JUN 21 1966

DEPARTMENT OF COMPUTER SCIENCE
UNIVERSITY OF ILLINOIS
URBANA, ILLINOIS

COO-1469-0003

MASTER

RELEASED FOR ANNOUNCEMENT
IN NUCLEAR SCIENCE ABSTRACTS

GENERAL DESCRIPTIONS OF A SATELLITE PROCESSOR FOR ILLIAC II*

by

C. E. Carter and R. E. Willard

LEGAL NOTICE

This report was prepared as an account of Government sponsored work. Neither the United States, nor the Commission, nor any person acting on behalf of the Commission:
A. Makes any warranty or representation, expressed or implied, with respect to the accuracy, completeness, apparatus, method, or process disclosed in this report, or that the use of any information, apparatus, method, or process disclosed in this report may not infringe privately owned rights; or
B. Assumes any liabilities with respect to the use of, or for damages resulting from the use of any information, apparatus, method, or process disclosed in this report.

As used in the above, "person acting on behalf of the Commission" includes any employee or contractor of the Commission, or employee of such contractor prepares, disseminates, or provides access to, any information pursuant to his employment or contract with the Commission, or his employment with such contractor.

File No. 654
May 12, 1965

*Reference File No. 652 and File No. 550-56.

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency Thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

DISCLAIMER

Portions of this document may be illegible in electronic image products. Images are produced from the best available original document.

1. General Requirements

The Department of Computer Science of the University of Illinois has constructed and put into service ILLIAC II, a large-scale digital computer. ILLIAC II is used both in computer systems research and in general University service.

One of the parts of ILLIAC II is Interplay, an independently operating input-output control. Interplay controls a general input-output bus with a maximum of 32 time-slotted channels. One or more of the channels of Interplay has been allocated for bilateral connection to a small general-purpose peripheral computer.

This small general-purpose computer will serve as a communications processor. Its role in the augmented ILLIAC II system will be threefold:

- (1) to supervise remote consoles of various types, speeds, and character sets;
- (2) to preprocess any and all remote console information; and
- (3) to communicate this information to and from ILLIAC II.

Investigation by various members of the Department of Computer Science has shown that the three functions may be ascribed to something possibly different from a small general-purpose digital computer. They are more closely properties of a computer communications control. In this light, it seems advisable to give the peripheral processor another name: Satellite Processor Communications System.

With these things in mind, it is easy to visualize a block diagram similar to Fig. 1, in which the box, "Satellite Processor Communications System" (SPCS) performs all three functions described above. ILLIAC II will have direct cable communications with the Satellite and the consoles will be made to think they are communicating directly with ILLIAC II.

FIGURE 1

Next, consider the interfaces. Figure 2 gives a slightly more detailed representation of that portion of Fig. 1 appearing between the two broken lines.

FIGURE 2. SATELLITE PROCESSOR COMMUNICATIONS SYSTEM

The left-hand interface region involves ILLIAC II (Interplay) only, and is to be defined by a close and active interaction with the supplier and ILLIAC II engineers. The right-hand interface, involving connection to many devices, has a multiplicity of properties, each depending upon the particular remote device, its character set, its speed of transmission, and its detailed connection. At present, we are considering Teletype 33 units on both a D.C. loop and data-phone connections as well as IBM 1050 units on data-phone connection.

If we carry the investigation of the SPCS a bit further, Fig. 2 can be expanded into Fig. 3. Figure 3 shows the major divisions in the SPCS. The Satellite Processor Unit (SPU) is the part loosely considered to be the small general-purpose digital computer in the original development. The Communications Unit (CU) is an important additional portion permitting control of the remote consoles and the transmission of their information to and from the SPU.

The one remaining step in developing the pattern is to show that most manufacturers make a further distinction in the chain of functions. It

FIGURE 3

seems to be common practice to separate the communications unit into parts which allow some modular construction. This is shown in Fig. 4.

FIGURE 4

Figure 4 shows a final breakdown into components which represents a commonly accepted picture of remote console control computers today. The Communications Control Unit (CCU) typically represents a major single item

which can be purchased with a large built-in expansion feature. This unit does not depend, in general, on the type of console being served but only on the maximum character rate to be served. The smaller box, called "GROUP," generally is specified by the type of console in use and sometimes is very specialized. The smallest box, "PORT," is the most specialized part of the system for it must respond to all of the peculiar features that distinguish the particular console attached to it. Figure 5 (the complete SPCS) is shown to contain additional input/output gear which though incidental to its primary function as a Communications Computer is essential for initial test as well as for general maintenance purposes. This input/output gear includes a printer (Teleprinter ≥ 10 cps), a reader (paper tape ≥ 300 cps), and punch (paper tape ≥ 60 cps).

FIGURE 5

2. Time Sharing

The purpose of appending to ILLIAC II a Satellite Processing Communications Systems is to facilitate a time-sharing, multi-user system as envisaged by Metze¹ and others. It is presumed that the system has a resident program in ILLIAC II core of at least one 256-word block, whose function is to converse with all active remote consoles. The minimum conversational unit with which ILLIAC II should be forced to deal directly is a line of 60 to 100 characters; smaller conversational units require too much red tape time for the central processor. At the opposite extreme, if the remote consoles are allowed to be at all remote, communications between the consoles and the physical environs of ILLIAC II must be handled serially. The role of the Satellite Processing Communications Systems is, then, to reconcile the two requisite units of information: the character and the line. The primary data-handling function of the SPCS is (for input) to deserialize incoming characters, buffer them into lines of 60 to 100 characters and forward them to ILLIAC II in bursts of one line each. A similar converse function exists for output.

A secondary goal of the time-sharing system, namely, diversity of remote console devices, points out a secondary function of the SPCS--dealing with different character sets. It is anticipated that, at least initially, the majority of remote consoles will be TTY 33 and IBM 1050 units. Since ILLIAC II is well equipped with software for dealing with one character set (BCD), it is proposed that all character sets be recoded to BCD by the SPCS before transmission to ILLIAC II and recoded from BCD by the SPCS in the converse case of output.

Additionally, it is proposed that the SPCS be capable of recognizing some data delimiter characters: End of Line, End of Message, Carriage Return, etc.

From the preceding consideration of the nature of the duties of the SPCS, it is possible to point out some necessary features of the Satellite Processing Unit:

¹Gernot Metze, "On-Line, Time-Sharing Use of ILLIAC II," File No. 554, Digital Computer Laboratory, University of Illinois, Urbana, Illinois, July 11, 1963.

- (1) It must be strongly character oriented to enable packing characters into words.
- (2) It must have a complete set of Boolean instructions to facilitate character recognition and recoding.
- (3) It must have a large, directly addressable core memory and indirect addressing to allow rapid access to line buffering areas.

3. Interface to ILLIAC II

It is proposed that data transfer between ILLIAC II and the SPCS be handled in a burst mode: once a data transfer has been initiated by the SPCS, it will continue at a rate of one SPU word per SPU core cycle until an entire line has been transmitted. This method has two distinct advantages: minimum time to complete a transfer and minimum extra-cost hardware within the SPU. Assuming a maximum line length of 128 characters, packed two per SPU word, and an SPU core cycle of 2 μ sec, the time to transfer a line is at most 130 μ sec; this amount of time is required to perform the transfer whether input/output is overlapped in the SPU or not. Since 130 μ sec of nonavailable processing time is small compared with the maximum allowable response time of the SPU, burst mode transfers are feasible.

In order to make burst mode transfer practical, it is proposed that the ILLIAC II Interplay portion of the interface be a bifurcated channel; it shall be viewed from ILLIAC II as two separate simplex channels, but from the SPU as a single half-duplex channel. Burst mode operation requires that the data transfers to/from the SPU core be controlled asynchronously by word by the channel. Signals required at the SPU-CHPU interface to permit burst mode transmission are listed below:

In-Out	CHPU $\rightarrow$ SPU	1
Data In	SPU $\rightarrow$ CHPU	12 or more
Data Out	CHPU $\rightarrow$ SPU	12 or more
Core Cycle Request	CHPU $\rightarrow$ SPU	1
Core Cycle Reply	SPU $\rightarrow$ CHPU	1
Strobe	SPU $\rightarrow$ CHPU	1 or more

BIPS: Bits in parallel.

BSY: Busy.

CHBSY: Channel busy.

CHEND: Channel end. This can interrupt the CHPU.

CH PRIM: Channel primed. This can interrupt the SPU when Channel Busy is turned on.

FIGURE 6. CHPU-SPU INTERFACE

If addresses related to the core cycle requests need to be furnished to the SPU by the CHPU during each core cycle request, there must also be an added path from the SPU to CHPU for setting up initial addresses and/or decrement counts. In that case, add to the above lists:

Address Out	CHPU → SPU	15 or more
Address In	SPU → CHPU	15 or more
Address in Strobe	SPU → CHPU	1 or 2

It is also necessary for each half of the bifurcated channel to be able to interrupt the SPU, in order that the SPU be made immediately cognizant of a newly created opportunity for data transfer in either direction. Since input and output queuing may become necessary, it will also be required that the SPU be capable of sensing the availability status of the two half-channels. It is preferable for these latter two sensing signals to be dc rather than pulsed; necessary strobing should be done by the SPU. Therefore, the final necessities of this interface are four signals from CHPU to SPU meaning: In Channel Busy, Out Channel Busy, In Channel Interrupt, and Out Channel Interrupt.

4. Interfaces to Remote Consoles

On communication lines of the type discussed herein, data is transmitted serially and each character represents a separate transmission. A character containing k bits of information is transmitted as a sequence of $k + 2$ impulses occupying $k + 2 + \epsilon$ units of time. For three devices of interest, these numbers are:

<u>Device</u>	<u>Code</u>	<u>k</u>	<u>k + 2 + ϵ</u>
IBM 1050	BCD + parity	7	9.1
TTY 33	ASCII	8	11.0
TTY 32	Baudot	5	7.42

The impulse sequence used is: start impulse; k data impulses; stop impulse, wherein the stop impulse is $1 + \epsilon$ units in length and all other impulses of unit length. The existence of a stop impulse of greater duration than the start or data impulses allows the receiving unit to resynchronize at the start of each character. The stop impulse is marking, which is the rest state of the line, and the start impulse is spacing, or the active state of the line; therefore, the communication line is guaranteed to be in its rest state at least one bit time between each character transmitted and the start of each character is denoted by this first marking-to-spacing transition. Due to this resynchronizing, the sending and receiving units may (under optimal conditions) differ in character rates by five to ten per cent without causing transmission errors.

As used here, the term "Communication Port" refers to the logical interface between a Communications Unit and an individual Communications Line with its associated buffer. The demand which an individual Communication Port places on a Communication Unit, then, is a function of the speed of the Remote Console which is controlling the communication and the type of buffering used by the Communication Port. Three methods of buffering will be described: the bit buffer, the shift buffer, and the character buffer.

A bit-buffered Communication Port samples the Communication Line (marking or spacing) once per bit time, and then interrupts the SPU to save the sample. Since the SPU must save the sample before the Communication Port extracts its next sample, the maximum allowable response time of the SPU is less than one bit time. A variance on this scheme is the line-sampling or unbuffered Communication Port in which the response of the Processing Unit to an interrupt is to sample the Communication Lines directly. In this case, any difference between the maximum and minimum response times of the Processing Unit will cause sampling to take place away from the center of the bit time and will directly degrade the service.

The next level of buffering is the shift register. The shift-buffered Communication Port samples the Communication Line once per bit time and retains these samples until it has a complete character. When the last bit of the character is obtained from the Communication Line, the Processing Unit is interrupted, and the maximum allowable response time is somewhat greater than the duration of the stop impulse. The maximum allowable response time of the shift-buffered scheme is not radically greater than that of the bit-buffered scheme, but the number of interrupts which must be processed has decreased by a factor of k , the number of bits per character.

The highest level of buffering is the character buffer. A character-buffered Communication Port consists of a shift-buffered Communication Port plus a one-character buffer. In this case, when the shift register is filled, its contents are gated into the character buffer and the Processing Unit is interrupted concurrently. This scheme creates the same number of interrupts as does the shift-buffered scheme, but the maximum allowable response time has been extended to one character time. An obvious extension of character buffering is the multiple-character buffer, which could be used to extend the permissible response time or decrease the number of interrupts or both.

In comparing these three methods of buffering as applied to a multi-port communication processing system, the criterion must be the traffic level at which the system saturates. The maximum allowable response time is the critical parameter under peak traffic conditions. Under these circumstances the SPU deals only with the Communication Ports and it must be ensured that no data is lost. Under average traffic conditions, the number of interrupts per character assumes more importance due to the large amount of red tape time in making the transition from dealing with Communications Ports to dealing with messages and vice versa. With respect to the bit-buffered scheme, the shift-buffered scheme extends the average traffic capability by a factor of 5 to 8, depending on the number of bits in a character; it extends the peak traffic capability by a factor of 2 to 4, depending on the ratio of stop-impulse time to data-impulse time. The character-buffered scheme further extends the peak traffic capability by a factor of 2 to 4, depending on all three of the preceding points; it does not increase the average traffic capability.

For purposes of comparison, bit times and character times for three devices are listed below:

<u>Device</u>	<u>Speed in cps</u>	<u>Baudot</u>	<u>Bit Time in msec</u>	<u>Character Time in msec</u>
IBM 1050	14.8	135	7.4	67.5
TTY 33	10.0	110	9.1	100
TTY 32	10.0	75	13.5	100