Title
AN ERROR-CORRECTING DATA LINK BETWEEN SMALL AND LARGE COMPUTERS

Authors
Andreae, Sypko W.
Lafore, Robert W.

Publication Date
1968-01-16

University of California

Ernest O. Lawrence Radiation Laboratory

AN ERROR-CORRECTING DATA LINK
BETWEEN SMALL AND LARGE COMPUTERS

Sypko W. Andreae and Robert W. Lafore Jr.

January 16, 1968

To be presented at the American Federation
of Information Processing Societies Spring
Joint Computer Conference, Atlantic City,
New Jersey, April 30–May 2, 1968 -
and Proceedings

UCRL-17999
Preprint

UNIVERSITY OF CALIFORNIA

Lawrence Radiation Laboratory
Berkeley, California

AEC Contract No. W-7405-eng-48

AN ERROR-CORRECTING DATA LINK BETWEEN SMALL AND LARGE COMPUTERS

Sypko W. Andreae and Robert W. Lafore Jr.

January 16, 1968

AN ERROR-CORRECTING DATA LINK BETWEEN SMALL AND LARGE COMPUTERS

Sypko W. Andreae and Robert W. Lafore, Jr.

Lawrence Radiation Laboratory
University of California
Berkeley, California

January 16, 1968

OPERATING ENVIRONMENT

The need for a data-link connecting small data-acquisition computers to a central computer with great analysis power arose in a particular context at the Lawrence Radiation Laboratory in Berkeley. Both the type of high-energy physics experiments being performed and the operation of the available large computer, a CDC 6600, posed unusual design problems.

Experimental Requiresments

A comparison of two important approaches to the recording of data from high-energy physics experiments is instructive in providing a background for the operation of the data link.

The bubble chamber approach uses a photographic process to record tracks made by particles in a nuclear event. Afterwards the photographs are examined by elaborate man-machine scanning systems to yield data in the appropriate digitized form. This data may then be analyzed by computer. Bubble chamber pictures in general contain much more information than can be abstracted from them during the first such analysis. It is thus possible to scan the same pictures many times to digitize the data relating to different phenomena.

By contrast, the counting approach uses a technique in which the data from the experiment is digitized directly. In the past, electronic counters were the principle devices used to record the data; today there is a variety of such direct-digitizing devices, including spark chambers and photomultiplier tubes. The success of a counting experiment depends largely on how correct the physicist is initially in assuming which phenomena are to be expected. The experiment is specifically aimed at one or perhaps a few phenomena, and if well-aimed, will provide the required data. But if the physicist's assumptions were not accurate, there is then no opportunity to reexamine the data for other phenomena, as with bubble-chamber photographs. The entire experiment must be rerun with the equipment set up to record a different phenomenon.

Computer Availability

The physicist therefore requires rapid feedback in the form of completely analyzed data to permit him to alter his experiment's configuration during a run if required. The capability of the small computers, such as the PDP-8, commonly used for on-line data acquisition at the experiment are too limited to allow this type of analysis, although they can perform simple checks to determine whether the experimental equipment is working normally and whether the data looks reasonable in a general sense. The only way the experimenter can use the large computer is to hand-carry magnetic tapes from the small to the large computer, a process resulting in a turnaround time of hours or even days.

To provide feedback within a useful time the experimenter thus requires sufficient computing power to provide him, on-line, with a detailed analysis of his data in terms of physics. A suitable computer is available at the Radiation Laboratory, a CDC 6600 located several thousand feet from the Bevatron and 184-inch cyclotron where high-energy physics experiments are conducted. Private telephone lines are also available, running more or less along the desired routes. The data link was conceived to provide a reliable high-speed transfer medium between the small and large computers.

The CDC 6600 provided several unusual problems in the design of the data link because of its construction and the way it is normally used at the Radiation Laboratory. The 6600 is designed to protect the central processing unit (CPU) as much as possible from the interference of input/output (I/O) devices. The CPU may be thought of as being surrounded by a protective layer of central memory (see Fig. 1), which is in turn surrounded by 10 peripheral processors (PP) which communicate with I/O devices via 12 data channels. The only way the CPU can communicate with the outside world is via the 131K core memory. Nearly all the PP's are involved in I/O communications. Their tasks are assigned by a controlling PP on the basis of availability, which makes for very efficient use of the PP's. In this sense there is a kind of "time sharing" within the system, an approach used in many areas of 6600 design. This kind of time sharing results in a very fast computer, but not one which is easily used for time sharing in the more usual sense of devices outside the computer. For example, there is no facility for interrupts, so that a PP is forced to continually check a flag from a particular device to determine when service is required. This is satisfactory for such devices as tape units. However, dedicating a PP to watch a flag from the data link would result in a prohibitively heavy demand on the pool of available PP's, since many minutes might elapse between calls for service from the experiment. The data link cannot therefore be treated as a normal I/O device.

Another possibility is to alter the operating system to assign a PP to check at infrequent intervals if the link is requesting service. However, any such modifications are difficult to make on the 6600 operating systems and would result in degradation of the batch processing throughput. It is therefore necessary to include the 6600 operator in the interrupt chain; he is in fact the only part of the 6600 system which can be interrupted.

DESIGN OBJECTIVES

From the situation described above, the following design objectives evolved for the data link:

1. No hardware modifications were to be made to the 6600. Modifications to the software operating system were to be minimized to avoid degradation of the normal batch processing.

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, error-detection capability 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.

DESIGN PHILOSOPHY AND EVOLUTION

The approach to the design of the data link was governed by one important consideration: since a relatively small amount of data handling would require use of the link, cost was to be minimized, and existing equipment was to be used as much as possible. This eliminated immediately such alternate approaches as using a medium-size computer on-line or substituting coaxial cable for the already existing phone lines.

Synchronization

The use of twisted-pair lines, however, raises fairly serious noise problems, since the route through which they run contains many powerful noise-producing devices, such as spark chambers, magnets, and the rf field of the cyclotron. Associated with the noise problem is one of synchronization: how to provide communication between two synchronous devices each running independently on its own clock.

The simplest solution to the synchronization problem is to transmit an echo, or "received your last word" signal from the receiver to the transmitter, and send the next word only when the echo is received. In this way either the receiver or the transmitter can halt the flow of data when the words can not be obtained from, or accepted by, the respective computer. The alternative approach, that of sending all words without interruption at a rate slow enough for the slowest computer to handle even during the worst case, would actually have required a slower transmission rate.

Error Correction

In order to match the format of both the PDP-8 and the PP, data is transmitted in words of 12 parallel bits, with a 13th bit indicating whether the word is data or function/status (see below). Various error-detection schemes were considered. The most common, the addition of a parity bit in parallel with the data word, has decreasing utility as the number of bits of the transmitted word increases. In this case, with 13 parallel bits and the possibility of powerful noise wiping out entire words, parity was considered too weak a system. Another common error-detection method, the checksum, suffers from a similar difficulty in that if the average error rate is high enough for several errors to occur in each block of data, many retransmissions of each block may be required before the record is received with the correct checksum. Also, either storage devices capable of holding an entire data block must be provided at each end, which is prohibitively expensive, or software in the two computers must intervene in the case of an error to cause the data block to be retransmitted.

A method which both solves the synchronization problem and provides a powerful error-detection technique is to echo each word in its entirety back to the transmitter for a complete bit-by-bit comparison before the next word is sent. This of course increases the time to transmit each word, but the increase is not as large as the propagation delay, since reading and writing data between the computers and the link may be overlapped with transmission. However, for the line lengths in use, the propagation delay is the limiting factor: one typical line of 4000 ft has a round-trip delay of about 12.5 μsec.

Line Transmission and Reception

Each individual bit of a word is transmitted over its own line. Each line is a twisted pair, transformer isolated from its associated transmitter and receiver. A full duplex approach is used, and therefore 26 twisted pairs are dedicated to word transfer. An additional 8 lines are used for the control signals, making for a total systems requirement of 34 twisted pairs. These are privately owned telephone lines.

In order to keep the cost low, it was desirable to design transmitters and receivers using simple circuitry but still capable of rejecting a significant amount of noise. Each transmitter consists of two power-nand integrated circuit gates which drive two transistors connected in a push-pull configuration. This produces a bipolar pulse approximating a square wave in the secondary of the output transformer. Two monostable vibrators control the width of the positive and negative areas of the bipolar pulse for all 13 data transmitters in one synchronizer.

Both areas of the bipolar pulse are, in general, equal in width. For reasons of convenience, the width is chosen in this design to maintain an amplitude attenuation of a factor of 4; for instance, 1.3 μsec for a 4000-ft line.

Our initial receiver design actually proved unnecessarily complex. It required the positive portion of the incoming pulse (now resembling a sine wave due to the filtering effect of the lines) to pass through a height window much like a single-channel analyzer. In tests over the actual phone lines, however, we found that, even with the window set quite wide, many more errors were caused by failure to detect an existing bit than by triggering on unwanted noise. The final receiver design therefore requires only that the positive portion of the incoming signal exceeds a certain threshold.

The maximum word-transfer repetitive rate of the link is 80,000 words per sec for a 4000-ft line (12.5-μsec round-trip delay). Obviously the lines of the link are operating with a bandwidth far in excess of the bandwidth of normal telephone lines. In contrast to the latter, the link lines are limited to a length of a few miles and do not run through switch circuits, line amplifiers, etc.

IMPLEMENTATION AND OPERATION

Levels of Communication

Communication over the link system actually takes place on several different levels, as indicated somewhat idealistically in Fig. 2. In a general sense the link may be thought of as a medium of communication between the experiment and the mathematical analysis of the experimental data. This data, and the results of the analysis, are transformed into 12-bit data words by the small computer program and the FORTRAN program in the CPU of the 6600.

The programs however are not restricted to sending data. They can also communicate instructions to each other. For example, the experimenter, via the small computer, can request different analysis approaches or output format, while the FORTRAN program may instruct the small computer to modify a display program depending on the results of analysis. Thus the experimenter has available to him in a limited way an extended on-line processing capability that includes some control over the analyzing process.

At a lower level, it is necessary for the I/O routine in the PDP-8 and the PP's I/O program to communicate certain information involving the timing and format of the data records to be sent. This is accomplished by the use of words distinguished from data by their 13th bit being set to 0, instead of 1, as with data. Borrowing from CDC terminology, these words are called function words if they originate at the 6600 end and travel toward the experiment, and status words if they originate at the PDP-8 end and travel toward the 6600. The 13th bit also permits detection of an all-zero data word, which would otherwise cause no transmission.

On a still lower level, the device synchronizer (DS), which is the portion of the data link at the experiment end of the lines, and the channel synchronizer (CS) at the 6600 end, must communicate information concerning the words being sent. The control signals which transmit this information are critical to the operation of the system, in that any error occurring in their transmission will result in irretrievable failure of the system. This is true because these control signals are responsible not only for initializing and terminating each record sent, but also because they convey information concerning any errors that may occur in the data record.

It is always possible to increase reliability at the expense of increased time, and this was the technique applied to the control signals. Instead of consisting of a single bipolar pulse, as the 13 parallel bits of the data words do, they consist of a train of 16 bipolar pulses. The receivers require that at least half of these pulses arrive: the remaining half may drop out anywhere along the pulse train.

Error Rates

In practice it was found that the error rate was lower than expected. On the 4000-ft line tested, errors (either the dropping of data bits or the addition of noise bits) occurred once in 10^7 transmitted words. However, this rate may well prove to be significantly worse in longer lines. Also, it is difficult to extrapolate the results of error rates from one location or time to another, since the addition of new equipment may result in unexpected large-scale increases in noise.

Tests were run in which the number of times the control signals were transferred was several orders of magnitude larger than would be transferred during the probable lifetime of the link. Not a single error was found during these tests.

The Link "Conversation"

Before an actual transfer of data can take place 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, and also specify the timing. This is done using function and status words as defined above.

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. 3). 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 to the PP.

Following the transfer to 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 at the beginning of each record the PP must inform the small computer 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.

The Carrousel

The four registers connected directly to the long-line receivers and transmitters are called the "carrousel" and constitute the heart of the error checking and correcting system. (See Fig. 4. The register names are arbitrary).

During normal data transmission (for example, during input to the 6600) a word flows out of the small computer memory buffer into the BR in the device synchronizer, then into the CR, and finally the DR, where it is both stored and transmitted to the channel synchronizer. In the CS it is received in the CR and then shifted to the DR, where it is stored and retransmitted to the DS. Arriving in the CR of the device synchronizer, 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. Once both CR and DR are cleared, it will be free to shift into the CR, and the sequence will be 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 echoing 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 will simply pause 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 in the transmitting synchronizer the comparison check on the echoed word should fail, only the CR is cleared, and a control signal is sent to the receiving synchronizer to warn that the previously received word was in error. This clears the DR in the receiving system 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 it should cycle for too long a time the link assumes that a serious fault has occurred and sends function and status words to this effect to the two I/O programs.

While the link is waiting for the computers or the 6600 operator to give it further instructions, it arranges that both the synchronizers are 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.)

CONCLUSION

In the special situation for which it was designed, the data link provides an effective technique for communication of data between computers. Its strong points are:

1. It is relatively inexpensive.
2. It provides a very strong error-correction capability, necessary in its noisy environment.

3. It requires very little modification to the existing large computer, either hardware or software.

On the other hand, the techniques used in the data-link system are not immediately applicable to data transmission over more than a few miles, because the round-trip propagation delay then becomes excessively large, or over commercial phone lines where the bandwidth is restricted to about 3kHz.

ACKNOWLEDGMENT

The authors wish to acknowledge the contribution of Alan Oakes to the design and development of the receivers and transmitters.

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

FIGURE LEGENDS

Fig. 1. CDC 6600 organization.

Fig. 2. Levels of communication.

Fig. 3. Conversation procedure.

Fig. 4. Register arrangement.

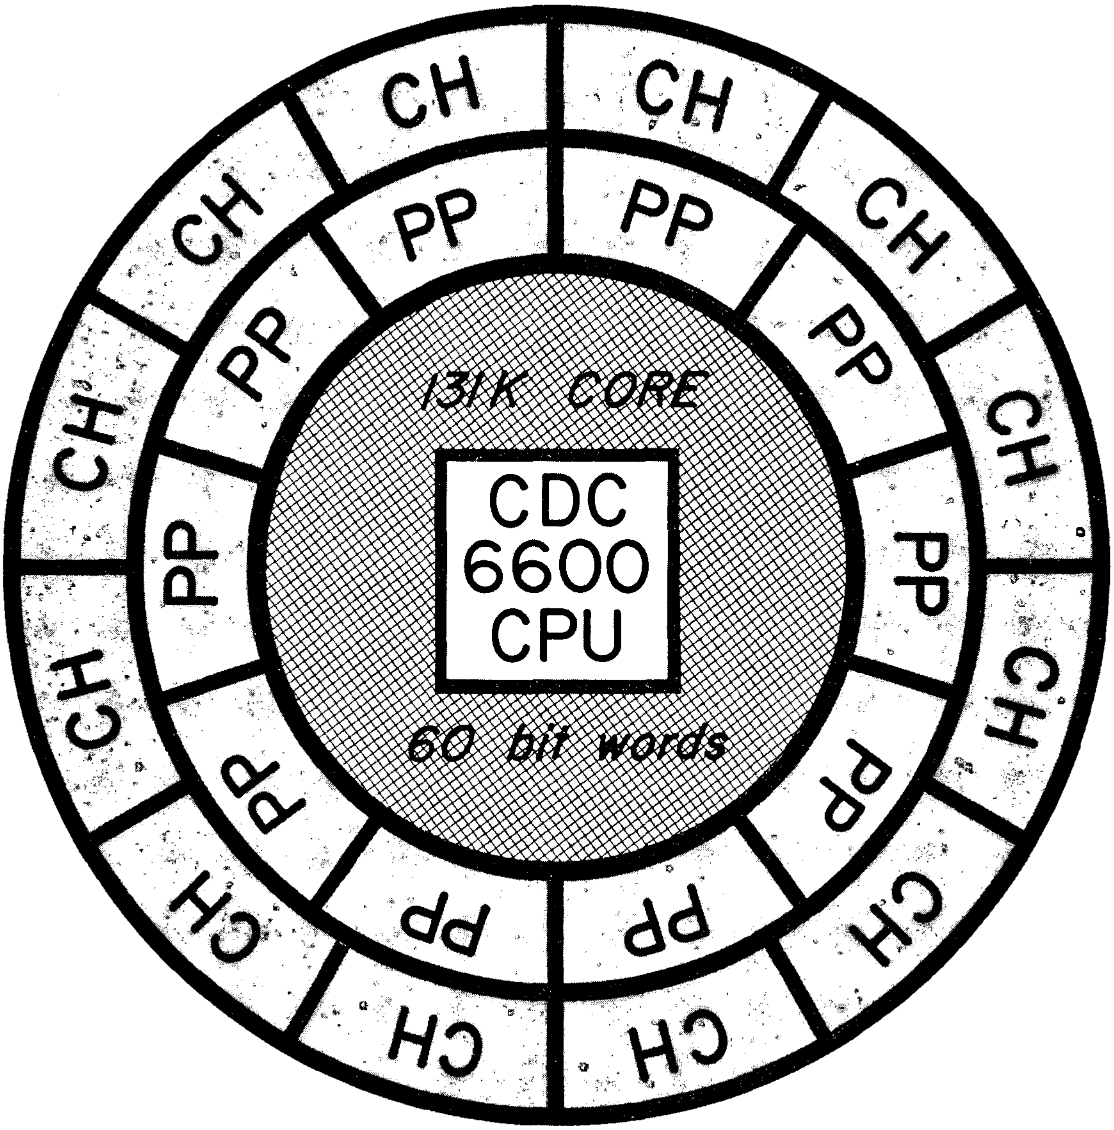

CBB 6710-6236

Fig. 1

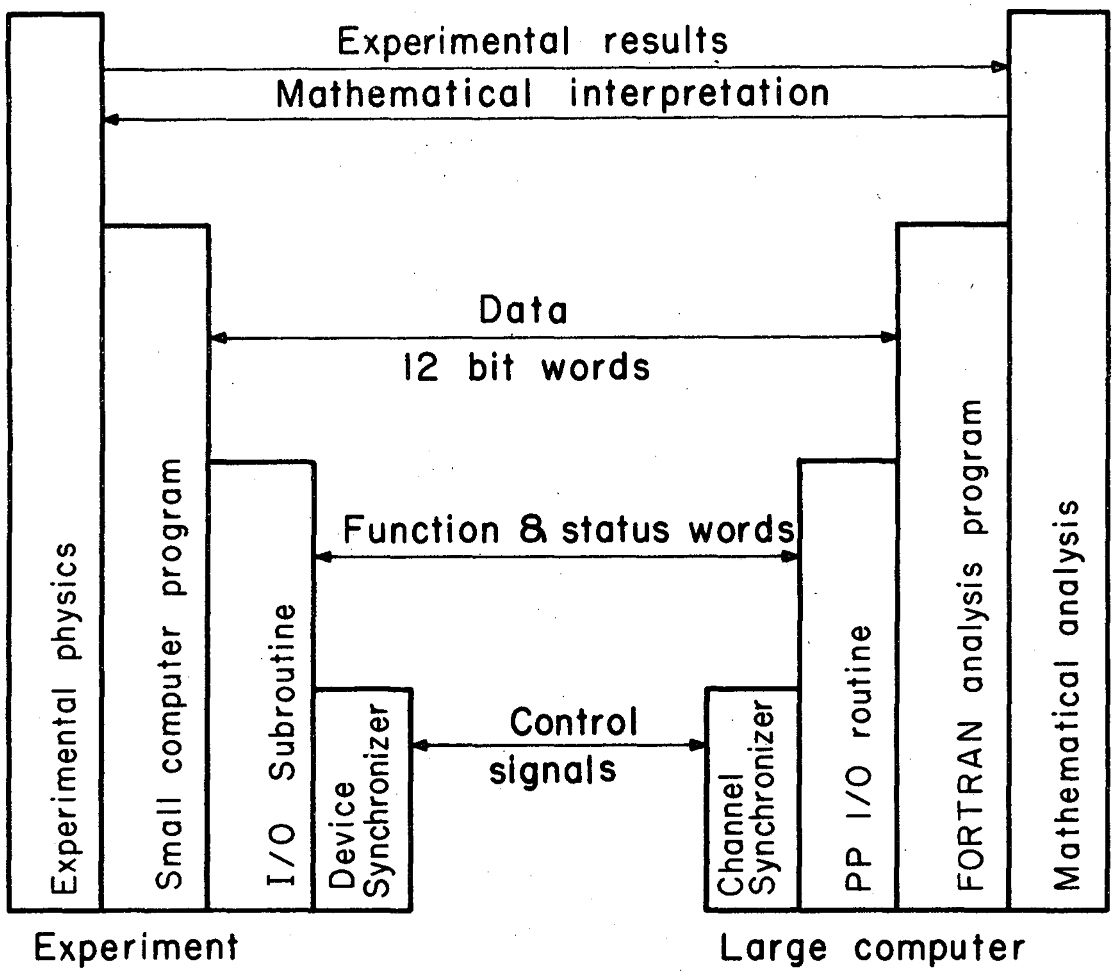

Levels of Communication
in the Data-Link System

XBL681-1572

Fig. 2

LINK CONVERSATION

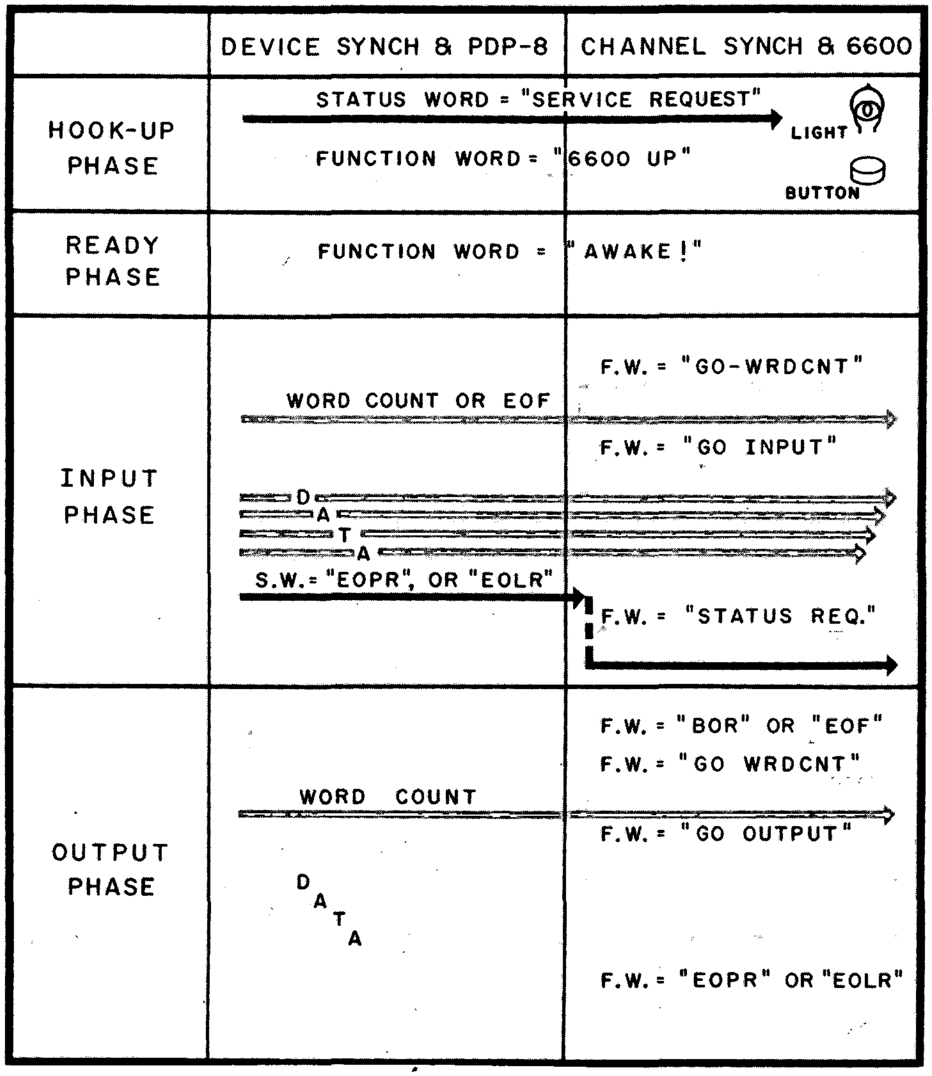

CBB 6710-6240

Fig. 3.

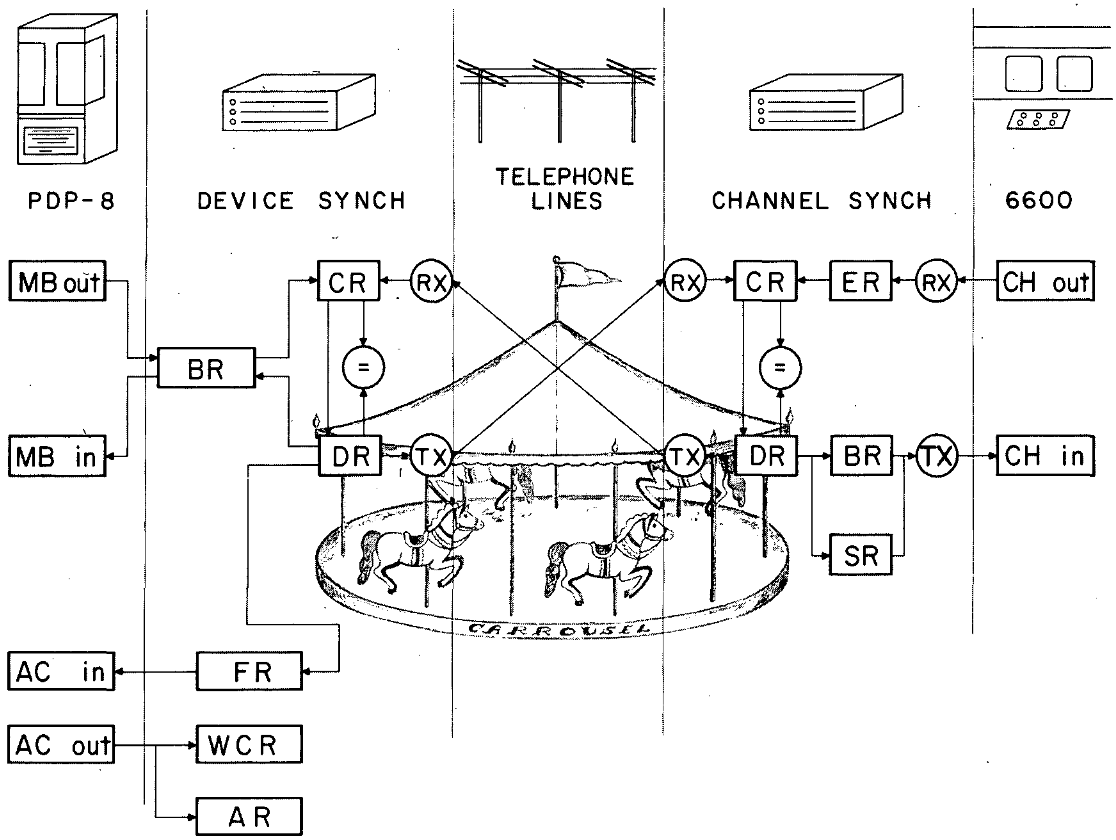

XBB 6710-6260

Fig. 4