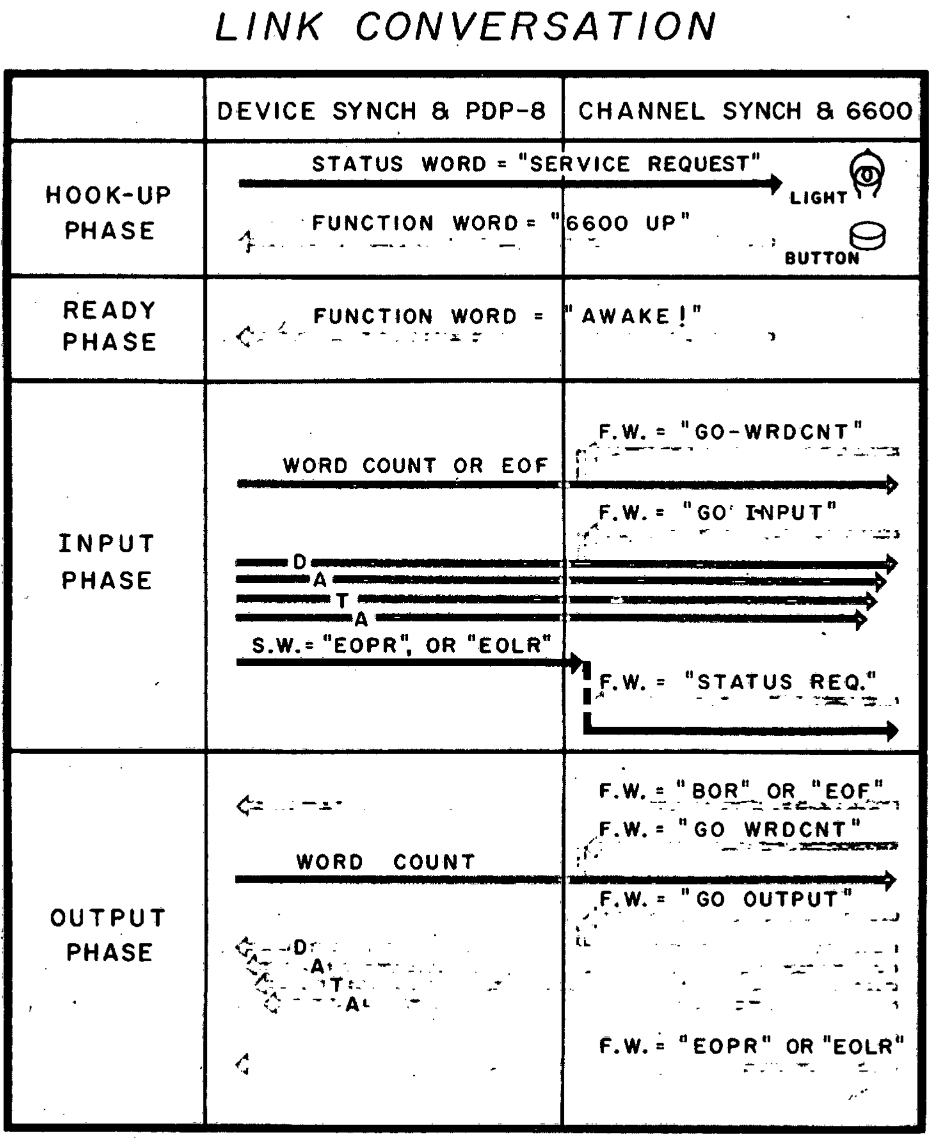

CBB 6710-6240

Fig. 2. Conversation procedure.

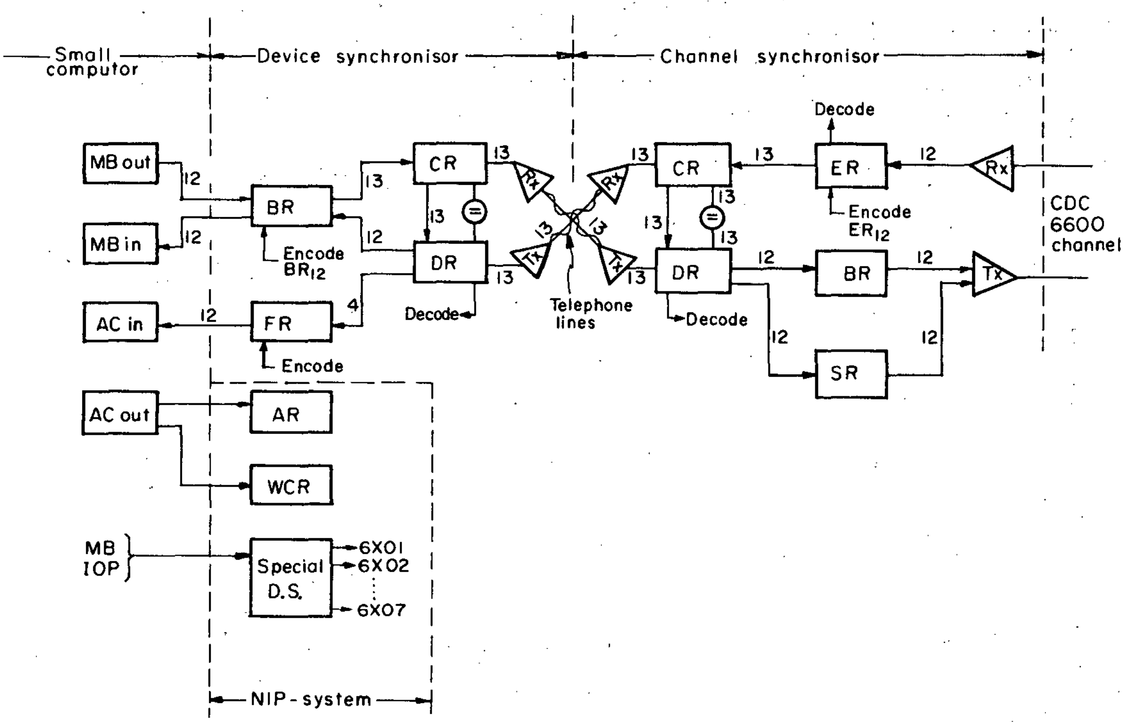

XBL6710-5378

Fig. 3. Data Link system.

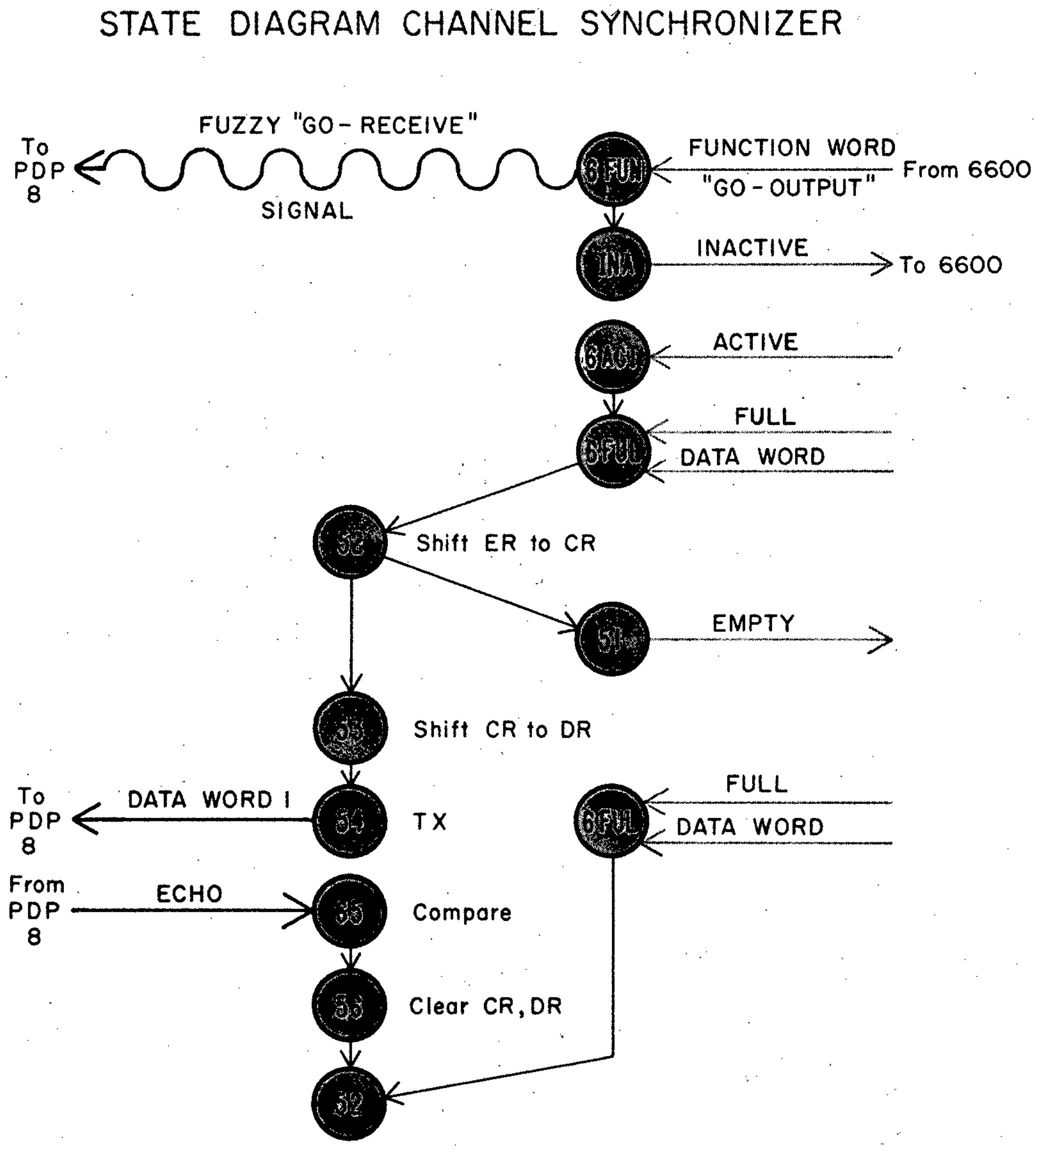

CBB 6710-6238

Fig. 4. Typical state diagram (channel synchronizer, data output).