


PROGRAMMERS' REFERENCE MANUAL




HONEYWELL 800
Transistorized Data Processing System



Litho in U. S. A.

DSI-31
2460
2660
1361



**PROGRAMMERS'
REFERENCE MANUAL**

Honeywell



Electronic Data Processing



Copyright 1960
Minneapolis-Honeywell
DATAmatic Division
Wellesley Hills 81, Massachusetts

TABLE OF CONTENTS

		Page
Section I.	Introduction	1
	Numbering Systems	2
	Decimal System	2
	Binary System	3
	Binary Codes	3
	Octal and Hexadecimal Systems	3
Section II.	The Honeywell 800 System	4
	The Equipment	4
	1. The Central Processor	4
	2. The Tape Unit and Tape Control Unit	5
	3. The Card Equipment	7
	a. Card Readers	7
	b. Card Punches	10
	4. The Printers	10
	a. Standard-Speed Printers	10
	b. High-Speed Printer	11
	5. The Console	12
	6. Off-Line Controls	12
	Traffic Control	13
	System Configurations	16
	Multiprogram Control	17
	Orthotronic Control and Checking	20
Section III.	The Honeywell 800 Word	24
	Data Words	24
	Special Register Words	26
	Instruction Words	26
	General Instructions	27
	Unmasked General Instructions	28
	Masked General Instructions	28
	Inherent Mask Instructions	29
	Peripheral and Print Instructions	30
	Simulator Instructions	30
	Special Words	31
Section IV.	Addressing	33
	Direct Memory Location Address	34
	Direct Special Register Address	35
	Indexed Memory Location Address	37
	Indexed Special Register Address	39
	Indirect Memory Location Address	40
	Indexed Indirect Memory Location Address	42
	Summary of Address Forms	43
	Significant Main Memory Addresses	43
	Stopper Address	45
	Inactive Addresses	46

TABLE OF CONTENTS (cont.)

		Page
Section V.	Special Registers	49
	Sequencing Counters	53
	History Registers	55
	Index Registers	55
	Mask Index Register	56
	General Purpose Registers	57
	Read-Write Counters	58
	Arithmetic Control Counters	59
	Unprogrammed Transfer Register	60
Section VI.	Arithmetic Instructions	62
	The Accumulator	62
	The Low-Order Product Register	64
	Binary Add, BA	64
	Decimal Add, DA	66
	Binary Subtract, BS	66
	Decimal Subtract, DS	67
	Word Add, WA	67
	Word Difference, WD	68
	Binary Accumulate, BT	68
	Decimal Accumulate, DT	70
	Binary Multiply, BM	70
	Decimal Multiply, DM	71
Section VII.	Logical Instructions	72
	Extract, EX	72
	Substitute, SS	74
	Half Add, HA	75
	Superimpose, SM	75
Section VIII.	Transfer Instructions	77
	Transfer A to C, TX	78
	Transfer A to B and Go to C, TS	78
	N-Word Transfer, TN	79
	Multiple Transfer, MT	80
	Record Transfer, RT	81
	Item Transfer, IT	82
Section IX.	Decision Instructions	83
	Inequality Comparison, Alphabetic, NA	84
	Less Than or Equal Comparison, Alphabetic, LA	85
	Inequality Comparison, Numeric, NN	85
Less Than or Equal Comparison, Numeric, LN	85	
Section X.	Shift Instructions	86
	Shift Preserving Sign and Substitute, SPS	88
	Shift Preserving Sign and Extract, SPE	89
	Shift Word and Substitute, SWS	89
	Shift Word and Extract, SWE	89
	Shift Word and Select, SSL	90

TABLE OF CONTENTS (cont.)

		Page
Section XI.	Peripheral Instructions	92
	Read Forward, RF	94
	Read Backward, RB	98
	Write Forward, WF	99
	Rewind, RW	103
Section XII.	Miscellaneous Instructions	105
	Print, PRA, PRD, PRO	105
	Control Program, MPC	108
	Proceed, PR	112
	Simulator, S	113
	Compute Orthocount, CC	114
	Check Parity, CP	116
Section XIII.	Summary of Instructions	119
	General Instructions, Unmasked or Masked	119
	General Instructions, Unmasked	120
	Inherent Mask Instructions	122
	Peripheral and Print Instructions	123
	Simulator Instructions	124
Appendix A.	Addition in the Honeywell 800	125
Appendix B.	Orthotronic Control	130
Appendix C.	Timing Summary	136
Appendix D.	Tables	143

LIST OF ILLUSTRATIONS

		Page
Figure	II-1. Conversion of Punch Positions to Honeywell 800 Character Positions	8
Figure	II-2. Card Reader Control Word	9
Figure	II-3. Stages of Traffic Control	14
Figure	III-1. Honeywell 800 Word Structure	25
Figure	III-2. Honeywell 800 Major Instruction Types	28
Figure	IV-1. Main Memory Address	33
Figure	IV-2. Special Register Address	33
Figure	IV-3. Direct Memory Location Address	35
Figure	IV-4. Direct Special Register Address	36
Figure	IV-5. Indexed Memory Location Address	38
Figure	IV-6. Indexed Special Register Address	39
Figure	IV-7. Indirect Memory Location Address	41
Figure	IV-8. Interpretation of Address Bit Structure	44
Figure	V-1. Special Register Names, Subaddresses, and Mnemonic Addresses	50
Figure	V-2. Mask Index Register	56
Figure	V-3. Generated Mask Address in Shift Instructions	56
Figure	V-4. Generated Mask Address in Field Instructions	57
Figure	V-5. Unprogrammed Transfers of Control	61
Figure	XII-1. Honeywell 800 Special Register Group Indicator Relationships	110
Figure	XII-2. B Address Group Function in Control Program Instruction	111
Figure	XII-3. Bit Layout -- Memory and Tape	117
Figure	A-1. Schematic Representation of Bus and Accumulator	126
Figure	B-1. Relation of Data Words to Orthowords	131
Figure	B-2. Computation of Orthowords	132
Figure	B-3. Orthotronic Check of Tape Error	133
Figure	B-4. Orthocount of Error Record	134
Figure	B-5. Correction of Tape Error -- First Method	135
Figure	B-6. Correction of Tape Error -- Second Method	135
Table	I. Honeywell 800 Coding and Punched or Printed Equivalents . .	143
Table	II. Honeywell 800 Command Codes	144

SECTION I

INTRODUCTION

The purpose of the Programmers' Reference Manual is to define the internal machine language of the Honeywell 800 Electronic Data Processing System and the manner in which that language is interpreted and manipulated by the system. This manual is not intended to be either an introduction to programming or an introduction to the Honeywell 800. Instead, it is written as a handbook for the experienced programmer who has completed the Honeywell 800 programming course.

Every Honeywell 800 user is furnished with a complete automatic programming package which eliminates many of the routine human tasks involved in program preparation, checkout, and execution. The Honeywell automatic programming package includes the following compatible elements:

1. ARGUS (Automatic Routine Generating and Updating System). ARGUS consists of a mnemonic and symbolic language for program preparation, an assembly routine which translates this notation into the internal language of the Honeywell 800 and inserts existing subroutines at points specified by the programmer, sort generators capable of tailoring a sort routine to a specific set of input data, and a program test system which automatically executes a series of programs to provide extensive information about their operation.
2. FACT (Fully Automatic Compiling Technique). FACT interprets a narrative English description of a data handling procedure and produces (or compiles) a complete machine-language program capable of handling all aspects of the business-oriented process, such as input editing, data sorting, file updating, and generation of printed reports.
3. Algebraic Compiler, the scientific counterpart of FACT. The Algebraic Compiler interprets a series of problem statements in a combined arithmetic and logical notation closely approximating the standard symbols of mathematics, and produces a complete machine-language program specifically directed to the solution of a problem in the scientific area.
4. A library of thoroughly tested subroutines for inclusion in compiled or assembled programs.
5. The Executive Routine, which directs the scheduling and operation of a multi-program run to facilitate maximum usage of the benefits of Honeywell parallel processing.

SECTION I. INTRODUCTION

The extent and power of this automatic programming package means that Honeywell 800 programs can normally be prepared without reference to the language set forth in the Programmers' Reference Manual. In other words, the ARGUS Manual of Assembly Language and the FACT and Algebraic Compiler Manuals are the standard programming documents for the Honeywell 800. Nevertheless, this manual is offered in recognition of the fact that the sophisticated programmer is naturally eager to consider subtle ways in which he can take advantage of the unusual and powerful features of the machine. When this level of experience is reached, an understanding of the internal configurations of instructions and data, as presented in the following sections, will prove invaluable.

Throughout the manual, the binary-digit structure of instructions and addresses is presented wherever it can enhance the explanation. Moreover, other numbering systems, such as octal and hexadecimal, are utilized wherever the subject matter requires them. For comparative purposes, the ARGUS format of many examples is also shown.

Numbering Systems

Modern numbering systems, including all systems used in this manual, are based on positional notation. This means that each system is based upon a root number, or "radix," and that each position within a number represents a specific power of the radix of the particular system being used. Positive and negative powers of the radix are separated by an indicator point, with the zeroth power of the radix appearing immediately to the left of the indicator point. Positive powers of the radix appear in successive positions to the left and negative powers in successive positions to the right of the indicator point. The radix of a numbering system is equal to the number of digits comprising that system; these digits cover the range from 0 to one less than the radix.

Decimal System

The familiar decimal system is based on a radix of 10 and uses 10 digits from 0 to 9. Each position in a decimal number represents a specific power of 10 and can have any of 10 values. The total value of a decimal number is computed by multiplying the value of each digit by the positional value (power of 10) of its position within the number and then summing all of these products. For example, the decimal number 356. has the value 3×10^2 plus 5×10^1 plus 6×10^0 , or $300 + 50 + 6$. The number

3.56 has the value 3×10^0 plus 5×10^{-1} plus 6×10^{-2} , or $3 + \frac{5}{10} + \frac{6}{100}$. When positional notation is understood in the familiar decimal context, the interpretation of any other positional system becomes clear.

Binary System

This system is based on a radix of 2 and uses the two binary digits (or bits) 0 and 1. Binary numbers are the common internal system for digital computation due to the relative simplicity of recording, storing, and recognizing variables of only two values. The value of a binary number is computed by multiplying the value of each digit by the corresponding power of 2 and summing all of the products. For example, the binary number 1001 has the value 1×2^3 plus 0×2^2 plus 0×2^1 plus 1×2^0 , or $8 + 0 + 0 + 1$, which equals 9. Where the system in use is not made clear by the context, its radix may be appended to the number as a subscript as, for example, to distinguish the binary number 1001_2 from the decimal number 1001_{10} .

Binary Codes

In addition to the use of "pure binary" numbers, as described in the preceding paragraph, binary digits may be grouped so that each group represents a decimal digit, alphabetic character, or other symbol. For example, bits may be manipulated in groups of four with each group representing a decimal digit (from 0000_2 to 1001_2). Similarly, groups of six bits may represent up to 64 digits, characters, or symbols. Such 4-bit and 6-bit codes are called "binary-coded decimal" and "alphanumeric," respectively. They facilitate the handling of the external decimal and alphabetic symbols by machine elements which recognize only variables of two values.

Octal and Hexadecimal Systems

These two systems, based on radices of 8 and 16, respectively, are useful as shorthand methods of writing pure binary numbers. If a binary number is divided into groups of three bits, proceeding in either direction from the indicator point, each group may be replaced directly by its octal equivalent, since a 3-bit group has a total of eight possible values. If the same number is divided into 4-bit groups in the same manner, each group may be replaced directly by its hexadecimal equivalent, since a 4-bit group has a total of 16 possible values. The 16 hex digits are represented in this manual by the symbols 0 - 9 and B - F, respectively.

SECTION II

THE HONEYWELL 800 SYSTEM

The Equipment

A Honeywell 800 Data Processing System consists of a central processor to which varying types and numbers of input and output units are attached. The programmer must know the system configuration with which he is to work, and an understanding of the function of each component and its relation to the system as a whole will make his task easier.

1. The Central Processor (801)

The basic central processor consists of a control unit, a control or special-register memory, an arithmetic unit, and two banks of main (or high-speed) memory, each capable of storing 2048 Honeywell 800 words. (A Honeywell 800 word is composed of 48 information bits and six checking bits.) To this basic unit, additional memory banks may be added in amounts of 4096 words up to a maximum of 32,768 words. Also available is an optional floating-point unit (801B) which adds floating-point instructions to the command repertoire.

The control unit with its control memory is the nerve center of the central processor. As the site of traffic control and multiprogram control, it monitors the time sharing of the entire system to achieve maximum efficiency. In addition to its multiplexing function, it is also the unit which selects, interprets, and directs the execution of instructions, and governs address selection in both control memory and main memory.

The control memory is a magnetic-core storage array providing storage for 256, 18-bit words. The read-restore cycle of the control memory is out of phase with that of the main memory in such a way that if reference must be made to the control memory between references to main memory, it is usually possible to make such reference without loss of a main-memory cycle. As discussed more fully in Section V, the control memory contains eight identical groups of special

registers such as sequencing counters, index registers, registers used for indirect addressing, etc., the contents of which are used to select a full Honeywell 800 word from the main memory. The offset cycle of control memory makes it possible to anticipate an address selection involving the contents of a special register and to prepare the address of a second operand while another unit is using the first operand. Because of this anticipatory technique, it is unnecessary in many cases to add memory cycles to an instruction for indexed or indirect addressing. Even when the contents of a special register are modified before they are restored, no extra memory cycle need be added, since the special register circuitry includes a separate adder, with complete and independent checking, used only for special register modification. This applies to both automatic modification, as when a sequencing counter is incremented after use, and program-controlled modification, as when an increment is specified in an address. However, while the two memory units are sufficiently out of phase to allow reading from the control memory prior to the start of a main memory cycle, a read-restore operation in which the result of an instruction is returned to a special register cannot overlap a main memory cycle; in this case, an extra cycle must be added to the instruction time.

The arithmetic unit is the portion of the central processor in which digits are combined to form new arrays in accordance with the logical rules of the command codes. The Honeywell 800 central processor has provision for both binary and decimal arithmetic, complete logical abilities, and competent internal checking. For the interested reader, a complete description of the addition logic can be found in Appendix A.

The optional floating-point unit is essentially a secondary control and arithmetic unit which manipulates Honeywell 800 words in floating-point form. A complete description of the operation of this unit will be found in a separate manual.

2. The Tape Unit (804) and Tape Control Unit (803)

The Honeywell 800 tape unit has been designed for reliability and accuracy. The recording surface of the tape is not touched except by the read-record head so that wear and damage to the surface are reduced to an absolute minimum. Vacuum is used to seat the reels on the hubs, to maintain loops in the loop chambers, and to provide contact with the capstans which cause the tape to move under the head,

SECTION II. THE HONEYWELL 800 SYSTEM

giving accurate control without the dangers inherent in mechanical techniques. The tape is edge guided along its entire path to protect the tape edges from damage, and avoid skew.

The system uses a 3/4-inch-wide tape with Mylar base and oxide coating. A full reel of tape consists of 2400 feet of oxide-coated tape plus two 50-foot, clear leaders used for control. Information is written on the tape in 10 longitudinal channels, eight for information bits from the word, one for a parity checking bit, and one for a clocking indicator. One array of bits across the tape is called a frame. Information from six frames makes up the 54-bit word (including six parity bits). The clocking channel is used to give positive indication of the frame location on tape.

Variable-length recording is a basic feature of the Honeywell 800, and records of any size may be read or written from tape, although for control purposes, a maximum size limit may be placed on records at the beginning or end of tape. Gaps between records are 0.67 inches long. Frames are recorded on tape at a nominal density of 400 per inch. During reading and writing operations, the tape is moved at 120 inches per second, giving instantaneous transfer rates of 8000 words or 96,000 decimal digits per second.

When rewinding, the tape moves at a speed of 360 inches per second. A small photoelectric device in the tape unit senses the presence of edge "windows" (clear Mylar) in the tape to provide beginning-of-tape and end-of-tape indications for the programmer. A physical slot in each clear leader is used to negate the vacuum and provide positive protection against pulling the tape off the reel.

When a metal ring is inserted in the front of a tape reel, writing is allowed to take place on the tape. When this ring is removed, the tape is protected and cannot be written upon. The ring may be installed or removed without rewinding the tape or removing the reel. There is, in addition, a manual switch on the tape unit panel which can be set to prohibit writing.

A tape control unit can control up to eight tapes. Each control unit has an input and an output channel so that one of the tape units attached to the control unit may be reading and another writing simultaneously.

3. The Card Equipment

a. Card Readers (823)

Two 80-column card readers are available with the Honeywell 800 System. From the programmer's point of view, these two readers differ only in that the 823-1 reads 240 cards per minute, while the 823-2 reads 650 cards per minute. In the manner in which they respond to instructions and transmit information to the central processor, they are identical.

The card reader with its control unit reads the card, converts the punch configurations to Honeywell 800 notation, and transmits the information to the main memory. Two modes of conversion may be selected by the operator by means of an external switch. The first is called the normal or alphanumeric mode. Conversion of a card in this mode results in the transmission of 10 Honeywell 800 words of eight characters each, where the high-order character of the first word corresponds to column 1 of the card and each successive 6-bit character in the next lower-order position corresponds to the next higher-numbered column. The correspondence between punched-card codes and Honeywell 800 codes is shown in Table I, page 143. In the normal conversion mode, an illegal punch check switch is set by the operator to define as legal either the entire 64-code set listed in Table I or a special 50-code subset, excluding the asterisked codes.

The second mode is called the transcription mode, in which the information from each card read is transmitted to the main memory in 20 Honeywell 800 words. The value of each of the 960 bits indicates the presence or absence of a punch in a specific punching position. Figure II-1 shows the correspondence between card format and memory format in these two modes.

Each of the card readers has two reading stations. The results of the two readings of each card are compared, and any discrepancy is noted. As the 80 columns of information are converted in the control unit, additional checking is done to insure correct conversion.

One extra word is appended to each card record as it is sent to memory, indicating the status of the error indicators at the completion of the reading

SECTION II. THE HONEYWELL 800 SYSTEM

ALPHABETIC MODE

Word 1	Column 1	Column 2	Column 3	Column 4	Column 5	Column 6	Column 7	Column 8
Word 2	Column 9	Column 10	Column 11	Column 12	Column 13	Column 14	Column 15	Column 16
.								
.								
.								
.								
Word 10	Column 73	Column 74	Column 75	Column 76	Column 77	Column 78	Column 79	Column 80

TRANSCRIPTION MODE

	Column 1	Column 2	Columns 3-4-5-6-7	Column 8
Word 1	← Row →	← Row →	← Row →	← Row →
	9 8 7 6 5 4	9 8 7 6 5 4	9 8 	9 8 7 6 5 4
.				
.				
.				
.				
.				
Word 10	Columns 73-74-75-76-77	Column 78	Column 79	Column 80
	← Row →	← Row →	← Row →	← Row →
	9 8 7 6 5 4	9 8 7 6 5 4	9 8 7 6 5 4
Word 11	Column 1	Column 2	Columns 3-4-5-6-7	Column 8
	← Row →	← Row →	← Row →	← Row →
	3 2 1 0 X R	3 2 1 0 X R	3 2 	3 2 1 0 X R
.				
.				
.				
.				
Word 20	Columns 73-74-75-76-77	Column 78	Column 79	Column 80
	← Row →	← Row →	← Row →	← Row →
	3 2 	3 2 1 0 X R	3 2 1 0 X R	3 2 1 0 X R

Figure II-1. Conversion of Punch Positions to Honeywell 800 Character Positions

and conversion operation for that card. The control word (see Figure II-2) may indicate:

1. that the card was punched, read and converted correctly;
2. that the card contained a punch combination defined as illegal;
3. that an error occurred on reading or conversion.

Thus, a standard card record is either 11 words (alphanumeric mode) or 21 words (transcription mode).

Since an anticipatory card reading technique is used (actually the card has been read and converted before the read instruction is received by the card reader), all errors except detection of a parity failure on transmission to the central processor must be indicated by the control word. Parity failure on transmission is indicated by an automatic transfer of control at the time of the next read instruction (see Section XI).

The operator may, by setting switches on the card reader control unit, specify the following options in the event of a reading error:

1. NORMAL - Do not eject card; do not stop reader.
2. EJECT - Eject card to reject stacker; do not stop reader.
3. STOP - Eject card to reject stacker; stop reader.

In these three options, the information from the card, properly tagged in the control word, is transferred to the central processor.

4. DISCARD - Eject card to reject stacker; do not stop reader; do not transfer information to central processor.

<u>Bits</u>	<u>Definition</u>
15-16	11 - Card was correctly punched, read and converted.
	01 - Card was illegally punched. Reading and conversion, however, were correct.
	10 - Card was either incorrectly read or converted.
All other bits of any card reader control word are zero.	
Note: In the event of both an illegal punch and an incorrect read or conversion, the latter signal takes precedence.	

Figure II-2. Card Reader Control Word

b. Card Punches (824)

The two card punches available with the Honeywell 800 System also appear identical to the programmer, differing only in speed. The 824-1 punches 100 cards per minute; the 824-2 punches 250 cards per minute. The punch is capable of punching in the same two modes in which the card reader may operate, and the relationship between the card format and the format of information within the Honeywell 800 is identical. (see Figure II-1). The selection of the mode to be used in punching an individual card, however, is made by the program. A control word precedes the actual information to be punched, sets the punch to accept either 10 or 20 additional words, and informs the punch control unit as to which conversion mode is desired. Only four bits of the punch control word are used by the punch control unit. The values of bits 16 and 17 specify the desired conversion mode and hence the number of words to be accepted. Bits 13 and 14, called the End-of-Run bits, can cause an indicator light on the punch control unit to glow, indicating to the operator that the program has completed the punching of a file.

4. The Printers (822)

a. Standard-Speed Printers

The standard-speed printers, models 822-1 and 822-2, are 407 and 408 tabulating machines, printing 150 lines per minute. There are 120 printing positions, and 47 characters are available. A 12-channel, punched paper carriage tape is used for vertical format control. When executing a write instruction which addresses a standard printer, the Honeywell 800 transmits a 16-word record to the printer control unit. The first word contains vertical-format information. The remaining 15 are considered as 120 alphabetic characters of which the high-order character of word 1 corresponds to the left-most printing position. The vertical-format word contains 12 bits (bits 19-30) which correspond with the 12 channels of the tabulator carriage tape. A one in any of these bit positions designates a controlling carriage-tape channel and the carriage is advanced until the next punch is sensed in the designated channel. Values of bit positions other than those specified above are ignored.

The tabulators retain most of the abilities normally provided through control panel wiring; in particular, the comparison abilities of the 408 are retained.

b. High-Speed Printer

The high-speed printer, model 822-3, has a rated speed of 900 lines per minute for continuous single-space printing and approximately 800 lines per minute for continuous double-space printing. There are 160 print positions of which any prescribed array of 120 may be made active during a given run. A total of 56 characters is available at each print position, including 26 alphabetic, 10 numeric, and 20 special symbols. The numeric font used on the printer is designed for automatic reading by optical scanning equipment.

Horizontal spacing is 10 characters to the inch. Vertical spacing is six lines to the inch with eight lines per inch available as a special option. Skipping speed in the non-printing mode is approximately 20 inches per second. Vertical format control is accomplished under program control in conjunction with a pre-punched, 6-channel, paper carriage tape. Both horizontal and vertical vernier adjustment of form position is standard.

Stock ranging from 8-pound single-part forms to 6 1/2-mil continuous card stock can be printed. Conversion adjustment for stock thickness is provided. With light-weight 8-pound stock, an original and up to four highly legible copies can be produced. Two-part heat transfer forms and Multilith master forms can be accepted. The width of the stock to be printed may vary from 4-3/4 inches minimum to 22 inches maximum, measured paper edge to paper edge.

For each printed line, 16 words are sent to the printer control unit. The first word contains vertical-format information, while the remaining 15 words are interpreted as 120 characters to be printed. The high-order character in the first information word corresponds to the first of the 120 print positions selected to be active. Line spacing is specified in the vertical-format word as follows:

- Bit 1 - Position at head-of-form, as indicated by a punch in channel 1 of carriage tape.
- Bit 6 - Inspect channel 6 of carriage tape for punch indicating end-of-form.

If bit 6 is a one, the machine senses for end-of-form as determined by a punch in channel 6 of the paper tape. When the punch is sensed, the carriage is advanced to

beginning-of-form. Normally, in a listing, bit 6 always has the value of one, so that the end-of-form channel is continually examined.

Bits 7-12 - Line Count - Specify a number of lines (0-63) to be spaced after the accompanying line has been printed.

Bits 13-14 - End-of-Run - Cause end-of-run light to glow on printer control unit to signal operator that run is complete.

Values of bit positions other than those specified above are ignored, so that it is possible to edit a record for listing on either standard-speed or high-speed printer.

5. The Console

The Honeywell 800 console is basically a part of the central processor and is multiplexed into the system via the multiprogram control unit. A monitor typewriter is used by the operator to communicate directly with the central processor. Manual operations on this typewriter can start and stop individual programs and interrogate Honeywell 800 storage. Under program control, the console typewriter can also print information useful to the operator. An additional typewriter, called a console slave typewriter, may be added to the system. No manual operations may be performed from the slave typewriter, but printing may be programmed to occur on either the slave typewriter or the console monitor typewriter. In addition to the typewriter(s), the console includes display lights to give the operator an at-a-glance summary of the number of active programs, their control centers, and their progress. An auxiliary board may be installed with indicators and displays for monitoring the status of tape units and peripheral devices. From 1 to 45 remote inquiry stations may be included in the system for direct interrogation of stored information and printout of results.

6. Off-Line Controls (815, 816, 817)

When a terminal unit is to be used with a tape unit independently of the central processor, as in a card-to-tape or tape-to-printer operation, an off-line auxiliary control must be placed between the peripheral control unit and the tape unit to provide control signals and power normally supplied by the central processor. If the terminal unit is an input device, then an off-line input auxiliary control (816) must be used. In like manner, an output device requires an off-line output auxiliary control (815) and a multiple-unit control calls for an off-line input-output auxiliary control (817).

When terminal units are used off-line, information is written on or read from tape in units corresponding to the record size of the device. In a card-to-tape operation, each tape record consists of one card's worth of information (10 or 20 data words) plus a control word, two orthotronic correction words, and an end-of-record word automatically generated within the off-line input control unit. Records read from tape for printing must each be 19 words in length (one vertical format word, 15 words of characters to be printed, two orthotronic words, and an end-of-record word), and each record for punching must contain 14 or 24 words as described for the card reader.

Traffic Control

Traffic control is the Honeywell 800 element which directs the time sharing of memory by the tape and peripheral units and the central processor. Multiprogram control is the element which directs the time sharing of the central processor by the active program control centers. A clear concept of both of these elements is basic to the understanding of parallel processing and allowable system configurations and is the key to a thorough knowledge of the Honeywell 800.

Traffic control has as its main object the efficient use of the entire system according to a set of priorities which derive directly from the nature of the equipment and are independent of the programs. For example, a tape unit reading at full speed assembles one Honeywell 800 word in a one-word buffer every 125 microseconds. If instant access is not provided to memory, a second word of buffer storage must be provided to retain this word. At the end of the next 125 microseconds, another word will have been read. If the first word has not yet been placed in main memory, another word of buffer storage must be provided. Since eventually one access to memory must be made for each word to be stored, it is obviously economic to store each word as it is assembled from tape and thus reduce the required buffer storage to a minimum. However, to keep the memory continuously available to the tape buffer during a read operation would be to introduce inefficiencies in the system, for only one memory cycle (six microseconds) is needed to store each word, and during the remaining 119 microseconds the entire system would lie idle.

SECTION II. THE HONEYWELL 800 SYSTEM

In the Honeywell 800, one access to memory every 125 microseconds is guaranteed each tape unit. When a word is assembled from six frames for storage, a demand signal is generated by the buffer for one access to memory and is honored within 125 microseconds, clearing the buffer. In this case, only two words of storage are needed for each active tape unit, and the memory is utilized by the input-output operation only $6n$ microseconds out of each 125 microseconds, where n equals the number of active units. To achieve this time-sharing, traffic control monitors the demand signals from the buffers and arranges access to the memory within the prescribed time for each buffer demand.

As its name implies, traffic control monitors the transmission of information to and from the main memory. Its operation is represented schematically in Figure II-3. The 17 divisions of the band are called "stages" and one stage is assigned to each of the eight output channels, each of the eight input channels, and the central processor.




Figure II-3. Stages of Traffic Control

The creation of a demand signal by any device is represented in Figure II-3 by the closing of the switch shown in the corresponding control stage. When any program has been turned on in the central processor, the switch corresponding to the central processor stage is continuously closed. Traffic control begins each scan at the left end of the band. It proceeds to the right, ignoring all stages which show no demand signal, until a demand stage is reached. This stage is allowed access to the main memory for one memory cycle only. Traffic control then returns to the left end of the band to begin the next scan. Because the control search is anticipatory, no system time is consumed in bypassing stages in which no demand exists.

If no input-output units are active but a program is running, the central processor stage will have complete use of the memory since each scan of the band will find no other demand signals. If a tape unit attached to channel 2 is reading, then once every 125 microseconds a demand signal will halt traffic control at the stage marked "input 10" for one memory cycle to allow transfer of the assembled word from buffer storage to memory. Thus, between each memory cycle of instruction execution in the central processor, 0 to 16 memory cycles may intervene, depending on the number of tapes and peripheral devices which require access to memory.

Since, in normal read-write operations, only one memory cycle may be allowed any stage before the next scan of the band, a maximum of 16 memory cycles or 96 microseconds will elapse between successive interrogations of any stage. As this is within the 125-microsecond maximum, eight input units and eight output units may be active simultaneously with the central processor without conflict in demands for access to memory. Furthermore, since the memory is made available to the central processor for any cycle in which it cannot be utilized by an input-output stage, no idle time is introduced as long as any program is active. Thus, traffic control insures that the system responds to input-output device demands as required without introducing idle cycles, and that as long as any program can proceed, useful work is being done.

SECTION II. THE HONEYWELL 800 SYSTEM

The rule of operation which states that only one memory cycle may be allowed any stage before the next scan has one exception. When a tape unit is executing a distributed read or distributed write operation, two consecutive memory cycles are allowed a stage for modifying the appropriate address counter on recognition of the demand signal accompanying an end-of-item word. Distributed read and write instructions are more fully discussed in Section XI. Two points in connection with this exception bear mentioning here. First, the difference between the normal 96-microsecond cycle and the maximum allowable 125-microsecond cycle allows servicing of four such demands without timing conflicts. Secondly, in the improbable case that all system channels are active and five or more channels are being used in distributed tape operations, and at least five of these each assemble an end-of-item word and create a demand signal within 125 microseconds of each other, an error signal will be generated for any unit whose demand is not serviced within the required time limit. This signal will appear to the program as a normal reading error signal indicating the need for a reread from the specified unit. Complete information is available to the program so that intelligent action can be taken under program control.

System Configurations

The central processor is the heart of any Honeywell 800 installation. Peripheral and tape units are attached to the eight input and eight output channels of the central processor. Input and output channels alike are numbered 1, 2, 3, 4, 5, 6, 7 and 0. The output channels are associated with stages 1-8 of traffic control; the input channels are associated with stages 9-16. Each channel has a priority according to the sequence in which the corresponding stage of traffic control is interrogated. In other words, output channel 1 has the highest priority and input channel 0 the lowest. Each special register group includes an input buffer counter and an output buffer counter associated, respectively, with an input and an output channel. Thus output and input channels 1 are associated with the buffer counters of special register group 1, while output and input channels 0 are associated with the buffer counters of special register group 0. There exists no relationship between the special register group whose buffer counters are associated with a particular device and the special register group controlling the program which uses that device.

Any device requiring both an input and an output channel, such as a tape control unit, must be assigned to channels whose associated buffer counters reside in the same

special register group. Otherwise, any input device may be assigned to any input channel and any output device to any output channel, subject only to the restriction that tape control units must be assigned to channels of higher priority than peripheral control units. Specifically, if an installation includes two tape control units, a card reader, and two printers, the first tape control unit is assigned to output and input channels 1 and the second tape control unit to output and input channels 2. The card reader may be assigned to any remaining input channel and the printers to any remaining output channels.

Each tape control unit may control up to eight tapes. The tape units connected to a tape control unit are normally assigned consecutive positions starting with position 1. However, since tape addresses are assignable by patchboard, the programmer will find this no restriction.

In most systems each terminal unit will have its own control unit. The exception is the use of a Model 811 Multiple Terminal Unit Control to control a card reader, a card punch, and a printer (or any two of these devices). The multiple control unit has one buffer to handle traffic in either direction. Thus, only one device, either input or output, may be used at a time. Selection of the device to be used is made by a manual switch on the Multiple Terminal Unit Control. If the multiple control unit has both a card reader and an output device attached, it must be assigned both an input and an output channel. Since only one device may be used at a time, however, the channels assigned need not be a corresponding pair.

Multiprogram Control

Multiprogram control directs the time sharing of the central processor by the active programs. Each of the eight groups of special registers may direct the execution of an independent program. After a program is loaded, one of the special register groups is activated to direct the selection of instructions. This special register group, and the program it directs, is said to be "on." When the program is completed, the directing special register group is inactivated (turned "off"). Special register groups may be turned on and off independently of each other, either from the console or by program control.

SECTION II. THE HONEYWELL 800 SYSTEM

Each time traffic control allows the central processor access to main memory, one memory cycle of one instruction is performed. If only one special register group is active, all cycles allowed the central processor are used in executing instructions from the active program. Since traffic control allows the central processor all available cycles except those needed to honor the intermittent demands of tape and peripheral devices, this case represents that of the conventional single-program computing machine with the ability to implement input-output operations simultaneously with computing.

When more than one special register group becomes active, central processor cycles must be shared between the several programs. The rules under which multiprogram control operates are as follows:

1. When an active special register group is selected, the next cycle allowed the central processor must be used to select the next instruction to be executed in the program controlled by that group.
2. All succeeding central processor cycles must be devoted to the execution of this instruction until it is completed. (This may range from two cycles upward.)
3. If the instruction is one which does not leave unstored information in the arithmetic or control units, its completion causes multiprogram control to scan or "hunt" for the next active special register group in sequence.
4. If the instruction does leave unstored information in the arithmetic or control unit, scanning or hunting is inhibited and the next instruction is selected from the same program. Hunting is not allowed if:
 - a. the instruction generates a 2-word result, of which only one word is stored in memory as the result of the instruction execution, e.g. multiply;
 - b. the execution of the instruction results in a sequence change; in the case of a conditional change, hunting is inhibited only if the condition is satisfied;
 - c. the instruction has an inactive C address;
 - d. an unprogrammed transfer takes place as the result of executing the instruction;
 - e. the instruction is capable of specifying that hunting shall not take place, and does so specify.

For example, if three programs are active, then one instruction is performed in turn from program 1, then 2, then 3, then 1, and so on, as long as all instructions selected allow hunting. Such alternation is temporarily held up if an instruction is selected which does not allow hunting, but it is resumed as soon as an instruction is selected

which does allow hunting. Thus, the central processor cycles are shared on a fairly equal basis between all active programs.

The true power of multiprogram control appears when an active program attempts to execute an instruction which cannot be implemented because of the unavailability of a system component. Suppose, for example, that a write instruction calling for tape 4 is selected from an active program. The central processor, in attempting to execute this instruction, finds that either tape 4 or its associated output channel is involved in completing the execution of another instruction from the same or some other program. Seeing that the instruction cannot be executed immediately and recognizing the reason therefor, multiprogram control places the special register group in a "stall" condition. This condition indicates to multiprogram control that a) this program, although still active, shall not be allowed any central processor cycles as long as the "stall" indication remains, and b) when the channel and/or the device involved completes its present task, the stall condition shall be automatically removed and the program restored to its full active status.

Thus, when an instruction cannot proceed because of input-output conflicts with either the same or another program, the central processor cycles which it would have used are used instead for the other active programs, causing them to proceed faster. The result of the operation of multiprogram control is that there is never an idle memory cycle in a Honeywell 800 system as long as there is any active program in which an instruction can be executed.

If more than one active program is stalled because of input-output conflict, multiprogram control remembers the sequence in which the programs were stalled and operates on a "first off - first on" basis, so that no one program can monopolize an input or an output channel in case of conflict.

Although it need not concern the programmer, the reader may be interested in a short discussion of the actual machine procedure in cases of program "stalls." The computer will have selected the instruction and entered the second cycle of execution before the unavailability of the input or output channel and/or device is discovered. No memory alteration will have been made, but the sequencing counter will have been advanced by one. When the stall condition is ascertained, the computer will subtract one from the sequencing counter and deactivate the program from multiprogram con-

SECTION II. THE HONEYWELL 800 SYSTEM

trol. Each time that any input-output channel or device terminates an operation, all programs in a stalled state are reactivated. Multiprogram control looks again at each reactivated program in the order that they were originally stalled, replacing in a stalled condition those programs whose input-output demands still conflict.

Multiprogram control also receives demand signals generated by the console and by inquiry stations. The console is regarded by multiprogram control as a ninth group of special registers, except that it takes precedence over any active program in the assignment of available central processor time. A demand signal from the console is serviced at the time of the next hunt, although a non-hunting sequence of instructions will not be interrupted.

When a console demand is recognized, the computer generates an instruction to implement the activity indicated by the console command. Sufficient memory cycles are then assigned to execute this generated instruction in the same fashion as if it had been selected from memory. When the instruction has been executed, the normal hunting process is resumed. If the system includes one or more remote inquiry stations, demand signals from these stations are recognized and implemented in a similar fashion. This technique allows the operator to communicate manually with the central processor without stopping the computer, and thus allows the system full-efficiency operation even during manual manipulation on one program.

Orthotronic Control and Checking

Orthotronic control is a powerful technique, exclusive with Honeywell-DATAmatic, which insures against loss of information from tape during writing, storage, or subsequent reading. Experienced data processing personnel know that long storage periods or inept operator handling can cause information to disappear from a tape even though the accuracy of the record was checked at the time the record was written. Even infrequent occurrences of this type can result in many man-hours and machine-hours spent in re-creation of the records. While no technique will ever completely eliminate information loss, the high reliability and accuracy of the Honeywell 800 tape unit, plus the presence of orthotronic control as a standard feature of every Honeywell 800 system, insures that such loss is eliminated as a practical problem.

Orthotronic control is based on studies of the types and extent of information losses which have occurred on tape systems. It is partly automatic and partly program-controlled. An instruction is provided which automatically creates two orthotronic words for a specified record. These words are a logical combination of all the words in the record such that only a highly unlikely periodicity of error can go undetected and uncorrected. The orthotronic words are automatically positioned to accompany the record as it is written. Read and write instructions assume the presence of the orthowords and automatically include them in the record, using them in an automatic first-level check of the correctness of the information handled. The instruction which generates the original orthowords may also be used to reconstruct missing information if loss is detected. A full discussion of orthotronic control can be found in Appendix B.

Orthotronic control is a checking device peculiar to the tape units. In the on-line mode of operation, the central processor must be programmed to reconstruct lost data from a garbled record. When tapes are read or written in an off-line configuration, however, the off-line control unit provides, in part, the central processor function. The off-line input control unit automatically generates the two orthowords which must accompany each record when it is read by the central processor, and performs the first-level check on the information which involves these two words. The output off-line control performs a check of every record read exactly as the central processor does, and is capable not only of detecting an error in the record but, in the majority of cases, of reconstructing the garbled information. The corrected information can then be printed or punched without stopping or repositioning the output device as would be necessary without such automatic error correction.

In addition to orthotronic control, and in some ways complementing it, a parity bit is written on tape accompanying each frame. The parity bit is read from tape together with the eight information bits of the frame and remains with these bits as frames are collected to form words. As each word of six frames is transmitted to memory, the accompanying parity bits are monitored to insure an error-free transmission. Each time a word is sent to or from main memory, a transmission check is performed using these six parity bits, and when the word is again written on tape, each bit accompanies its corresponding frame.

When a word is brought to the arithmetic unit, the computer generates a modulo-3 check on each frame pair (16 information bits) for use in checking arithmetic opera-

SECTION II. THE HONEYWELL 800 SYSTEM

tions. The value of the 2-bit mod-3 check digit is the remainder (a value from 0 to 3, where either 0 or 3 may represent no remainder) which results when the decimal equivalent of the 16 bits is divided by 3. After the mod-3 check is generated, the parity bits are checked and then replaced by the six mod-3 checking bits. When arithmetic operations have been completed and the mod-3 check has been performed to insure that they have been completed correctly, the parity bits are again generated, replacing the mod-3 bits, and used to check the correct transmission of the result to memory.

The control unit of the central processor checks the interpretation and execution of the program instructions. Selection of instructions and operand locations is checked. The checking process of an add instruction illustrates the thoroughness of the Honeywell 800 checking system.

1. The selection of the instruction location is verified.
2. The instruction itself is verified for proper parity.
3. During the processing of the A address:
 - a. a mod-3 check group is attached to the address, then independently recalculated and compared with the original when this information is transferred to the memory selection circuits;
 - b. the selection operation is verified by comparing with the mod-3 check for the original address the special check digits delivered with the operand;
 - c. the operand itself is checked for proper parity when read from memory;
 - d. three mod-3 check digits associated with the operand are generated and stored.
4. During the processing of the B address (when the contents of the B address are added to those of the A address), steps a, b, and c are repeated for the B address. Also, three mod-3 check digits are generated as in 3d, but are added (mod-3) to the check digits previously stored.
5. As the result of the addition is transferred to memory, the C address memory selection is verified as in 3a, b, and c, and a new set of mod-3 check digits is formed from the computed sum and compared for equality with the check digits sum formed in (4) above. If the two sets of digits are equal, then the add instruction has been processed properly.

6. An example of the mod-3 arithmetic follows:

Number in A address	8426	9721	4075
The associated mod-3 check digits	2	1	1
Number in B address	1276	0216	4925
The mod-3 check digits	1	0	2
The sum of A and B is	9702	9937	9000
The mod-3 check digit	0	1	0
The mod-3 sum of the check digits is	0	1	0

In the general case where carries might occur between the operand groups, corrections of +1 and +2 are added to the appropriate check digits. Since two groups are simultaneously affected by a carry correction, any error in the addition, including the carry generation or correction process, is automatically detected.

Each of the special registers retains a mod-3 check on the 16 information bits it contains, which is used to check transmissions and arithmetic operations within the control unit. When the contents of a special register are transferred to the arithmetic unit or to main memory, they are expanded to full-word form and the mod-3 check is replaced by parity bits.

Card reading is checked not only for correct reading by the equipment, but, in the alphanumeric mode, for correctness of conversion and proper keypunching also. Card punching is checked by echo pulses sent from the punch to be compared with the punching image. Printing is also checked by comparison of echo pulses generated by the printer with the print image. Up to 30 columns of double-punch, blank-column detection is also available as an option on the punch.

SECTION III

THE HONEYWELL 800 WORD

The basic unit of information in the Honeywell 800 System is a fixed-length word consisting of 54 binary digits, of which six are parity bits used by the automatic checking circuitry and 48 are information bits. Each main memory location is capable of storing one such word, and each arithmetic register is one word in length. A main memory word may represent a machine instruction or one or more pieces of data. In addition to the main memory, the central processor includes the control memory of 256 special registers, used primarily for control purposes and address modification. A special register has the capacity to store a partial word consisting of 16 information bits and two checking bits.

The check bits of the main memory and special register words are not directly available to the programmer, nor are their values subject to program control. Subsequent discussions of the Honeywell 800 word, therefore, will refer only to the information bits, unless otherwise noted.

Data Words

A computer program generally manipulates data in one or more different forms: decimal, alphanumeric, binary, or a combination of these. The Honeywell 800 is capable of handling all these types of information. It may interpret the 48 bits of a word in groups of four for the purpose of binary-coded-decimal operation, in groups of six for alphanumeric operation, or as individual units of information for pure binary operation. Figure III-1 illustrates the structures of these different words.

A decimal word in the Honeywell 800 contains either 11 decimal digits with a sign, or 12 decimal digits without sign. The decimal arithmetic instructions interpret all operands as a sign and 11 digits. The sign consists of four bits which may represent either the sign of the entire word or individual, 1-bit signs for as many as four different pieces of information within the word. Although a positive sign is normally represented by four binary ones and a negative sign by four binary zeros, a non-standard

SECTION III. THE HONEYWELL 800 WORD

BIT POSITION	1	5	9	13	17	21	25	29	33	37	41	45	
DECIMAL	±	1	2	3	4	5	6	7	8	9	0	1	
ALPHANUMERIC	R	O	B	I	N	S	O	N					
ALPHANUMERIC COMPRESSED	C	.	W	E	B	B	1	7	4				
BINARY	±	(44 Binary Digits)											
INSTRUCTION	OPERATION CODE			ADDRESS A			ADDRESS B			ADDRESS C			
SPECIAL REGISTER										±	(15 Binary Digits)		

Figure III-1. Honeywell 800 Word Structure

configuration is perfectly acceptable as input to the arithmetic unit, which interprets any combination of bits except four binary zeros as a positive sign. The sign supplied with the result of an arithmetic operation, however, is always one of the two standard conventions, either four binary ones or four binary zeros. A more detailed discussion of sign conventions can be found in Section VI.

An alphanumeric word in the Honeywell 800 consists of eight 6-bit groups. Each 6-bit configuration may represent any one of 26 alphabetic characters, 10 decimal digits, or 20 special symbols, such as punctuation marks (see Table I, page 143). The alphanumeric mode is always used for communication between the central processor and card readers, card punches, or printers. Although a number may be stored in alphanumeric form, the arithmetic unit is not designed to handle information stored as 6-bit alphanumeric codes.

The 48 binary digits of a word may also represent a pure binary number, which may be stored as a sign and 44 bits, or as 48 unsigned bits. With the exception of the instructions word add and word difference, which treat their operands as 48-bit unsigned numbers, the binary arithmetic instructions interpret operands as signed 44-bit numbers. The sign convention in binary arithmetic is identical to that described for decimal words.

SECTION III. THE HONEYWELL 800 WORD

The data words described above are identified in ARGUS language by the following constant codes: DEC, decimal number, signed or unsigned; ALF, alphanumeric word; FXBIN, pure binary number; and M (mixed constant), compressed alphanumeric word. In addition, ARGUS recognizes an octal word identified by the constant code OCT. This word contains 16 unsigned or 15 signed octal digits. If 15 signed digits are specified, the most significant digit must be less than four, since a sign is represented by four bits, leaving only two bits for the high-order octal digit.

Several differences should be noted between ARGUS notation for data words and the format shown in Figure III-1. When ARGUS notation is used for decimal words, high-order zeros in signed decimal numbers and low-order zeros in unsigned decimal numbers need not be expressed. For example, ARGUS converts the number +125 to the signed 11-digit number +0000000125 and the unsigned number 32 to the 12-digit number 320000000000. A binary word in ARGUS notation is not expressed as a 44- or 48-bit binary number, but as the decimal equivalent of the desired binary information bits. As such, a binary word in ARGUS may contain up to 14 decimal digits and a sign. For complete details on the specification of data words in ARGUS language, reference should be made to Section IX of the ARGUS Manual of Assembly Language.

Special Register Words

As previously noted, a special register can store 16 information bits, or one-third of a full Honeywell 800 word. When these bits are manipulated within the special register circuitry, the high-order bit is interpreted as a sign (1 = plus, 0 = minus). Depending upon the type of addressing used, the remaining 15 bits of a special register word may be interpreted as a main memory address, consisting of a bank indicator and subaddress, or as a special register address, consisting of a group indicator and subaddress (see Figures IV-1 and IV-2). When a special register word is modified arithmetically within the special register circuitry, the value of the sign bit determines whether it is incremented or decremented. A special register word is identified in ARGUS language by the constant code SPEC.

Instruction Words

The 48 bits of a Honeywell 800 instruction word are interpreted as four groups of 12 bits each. Bits 1-12 represent the command code; bits 13-24, 25-36, and 37-48

are designated as the A address group, B address group, and C address group, respectively. The address portions of instructions normally are used to designate the locations of operands and results, but in certain instructions they may contain special information such as the number of words to be moved, the number of bits to be shifted, a change of sequence counter, and so forth. A detailed discussion of addressing in the Honeywell 800 will be found in Section IV.

Excluding the scientific instructions, to be found in another manual, machine instructions fall into four major categories: general instructions, unmasked and masked; inherent mask instructions; peripheral and print instructions, and simulator instructions. The masked general instructions and the peripheral and print instructions are uniquely designated by six bits -- bits 7 through 12 of the instruction word. The unmasked general instructions and the inherent mask instructions are uniquely designated by eight bits -- bits 7 through 12, plus bits 2 and 3. The simulator instructions are uniquely defined by only three bits -- bits 10 through 12. These groups of bits which uniquely specify the operation to be performed are called the operation code. The bits of the command code which are not used for the operation code serve various other purposes which will be described as the instruction types are discussed. A graphic summary of the format of the major command code types appears in Figure III-2. The command codes for the individual instructions, together with their mnemonic operation codes in ARGUS language, are set forth by major instruction type in Table II, page 144.

General Instructions

General instructions include the arithmetic operations, logical operations, decisions, and information transfers. As noted in Table II, certain of these instructions may only be performed without masks; others may be performed either with or without masks. Regardless of masking, bit 1 of a general instruction command code, called the bisequence bit, always specifies the source of the next instruction (see Figure III-2). If bit 1 is a zero, the next instruction will be taken from the sequence counter; if bit 1 is a one, the next instruction will be selected from the cosequence counter. In ARGUS language, the source of the next instruction is specified in column 23 of the ARGUS input card by an "S" or blank for sequence counter or by a "C" for cosequence counter.

SECTION III. THE HONEYWELL 800 WORD

BITS	1	2	3	4	5	6	7	8	9	10	11	12
General Instructions Unmasked	S/C	Op Code	Memory Designator			Operation Code						
			A	B	C							
General Instructions Masked	S/C	Partial Mask Address					Operation Code					
Inherent Mask Instructions	S/C	Op Code	Memory Designator			Operation Code						
			A	B	C							
Peripheral Instructions	Peripheral Address					Operation Code						
	I/O Channel		Device									
Simulator Instructions	D/I	Remainder of Address								1	1	1

Notes: S/C = Sequence or Cosequence Counter Address Used
 I/O = Input or Output
 D/I = Direct or Indexed

Figure III-2. Honeywell 800 Major Instruction Types

Unmasked General Instructions

Unmasked general operation codes are specified by command code bits 2, 3, and 7 through 12. Bits 4, 5, and 6 designate whether the A, B, and C addresses, respectively, refer to a main memory or a control memory location. If a memory designator bit is zero, then the corresponding address refers to main memory; a designator bit of one denotes a control memory address. In ARGUS language, the memory designator bit is not explicitly stated but is implied by the type of addressing used (see Section IV).

Masked General Instructions

When general instructions are performed under the control of masks, they usually designate partial words as operands and results. For this reason, they are frequently referred to as "field" instructions. When a field instruction is performed, the same mask is applied to operands and result. Only those bit positions of the operands which correspond to binary ones in the mask word are used in the operation. The positions of the result location which do not correspond to binary ones in the mask are not altered by the operation. The location of the mask used in a field instruction is speci-

fied by bits 2 through 6 of the command code, in conjunction with bits 2 through 5 and 11 through 16 of a special register called the mask index register (MXR). A complete description of the way in which the five command code bits, called the partial mask address, are united with the ten bits of the MXR to designate the location of the mask will be found in Section V under the discussion of the mask index register.

In ARGUS language the location of the mask is specified by writing its symbolic tag in the command code field, following the operation code and separated from it by a comma. Thus, the instruction

```
DS, MASK2 WAGES DEDUCTNS WEEKSPAY
```

is performed under control of the mask stored in the memory location assigned by ARGUS to the symbolic tag MASK2. Since field instructions use the memory designator bit positions in the partial mask address, it is impossible for these instructions to address control memory.

Inherent Mask Instructions

The use of masks is not restricted to field instructions, but extends to the inherent mask instructions. These instructions, which have the same command code format as the unmasked general category, include five shift instructions, a substitute, and an extract instruction. The chief distinction between the two types of masked instruction lies in the fact that the inherent mask instructions use bits from the B address group rather than from the command code to specify the location of the mask. For the shift instructions, the low-order six bits of the B address group are used in conjunction with bits 2 through 5 and 6 through 10 of the mask index register to locate the mask. In the substitute and extract instructions, the entire B address group is used to specify the location of the mask, without reference to the MXR.

A further difference between inherent mask and field instructions is that the latter always operate in the "protected" mode; in other words, the portions of the result location corresponding to binary zeros in the mask are preserved during the operation. The inherent mask group, on the other hand, includes three instructions which operate in the "unprotected" mode, in which the unmasked portions of the result location are cleared to zeros. In ARGUS language, the location of the mask for a shift instruction

is specified in the same way as for field instructions: by writing its symbolic tag in the command code field following the operation code.

Peripheral and Print Instructions

Every instruction in the peripheral group performs some function involving a magnetic tape unit or a peripheral device. The high-order six bits of the peripheral instruction command codes are used to specify a magnetic tape or peripheral address. Thus, these instructions cannot specify the source of the next instruction or address the control memory. The peripheral address bits are divided into two groups of three bits each. Bits 1 through 3 specify one of eight input or output channels (the operation code itself defines whether the channel is input or output) and bits 4 through 6 specify one of the devices attached to this channel. A more detailed explanation of the assignment of peripheral address bits will be found in Section XI.

The print instruction involves the use of a console or inquiry station typewriter. In this instruction, the high-order six bits of the command code are used as follows: bit 1 designates the sequence or cosequence counter as the source of the next instruction; bits 2 and 3 are irrelevant; and bits 4, 5, and 6 serve as A, B, and C address memory designators, respectively. Thus, this instruction can specify the source of the next instruction and address the control memory.

Several differences should be noted between the machine command codes and ARGUS notation for these instructions. First, the peripheral command codes in ARGUS language are reversed in terms of machine language. In other words the mnemonic operation code is written first, followed by the device address expressed as an alphabetic code from AA to HH. Secondly, although there is but one machine instruction for the print function, ARGUS recognizes three distinct mnemonic codes to indicate alphanumeric (PRA), hexadecimal (PRD), or octal (PRO) print format. In machine language, the type of print format is specified by bits 5 and 6 of the B address group.

Simulator Instructions

Any instruction in which command code bits 10 through 12 are all ones is called a simulator instruction, since it permits the programmer to represent with a single instruction any function not built into the equipment logic, such as a machine instruction

for some other data processing system. Each such instruction provides an entry to a simulator routine which is coded by the programmer and stored beginning with the next memory location after the address specified by command code bits 2 through 12. When the instruction is performed, it is transferred to the memory location specified by command code bits 2 through 12, the cosequence counter is set to the next higher address, and the next instruction is taken from the cosequence counter. If bit 1 of the command code is zero, bits 2 through 12 are interpreted as a main memory subaddress. If bit 1 is one, bits 2 through 12 are interpreted as a 3-bit index register designator, and an 8-bit augments (see Section IV). The address portions of a simulator instruction have no assigned function and may be used to store parameters used by the simulator routine. The command code for an ARGUS simulator instruction is S, followed by a comma and an address designated by a symbolic tag or by an index register designator with an augments of seven.

Special Words

Two special Honeywell 800 words -- the end-of-record word and the end-of-item word -- deserve separate mention in this section. The end-of-record word is a word whose 48 information bits are:

1010 1010 0000 0000 1110 1110 1110 1110 1101 1101 1101 1101

This word is used to designate the end of a group of words constituting a single record. When records are being written on tape, the write operation stops only when an end-of-record word is sensed in memory. End-of-record words are automatically generated in the central processor during execution of compute orthocount and record transfer instructions. Their function will be detailed further as these instructions are discussed.

The end-of-item word is a word whose high-order 32 bits are identical to the high-order 32 bits of the end-of-record word. The low-order 16 bits are irrelevant for purposes of identification. As the name implies, an end-of-item word is used to designate the end of a group of words constituting a single item within a record. A record may contain an unspecified number of items, each of which is followed by an end-of-item word. (The end-of-record word is also an end-of-item word.) End-of-item words

SECTION III. THE HONEYWELL 800 WORD

are automatically generated in the central processor during execution of an item transfer instruction, while certain other instructions sense for these words during execution. Their function will be highlighted in the discussion of these instructions.

SECTION IV ADDRESSING

The Honeywell 800 main memory consists of two, four, six, or eight banks of 2048 words each, providing up to 16,384 words, and may be increased by an additional 16,384 to provide a maximum of 32,768 words. Each main memory location is directly addressable and is uniquely designated by a 15-bit configuration. This array of bits may also be thought of as an 11-bit subaddress to specify, in binary, one of the 2048 locations in a bank and a 4-bit bank indicator to specify a memory bank. Since the more typical systems do not include the additional 16,384 words of main memory, discussions of addressing in this manual assume a system with no more than eight banks, unless otherwise noted. In such a system, the high-order bit of the bank indicator is always zero. The bit structure of a main memory address is shown in Figure IV-1. (It is sometimes convenient to express the value of such a 15-bit array as five octal digits. In this notation, the memory addresses of a 16,384-word system range from 00000 to 37777.)




Figure IV-1. Main Memory Address

The control memory consists of 256 special registers divided into eight groups of 32 registers each. Every special register is directly addressable by a unique 8-bit configuration. This array of bits may also be thought of as a 3-bit group indicator designating one of the eight special register groups and a 5-bit subaddress specifying one of the 32 registers in a group (see Figure IV-2).




Figure IV-2. Special Register Address

SECTION IV. ADDRESSING

Since both the main and control memories are directly addressable, some means must be provided to specify which memory is being addressed in each of the three address groups of an instruction. This is accomplished by defining one memory designator bit position in the command code of the instruction for each of the three address groups. A zero designator bit indicates that the respective address refers to main memory; a designator bit of one indicates that the address refers to a control memory location (special register). For those command codes which do not provide the memory designator bit positions, designators of zero are always implied and the respective addresses are always interpreted as main memory addresses. During the execution of an instruction, the designator bit does not appear in the address selectors, but is stored in a separate unit of the control circuitry.

Since an instruction address includes 13 bits (12 bits in the address group plus a memory designator bit, explicit or implied), it does not precisely specify a complete main memory address (15 bits) or a complete control memory address (8 bits). The instruction address bits may be interpreted by the central processor in a number of different ways to form a complete main or control memory address. For example, direct addressing is the explicit statement of the desired main or control memory sub-address in the instruction address. Indexed addressing refers to the technique of augmenting a main or control memory address stored in an index register to form the desired address. Indirect addressing refers to the technique of stating the address of a special register in which the desired main memory address is stored. Main memory locations can be addressed in any of these ways; special registers are addressed either directly or by indexing.

As noted in Section III, the main memory may be addressed by all instructions in any of the four major categories. Special registers, on the other hand, may be addressed only by general unmasked instructions, inherent mask instructions and the print instruction, since these are the only types of instruction which provide for the memory designator bits in the command code.

Direct Memory Location Address

Each address group in a Honeywell 800 instruction word consists of 12 binary digits. Bit 1 of this 12-bit configuration specifies whether the address is direct or indexed. If bit 1 is zero, the address is direct, if bit 1 is one, the address is indexed.

If bit 1 of an address group is zero (direct) and the corresponding memory designator bit is zero (main memory), then the remaining 11 bits of the address group are interpreted as a subaddress designating one of the 2048 locations in a bank of memory. The bank indicator stored in the sequencing counter which selected the instruction is appended to this subaddress to form a complete 15-bit address (see Figure IV-3). Thus, every direct memory location address in an instruction always refers to the bank in which the instruction was stored.

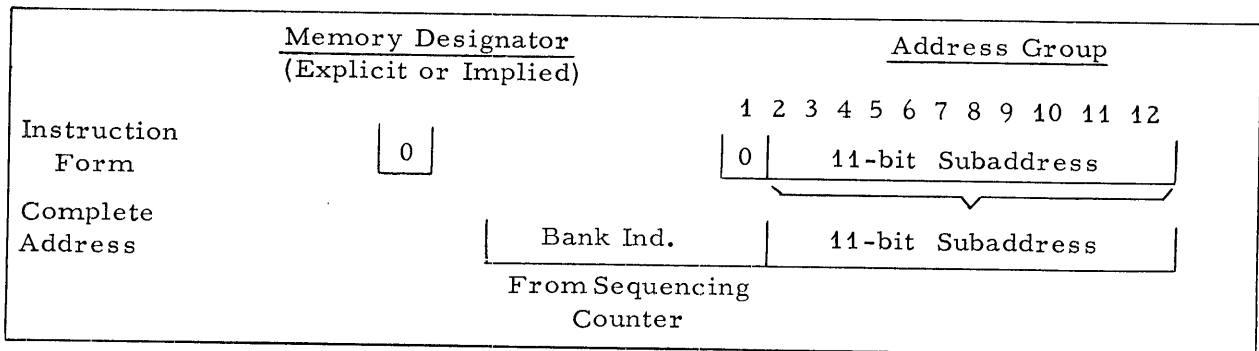


Figure IV-3. Direct Memory Location Address

In an ARGUS instruction, a direct memory location address is specified by a symbolic tag or by address arithmetic, in which the address is designated according to its relative position with reference to the instruction in which it appears or with reference to a symbolic tag. These three types of direct memory location addressing are illustrated in the ARGUS instruction:

DS SALARY C, + 20 SALARY - 2.

SALARY is the symbolic tag of a main memory location; C, + 20 represents the location 20 after the location of the instruction itself; and SALARY - 2 represents the location two before that tagged SALARY.

Direct Special Register Address

If bit 1 of the address group is zero (direct) and the memory designator bit is one (control memory), bits 8-12 of the address group are interpreted as the subaddress of one of the 32 special registers in the group which includes the sequencing counter that selected the instruction. The central processor attaches to this subaddress the group indicator associated with the sequencing counter to form a complete 8-bit special reg-

SECTION IV. ADDRESSING

ister address (see Figure IV-2). If bit 7 of the address group (called the tabular bit) is zero, then the 8-bit array is interpreted as a direct special register address; that is, the specified register is used as an operand location or as a result location. Bits 2 through 6 of the address group specify an increment in the range 0 through 31 (in binary), which may be added, under control of the special register sign bit, to the low-order bits of the special register after use, thereby altering them permanently. If the special register sign bit is positive, the value of the increment is added to the contents of the special register, and the contents are said to be incremented. If the sign is negative, the value of the increment is subtracted from the contents of the special register and the contents are said to be decremented. Incrementing (or decrementing) always occurs when the special register is addressed as the source of an operand, never when the special register is addressed as a result location. The formation of a direct special register address is illustrated in Figure IV-4.

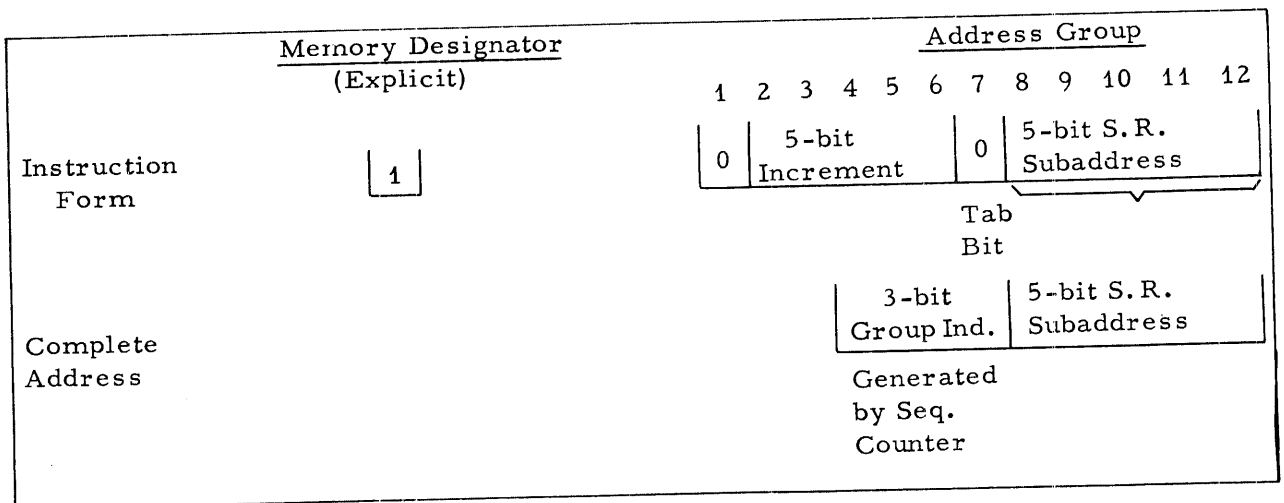


Figure IV-4. Direct Special Register Address

Since the group indicator attached to the special register subaddress is always that of the sequencing counter which selected the instruction, a direct special register address always refers to the special register group containing the sequencing counter which selected the instruction.

In ARGUS language, the direct address of a special register is indicated by the letter Z, followed by a special register designation and an unsigned increment from 0 to 31, all separated by commas. The letter Z, in effect, represents a one in the

memory designator bit position of the command code, a zero in the first bit position of the address group, and a zero in the tabular bit position of the address group. The special register designation may be either the numeric subaddress or the mnemonic subaddress of the desired register, as shown in Figure V-1 (page 50). If the contents of the special register are not to be altered, the programmer may specify an increment of zero or may omit the increment entirely. The ARGUS address

Z, R1, 10

indicates that general purpose register 1 is addressed directly as the source of an operand or as a result location; if addressed as an operand source, its contents are to be incremented (or decremented) by 10 after use. The address

Z, X0

indicates that index register 0 is directly addressed and that no incrementing is to take place.

It should be noted that the relative positions of the special register subaddress and the increment are reversed in ARGUS language from their machine language arrangement. In other words, the increment appears in the low-order position of the ARGUS address, but in the high-order bits of the machine address group.

Indexed Memory Location Address

Each special register group includes eight index registers. An indexed address refers to one of these registers in the special register group of the sequencing counter which selected the instruction and is defined by a one in bit 1 of the address group. The remaining 11 bits of the address group are interpreted as an index register number and an augments to be added to the contents of that index register before use. Bits 2 through 4 designate one of the eight registers in the group, while bits 5 through 12 specify, in binary, a number from 0 to 255 which augments the low-order bits of the contents of the index register. (Note that an indexed address specifying index register 7 with an augments of 255 is interpreted as an inactive address, see page 46). It should be emphasized that in indexed addressing the index register is not identified by its 5-bit subaddress, but by only three bits which designate its position

SECTION IV. ADDRESSING

within the group of eight index registers. Whenever a special register is addressed in an instruction by its full 5-bit subaddress, it is said to be explicitly addressed. When it is referenced in any other way, it is said to be implicitly addressed or referenced. Since the index register in an indexed address is denoted by only three bits, an index register is always referenced implicitly in indexed addressing.

Like all special registers, an index register has the capacity to store 16 information bits of which bit 1 is a sign bit. If the memory designator bit in the command code of the instruction is zero (either explicit or implied), the low-order 15 bits stored in the referenced index register are interpreted as a bank indicator and an 11-bit main memory subaddress. When the instruction is performed, the 8-bit augmenter is added to the stored 15-bit address, under control of the index register sign, to form the desired main memory address. This process has no effect upon the contents of the index register, which retains the unaugmented address. The interpretation of an indexed memory location address is shown in Figure IV-5.

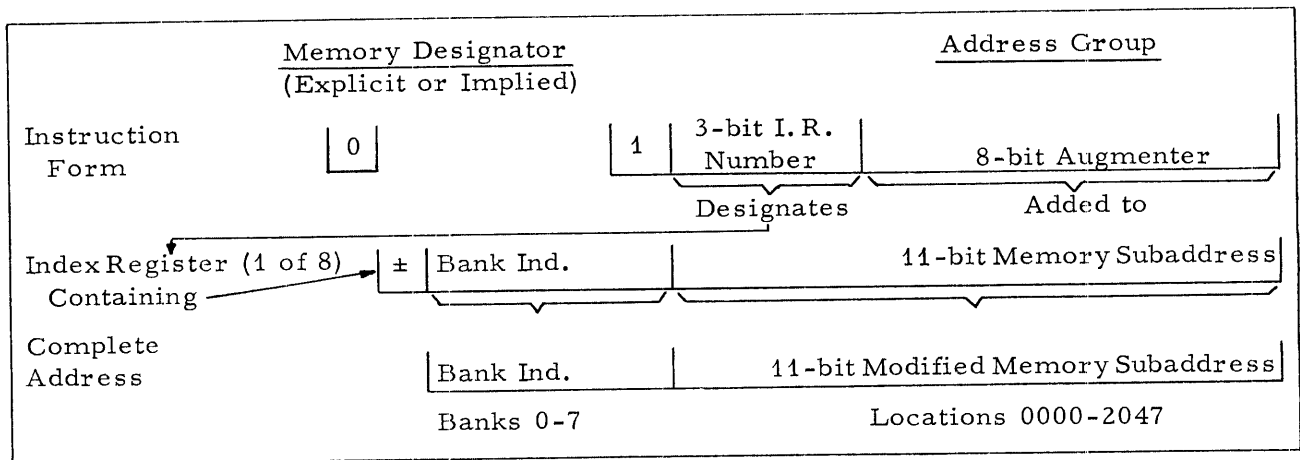


Figure IV-5. Indexed Memory Location Address

Since the index register contains a full 15-bit memory address, indexed addressing, unlike direct addressing, permits the programmer to address locations in any main memory bank, regardless of the bank indicator stored in the controlling sequencing counter. This type of addressing is also useful in processing multi-word items or in referring to a stored table, where the address of the first word of the item or table is stored in an index register and all references to the item or table are made using the index register with appropriate augmenter. It must be remembered, however, that positive augmentation occurs only if the index register sign is positive. If

the sum of the augments plus the stored subaddress exceeds 2047, a carry occurs into the bank indicator, and the resulting address designates a location in a different bank from the address stored in the index register.

In ARGUS language, an indexed memory location address is indicated by writing an index register number (from 0 to 7) followed by a comma and a number from 0 to 255 or a symbolic tag to represent the augments. Thus the address

5, 10

specifies that the contents of index register 5 in the related special register group are to be augmented by 10 to form the complete memory address of an operand or result location. Reference should be made to Section VI of the ARGUS Manual of Assembly Language for details on the use of symbolic tags to represent the augments.

Indexed Special Register Address

If bit 1 of the address group is one (indexed) and the memory designator bit is one (control memory), the address group is interpreted as an index register number and an augments, but the augmented contents of the referenced index register are interpreted as a special register address rather than a main memory address. When the augments has been added to the low-order eight bits of the index register, the resulting configuration is interpreted as shown in Figure IV-6.

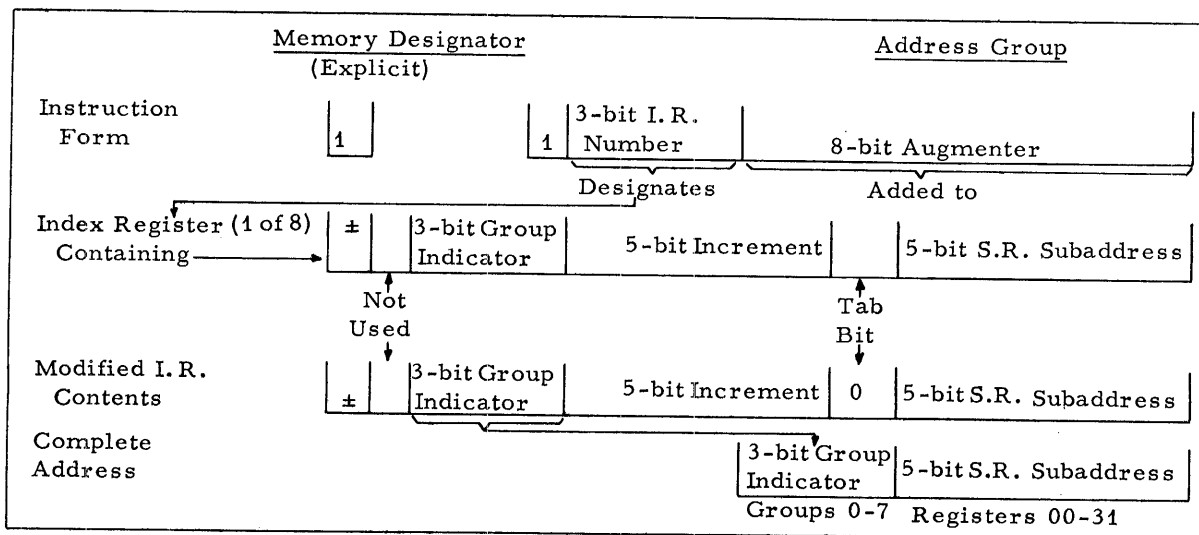


Figure IV-6. Indexed Special Register Address

SECTION IV. ADDRESSING

Two facts illustrated by Figure IV-6 should be particularly noted. Since the augmented contents of the index register are interpreted as a special register address complete with group indicator, this type of addressing, unlike direct special register addressing, permits the programmer to address special registers in any group. As noted in Section V, this facility has particular significance in connection with access to certain special registers involved in reading and writing operations.

The second point involves the value of the tabular bit. Depending upon the original contents of the index register and the value of the augments, the tabular bit in the modified index register contents may be zero or one. If this bit is zero, then the special register is directly addressed, as defined under the discussion of direct special register addressing. If the tab bit is one, on the other hand, the type of addressing is indirect, as described below under the discussion of indirect addressing. Regardless of the value of the tabular bit, the contents of the special register will be permanently modified, after use, by the value of the 5-bit increment, under control of the sign of the special register itself, provided that the special register is not addressed as a result location. As always, the contents of the index register are not altered by the indexing process.

In ARGUS language, an indexed special register address takes the form

Index Register Designator, Z, Special Register Designator, Increment.

The index register designator is a number from 0 to 7 which specifies one of the eight index registers related to the controlling sequencing counter. The special register designator may be a number from 0 to 31 or it may be mnemonic (see Figure V-1, page 50). The increment may be a number from 0 to 3 or it may be omitted. The manner in which these numbers are used to modify the index register contents and form a special register address is discussed in detail in Section VI of the ARGUS Manual of Assembly Language.

Indirect Memory Location Address

In some instances, it is useful to be able to specify in the address group of the instruction the address of a special register where the main memory address of the desired operand is stored, rather than to specify the location of the operand directly. This method of locating an operand is called indirect memory location addressing.

In this type of addressing, the bit configuration of the address group is identical to that described under direct special register addressing with the exception that the tabular bit (bit 7) in the address group has the value of one rather than zero. Since the memory designator bit must also have the value of one (control memory), this type of addressing may be used only in unmasked general instructions, inherent mask instructions, or print instructions. The special register whose address is generated, as shown in Figure IV-7, contains not the operand for the instruction, but the address of the operand, in main memory.

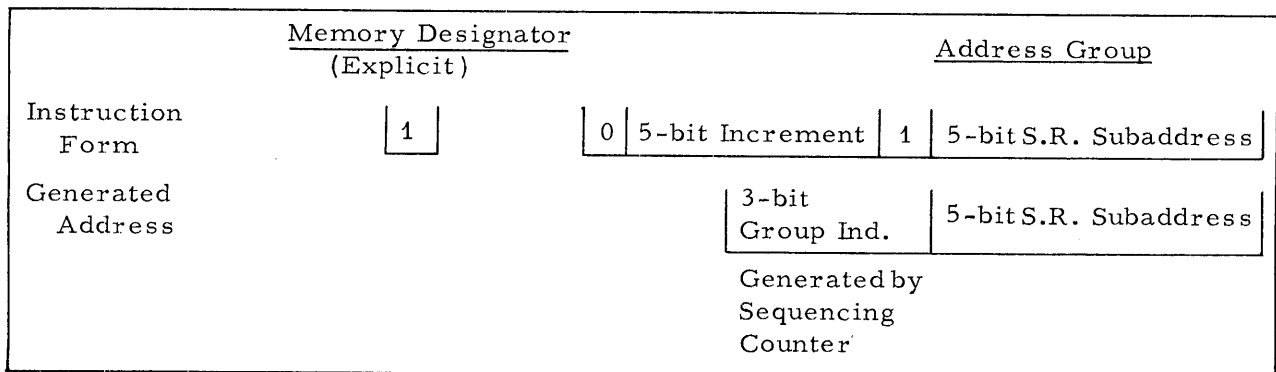


Figure IV-7. Indirect Memory Location Address

The address generated is that of a special register in the same group as the sequencing counter which selected the instruction. The contents of this special register are interpreted as a sign, followed by a 4-bit bank indicator and an 11-bit subaddress designating the main memory location (in any bank) where the operand will be found. Thus, indirect memory location addressing, like indexed memory location addressing, provides access to operands in any bank of memory regardless of the bank indicator stored in the sequencing counter. After the contents of the special register have been used to locate the desired operand, the low-order bits of the stored contents are permanently modified by the increment specified in the address group of the instruction, under control of the special register sign. This incrementing (or decrementing) takes place regardless of whether the memory location addressed is an operand or a result location.

Indirect addressing is a useful tool for stepping sequentially through an array of items, processing the N^{th} word of each item. The location of word N of the first item is stored in a special register. When this location is addressed indirectly, the use of

SECTION IV. ADDRESSING

the proper increment (equal to the item size) sets the contents of the special register to the location of word N of the second item, and so forth, until word N of each item has been processed.

ARGUS recognizes the use of indirect memory location addressing by the notation

N, Special Register Designator, Increment

where N represents a memory designator bit of one in the command code, a zero in the first bit position of the address group, and a one in the tabular bit position of the address group. The special register designator specifies one of the registers in the related group, either numerically or mnemonically (see Figure V-1, page 50), and the increment is a number from 0 to 31. The computer interprets the contents of the specified register as the bank indicator and subaddress of a memory location in any bank. The increment is added to the low-order bits of the contents of the special register after use, permanently altering them. Thus, the address

N, R1, 10

designates the contents of the related special register R1, which are interpreted as the location of an operand in main memory. After use, the contents of R1 are incremented by 10.

Indexed Indirect Memory Location Address

As noted in the discussion of indexed special register addressing, the contents of an index register may be interpreted as a special register group indicator and subaddress, a tabular bit, and an increment. If the tabular bit position of the augmented index register contents has the value of one, then the contents of the special register (designated by the augmented index register contents) are used to locate the operand in main memory. This type of addressing is called indexed indirect memory location addressing. The generation of an indexed indirect memory location address is identical to that shown in Figure IV-6. However, the tabular bit position in the augmented contents of the index register has the value of one for an indexed indirect memory location address whereas it has the value of zero for indexed special register addressing.

Indexed indirect memory location addressing makes it possible to use any of the 256 special registers in the system to address any available memory location indirectly. The retained contents of the special register are always modified after use by the amount of the increment, under control of the special register sign bit.

ARGUS notation for an indexed indirect memory location address resembles that for an indexed special register address, taking the form

Index Register Designator, N, Special Register Designator, Increment.

The comments made with respect to ARGUS notation for an indexed special register address are also applicable to an ARGUS indexed indirect memory location address.

Summary of Address Forms

The binary forms of six different address types are described in the preceding pages and illustrated in Figures IV-3 through IV-7. Figure IV-8 suggests a method with which the reader may determine by inspection the address type of any binary address configuration. This figure is not illustrative of the steps taken by the machine in interpreting addresses.

Significant Main Memory Addresses

Regardless of the amount of main memory available, the first memory bank of every Honeywell 800 system includes certain locations whose use by the programmer is restricted. These locations are automatically involved in certain central processor functions described below. Although reserved for these functions, they are nevertheless directly addressable by the programmer and may be used with caution by a person familiar with the situations in which the central processor uses them automatically. (It is recommended that they not be used when processing is done in parallel.) The complete addresses, in octal notation, for these "reserved" locations are as follows:

00000-00017:	Used for automatic access in the multiply instructions
00021	: Main console typewriter buffer
	and
	Console slave typewriter buffer
00023-00077:	Buffers for inquiry station typewriters.

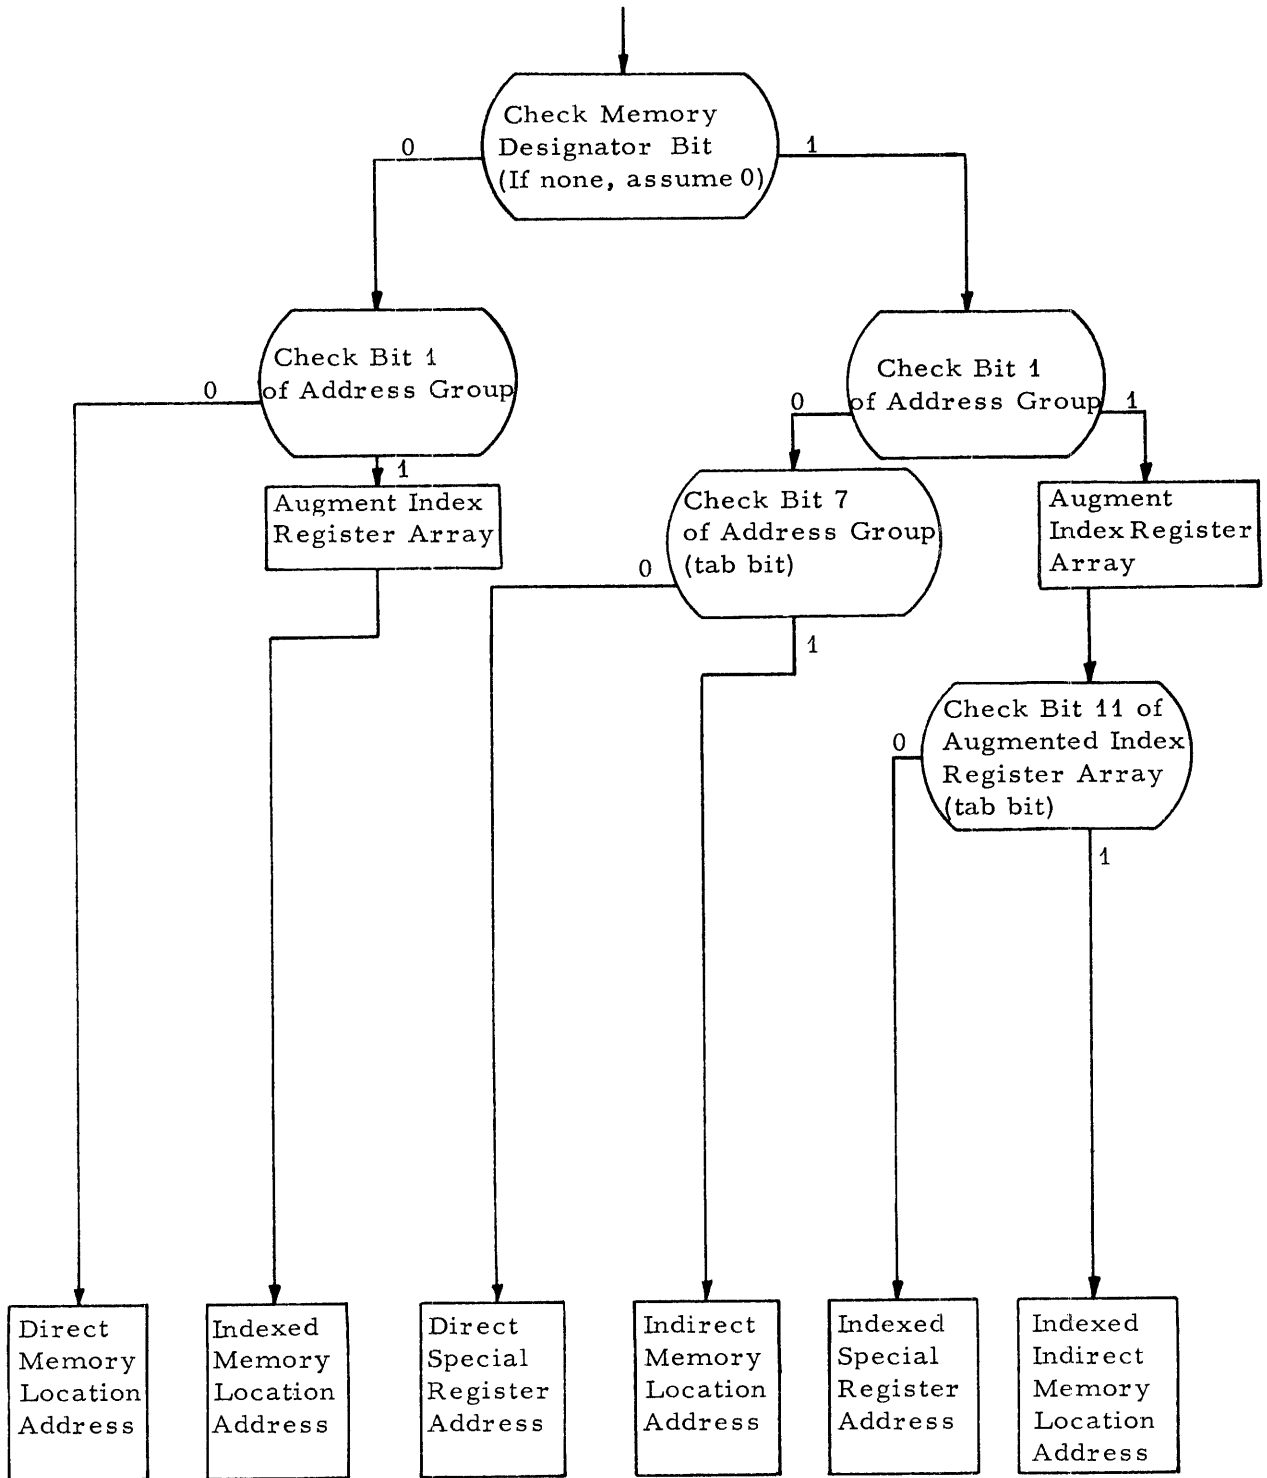


Figure IV-8. Interpretation of Address Bit Structure

The multiply instructions (see Section VI) in the Honeywell 800 generate a set of multiples of the multiplicand which are stored in locations 00000-00017, destroying any information previously stored in these locations by the programmer. Since 16 multiples of the multiplicand are generated during execution of a binary multiply, this instruction involves all 16 locations. The decimal multiply instruction, however, requires only 10 partial products, so that only locations 00000-00011 inclusive are affected by this instruction. It must be remembered that every program running in parallel uses these same locations whenever a multiply instruction is executed. Thus, a programmer who hopes to use these locations with impunity must consider the requirements of other programs which may be run at the same time as his own. It should also be noted that the multiples stored in these locations include modulo -3 check bits stored in the parity bit positions, with the result that these locations will generally contain words which the parity checking circuits will find invalid.

Since the console typewriter is used at least to a limited degree in every system, its buffer location should not be used by the programmer for storage. The buffer locations for inquiry station typewriters not included in a particular system, however, may be used freely.

Stopper Address

When the contents of a special register, interpreted as a main memory address, are modified by incrementing or augmenting, a carry may occur from the 11-bit subaddress into the bank indicator bits. Thus, a sequencing counter can be stepped through successive memory banks, and a single peripheral or transfer instruction can handle a record which is not stored entirely within one memory bank. There is one address, however, which by definition is neither incremented nor decremented when it appears in a special register. This address, called a stopper address, represents the highest-numbered location in the memory of a given Honeywell 800 system, regardless of the number of banks in the system. Its 11-bit subaddress, therefore, represents location 2047 in some memory bank. Its bank indicator is the highest such indicator in the particular system and varies from installation to installation. The largest possible value, in octal, for a complete stopper address is 77777 (15 binary ones). This occurs only in a system having 32,768 words of main memory.

SECTION IV. ADDRESSING

The stopper location can be addressed directly only by a program controlled by a sequencing counter which contains the highest bank indicator in the system. Other programs must address the stopper either by indexing or indirectly through a special register containing the highest bank indicator and a subaddress consisting of 11 binary ones. By addressing the stopper in the A address group of a read instruction, it is possible to move tape without disturbing any memory locations except the stopper. Similarly, it is possible to read only part of a record into memory and discard the balance by causing the first unwanted word to fall in the stopper location.

Although the stopper address cannot be incremented or augmented, it is possible for an address in a special register to receive an increment or augments greater than the difference between its initial value and the stopper or to be decremented by an amount greater than its initial value. The effects of such operations, however, should be carefully noted. These are summarized below:

1. If a word in the control memory receives an increment greater than the difference between its initial value and 77777 or a decrement greater than its initial value, the result restored to the special register contains invalid parity bits. The next attempt to read out this special register will result in a control error in the system. However, it is possible to write into this special register without error.
2. Whenever the stopper address for a given installation is less than 77777, and a word in the control memory receives an increment greater than the difference between its initial value and the stopper (but not greater than the difference between its initial value and 77777), a legal special register word is created and direct addressing of this special register may take place without error. The resulting special register word, however, represents the address of a non-existent main memory location for this installation. Thus, if this special register is referenced as the source of a main memory address (indirect memory location addressing), a control error will result.
3. If a legal special register word representing a memory location with an address greater than the stopper and having a negative sign bit appears in an index register, this index register may be used for indexed addressing without error, provided that the result of indexing is the address of an existing memory location.

Inactive Addresses

The Honeywell 800 central processor contains three arithmetic registers which have no addresses: the accumulator, the mask register, and the low-order product register. Programmer access to these registers is provided by the technique of in-

active addressing with certain specified instructions. Whenever an address group in an instruction has the octal value 7777 (12 binary ones), that address is said to be inactive. This definition is independent of the value of the corresponding memory designator bit in the command code.

Access to the accumulator is provided by the proper use of inactive addressing in conjunction with the add instructions. Inactive addressing with the extract instruction provides access to the mask register. Access to the low-order product register is made possible by inactive addressing with the instruction "transfer A to B, go to C" (TS in ARGUS language).

The behavior of the accumulator when inactive addressing is used in the binary add, decimal add, or word add is specified as follows:

1. If address A is inactive, the previous contents of the accumulator are used as if they were the contents of A. However, if the contents of the accumulator have already been delivered to a memory location by a previous instruction, a control error will result.
2. If address B is inactive, the accumulator (which contains the contents of A if A is active or the previous contents of the accumulator if A is inactive) is left undisturbed.
3. If address C is inactive, the normal process of hunting for the next sequence counter in demand is inhibited, and the result remains in the accumulator at the conclusion of the instruction.

Thus, if the B and C addresses are inactive, the effect of the instruction is to transfer the contents of address A to the accumulator. If the A and B addresses are inactive, on the other hand, the effect is to transfer the contents of the accumulator to the location specified in the C address. Certain restrictions should be noted with respect to the sequencing of these instructions when they contain inactive addresses. If a decimal or binary add instruction is used to place a word in the accumulator, the same type of instruction should be used to transfer the contents of the accumulator, with proper sign, to memory. Similarly, if a word has been placed in the accumulator by a word add instruction, the word add instruction must be used to deliver the contents of the accumulator to memory. The explanation for this restriction is reserved for the section discussing the arithmetic instructions in detail (Section VI).

The mask register, not to be confused with the mask index registers in the con-

trol memory, is a full-word register in the central processor which stores the mask during the execution of a masked instruction. In the extract instruction, the location of the mask is specified in the B address. At the conclusion of the execution of an extract instruction with all addresses active, the original contents of the B address are left in the mask register. Thus, a word may be loaded into the mask register, without disturbing memory, by using an extract instruction with an inactive C address. If address B is inactive, the previous contents of the mask register will be used as the mask. The contents of the mask register are transferred to the location specified in address C of an extract instruction if address B is inactive and the contents of address A consist of all (48) binary ones.

The low-order product register is of interest to the programmer primarily because it stores the low-order portion of the result of a multiply instruction. In order to obtain this result, the programmer may use the instruction "transfer A to B, go to C" (TS) with an inactive A address, which transfers the contents of the low-order product register to the location specified by the B address. If address A of this instruction is active and address B is inactive, the contents of A are transferred to the low-order product register and to the accumulator.

In addition to its use in providing access to these three arithmetic registers, the technique of inactive addressing has special significance in the peripheral read, write, and rewind instructions. Discussion of these features, however, is postponed until the discussion of the instructions themselves (Section XI).

Two other general situations in which the effect of inactive addressing has been specified should be noted:

1. After execution of any instruction whose C address normally calls for a change in the sequencing counter, the counter will not be changed if the C address is inactive.
2. Whenever an inactive C address occurs in any instruction for which an inactive C address is allowed, the normal process of hunting for the next sequencing counter in demand is omitted.

In all cases of inactive addressing not specifically mentioned in this section, the behavior of the Honeywell 800 is presently unspecified.

In ARGUS language, an inactive address is specified by a hyphen (-) in the address field.

SECTION V

SPECIAL REGISTERS

Each of the 256 special registers in the Honeywell 800 control memory is uniquely designated by eight bits, consisting of a 3-bit group indicator to specify one of eight special register groups and a 5-bit subaddress to designate one of 32 special registers in a group. Figure V-1 lists the 32 registers associated with each group, together with their numerical subaddresses from 0 to 31 and their mnemonic designations in ARGUS language. The function of each of these registers is detailed in this section.

Each special register has the capacity to store 16 bits of information, plus two checking bits. The information bits consist of a sign bit ("one" for plus, "zero" for minus) and 15 bits which usually represent the bank indicator and subaddress of a main memory location. When the contents of a special register are modified arithmetically within the special register addition circuitry, the sign bit determines whether they are incremented or decremented. When a 16-bit special register word is transferred to the accumulator or to a main memory location, it is stored in the low-order 16 bits (bit positions 33 to 48) of the specified location. The high-order 32 bit positions of the location are all cleared to zero. Thus, when a special register word is manipulated within the arithmetic unit by an instruction which treats its operands as full-word signed numbers (see Section VI), the word appears to be negative, since it contains four zeros in the sign bit positions. When information is transferred from the accumulator or from a main memory location to control memory, only the low-order 16 bits of the word are stored in the special register; the high-order 32 bits are discarded.

The discussion of indexed memory location addressing in Section IV states that whenever a special register is addressed by its full 5-bit subaddress in an instruction address group, it is said to be explicitly addressed; whenever a special register is referenced in any other way, it is said to be implicitly addressed or referenced. It is also stated that an index register is implicitly referenced when designated by the 3-bit index register number in an indexed address. Implicit addressing is further illustrated by the following examples:

SECTION V. SPECIAL REGISTERS

1. A sequencing counter is referenced implicitly every time an instruction is selected and executed;
2. A read address or write address counter is referenced implicitly whenever a peripheral read or write instruction is executed;
3. Two AU-CU counters are implicitly referenced whenever an N-word transfer instruction is executed.

It will be noted that in each of these examples the implicitly referenced special register is called a counter. The following rule may be stated: those special registers designated in Figure V-1 as counters are always automatically incremented (if the special register sign is positive) or decremented (if the special register sign is negative) by one each time they are referenced implicitly. When these counters are

<u>Subaddress</u>	<u>Mnemonic Address</u>	<u>Name</u>
00	AU1	AU-CU Counter No. 1
01	AU2	AU-CU Counter No. 2
02	SC	Sequence Counter
03	CSC	Cosequence Counter
04	SH	Sequence History Register
05	CSH	Cosequence History Register
06	UTR	Unprogrammed Transfer Register
07	MXR	Mask Index Register
08-15	X0-X7	Index Registers
16-23	R0-R7	General Purpose Registers
24-31	S0-S7*	General Purpose Registers
28	RAC	Read Address Counter
29	DRAC	Distributed Read Address Counter
30	WAC	Write Address Counter
31	DWAC	Distributed Write Address Counter

* In those special register groups associated with active input and/or output channels, S4-S7 are replaced by RAC, DRAC, and/or WAC, and DWAC.

Figure V-1. Special Register Names, Subaddresses, and Mnemonic Addresses

addressed explicitly in an instruction address group, however, incrementing is not automatic but occurs only under program control, if specified in the address group. The value of the sign bit of a special register is never changed except through explicit addressing.

Under program control, the contents of a special register may be arithmetically modified within the special register circuitry in one of two ways: by augmenting an index register or by incrementing an explicitly addressed special register. Augmenting an index register does not alter the retained contents of the register, since the augments are actually added to the contents of the index register after they have been read out. One and only one address selection occurs each time the index register contents are augmented. Thus, even if the same index register is referenced in two successive instructions or twice within the same instruction, the same base address is used for each address selection. Incrementing, on the other hand, alters the retained contents of the special register. After the special register has been selected and its contents used, the increment specified in the address group is actually added to those contents, and the result is returned to the special register before the next address is selected. Thus, when the same special register is addressed twice within an instruction, the contents of the register at the second addressing are different from the contents at the first addressing, unless, of course, a zero increment is specified in the first instance. Whether done automatically or controlled by the programmer, incrementing (and augmenting) is a checked operation which takes place in addition circuitry peculiar to the special registers.

It is important to emphasize a few basic rules which govern programmed incrementing of explicitly addressed special registers. If a special register is addressed directly as the source of an operand, incrementing takes place. If a special register is addressed directly as a result location, on the other hand, no incrementing takes place even if programmed. Whenever the contents of a special register are used to locate either a source or result location in memory (indirect memory location address), incrementing always takes place except for the unusual case in which the register so used is the sequencing counter from which the next instruction will be selected. More specifically, if a sequencing counter is used to address a memory location

SECTION V. SPECIAL REGISTERS

indirectly in the C address and the instruction is one which would normally change the contents of that counter to the memory location address specified by C, the contents of the sequencing counter will not be incremented even though an increment is specified in the address group.

The rules of incrementing which apply under indirect memory location addressing are illustrated by the instruction "transfer A to B, go to C" shown below (in ARGUS format):

```
TS      ITEMA      N, R1, 1      N, SC, 5
```

When this instruction is executed, the word at ITEMA is transferred to the memory location designated by the address stored in special register R1. The contents of R1 are incremented by one after use and replaced in R1. Since the increment of five is not added to the contents of the sequence counter, the result is the same as if the C address had been inactive.

Any instruction which can explicitly address a special register may operate on its contents. This means that a special register word may be shifted, may be operated upon arithmetically, may be compared, and may be moved around in either main or control memory. When a special register word is brought into the accumulator, bits 1 through 32 are filled with zeros. Since four zero bits in the sign position define a negative number, a control memory word is always negative when manipulated in the accumulator, regardless of the value of the special register sign bit.

The peripheral read and write command codes do not provide memory designator bits for explicit addressing of the control memory. It is therefore impossible to read directly into a special register from a peripheral device or to deliver the contents of a special register directly to a peripheral device. For the same reason, it is also impossible to address a special register in a masked general instruction.

Each group of special registers forms a control center for a single program. Thus, as many as eight independent programs may be active at the same time. Each program proceeds under control of the sequence or cosequence counter in its own special register group and references the other special registers (index registers, mask

index register, and so forth) in this group. Direct memory location addressing allows the programmer to address only those 2048 memory locations within the bank specified by the bank indicator of the sequencing counter which referenced the instruction. When the main memory is addressed through the special registers, however, the program may have access to a location in any bank. Furthermore, when any of the counters in the control memory is incremented, any resulting carry may propagate throughout the full 15-bit address, with the result that the main memory is completely continuous when referenced through these counters. Thus, sequencing of control, reading, writing, and transfer of information may all proceed without regard to bank designation.

Sequencing Counters

Each special register group contains two sequencing counters called the sequence counter (SC) and the cosequence counter (CSC). Except in the case of simulator instructions, the programmer may use either of these counters to sequence his program. Furthermore, in any instruction except the simulator, proceed, and peripheral instructions, he may specify which counter will select the next instruction, with the result that he may change control between the two with complete freedom. The use of two counters in this way is called the bisequence operation mode. Since the behavior of the two counters is identical, the following description of the sequence counter is also applicable to the cosequence counter.

The sequence counter contains a sign and 15 bits which are interpreted by the control circuitry as a bank indicator and a subaddress. These 15 bits represent the complete address of a main memory location from which an instruction is to be selected. Each time the sequence counter is implicitly referenced for the selection of an instruction, its contents are automatically incremented or decremented by one (according to the value of the sign bit) and immediately replaced in the counter. During the execution of an instruction selected from location N , for example, the sequence counter contains the quantity $N + 1$ if the sign bit is positive. In this case, therefore, instructions are taken from successively higher memory locations. If the sign bit is negative, on the other hand, the instructions will be selected from successively lower memory locations.

As previously noted, carries may propagate across the entire 15 bits of a sequencing counter during incrementation. Instruction sequences can therefore pass freely

SECTION V. SPECIAL REGISTERS

from one memory bank to another. If an attempt is made to sequence the counter beyond the highest memory address included in a particular system, however, a control error will result and the machine will stop. The same result will occur if a sequencing counter containing a negative sign and 15 binary zeros is implicitly referenced.

The initial setting of the sequence counter is normally made by transferring a word from main memory whose low-order 16 bits represent the desired sign and the memory address from which the first instruction is to be selected. Alternatively, the starting address may be entered directly into the counter from the console typewriter. Once the sequence counter is set and referenced, it continues to select instructions from successive locations until an instruction is executed which specifies the alternate counter as the source of the next instruction or which changes the contents of the counter itself through explicit or implicit addressing. When an instruction specifies a change of counter but not a change in the contents of a counter, the only change which occurs in the two sequencing counters is the normal incrementation of the counter which selected the instruction. An instruction which explicitly addresses a sequencing counter as a result location simply causes the contents of that counter to be replaced. For example, the instruction

TX AU1 - Z, SC

merely replaces the contents of the sequence counter with the contents of AU1, so that the next instruction is selected from the location whose address is stored in AU1. No record of such a sequence change is retained by the machine. An instruction which changes the contents of a counter by implicit reference, on the other hand, alters the contents of the counter specified as the source of the next instruction and stores the contents of the counter which selected this instruction in the history register (see below) associated with the counter whose contents are changed. Thus a record is always available internally of the last implicit sequence change which affected the contents of either counter. For example, an instruction selected under control of the sequence counter specifies the cosequence counter as the source of the next instruction and directs the program to transfer a word from A to B and select the next instruction from the location specified by C:

TS C RECORD OUTPUT SECTIONA

This instruction puts the address tagged SECTIONA in the cosequence counter, where it is selected as the address of the next instruction, and stores the incremented contents of the sequence counter in the cosequence history register.

History Registers

For each sequencing counter in the system, there is a corresponding history register called the sequence history register (SH) or the cosequence history register (CSH). These registers are used to store the contents of a sequencing counter whenever the counter is implicitly addressed by an instruction specifying a change in its contents. If an instruction selected by the sequence counter from location M specifies a sequence change to location N, and the next instruction is also to be selected by this counter, then the address $M + 1$ is stored in the sequence history register and the sequence counter itself is set to the address N. If, however, the alternate counter is specified as the source of the next instruction, then the cosequence counter is set to N and $M + 1$ is stored in the cosequence history register. In other words, the contents of a history register always represent the incremented address of the instruction which last changed the contents of the associated sequencing counter by implicit reference. As previously noted, no change in the history register occurs if the contents of a sequencing counter are changed by explicit addressing.

Index Registers

The Honeywell 800 contains a total of 64 index registers, of which eight (designated X0-X7) are located in each special register group. Like the other special registers, they contain 16 bits normally interpreted as a sign and a main memory or special register address. Although the index registers must always be loaded and unloaded by the use of an explicit address, they are always implicitly addressed by a 3-bit number when used for their intended purpose in indexed addressing. As explained in Section IV, an indexed address group includes a 3-bit index register number and an 8-bit augmenter to be added to the low-order contents of this register. The retained contents themselves are not modified; the special register addition circuitry merely uses the contents, together with the augmenter, to generate the main memory or special register address of an operand or result location. Since the sign of the register may be positive or negative, at the programmer's option, the generated address may be higher or lower than the base address stored in the register. The central processor accepts augmenters valued from 0 to 255.

SECTION V. SPECIAL REGISTERS

Mask Index Register

Each special register group contains a mask index register (MXR) which is implicitly referenced whenever a field instruction or a shift instruction is executed. The 16 bits of this register are interpreted as a sign, a bank indicator, and the high-order portions of two different subaddresses. Bits 6 through 10 specify a partial address for masks used with the shift instructions; bits 11 through 16 serve the same purpose for masks used in field instructions (see Figure V-2).




Figure V-2. Mask Index Register

The low-order portions of the mask addresses are found in the instructions themselves, as explained in Section III (see pages 28 and 29). When a shift instruction is executed, the central processor unites the low-order six bits of the B address group with the bank indicator and bits 6 through 10 from the mask index register to form the complete main memory address of the mask. This process is illustrated in Figure V-3.

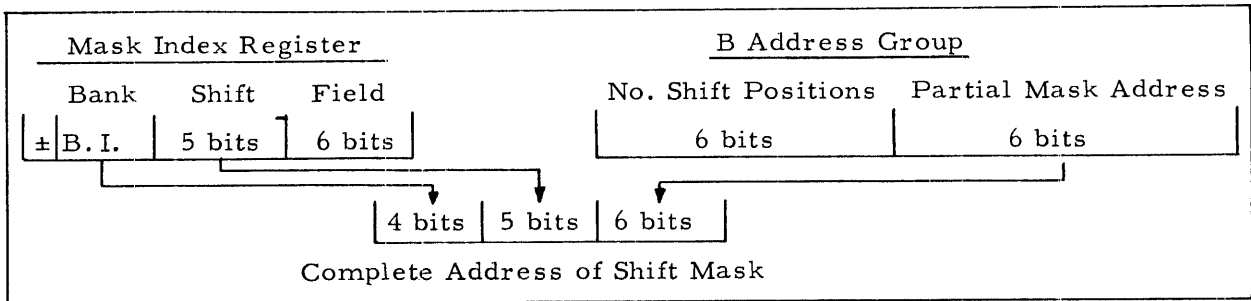


Figure V-3. Generated Mask Address in Shift Instructions

When a field instruction is executed, the central processor attaches the 5-bit partial mask address from the instruction command code to the bank indicator and bits 11 through 16 from the mask index register to form the complete main memory address of the mask. This process is illustrated in Figure V-4.

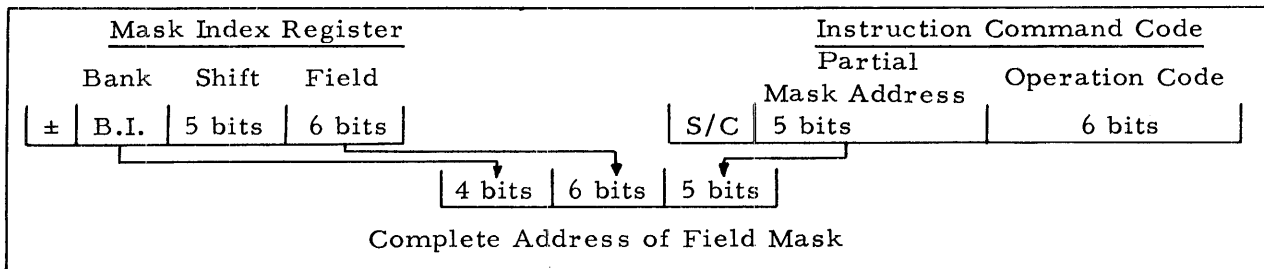


Figure V-4. Generated Mask Address in Field Instructions

Since the mask index register contains a single bank indicator, both shift masks and field masks are stored in the same memory bank. The value of the sign bit is not relevant in locating masks.

The mask index register is set by explicit addressing. Each time the programmer loads the register, he designates 96 memory locations as mask addresses, 64 for shift masks and 32 for field masks. Since the programmer may change the contents of the mask index register whenever he wishes, the number of masks available for his use is virtually unlimited.

General Purpose Registers

Each special register group contains a minimum of 12 general purpose registers (R0-R7, S0-S3), and some groups may contain 14 or 16, depending upon the assignment of input-output channels (see Read-Write Counters below). Like the index registers, general purpose registers are used primarily for address modification. Their use differs from that of index registers, however, in several respects. First, they are always addressed explicitly. Secondly, the specified increment, which has an upper limit of 31, alters the retained contents of the register after use. As in the case of index registers, the address of a memory location generated by adding the increment to the original contents of the register may be higher or lower than the address originally contained in the register, according to the value of the sign bit. These registers are used mainly in the indirect addressing mode to address an operand or a result location in any bank of memory, but they may also be used as programmed counters, as temporary storage for the contents of other special registers, and for any other purpose the programmer may devise.

SECTION V. SPECIAL REGISTERS

Read-Write Counters

Every Honeywell 800 system includes eight channels for entry of information into the central processor from peripheral units and eight channels for output of information from the central processor to peripheral units. Two special registers known as the read address counter (RAC) and the distributed read address counter (DRAC) are associated with each input channel. For each output channel there are two similar counters known as the write address counter (WAC) and the distributed write address counter (DWAC).

Each of the eight special register groups contains a pair of read counters and a pair of write counters. The counters corresponding to the first input and output channels are located in special register group 1; those corresponding to the second input and output channels are located in special register group 2, and so forth up to the counters corresponding to the last channels, which are located in special register group 0. The read-write counters are, therefore, exceptions to the rule that an implicitly addressed special register always belongs to the group associated with the sequencing counter which selected the instruction. An implicitly addressed read or write counter is always one of the pair corresponding to the input or output channel connected to the device addressed. For example, if a program controlled by special register group 3 addresses a peripheral device attached to the first input channel, then the read or write counters in special register group 1 are activated, rather than the counters in special register group 3.

When peripheral equipment is attached to a given channel, the read or write address counter corresponding to that channel is implicitly referenced whenever the peripheral devices are addressed or are operative. In certain cases, the distributed read-write counters are also implicitly referenced. When no hardware is attached to a specific channel, or when the hardware is not being used by any active program, the associated read-write counters may be used as general purpose registers. Whenever these special registers are addressed explicitly, they lose their identity as automatically incremented counters.

When used to control input-output functions, those read-write counters used by a peripheral instruction are automatically loaded during execution of the instruction. The 16 bits initially loaded into the read or write address counter (RAC or WAC)

always represent a sign bit, plus 15 bits (generated from the A address group) which specify the main memory location into which the first word will be read or from which the first word will be written. When a tape is read in the forward direction or written, the sign of the counter is automatically made positive. When a tape is read backward, the sign bit is automatically made negative. The distributed read or write counter (DRAC or DWAC) is automatically loaded during execution of a tape instruction which senses for end-of-item words. Their initial setting represents the address (generated from the B address group) of a main memory location which contains the first entry in a table of addresses used to specify the starting location for each item to be read or written, beginning with the second item. (The starting address for the first item is obtained from the A address group and stored in RAC or WAC.) The sign bit in the distributed counters follows the same convention as that of the read and write address counters. Further details on the functions of the read-write counters are found in Section XI under the discussion of peripheral instructions.

Arithmetic Control Counters

Each special register group contains two arithmetic control counters known mne- monically as AU1 and AU2. One or both of these counters are implicitly referenced, loaded, and automatically incremented during execution of an N-word transfer, item transfer, record transfer, binary or decimal multiply, compute orthocount, or sim- ulator instruction. As an example, during execution of an N-word transfer of 10 words, the initial setting of AU1 represents the main memory or special register address from which the first word is to be transferred, while the initial contents of AU2 rep- resent the location to which this word will be delivered. As successive words are transferred, the counters are automatically incremented to specify a source and re- sult address for each word transferred. At the completion of the instruction, the counters contain addresses equal to their initial settings plus ten.

Since carries may propagate across the low-order 15 bits of the counter, a re- cord which is divided between two memory banks may be transferred as easily as one contained entirely in one bank. It should be noted, however, that when the contents of an arithmetic control counter are interpreted as a special register address, a subad- dress overflow will not change the group indicator but instead will change the value of the tabular bit (bit 11) from zero to one.

SECTION V. SPECIAL REGISTERS

Like other special registers, the arithmetic control counters may be addressed explicitly in order to transfer their contents to main memory or to use them as general purpose registers. The programmer who uses them thus, of course, must remember that information stored in one or both of these registers will be destroyed by the execution of certain instructions. Whenever these special registers are addressed explicitly, they lose their identity as automatically incremented counters. The arithmetic control counters are described more fully in connection with the instructions which use them.

Unprogrammed Transfer Register

The Honeywell 800 is so designed that the occurrence of certain unusual events during execution of a program does not stop the machine but rather effects a transfer of control out of the normal sequence of the program to initiate appropriate action as specified in a programmed subroutine. The unprogrammed transfer register (UTR) in the special register group controlling the program is the key to the location of these subroutines designed to handle the seven different types of conditions which may cause an unprogrammed transfer.

The unprogrammed transfer register is initially loaded, by direct addressing, with 16 bits which represent a sign, a bank indicator, and an 11-bit main memory sub-address. The address thus loaded, called U, must be even or a control error will result when an unprogrammed transfer is attempted. When an unprogrammed transfer situation arises, the control circuitry reads out the contents of the UTR and inserts a one into bit 16 if the instruction causing the unprogrammed transfer was selected from the cosequence counter. The instruction itself is stored in the address thus generated. Thus, the instruction causing the transfer will be stored in U if it was selected by the sequence counter or in $U + 1$ if it was selected by the cosequence counter.

The conditions which result in unprogrammed transfers are listed in Figure V-5. Each of these events causes the execution of one instruction whose address is formed by adding a constant to the contents of the UTR, under control of the UTR sign bit. If the instruction causing the transfer was selected by the sequence counter, the constant is an even number from 2 to 14; if by the cosequence counter, the constant is an odd number from 3 to 15. The value of the constant depends upon the event which caused the transfer and the counter which selected the instruction, as listed in Figure V-5.

The execution of an unprogrammed transfer does not alter the contents of either sequencing counter. Thus, control returns immediately to the normal sequencing of the program unless the unprogrammed transfer instruction itself changes the contents of the sequencing counter from which the next instruction is selected. Note, however, that the bi-sequence bit in the instruction causing the transfer is ignored.

By inserting the appropriate address into the unprogrammed transfer register, the programmer may select any area of memory for use as a corrective routine selection table. He may also change the contents of the UTR at any point in the program or he may change the contents of any entry in the table to correspond with the particular portion of the program being executed at the time. Finally, the unprogrammed transfer register may also be used in connection with fixed starts from the peripheral equipment under conditions which will be discussed with the peripheral equipment in a separate manual.

<u>Event Causing Unprogrammed Transfer</u>	<u>Instruction Stored In</u>	<u>Next Instruction From</u>
Parity Failure		
Sequence Counter	U	$U \pm 2$
Cosequence Counter	$U + 1$	$U \pm 3$
Beginning or End of Tape		
Sequence Counter	U	$U \pm 4$
Cosequence Counter	$U + 1$	$U \pm 5$
Read or Write Error		
Sequence Counter	U	$U \pm 6$
Cosequence Counter	$U + 1$	$U \pm 7$
Addition or Subtraction Overflow		
Sequence Counter	U	$U \pm 8$
Cosequence Counter	$U + 1$	$U \pm 9$
*Division Overcapacity		
Sequence Counter	U	$U \pm 10$
Cosequence Counter	$U + 1$	$U \pm 11$
*Exponential Underflow		
Sequence Counter	U	$U \pm 12$
Cosequence Counter	$U + 1$	$U \pm 13$
*Exponential Overflow		
Sequence Counter	U	$U \pm 14$
Cosequence Counter	$U + 1$	$U \pm 15$

* These unprogrammed transfers apply only to the scientific instructions, to be found in another manual.

Figure V-5. Unprogrammed Transfers of Control

SECTION VI

ARITHMETIC INSTRUCTIONS

Arithmetic instructions in the Honeywell 800 involve the use of two arithmetic registers called the accumulator (AC) and the low-order product register (LOP). As discussed in Section IV, these registers are accessible to the programmer by the technique of inactive addressing with certain specified instructions. The accumulator is used in all the instructions described below. The low-order product register is used in the multiply instructions and in the masked arithmetic instructions.

The Accumulator

The accumulator consists of 48 flip-flops, each capable of storing a single binary digit. Used in conjunction with the accumulator is a single flip-flop called the sign flip-flop. Since the operands handled in most arithmetic operations are treated as 44-bit numbers with 4-bit signs, most arithmetic instructions use the sign flip-flop together with the low-order 44 bits of the accumulator. The exceptional instructions which handle unsigned 48-bit numbers are so noted as they are described.

For all arithmetic operations on signed numbers, the sign flip-flop, originally set to the value of zero, is changed to the value of one if any of the four sign bits in the A operand is a one. (This is the logical OR function.) The setting of the sign flip-flop, the sign of the B operand, and the operation code determine the sign of the result. When the result is read out of the accumulator, a sign consisting of four bits identical in value (zero or one) to the final setting of the sign flip-flop is attached to the low-order 44 bits of the accumulator. Thus only two sign configurations may be obtained as the result of an arithmetic operation: four binary zeros indicating a negative result or four binary ones indicating a positive result. If the result of an add or subtract instruction is zero, the result takes a sign based on the sign of the A operand.

When addition is performed on signed numbers, bits 1 through 4 of the accumulator are automatically filled with binary ones so that if overflow occurs it may be sensed in bit 1, the same position in which it is sensed during the addition of unsigned 48-bit

operands. Similarly, bits 1 through 4 are filled with binary zeros when subtracting signed numbers in order that borrows may be sensed at the same position for both signed and unsigned numbers. If overflow occurs, the instruction is completed and the low-order 44 bits of the accumulator, plus a 4-bit sign based on the value of the sign flip-flop, are stored in the location specified by the C address. An unprogrammed transfer is then made to $U + 8$ or $U + 9$, where the programmer should have stored the entry to a subroutine to handle this condition. The instruction which resulted in the overflow is stored in U or $U + 1$ (see Figure V-5, page 64).

The arithmetic operations are not actually performed in the accumulator but in an adder consisting of 12 gate buffer amplifiers. Both binary and decimal arithmetic are performed in the same adder, which handles three 4-bit groups at a time. For binary instructions, these 4-bit groups are considered as hexadecimal digits, with carry occurring after a group reaches the value of 15. This results in a pure binary operation. For decimal instructions, the adder is made to carry when a 4-bit group reaches the value of nine, so that decimal arithmetic is performed. However, if a decimal addition instruction involves an operand which contains hexadecimal digits, a variant on normal addition occurs in accordance with the following rules:

1. If the hexadecimal digit appears in the A operand, the corresponding digit in the B operand is added in hexadecimal fashion. In other words, $14 + 1$ becomes 15, $14 + 3$ becomes 1 with a carry of 1.
2. If the A operand contains decimal information and a hexadecimal digit occurs in the B operand, then the result is decimalized, as follows: $3 + 14$ becomes 7 with a carry of 1. However, if the result is greater than 19, a control error may occur.

The contents of the operands are inspected digit by digit. Therefore, the result to be obtained by adding two words having both hexadecimal and decimal digits must be ascertained on a digit-by-digit basis. Since this condition is not considered an error by the central processor, except in the circumstance noted above, the programmer will receive no indication of the existence of a hexadecimal digit in an operand handled by a decimal instruction.

Addition and subtraction of signed numbers conforms to normal algebraic rules. Thus, an add instruction causes operands with like signs to be added and operands with unlike signs to be subtracted. A subtract instruction causes operands with un-

SECTION VI. ARITHMETIC INSTRUCTIONS

like signs to be added and operands with like signs to be subtracted. A more detailed discussion of addition in the Honeywell 800 will be found in Appendix A.

Several precautions must be observed in working with the accumulator. In the discussion of inactive addressing (Section IV, page 46ff), it is pointed out that the result of an addition may be left in the accumulator by using an add instruction with an inactive C address and that the contents of the accumulator may be stored in memory by using an add instruction with inactive A and B addresses. It should also be noted that if the contents of the accumulator were formed by an instruction which treated the operands as signed 44-bit numbers, such an instruction must be used to store the contents of the accumulator in order to guarantee them the proper sign. Otherwise, the entire 48-bit contents will be stored rather than the low-order 44 bits with a 4-bit sign determined from the value of the sign flip-flop. Similarly, if the contents of the accumulator were created by an instruction which treated the operands as unsigned 48-bit words, such an instruction must be used to transfer the entire contents of the accumulator to memory without reference to the sign flip-flop.

Another precaution involves the condition of the accumulator after its contents have been delivered to memory. After a result formed in the accumulator has been transferred to memory, the contents of the accumulator are invalid. A second attempt to transfer the result, therefore, will cause a control error and the machine will stop. Since a hunt for the next program demand will have occurred immediately after the first transfer, this behavior imposes no real restriction on the use of the accumulator.

The Low-Order Product Register

The low-order product register is a 48-bit register similar to the accumulator. As its name implies, the register is used to store the low-order portion of the result of a multiply instruction. The contents of the register are interpreted as a sign and a 44-bit number.

Binary Add, BA

The binary add instruction causes the contents of the location specified by A to be added algebraically to the contents of the location specified by B and the result of the operation to be stored in the location specified by C. The contents of both A and B are regarded as 44-bit numbers with 4-bit signs. If any sign bit for the A or B oper-

and is a one, then the corresponding operand is considered positive. After the addition is complete, the low-order 44 bits of the accumulator, plus a 4-bit sign (1111 or 0000) corresponding to the value (1 or 0) of the sign flip-flop, are stored in the location specified by C. If overflow occurs, the instruction is stored in U if the sequence counter selected it or in U + 1 if the cosequence counter selected it, and the next instruction is taken from U + 8 or U + 9.

As an illustration, assume that the following information bits are stored in memory locations tagged MONTHDAY and CONSTNT1:

MONTHDAY	00110.....01110000
CONSTNT1	10000.....00011001

When the instruction

BA MONTHDAY CONSTNT1 TESTAREA

is executed, the following result will be stored in TESTAREA:

TESTAREA	11110.....010001001
----------	---------------------

In the unmasked version of the binary add, all three addresses may be direct, indexed, or indirect and may refer to main memory locations or special registers. When the instruction is masked to permit operations on partial words, however, the A, B, and C addresses may refer only to the main memory in the direct or the indexed mode.

If the A address of the instruction is active and B and C are inactive, the contents of A are placed in the accumulator. If the B address is active and A and C are inactive, then the contents of B are added to the contents of the accumulator. If the C address is active and A and B are inactive, the low-order 44 bits of the accumulator are stored in C with the sign bits 1111 or 0000, depending upon the value of the sign flip-flop.

The time required to execute an unmasked binary add instruction whose A, B, and

SECTION VI. ARITHMETIC INSTRUCTIONS

Addresses directly specify main memory locations is generally four memory cycles. Occasionally, one or two extra memory cycles may be required, as explained in Appendix A. The effect on timing when the instruction is masked or when it uses indexed operands or special registers is summarized in Appendix C.

Decimal Add, DA

The decimal add instruction causes the contents of A to be added algebraically to the contents of B and the result of the operation to be stored in C. The instruction differs from binary add only in the fact that the contents of each operand are handled as eleven 4-bit groups with a 4-bit sign. The sign conventions are identical, and the result stored in memory consists of the low-order 44 bits of the accumulator and a 4-bit sign corresponding to the value of the sign flip-flop.

If memory locations tagged REGPAY and OVERTIME contain the following information

REGPAY	+00000012500
OVERTIME	+00000001750

and the instruction

DA REGPAY OVERTIME WEEKPAY

is executed, then the following result will be stored in WEEKPAY:

WEEKPAY	+00000014250
---------	--------------

The comments on the binary add instruction with respect to overflow, addressing, masking, and timing are equally applicable to the decimal add instruction.

Binary Subtract, BS

This instruction causes the contents of B to be subtracted algebraically from the contents of A and the result to be stored in C. The contents of both A and B are regarded as 44-bit numbers with 4-bit signs. The sign and overflow conventions followed are the same as those described for the binary add instruction.

The remarks about addressing and masking with reference to the binary add instruction also apply to binary subtract. The time required to execute an unmasked binary subtract instruction whose A, B, and C addresses directly specify main memory locations is ordinarily four memory cycles. Under certain unusual conditions described in Appendix A, however, one or two additional memory cycles may be required. The timing effect of masking, indexing, and the use of special registers is summarized in Appendix C.

Decimal Subtract, DS

The decimal subtract instruction differs from binary subtract only in the fact that both operands are regarded as eleven 4-bit groups with 4-bit signs. In all other respects, the description of binary subtract is completely applicable to decimal subtract.

Word Add, WA

Word add is one of the two arithmetic instructions which regard operands as unsigned 48-bit numbers. The instruction adds the absolute values of the entire 48-bit contents of A and B in binary and stores the entire 48 bits of the accumulator in C, making no reference to the sign flip-flop. In contrast to the example cited under the discussion of binary add, when the instruction

WA MONTHDAY CONSTNT1 TESTAREA

is executed with the same operands:

MONTHDAY	00110.....01110000
CONSTNT1	10000.....011001

then the following result will be stored in location TESTAREA:

TESTAREA	10110.....010001001
----------	---------------------

Although the low-order 44 bits of the result stored in TESTAREA are identical for the two instructions, the high-order four bits are different since binary add inserts in these four positions a sign based on the value of the sign flip-flop whereas word add

SECTION VI. ARITHMETIC INSTRUCTIONS

stores the entire contents of the accumulator, without regard for the sign flip-flop.

Overflow is sensed in bit 1. If overflow occurs, the instruction is completed, and the overflow conventions set forth under the description of binary add are followed. With respect to addressing, masking, and timing, the word add instruction is identical to binary add.

Word Difference, WD

The second arithmetic instruction which treats its operands as unsigned 48-bit numbers, word difference, causes the entire 48-bit contents of B to be subtracted in binary from the entire 48-bit contents of A. The entire 48 bits of the accumulator are stored in C. If the absolute value of the contents of B is greater than the absolute value of the contents of A, then overflow occurs, and the result stored in C is the difference of the absolute values of the words. In all other respects, the word difference instruction is identical to word add.

Binary Accumulate, BT

The binary accumulate instruction totals the absolute value of the contents of A the number of times specified by the high-order six bits of the B address group, a number which ranges from 0 through 63. Although the words added are treated as signed 44-bit numbers, only their absolute values are added. The accumulator is not cleared between successive additions except for the high-order four bits. At the conclusion of the series of additions, the 44 low-order bits of the accumulator are stored in C, together with a sign (four binary ones or four binary zeros) based on the value of the sign flip-flop. This value represents the sign of the first word added (the contents of the location originally specified by the A address group).

Step by step, the instruction functions as follows. The low-order 44 bits of the A operand are transferred to the accumulator and the high-order four bits of the accumulator are set to one. If A contains a special register subaddress, incrementing is performed as specified. The high-order four bits of the accumulator are again replaced with ones, and the low-order 44 bits of A are added in binary to the contents of the accumulator. (Note well that the location now specified by A will be different from the original A if incrementing took place.) The specified incrementing is again per-

formed (if A contains a special register subaddress), the high-order four bits of the accumulator are replaced by ones and the low-order 44 bits of A are again added to the accumulator. This process is performed the number of times specified by the high-order six bits of the B address group. If the value of these bits is zero, no information is transferred to the accumulator, the instruction is not executed, and the next instruction is selected from the sequencing counter specified by bit 1 of the command code. The low-order six bits of the B address group are ignored.

As an example, consider the instruction

```
BT    N, R1, 1    4    TOTAL
```

in the case where R1 contains a main memory address tagged OPERAND. The high-order four bits of the accumulator are set to ones, and the low-order 44 bits of OPERAND are transferred to the accumulator. R1 is incremented by one, the high order four bits of the accumulator are replaced by ones, and the low-order 44 bits of OPERAND + 1 are added in binary to the contents of the accumulator. This process is repeated until the contents of OPERAND + 3 have been added to the total in the accumulator. The accumulator now contains the sum of the absolute values of the low-order 44 bits of OPERAND, OPERAND + 1, OPERAND + 2, and OPERAND + 3. This sum, with a sign based on the sign of OPERAND, is then stored in the location tagged TOTAL. Register R1 contains the address of location OPERAND + 4 at the conclusion of the instruction.

Overflow in the accumulator is sensed out of bit 1. If overflow is sensed, the instruction is completed, and the normal overflow procedure (described under the binary add instruction) is performed. Since the high-order four bits of the accumulator are replaced by ones between successive additions, the contents of these four positions are not available to indicate the number of overflows. Thus, unless the programmer knows from the logic of his problem precisely how many overflows may have occurred and at which points, he must repeat the addition process in pairs.

The time required to execute a binary accumulate instruction with direct and/or indirect memory location addresses is three memory cycles plus one memory cycle for each word accumulated. Masking is not permitted with the accumulate instructions.

SECTION VI. ARITHMETIC INSTRUCTIONS

Decimal Accumulate, DT

This instruction is implemented in precisely the same way as binary accumulate, with the exception that the words added are regarded as 11-digit decimal numbers and are added according to the rules of decimal arithmetic. If hexadecimal digits appear in the operands, they are added in the same fashion described under the decimal add instruction. The remarks on overflow, timing, and masking made with reference to binary accumulate also apply to decimal accumulate.

Binary Multiply, BM

The binary multiplication instruction in the Honeywell 800 stores a set of 16 multiples of the multiplicand (or A operand) in the first 16 locations of memory bank 0 (see page 43ff). Any information previously stored in these 16 locations by the programmer will be destroyed during execution of a multiply instruction. Since the multiples stored in these locations use the parity bit positions for modulo -3 check digits, these locations, in general, will contain words which the parity checking circuits will find invalid.

At the beginning of the instruction, the address 00000 (in octal) is placed, with a positive sign bit, in AU-CU counter 1 (AU1). The number zero (zero times the A operand) is stored in location 00000, and AU1 is incremented by one to specify the storage location of the next partial product. The A operand is then placed in location 00001; AU1 is stepped and twice the A operand is placed in location 00002; the counter is stepped again and three times the A operand is placed in location 00003; and so forth until the 16 multiples possible in the hexadecimal system are stored. These multiples are then selected and added in accordance with the value of the digits of the multiplier (or B operand). At the conclusion of the instruction the high-order product found in the low-order 44 bits of the accumulator is stored, together with a sign corresponding to the value of the sign flip-flop, in the location specified by the C address group. The low-order product remains in the low-order product register. AU1 contains the address 00020 (in octal).

Since hunting for the next sequencing counter in demand is not allowed at the conclusion of a multiply instruction, the programmer has the opportunity to store the contents of the low-order product register before an instruction from another program destroys them. As explained under the discussion of inactive addressing (Section IV,

page 46ff), this may be accomplished by performing a "transfer A to B, go to C" instruction with an inactive A address. If both halves of the product are stored, they will have the same sign. It should be noted that the high-order product stored in C is unrounded.

Several properties of the multiply instructions deserve particular attention. The contents of the mask register are destroyed, together with the contents of those locations used to store the partial products. Furthermore, if the C address of a multiply instruction is a direct or indexed special register address, a control error will occur and the machine will stop.

Masking is not permitted with the multiply instructions. Except for the restriction already stated with respect to the C address, the A, B, and C addresses may be direct, indexed or indirect. The time required to execute a binary multiply with direct memory location addresses is 33 memory cycles.

Decimal Multiply, DM

Decimal multiply is implemented exactly as the binary multiply instruction, with the exception that only 10 partial products are generated instead of 16. After the multiples of the A operand are generated, the proper multiples are selected and added in accordance with a digit-by-digit inspection of the multiplier, or B operand. If hexadecimal digits appear in the operands, an erroneous product will be generated, and a control error may be indicated.

Like binary multiply, the decimal multiply instruction produces a 2-word result. The high-order product, consisting of 11 decimal digits, is stored in the location specified by the C address group with a sign determined by the value of the sign flip-flop. The low-order result appears in the low-order product register and may be stored by the programmer in the manner described for binary multiply. At the conclusion of the instruction, AU1 contains the address 00012 (in octal).

Since only 10 multiples of the multiplicand are stored in a decimal multiply, compared with 16 multiples for a binary multiply, only memory locations 00000 through 00011 will be affected, and the time required to execute the decimal instruction with direct memory location addresses is six memory cycles less than for the binary instruction, or 27 memory cycles.

SECTION VII

LOGICAL INSTRUCTIONS

The four instructions which make up the logical group manipulate words on an individual-bit basis, combining bits from two words to form a third. The rules by which the bits are combined are similar to the rules under which the logical elements of a computer operate; hence the name logical instructions. All operands are regarded as 48-bit words in which each bit is an individual unit of information unrelated to any other bit.

The four logical instructions are extract, substitute, half add, and superimpose. Extract and substitute have command codes of the "inherent mask" format (see Section III) and utilize the entire B address group to specify the location of a mask. Half add and superimpose have command codes of the "general, masked or unmasked" format. The time required to execute an extract or substitute instruction with direct memory location addresses is five memory cycles; for an unmasked half add or superimpose instruction with direct memory location addresses, it is four memory cycles.

The logical instructions may address special registers in one or more address groups. The result of such an operation may be determined by applying the rules governing the transfer of a special register word to a 48-bit register and the transfer of a 48-bit word to a special register.

Extract, EX

The extract instruction places the A operand in the location specified by the C address group, using the B operand as a mask and not protecting the unmasked portions of C. (The mask index register is not used in locating the mask.) This is equivalent to combining the corresponding bits of the A and B operands in accordance with the following rule:

If corresponding bit positions in the word at A and the word at B both contain ones, the result shall contain a one in this position. In all other cases, the result shall contain a zero. This is the "logical AND" function.

SECTION VII. LOGICAL INSTRUCTIONS

The execution of an extract (or logical AND) instruction utilizes the accumulator, the mask register, and the low-order product register, and takes place in the following steps:

1. The 48-bit A operand is placed in the accumulator and the 48-bit B operand is placed in the mask register.
2. The contents of these two registers are then shifted right one bit position. The value of the bit shifted from the mask register is examined by a logical computer element called a "gate." If this bit is a one, the gate allows the value of the bit shifted from the accumulator to be placed in the high-order bit position of the low-order product register. If the bit from the mask register is zero, a generated zero bit is introduced into the corresponding position of the low-order product register. Actually, both the value of the accumulator bit and the generated zero bit are brought to the gate, which has been set according to the value of the mask register bit to allow one or the other of these two values to reach the low-order product register.
3. Step 2 is repeated 47 times, shifting the contents of the low-order product register as well. At the conclusion of 48 one-bit shifts, each of the 48 pairs of operand bits has been examined and each bit position in the low-order product register contains one or the other of the corresponding pair of bit values.
4. The 48-bit result generated in the low-order product register is then placed in the location specified by the C address group.

Whenever a masking operation is performed, whether in a logical, a general, or a shift instruction, a logical gate is set by the value of each bit from the mask register to open one transmission path and close another. In unprotected masking, one of the paths transmits the word to be masked while the other transmits generated zero bits. In protected masking, the generated zeros are replaced by the contents of the location specified by the C address group.

For example, the following locations contain the words shown:

OPERAND	110100100010-----
EXMASK	000011110000-----
RESULT	101001110110-----

when the instruction

EX OPERAND EXMASK RESULT

SECTION VII. LOGICAL INSTRUCTIONS

is executed. As a result of the instruction, the contents of location RESULT will be:

RESULT 000000100000-----

As discussed in Section IV (see page 46ff), the use of inactive addressing with the extract instruction provides access to the 48-bit arithmetic register called the mask register. When a mask is specified in an instruction, either in the command code or in the B address group, the contents of the specified location are placed in the mask register, where they remain until a subsequent instruction calls for a mask (or until they are destroyed by the execution of a multiply instruction). Therefore, the contents of the mask register are unrelated to the instruction being performed if this instruction does not specify a mask. Since the mask register has no address, the programmer cannot transfer its contents directly to memory. However, by executing an extract instruction which specifies the address of a word of 48 binary ones in the A address group and an inactive address in B, the programmer can store in the location specified by C a word guaranteed to be identical to the contents of the mask register. This guarantee can be verified by an inspection of the above rule of operation. If the programmer wishes to insert a full word into the mask register without disturbing any memory location, he may perform an extract instruction with an inactive C address.

Substitute, SS

The substitute instruction performs the same general function as extract except that the contents of the location specified by the C address group are protected. When this instruction is executed, therefore, the computer places the 48-bit contents of A in the accumulator and the 48-bit contents of B in the mask register and then forms a new word in the low-order product register according to the following rule:

Wherever the B operand contains a one bit, the corresponding bit of the A operand is stored in the corresponding position of the low-order product register. Wherever the B operand contains a zero bit, the corresponding bit of the word at C is stored in the corresponding position of the low-order product register.

Finally, the word formed in the low-order product register is stored in the location specified by the C address group. The result of the operation may differ from the result of an extract having the same operands only in those bit positions for which the B operand has a value of zero. The behavior of this instruction with one or more inactive addresses is unspecified.

For example, if the instruction

```
SS  OPERAND  EXMASK  RESULT
```

is executed where the contents of the three locations specified are as shown under the discussion of extract, the contents of location RESULT will be:

```
RESULT          101000100110-----
```

Half Add, HA

The half add instruction performs a binary addition of the 48-bit A and B operands in the accumulator, discarding all carries, and stores the result in the location specified by the C address group. This is equivalent to combining the corresponding bits of the A and B operands in accordance with the following rule:

If the corresponding bit positions in the A and B operands have the same value, the result shall contain a zero in this position. In all other cases, the result shall contain a one. This is the "logical exclusive OR" function.

The behavior of this instruction with one or more inactive addresses is unspecified.

The half add instruction is illustrated in terms of the same example used to explain extract. If the instruction

```
HA  OPERAND  EXMASK  RESULT
```

is executed and the operands are the same as in the preceding examples, the contents of location RESULT will be:

```
RESULT          110111010010
```

Superimpose, SM

The superimpose instruction performs a superposition of the 48-bit A and B operands in such a way that the result contains a one in every position in which a one existed in either A or B or both, and stores the result in the location specified by the C address group. This is equivalent to combining the corresponding bits of the A and B operands in accordance with the following rule:

SECTION VII. LOGICAL INSTRUCTIONS

If corresponding bit positions of the A and B operands are both zero, the result shall contain a zero in this position. In all other cases, the result shall contain a one. This is the "logical inclusive OR" function.

The behavior of this instruction with one or more inactive addresses is unspecified.

For example, if the above operands are manipulated by the instruction

SM	OPERAND	EXMASK	RESULT
----	---------	--------	--------

the contents of location RESULT will be:

RESULT	110111110010-----
--------	-------------------

SECTION VIII

TRANSFER INSTRUCTIONS

The logic of the Honeywell 800 includes six instructions designed to transfer data within the main memory, within the control memory, or from one memory to another. Using these instructions, called transfer instructions, the programmer may move any desired quantity of information, from a single bit to an entire record. Two of the instructions move single words only, or, when masked, fields within single words. The other four instructions move groups of words and cannot be masked.

The command codes for all six instructions provide the ability to address special registers and to designate either sequencing counter as the source of the next instruction. In instructions which transfer groups of words, the A address group indicates the location of the first word to be transferred, while the C address designates the location to which this word is delivered. Except for the multiple transfer (MT) instruction, transfers involving groups of words occur under control of special registers AU1 and AU2, which initially contain one of the following bit configurations, depending upon the type of addressing used in the A and C address groups, respectively.

1. Direct Memory Location Address: A positive sign bit, the bank indicator taken from the sequencing counter which selected the instruction, and the low-order 14 bits of the address group.
2. Direct Special Register Address: A positive sign bit, the group indicator associated with the sequencing counter which selected the instruction, and the low-order 14 bits of the address group.
3. Indexed Memory Location Address: A positive sign bit and the augmented low-order 15 bits of the index register referenced in the address group.
4. Indexed Special Register Address: A positive sign bit and the augmented low-order 15 bits of the index register referenced in the address group.
5. Indirect Memory Location Address: The sign bit (positive or negative) and 15-bit main memory address stored in the special register designated in the address group.
6. Indexed Indirect Memory Location Address: The sign bit (positive or negative) and 15-bit main memory address stored in the special register whose address is obtained by augmenting the

contents of the index register referenced in the address group.

As successive words are transferred, the AU counters are automatically incremented (or decremented if the sign bit is negative) by one to specify a source and a result address for each word transferred. At the completion of the instruction, the counters contain addresses equal to their initial settings plus (or minus) the number of words transferred, i. e., the address of the last word transferred plus (or minus) one.

Transfer A to C, TX

This instruction transfers one word from the location specified by the A address group to the location specified by the C address group. The B address group is ignored. Any of the six types of addressing previously discussed may be used in either the A or C address group, provided the instruction is used in its unmasked version. When the instruction is used with a mask, partial words (or fields) may be transferred from one memory location to another, protecting the unmasked portion of the result location. The masked version of the instruction, of course, cannot address special registers.

The time required to execute an unmasked instruction with direct memory location addresses is three memory cycles. The timing effect of masking, indexing, and the use of special registers is summarized in Appendix C.

Transfer A to B and Go to C, TS

This instruction transfers one word from the location specified by the A address group to the location designated by the B address group and changes the setting of the specified sequencing counter so that the next instruction is selected from the main memory location designated by the C address group. It is also used, with inactive addressing, to provide access to the low-order product register (see Inactive Addressing, Section IV). The instruction is identical to the TX instruction with respect to masking.

The A and B addresses may use any of the six types of addressing previously discussed. If the C address is active, it may specify a direct memory location address, an indirect memory location address, an indexed memory location address, or an indexed indirect memory location address. The behavior of the system for each of these cases is described below.

1. Direct Memory Location Address. The low-order 11 bits of the C address group are placed in the specified sequencing counter, protecting the bank indicator in that counter.
2. Indirect Memory Location Address. The complete contents of the addressed special register, including the sign and bank indicator, are transferred to the specified sequencing counter. The contents of the addressed special register are incremented in the usual fashion after use, unless the special register addressed is the counter to be changed (either sequence or cosequence counter), in which case incrementing does not take place. The tabular bit in the address group is ignored.
3. Indexed Memory Location Address. The contents of the referenced index register are augmented, and the complete 15-bit address thus formed is inserted, with a positive sign, into the specified sequencing counter.
4. Indexed Indirect Memory Location Address. A complete special register address is generated by augmenting the contents of the referenced index register in the usual fashion. The generated address specifies a special register whose entire 16-bit contents are transferred to the specified sequencing counter. The contents of the addressed special register are incremented in the usual fashion after use, unless this special register is the counter to be changed (either sequence or cosequence counter). The tabular bit position in the address group is ignored.

The time required to execute an unmasked TS instruction with three direct memory location addresses is four memory cycles.

N-Word Transfer, TN

This instruction transfers the number of words specified by the high-order six bits of the B address group from consecutive locations starting at A to consecutive locations starting at C. The number of words to be transferred can range from 0 to 63. If the B address group is zero, no information is transferred. The low-order six bits of B are ignored. The transfer of information occurs under control of special registers AU1 and AU2, according to the conventions outlined at the beginning of this section.

It should be noted that if a special register is directly addressed in the A address of an N-word transfer instruction and an increment other than zero appears in the address group, then this increment is applied after transfer to the contents of each special register thus addressed. For example, consider the instruction

SECTION VIII. TRANSFER INSTRUCTIONS

TN Z, X0, 10 5 Z, R0

This instruction causes the contents of index registers X0 through X4 to be transferred to special registers R0 through R4. At the conclusion of the instruction, the contents of index registers X0 through X4 will have been incremented by 10, and the low-order five bits of AU1 and AU2 will contain the subaddresses of X5 and R5, respectively.

When the instruction

TN N, X0, 10 5 N, R0

is executed, on the other hand, only the contents of X0 are incremented by 10 after use, since no other special register is referenced by the A address group. If X0 initially contained the main memory address tagged TRANS1 and R0 contained the main memory address tagged LOC1, then at the conclusion of the instruction AU1 will contain the address TRANS1 + 5 and AU2 will contain the address LOC1 + 5.

The time required to execute an N-word transfer with either direct or indirect memory location addresses is $5 + 2n$ memory cycles, where n equals the number of words transferred.

Multiple Transfer, MT

The multiple transfer instruction transfers the contents of the location specified by the A address group to the location specified by the C address group, repeating the transfer the number of times specified by the high-order six bits of the B address group. This number may range from 0 to 63. If B is zero, no transfer of information takes place. The low-order six bits of B are ignored.

Although all types of addressing are permitted with this instruction, the instruction is most meaningful when used with indirect addressing. For example, an area of 20 words in memory may be cleared to zeros by storing a constant of zeros in the location tagged ALLZEROS, setting general purpose register R0 to the address of the first location to be cleared, and executing the instruction

MT ALLZEROS 20 N, R0, 1

Zeros will be transferred to the 20 main memory locations starting with the address initially contained in R0, and the contents of R0 at the completion of the instruction will be equal to the initial contents plus 20. The AU counters are not involved in the execution of this instruction.

As a second example, consider the instruction

```
MT   N, X0, 10   5   N, R0, 1
```

If X0 initially contains a memory address tagged ONESTORE and R0 contains a main memory address tagged WORKAREA, execution of this instruction causes the words from locations ONESTORE, ONESTORE + 10, ONESTORE + 20, etc. to be transferred to WORKAREA, WORKAREA + 1, etc. At the conclusion of the instruction, X0 contains the address ONESTORE + 50, while R0 contains the address WORKAREA + 5.

The time required to execute a multiple transfer instruction with either direct or indirect memory location addresses is $1 + 2n$ memory cycles, where n equals the number of times the transfer is performed.

Record Transfer, RT

The record transfer instruction is used to move a group of related words from one location to another. Such a group does not necessarily constitute a record. Although the record transfer instruction is most frequently used to manipulate "records," it may also be used to advantage whenever the number of words to be transferred is greater than 63, the maximum number which can be moved with the N-word transfer. In fact, the only restriction on the number of words which can be transferred is the practical limitation of available storage space in the memory.

When a record transfer instruction is executed, an end-of-record word¹ is stored in the location specified by the B address group. Consecutive words are then transferred from the location starting with A to consecutive locations starting with C, until an end-of-record word is transferred. The behavior of the system is unspecified when the A address is a direct or indexed special register address.

1. As previously stated, an end-of-record word is a word whose 48 information bits are:
1010 1010 0000 0000 1110 1110 1110 1110 1101 1101 1101 1101

SECTION VIII. TRANSFER INSTRUCTIONS

Note that the transfer of information is stopped whenever an end-of-record word is transferred, regardless of where this word was stored. In other words, the word stopping the transfer is not necessarily the same word which the instruction has stored in the location specified by the B address. At the completion of the instruction AU1 will contain the complete address generated from the A address group plus the number of words actually transferred; AU2 will contain the complete address generated from the C address group plus the number of words transferred.

The time required to execute a record transfer instruction with direct or indirect memory location addresses is $7 + 2n$ memory cycles, where n is the number of words transferred.

Item Transfer, IT

The item transfer instruction operates exactly as the record transfer, with the exception that instead of storing an end-of-record word, an end-of-item symbol is substituted for the high-order 32 bits of the contents of B, protecting the low-order 16 bits. The transfer stops with the transfer of an end-of-item or an end-of-record word.

As described in Section III, an end-of-item word is any word whose high-order 32 bits are identical to those of an end-of-record word. The comments made with respect to the record transfer instruction are equally applicable to item transfer. The time required to execute the instruction is also the same.

SECTION IX

DECISION INSTRUCTIONS

The Honeywell 800 logic includes four decision instructions which are used to branch to an alternate path of control in response to a precisely defined condition. Each of these instructions looks for a special condition based upon the relationship between the contents of the location specified by the A address group and the contents of the location specified by the B address group. If the condition is met, the sequencing counter designated as the source of the next instruction is changed so that the next instruction is selected from the main memory location specified by the C address group. If the condition fails, the current setting of the designated sequencing counter selects the next instruction.

Any type of addressing may be used in the A and B address groups of these instructions unless they are masked. The C address group may contain a direct memory location address, an indirect memory location address, an indexed memory location address, or an indexed indirect memory location address. If the condition is met and the sequencing counter is changed, the behavior of the system for each of the four possible types of C address is the same as that described for the instruction "transfer A to B, go to C" (see Section VIII, pages 78 and 79). If the C address is inactive, no change is made in either sequencing counter and hunting for the next program in demand is inhibited. If the C address is active and the condition is met, hunting is also inhibited. If the condition is not met, a normal hunt is made for the next program in demand.

The accumulator is used in executing the decision instructions. Following the completion of such an instruction, the word in the accumulator is invalid. Any attempt to deliver this word to the memory will therefore cause a control error.

Two of the four decision instructions are inequality comparisons (one alphabetic, one numeric) in which the two operands are examined for inequality (\neq). If they are not equal, the next instruction is selected from the location specified by the C address

SECTION IX. DECISION INSTRUCTIONS

group. The other two instructions (one alphabetic, one numeric) are called "less than" comparisons, since the two operands are inspected to determine whether the contents of A are less than or equal to the contents of B (\leq). If so, the next instruction is selected from the location specified by C. All four instructions may compare either the entire contents of both words (unmasked version) or respective fields of the words, as defined by a mask. The location of the mask, if any, is partially specified by bits of the command code.

The alphabetic comparisons make a bit-by-bit comparison of all 48 information bits of the words at locations A and B. If the word at A has a zero bit in the first position where they differ, reading from left to right, it is considered "alphabetically less" than the word at B. The order of the alphanumeric characters is shown in Table I, page 143. When the instructions are masked, the decision is based upon a bit-by-bit comparison of the selected fields of the operands.

The numeric comparisons treat both operands as signed 11-digit numbers and make a bit-by-bit comparison of information bits 5 through 48 of the contents of A and B. Thus, the 11 digits of the signed numeric words are compared as in the alphabetic comparisons. However, the sign bits (bits 1 through 4) of both words are checked before a final decision is made. If any of the four bits is a one, the sign is considered to be positive; if all four bits are zeros, the sign is negative. Any positive number is larger than any negative number, with the exception that two words with 11 low-order zero digits are considered equal regardless of sign. When two negative numbers are compared, the one with absolute value closer to zero is considered greater. Thus, the numeric comparisons are true algebraic comparisons of signed 11-digit numbers.

The time required to execute an unmasked decision instruction with direct memory location addresses is four memory cycles, regardless of whether or not the condition is met. Reference should be made to the notes in Appendix C for the effect on timing of masking, indexing, and special register addressing.

Inequality Comparison, Alphabetic, NA

This instruction compares for inequality the 48-bit contents of the word at A and the 48-bit contents of the word at B. If the two words are not identical, the sequencing counter specified as the source of the next instruction is changed so that the next in-

instruction is selected from the location designated by the C address group. If they are equal, the next instruction is selected from the location specified by the current contents of the designated sequencing counter. Plus zero is not equal to minus zero.

Less Than or Equal Comparison, Alphabetic, LA

This instruction also makes a bit-by-bit comparison of the 48-bit word at A and the 48-bit word at B. If the contents of the location specified by the A address are less than or equal to the contents of the location specified by B, the designated sequencing counter is changed so that the next instruction is selected from the location specified by the C address group. If the word at address A is greater than the word at address B, the next instruction is taken from the location specified by the current contents of the designated sequencing counter. Plus zero is greater than minus zero.

Inequality Comparison, Numeric, NN

The numeric inequality comparison examines the contents of the locations specified by A and B, regarded as signed words, to determine whether they are algebraically unequal. If they are unequal, the designated sequencing counter is changed to select the next instruction from the location specified by the C address group. If the words are equal, the address of the next instruction is selected from the location specified by the current contents of the designated sequencing counter. In a numeric inequality comparison, plus zero equals minus zero.

Less Than or Equal Comparison, Numeric, LN

When this instruction is executed, the contents of the locations specified by the A and B address groups are regarded as signed words and compared to determine whether the word at A is algebraically less than or equal to the word at B. If the A operand is less than or equal to the B operand, the designated sequencing counter is changed to select the next instruction from the location specified by the C address group. If the contents of A are algebraically greater than the contents of B, the next instruction is selected from the location specified by the current contents of the designated sequencing counter. Plus zero equals minus zero in a numeric less than or equal comparison.

SECTION X

SHIFT INSTRUCTIONS

The ability to pack as many as four separate pieces of signed information (or a greater number of pieces of unsigned information) into a single Honeywell 800 word would have little practical value without the facility for manipulating each piece of information (or field) independently. In part, this facility is provided by the masking feature of the system, which permits certain instructions to operate on fields in pre-assigned positions within words. In order to achieve maximum flexibility in handling packed words, however, it is also necessary to provide the means for rearranging the positions of fields within words. This flexibility has been achieved in the Honeywell 800 by the inclusion of five shift instructions. Two of these instructions shift the low-order 44 bits of a word, preserving sign bit positions 1 through 4. The remaining three shift the entire 48 bits of the word.

The shift instructions cause the contents of the location specified by the A address group to be placed in the accumulator, where they are shifted right the number of bit positions specified by the high-order six bits of the B address group. The direction of the shift is always right end around, with the result that bits shifted out of the low-order position recirculate to the high-order position of the word. A 48-bit full word shift, therefore, produces the same result as a 0-bit shift. The effective number of bit positions by which a word may be shifted ranges from 0 to 44 for shifts protecting the sign of the operand, and from 0 to 48 for shifts of the entire word. If the number of bit positions specified to be shifted exceeds the maximum effective shift (it is possible to specify a 63-bit shift in the high-order six bits of B), the machine will actually perform the specified shift. The net number of places shifted as a result of the operation, however, will be the number specified minus 44 in the case of shifts protecting sign or the number specified minus 48 in the case of shifts not protecting sign.

The shift instructions are inherent mask instructions; that is, they are always performed with a mask. The mask is applied to the shifted operand before delivery to the result location (specified by the C address group). If the programmer wishes to

deliver the entire word, therefore, he must specify a mask consisting of 48 binary ones. The location of the mask is designated by the low-order six bits of the B address group, in conjunction with bits 2 through 5 and 6 through 10 of the mask index register (see Figure V-3, page 56). The two shift and substitute instructions and the shift and select instruction perform protected masking. This means that the values of the unmasked bit positions in the result location (those which do not correspond to binary ones in the mask) are preserved during the operation. The two shift and extract instructions perform unprotected masking; that is, the positions in the result location which do not correspond to binary ones in the mask are cleared to zeros. Protected masking is accomplished by gating the shifted contents of the accumulator and the original contents of the location specified by the C address group into the low-order product register under control of the mask in the mask register¹. Unprotected masking is accomplished by gating the contents of the accumulator and a generated word of all zeros into the low-order product register under control of the mask in the mask register. In either case, the contents of the low-order product register are then delivered to the location specified by the C address group.

Any type of addressing may be used in the A and C address groups of the shift instructions, with the exception that certain restrictions apply to the C address of a shift and select instruction. These will be discussed under the description of that instruction. If the B address group of a shift instruction is all zeros, then no shifting is performed and the word at A is transferred to the location specified by C under control of a mask located by attaching six low-order zero bits to the partial address stored in the mask index register. If the B address group is inactive, the behavior of the system is unspecified.

ARGUS notation for the shift instructions is designed to simplify the specification of masks and shifts. If the mask is designated by the programmer, its symbolic tag is written in the command code field (rather than in the B address field), separated from the operation code by a comma. Reference should be made to Section VII of the ARGUS Manual of Assembly Language for details on the generation of masks by ARGUS. The nature and extent of the shift is specified by three items of information in the B address field, all separated by commas:

1. This process is described more fully in Section VII, see page 73.

SECTION X. SHIFT INSTRUCTIONS

1. A character to designate the type of characters shifted, i. e., A, alphanumeric; D, hexadecimal; B or blank, binary.
2. A number to indicate the number of positions to be shifted (0 to 8, alphanumeric; 0 to 12 hexadecimal; 0 to 48 binary).
3. A character to designate the direction of the shift as left, L, or right, R or blank.

Any valid ARGUS address format may be used in the A or C address of a shift instruction, subject only to the restrictions noted in connection with the C address of a shift and select instruction.

Shift and substitute and shift and extract instructions with direct and/or indirect memory location A and C addresses require a basic execution time of five memory cycles plus a variable number of memory cycles required for the actual shifting operation. The shift logic in the system is so designed that the machine shifts a word in multiples of 1, 4, or 16 bits. Each of these shifts requires the same amount of time; i. e., a 16-bit shift takes no longer than a 1-bit shift. Two such shifts may be performed in a memory cycle. Thus, a 32-bit shift requires a single memory cycle, as does a 20-bit shift or a 5-bit shift. The total time required to shift the operand the specified number of bits, therefore, is a function of the number of individual shifts which the machine must perform. On this basis, the number of memory cycles required for actual shifting of an operand ranges from 0 to 4, so that the time required for complete execution of the instruction varies from a minimum of five to a maximum of nine memory cycles.

The shift and select instruction requires a basic execution time of six memory cycles plus the number of cycles required for the shift itself. The total time to execute the instruction, therefore, varies from six to ten memory cycles.

The exact times required for each number of positions shifted, as well as variations resulting from indexing or direct special register addressing, are listed in Appendix C.

Shift Preserving Sign and Substitute, SPS

This instruction directs the machine to shift the word at A, excluding sign bits

1 through 4, right end around, the number of bit positions specified by the high-order six bits of the B address group. Bits shifted out of position 48 recirculate to bit position 5. The result is masked using a mask whose address is generated from the low-order six bits of B and the partial address stored in the mask index register, as illustrated in Figure V-3. The masked result is delivered to the location specified by C, protecting the bit positions of the contents of C which do not correspond to binary ones in the mask.

A shift of n bits to the left is specified in machine language by setting the high-order six bits of B equal to $44-n$ (in binary). Thus, a 12-bit shift to the left is effected by specifying a 32-bit shift (44 minus 12) to the right. In ARGUS language, as previously noted, the programmer may specify a left shift directly.

Shift Preserving Sign and Extract, SPE

This instruction is identical to shift preserving sign and substitute, with the exception that the bit positions in the location specified by C which do not correspond to binary ones in the mask are not protected but are cleared to zeros when the result is stored.

Shift Word and Substitute, SWS

This instruction directs the machine to shift the entire word at A, including sign, right end around, the number of bit positions specified by the high-order six bits of B. The result is masked using a mask whose address is generated as in shift preserving sign and substitute. The masked result is delivered to the location specified by C, protecting the bit positions of the contents of C which do not correspond to binary ones in the mask.

A shift of n bits to the left is specified in machine language by setting the high-order six bits of B equal to $48-n$ (in binary). Again, a left shift may be specified directly in ARGUS language.

Shift Word and Extract, SWE

This instruction is identical to shift word and substitute, with the exception that the bit positions in the location specified by C which do not correspond to binary ones

SECTION X. SHIFT INSTRUCTIONS

in the mask are not protected but are cleared to zeros when the result is stored.

Shift Word and Select, SSL

The shift word and select instruction causes the machine to shift the entire word at A, right end around, the number of bits specified by the high-order six bits of B. The result of the shift is masked, using a mask whose address is generated as in shift preserving sign and substitute, and this masked result, with a positive sign implied, is added in binary to the address specified in C to form a new address. The sequencing counter designated as the source of the next instruction is then changed so that the next instruction is selected from the location specified by the modified C address group. Regardless of the mask used, the machine never adds more than 11 low-order bits to the C address.

The C address of a shift and select instruction may be a direct memory location address, an indirect memory location address, an indexed memory location address, or an indexed indirect memory location address. For these four types of address, the masked result, called C', is used with the C address as described below:

1. Direct Memory Location Address. C' is added to the low-order 11 bits of C, discarding the end carry, if any. The resulting 11 bits are placed in the specified sequencing counter, ~~protecting with~~ the bank indicator ~~in that counter.~~ *taken from the counter the previous instruction.*
2. Indirect Memory Location Address. C' is added to the low-order 11 bits of C, discarding the end carry, if any. The result is used to select a special register whose contents, including the sign and bank indicator, are transferred to the specified sequencing counter. The tab bit of the modified C address is ignored. The contents of the addressed special register are incremented in the usual fashion after use, unless the special register addressed is the counter to be changed (either sequence or cosequence counter), in which case incrementing does not take place.
3. Indexed Memory Location Address. The contents of the referenced index register are augmented in the usual fashion to form a complete 15-bit address. C' is then added to the low-order 11 bits of this complete address, discarding the end carry, if any. The resulting complete 15-bit address is transferred, with a positive sign, to the specified sequencing counter.
4. Indexed Indirect Memory Location Address. A complete special register address is generated by augmenting the contents of the referenced index register in the usual fashion. C' is added to the low-order 11 bits of this generated address, discarding the end carry, if any. The result of this addition is used to select a special register whose

contents, including the sign and bank indicator, are transferred to the specified sequencing counter. The tab bit in the generated special register address is ignored. The contents of the special register selected are incremented in the usual fashion after use, unless this special register is the counter to be changed (either sequence or cosequence counter).

Since the shift and select instruction permits the addition of as many as 11 bits to the low-order bits of the C address, it provides the ability to transfer control to any one of 2048 memory locations.

SECTION XI

PERIPHERAL INSTRUCTIONS

One of the unusual features of the Honeywell 800 is that all types of peripheral devices -- magnetic tape units, card readers, punches, and printers -- look the same to the central processor. Thus, the same 6-bit operation code is used to read from either a tape or a card reader, and another 6-bit operation code defines the instruction to write on a tape, to print, or to punch. These two instructions are called read forward (RF) and write forward (WF), respectively. The two remaining peripheral instructions are the read backward (RB) and rewind (RW) instructions, which are meaningful with tape units only.

In machine language, a peripheral operation code is specified by the low-order six bits of the command code, while the high-order six bits designate one of 64 possible devices to be addressed. More exactly, the high-order three bits designate a channel through which the information will travel (one of eight input or eight output channels, the direction determined by the operation code), while the following three bits specify a particular device attached to this channel from which information will be read or upon which it will be written. The assignment of peripheral codes to magnetic tape units and other input-output devices is established individually at each Honeywell 800 installation, subject to the restrictions discussed under System Configurations in Section II. From the table of peripheral assignments, the programmer must select a device which can perform the specified operation. For example, a write instruction cannot be executed by a card reader. Neither can a printer read backward nor a punch execute a rewind.

The peripheral commands, in ARGUS language, are reversed in terms of machine language. In other words, the operation code is written first, followed by the device address. This address is expressed as an alphabetic code from AA to HH. Thus the instruction

WF, CD OUTPUT - -

instructs the machine to write one record, stored in memory beginning at location OUTPUT, on device CD, which may be a magnetic tape unit, a printer, or a card punch, depending upon the assignment of the code CD at the particular installation.

Since the entire 12 bits of the command code are used to specify the operation code and the device address, a peripheral instruction cannot address a special register explicitly or specify an alternate sequencing counter as the source of the next instruction. A peripheral instruction, therefore, is always followed by an instruction from the same sequencing counter which selected the peripheral instruction. The A address group of a read or write instruction specifies the memory location into which the first word of a record is to be read or from which the first word of a record is to be written. The B address group is used only when the instruction is to sense for end-of-item words in a tape operation and read the separate items into or write them from non-sequential areas in memory. In this case, called distributed reading or writing, the B address group specifies the memory location of the first entry in a table of addresses which denotes the starting locations for the second item and all subsequent items to be read or written. Finally, the C address group may be used, at the programmer's option, to specify a new setting for the sequencing counter which selected the current instruction. All three address groups may specify either direct or indexed main memory addresses. If the C address is active, the behavior of the system follows that described for direct memory location or indexed memory location addressing with the C address of the instruction "transfer A to B and go to C" (see Section VIII, page 79).

If address A of a read or write instruction is inactive, the addressed device is tested for interlocks and errors (see below), but the device is not activated nor does any transfer of information take place. If address B is inactive, end-of-item words are not sensed and the record is read into or written from successive memory locations. If the C address is inactive, the contents of the sequencing counter are not changed, and no hunt is made for another active special register group. Since a hunt for another program demand is never made after an instruction which implicitly changes the contents of the sequencing counter, it follows that under no condition does a hunt occur after execution of a peripheral read or write.

SECTION XI. PERIPHERAL INSTRUCTIONS

The time required by the central processor to handle a read or write instruction to any device is determined in the following fashion. Three memory cycles are required to interpret a peripheral instruction whose A and B addresses are not indexed and whose C address is inactive. If the C address is active but not indexed, two additional memory cycles are required. Thus, interpretation of a peripheral instruction with direct memory location A, B, and C addresses requires five memory cycles. Complete details of the interpretation times for peripheral instructions are shown in Appendix C. In addition to the time required to interpret the instruction, one memory cycle is used for each word transferred from device to storage or from storage to device and an additional cycle is required for each item handled in distributed item reading or writing. The memory cycles used for information transfer are not consecutive but are allotted one at a time by traffic control, as described in Section II.

Read Forward, RF

When a read forward instruction is interpreted by the central processor, the main memory address generated from the A address group is sent, with a positive sign, to the read address counter associated with the input channel to which the addressed device is attached. The unit itself is signalled in order to initiate the mechanical operations, and then the buffer in the tape control unit or the peripheral control unit starts to receive information. From the time the unit is signalled until all the information in the record or card has been transferred to memory, the corresponding input buffer interlock bit in the program control register is set to one (see Section XII, page 108). When a word in the buffer is ready to be transmitted to memory, the appropriate traffic control demand is turned on to indicate that a word may be brought in through this channel. As successive words are transmitted to memory, the read address counter is stepped after each transfer so that it always specifies the location of the next word.

Once a read instruction is started, nothing will stop it except sensing a record gap on magnetic tape or the end of a card's worth of information from the card reader. There are several conditions, however, which will inhibit the instruction so that it is not executed. If for any reason the device is not available to the machine, including the condition in which either the device or the buffer is busy, or if an error occurred on the preceding read from that device, no part of the read instruction is executed. In other words, neither tape nor card is moved, no information is transferred, and if

the instruction calls for a change of sequence, this change is not made.

The first of these conditions requires no further elaboration. Sensing an error during reading will not prevent completion of the instruction (except during a distributed read, as explained later). Instead, when an error is sensed in information coming from either a card or a tape record, a parity error flip-flop is set and the read continues normally until all the words of the record have been stored in memory. When the next read instruction directed to this device tests the setting of the flip-flop and finds that an error was sensed during the previous read, the flip-flop is reset and execution of the instruction is inhibited. The instruction is stored in U or $U + 1$, and an unprogrammed transfer is made to $U + 6$ or $U + 7$, depending upon whether the instruction was selected from the sequence or cosequence counter (see Figure V-5, page 61). The instructions stored in $U + 6$ and $U + 7$ serve as entries to a subroutine designed by the programmer to handle the error condition. Since the location to which the unprogrammed transfer is made reflects the sequencing counter which selected the instruction and since the instruction itself is stored (in U or $U + 1$), the programmer can determine where he is in his program and on which device the error was detected. From this information he may determine how to proceed, in accordance with the various techniques discussed under the instructions compute orthocount and check parity in Section XII and in Appendix B.

Another type of unprogrammed transfer occurs as the result of reading the first end-of-tape record. The end of tape is physically marked by two optical windows, 32 inches apart. When the first window is sensed during tape recording, the record being written is completed. The beginning of the next record, called the first end-of-tape record, is written only after the second window has been sensed. When the tape is read forward, an unprogrammed transfer to $U + 4$ or $U + 5$ occurs on the instruction which reads the first end-of-tape record. The instruction itself is correctly executed and stored in U or $U + 1$. Every read forward instruction occurring beyond the first end-of-tape record also creates an unprogrammed transfer to $U + 4$ or $U + 5$. Each such instruction is also executed correctly and stored in U or $U + 1$.

If any record in the neighborhood of the first end-of-tape record contains more than 2048 words or fewer than seven words, an unprogrammed transfer may occur

SECTION XI. PERIPHERAL INSTRUCTIONS

one record earlier or one record later than the first end-of-tape record. The number of end-of-tape records which may be written, and hence read, after the first end-of-tape record is limited only by the length of the oxide trailer (see description of write forward instruction). If an error is sensed while reading any of these records, the next read instruction is inhibited and an error unprogrammed transfer occurs instead of a transfer to $U + 4$ or $U + 5$.

The preceding discussion is based on a normal read instruction in which the B address is inactive and a card or a tape record is read into successive locations in memory. When tape is being read, an active B address specifies a distributed read, in which case individual items of a record are read into non-sequential areas of memory. The central processor performs an additional operation in setting up such an instruction. As usual, the appropriate read address counter (RAC) is set to the main memory address generated from the A address group. The main memory address generated from the B address group is then sent, with a positive sign, to the corresponding distributed read address counter (DRAC). As the information is transmitted to memory, each successive word is placed in the next higher memory location. During this transmission, the central processor senses the input for an end-of-item word -- a special symbol having the high-order 32 bits identical to the high-order 32 bits of an end-of-record word. When such a word is sensed, it is placed in proper sequence relative to the words which preceded it. The contents of the RAC are then replaced with the contents of the memory location specified by the DRAC, the contents of the DRAC are incremented by one, and transmission of information proceeds with the first word of the next item placed in the location specified by the new contents of the RAC. It is the programmer's responsibility to see that each of the memory locations whose addresses appear successively in the DRAC is loaded with a constant whose low-order 16 bits represent the desired sign and main memory address.

As an example, a record on tape CD consists of five items to be loaded into discrete sections of the memory, using the following instruction:

```
RF, CD  ITEMA  TABLE  -
```

The locations specified by TABLE, TABLE + 1, TABLE + 2, and TABLE + 3 contain

a series of addresses (tagged ITEMB, ITEM C, ITEM D, and ITEM E), each with a positive sign bit, which designate the starting location in memory for each of the last four items. When the instruction is interpreted, the RAC is set initially to the memory location tagged ITEM A and the DRAC is set to the address tagged TABLE. The first item, including the end-of-item word, is read into the successive locations ITEM A, ITEM A + 1, ITEM A + 2, etc. When the end-of-item word is sensed, the contents of TABLE are placed in the RAC and the DRAC is stepped to TABLE + 1. The next item is then loaded into ITEMB, ITEMB + 1, etc. When the end-of-item word is sensed, the contents of TABLE + 1 are transferred to the RAC and the DRAC is incremented to TABLE + 2. This process continues until the entire record has been read into memory. Thus, the third item is stored in successive locations starting with ITEM C while the last item is stored in successive locations starting with ITEM E.

When an error is detected during a normal read, the read is completed before the parity error flip-flop is set. In a distributed read operation, however, no further information is transmitted to memory after an error is sensed, although the tape is allowed to move to the end of the record. As with the normal read, the next read instruction to the tape results in an unprogrammed transfer. If the programmer then wishes to reconstruct the record using the orthocorrection technique, he must reread it by means of a normal (non-distributed) read.

Whether the programmer uses a normal or a distributed read, he must remember that reading from any peripheral device proceeds concurrently with computation. This means that he must be certain a record has indeed been read into memory before he begins to operate upon the information. Several techniques are available for ascertaining that a read operation has been completed. The most obvious, perhaps, is to read records alternately into one of two buffer areas in memory. Thus, while information is being read into one area, information in the other area may be processed.

A shortcoming of this technique, especially if records are lengthy or if several tapes are being handled (as in a sort routine), may be the amount of memory required. If memory must be conserved, other options are available for checking the completion of a read operation. One of these is to give a read instruction with an inactive A address. This technique, which results in no tape movement and no transfer of information, not only insures completion of the previous read instruction but also guarantees

SECTION XI. PERIPHERAL INSTRUCTIONS

(if no unprogrammed transfer occurs) that the information has been correctly read. Another technique by which the programmer may test the status of a peripheral operation is by sensing the interlock bit in the program control register to see whether the corresponding buffer is busy. This technique is safe, however, only if the pertinent buffer is handling a single device during the current run. A third approach involves the use of the read address counter. If the programmer is working with fixed-length records, he knows what address should be found in the RAC when the entire record has been brought into memory. By addressing the RAC explicitly, he can compare its contents repeatedly with a constant representing its final contents to determine when the instruction has been completed. These techniques are not equally appropriate for all situations. The nature of the problem and the ingenuity of the programmer, therefore, determine which approach will be most satisfactory.

Read Backward, RB

The read backward instruction, as its name implies, causes a tape to be read in the reverse direction. The instruction is very similar to the read forward instruction, hence this discussion is limited to the dissimilar features.

The normal read backward instruction is implemented so that a record read in the reverse direction is stored in memory exactly as if it had been read forward, provided the A address group of the instruction specifies the memory location into which the last word of the record would have been placed by the read forward. This is accomplished as follows. When the instruction is interpreted, the main memory address generated from the A address group is inserted, with a negative sign, into the read address counter. As the counter is stepped with each successive word transfer, the effect is to decrement the counter so that each word is read into the next lower memory location until the record gap is reached.

When a distributed read backward instruction is executed, the main memory address generated from the B address group is placed in the distributed read address counter with a negative sign, so that this counter too is effectively decremented as successive items are handled. Thus, the machine not only stores successive words of an item in reverse order, but also sequences backward through the table set up for the distributed read. Since the contents of the locations listed in the table are transferred

to the RAC, the addresses stored in these locations must contain negative sign bits if the words of each successive item are to be placed in memory in reverse order. It should be noted that the end-of-record word, the first word to be brought into memory by a read backward instruction, does not effect a change in the contents of the RAC despite the fact that its high-order 32 bits have the same configuration as an end-of-item word. Instead, the end-of-record word, the two orthotronic words, and the last item of the record (through the end-of-item word for the penultimate item) are stored as one item.

When an error is sensed during a read backward (normal or distributed), the unprogrammed transfer conditions are identical to those described for the read forward instruction. The unprogrammed transfer convention associated with the beginning-of-tape condition, however, is slightly different from the end-of-tape convention which applies when tapes are read forward. When a tape is read backward, a beginning-of-tape unprogrammed transfer to $U + 4$ or $U + 5$ occurs on the instruction which reads the first record written on the tape, and the instruction is correctly executed. Any subsequent read backward instruction to this tape neither moves the tape nor alters the contents of main memory, and no unprogrammed transfer takes place. The actions corresponding to the B and C address groups, however, are carried out.

When a write instruction is given to a tape in rewound condition, the first record is written immediately following the clear leader. The beginning of the second record is not written until the head has reached an optical mark a fixed distance from the leader. If the first record written contains more than 2048 words, the beginning of tape unprogrammed transfer to $U + 4$ or $U + 5$ may fail to occur when this record is read backward. Furthermore, if the second record contains fewer than seven words, the unprogrammed transfer may occur upon reading the second record backward rather than the first.

Write Forward, WF

When a write forward instruction is interpreted by the central processor, the main memory address generated from the A address group is inserted, with a positive sign, in the write address counter (WAC) associated with the output channel to which the addressed device is attached. The unit is signalled, the output buffer interlock

SECTION XI. PERIPHERAL INSTRUCTIONS

bit is set to one, and the first word is selected from the location specified by the WAC and sent to the buffer area in the tape control unit or the peripheral control unit. The contents of the WAC are incremented, and when the buffer is empty, the traffic control demand for this channel is turned on so that the next word may be transmitted to the buffer.

Once the write instruction gets under way, it is stopped by one of three conditions only. If the instruction is directed to a tape unit, the write operation stops when an end-of-record word is sensed by the central processor. In other words, an instruction to write on tape starts with the word in the location specified by the A address group and continues through successively higher memory locations until an end-of-record word is encountered. If the write instruction addresses a printer or punch, on the other hand, the writing stops when 16 words have been delivered to the printer or when 11 words or 21 words (in transcription mode) have been sent to the punch, regardless of the presence or absence of an end-of-record word.

A continuous initial segment of tape is formed or extended by writing a sequence of records, uninterrupted by read or rewind instructions to the tape being written, and starting from a proper initial condition. A proper initial condition is defined as:

1. Following a rewind instruction, or;
2. Following a read backward instruction over a record which is already part of a continuous initial segment.

Thus, a continuous initial segment always consists of an unbroken sequence of records, beginning with the first record on tape. Information on a continuous initial segment of tape may be freely read forward or backward, without restriction.

Information previously recorded beyond the continuous initial segment of a tape is not necessarily recoverable. In particular, any attempt to read the first record following a continuous initial segment is likely to result in an error. One or more records of information previously recorded in the region beyond the continuous initial segment may have been destroyed, and spurious information may have been recorded. However, once the region of erroneous information adjacent to the continuous initial segment has been passed, information is recoverable up to the end point of the previously existing continuous initial segment.

Part of the checking process which occurs during tape reading and writing involves a longitudinal or channel check predicated on the assumption that two correct orthowords are included in every record read or written. If these words are not included with the record, the machine will in most cases indicate an error when the record is read or written. Before writing a record on tape, therefore, it is customary to compute the two orthowords to be included with the record, using the compute orthocount instruction described in Section XII. This computation is omitted at the risk of an error indication, unless the programmer is certain that the record was read into memory with correct orthocount and has not since been altered in any way. When a record is printed or punched on line, however, the question of orthowords is irrelevant, since no orthocheck is performed.

The same conditions which prevent execution of a read instruction also inhibit the write instruction. If either the device or the buffer is busy, or if an error occurred on the previous write instruction to that device, no part of the write instruction is performed. When an error is detected in information going from memory to tape, printer, or punch, a parity error flip-flop is set and the write is completed. The next write instruction to the same device is stored in U or U + 1 and initiates an unprogrammed transfer to U + 6 or U + 7. Since the instruction itself is stored, the programmer will have no difficulty in distinguishing between a write error and a read error, even though the unprogrammed transfer is made to the same location for both types of error. Note that because of the timing of the punch operation, an additional card will have been punched after the erroneously punched card before this unprogrammed transfer is made. A corrective error routine for punched output, therefore, should insure that two cards are repunched.

A special condition may result in an unprogrammed transfer to U + 6 or U + 7 during an instruction to write on tape. Each reel of tape has provision for manually inserting a ring in the hub; unless this ring is in place, the tape is protected from any attempt to write on it. The tape may also be protected, even when the ring is in place, by setting a switch on the tape drive to the position "Run Protected." If an attempt is made to write on a tape protected in either of these ways, the instruction is completely executed, the tape is moved normally, but no writing occurs. An unprogrammed transfer to U + 6 or U + 7 results on the following write instruction to this tape.

SECTION XI. PERIPHERAL INSTRUCTIONS

The end-of-tape unprogrammed transfer which occurs when writing is the same as when reading forward. The instruction which causes the first end-of-tape record to be written also causes an unprogrammed transfer to $U + 4$ or $U + 5$. A write instruction occurring beyond this record is executed correctly and also results in an unprogrammed transfer of control to $U + 4$ or $U + 5$. The number of records which may be correctly written is limited only by the length of the oxide trailer. If writing is attempted beyond the end of the trailer, the write instructions are executed in the same fashion, but the information is lost. A hardware interlock makes it impossible to pull the tape off the reel. When this interlock is activated, the tape unit presents a busy signal to the central processor so that any attempt to write beyond this point causes the program to stall.

If the B address group of a write forward instruction to tape is active, a distributed write operation is performed. Whereas the distributed read places in noncontiguous areas of the memory items which are part of a single tape record, the distributed write gathers up items scattered in the memory and writes them consecutively on tape as part of the same record. When the instruction is interpreted, the main memory address generated from the A address group is placed, with a positive sign, in the write address counter (WAC) and the main memory address generated from the B address group is inserted, with a positive sign, in the distributed write address counter (DWAC). As each word of an item is written out, the WAC is incremented by one to represent the location from which the next word will be written. When an end-of-item word is sensed, it is written on tape, the contents of the location specified by the DWAC are transferred to the WAC, and the contents of the DWAC are incremented by one. This process continues until an end-of-record word is sensed and the entire record has been written on tape. Since the compute orthocount instruction (see Section XII) can sense for end-of-item words and change control for distributed item handling, it is possible to provide the orthotronic control words required for file protection without disturbing the arrangement of distributed items in memory. If an error is sensed while performing a distributed write, the record is completely written and the next write instruction to this tape results in an unprogrammed transfer to $U + 6$ or $U + 7$.

Since writing on any peripheral device may proceed concurrently with computation, the programmer must take care that he does not begin to obliterate a stored record after

a write instruction until the record has been completely written. The techniques discussed under the read forward instruction for insuring that a read operation has been completed are equally useful, with appropriate modification, in guaranteeing the completion of a write instruction.

Rewind, RW

Tape rewind instructions are executed at three times normal tape speed, or 360 inches per second. The tape starts to accelerate to full rewinding speed immediately upon receiving the command; retraction of the head takes place during the acceleration period. Rewinding continues at high speed until the leader is reached, when the tape is stopped and the head restored to operating position. The device interlock is removed when the head has reached operating position. Including the acceleration-deceleration factor, the time to rewind a full reel of tape is 1-1/2 minutes.

Like the other peripheral instructions, the rewind instruction specifies in the high-order six bits of its command code the device to be rewound. If the specified device is not a tape unit, the instruction is recognized as a proceed instruction (see Section XII). If address A is active, an interlock is set after the tape is rewound. Until this interlock has been removed manually at the tape unit, the central processor cannot activate the device. This feature is particularly useful in file maintenance operations when a new file tape has been written and it is desired to insure that the tape is removed and replaced before another instruction with this tape address is processed. If the A address is inactive, no interlock is set. The contents of the B and C address groups are irrelevant to the execution of the instruction. The programmer may therefore use these address groups for the storage of information. Whether the C address is active or inactive, a hunt for the next sequencing counter in demand always occurs after implementation of the instruction.

Only two conditions inhibit execution of a rewind instruction. If the tape is already rewound at the time the instruction is given, the tape is not moved and the instruction is treated as a proceed. If the instruction is directed to a device upon which an error was detected during the previous read or write, the instruction is not executed, and an unprogrammed transfer to U + 6 or U + 7 occurs. It should be emphasized that the rewind instruction is unique in this ability to detect either a read or a write error,

SECTION XI. PERIPHERAL INSTRUCTIONS

since a peripheral read instruction checks a previous read only and a peripheral write instruction checks a previous write only. However, any instruction to a peripheral device resets the error flip-flop for this device.

The time required to execute a rewind instruction is two memory cycles.

SECTION XII

MISCELLANEOUS INSTRUCTIONS

Print, PRA, PRD, PRO

A Honeywell 800 system may include as many as 47 automatic typewriters. The standard unit, located at the central console, is known as the console typewriter. An optional unit, called the slave typewriter, may also be located near the console. Provision of a slave typewriter allows program printouts to be physically separated from the console input information or permits two programs operating in parallel to produce printout information on separate typewriters.

In addition to the console and slave typewriters, a system may contain up to 45 inquiry station typewriters for direct interrogation of information in the central processor and for direct printing of results. Interrogation information entered into the central processor through the inquiry station keyboard activates a stored inquiry program which performs the search of memory, magnetic files, or other connected units and controls the editing and delivery of the information to the inquiry station typewriter. Standard inquiry station typewriters can be located up to 1000 feet from the central processor.

Programmed communication between the central processor and the various typewriters is made possible by the print instruction, which is defined as follows: print the contents of the location specified by the A address group on the typewriter and in the format specified by the B address, and change the designated sequencing counter so that the next instruction is selected from the location specified by the C address. The command code for the print instruction permits special register addressing and change of sequencing counter. Masking is not possible with this instruction.

The B address group of the print instruction contains special information designating the mode and format of the printout and the typewriter to be used. Memory designator bit B is ignored. The significance of each of the 12 bits in the address group is described below:

SECTION XII. MISCELLANEOUS INSTRUCTIONS

- Bit 1: This bit must always be zero.
- Bit 2: Unspecified.
- Bit 3: This bit, called the carriage-return bit, is zero to specify carriage return after the word has been printed, one to specify no carriage return. However, if bit 4 is zero, the carriage will be returned regardless of the value of bit 3.
- Bit 4: This bit, called the more-to-follow bit, is zero to specify that no more information is to follow from this program, one to specify that more information is to follow. The programmer must set this bit to one whenever he wishes to insure that a message of more than one word will not be interrupted by printouts from other programs. Whenever this bit is one, the designated typewriter is interlocked against all other programs until another word is printed from the same program.
- Bits 5-6: These two bits specify the mode in which the information is to be printed, i. e.,
- 01 octal
 - 11 hexadecimal
 - 10 alphanumeric
 - 00 this combination is not defined.
- Bits 7-12: These bits designate one of 47 typewriters on which the information is to be printed. Consider as octal numbers, these bits must lie in the range 00-56. Address 00 designates the console typewriter, address 01 represents the slave typewriter, and addresses 02-56 are available for inquiry station typewriters. If the typewriter address is greater than 56 or does not specify a typewriter included in the system, the instruction is not defined.

Associated with the automatic typewriters are a group of typewriter buffer locations in main memory bank 0 (see Section IV, page 43). With one exception, the address of any typewriter buffer is generated from the corresponding typewriter address by adding the octal number 21. The exception is the slave typewriter which uses the same memory buffer (location 00021) as the console typewriter. Buffer addresses for the inquiry station typewriters lie in the range 00023-00077. Address 00022 has no special function.

If the A address of a print instruction is inactive, the contents of the accumulator are printed as specified by the B address. In this one case, a control error does not occur if the contents of the accumulator have already been delivered to a result location. If the B address is inactive, none of the typewriters is active. If the C address is inactive, no sequence change is made. If the C address is active, the specified sequence counter is set in the same manner as described for a "transfer A to B, go to

SECTION XII. MISCELLANEOUS INSTRUCTIONS

C" instruction having an active C address (see Section VIII, pages 78 and 79). Hunting for the next program in demand is always inhibited.

When printing under program control, the typewriter prints each word specified in the indicated mode. In the alphabetic mode, if the code for carriage return or tabulate is encountered, the corresponding symbol is printed but the function is not performed. When printing hexadecimal words, a space is inserted after every three characters, while a space appears between every four characters in an octal printout. Spaces appear in the alphabetic print mode only where they occur as part of the data. Under program control, the carriage return function is activated only by reaching the physical end of the carriage or by specifying "carriage return" or "no more to follow" in the B address group of the print instruction.

Whenever the preceding print instruction did not establish the "more to follow" condition, the group indicator associated with the program causing the current print-out appears to the left of the printed data, separated by spaces from the information. The values printed for the group indicator are 0 through 7.

If a parity error is detected in a word being printed on the typewriter, the word is followed by a space and two octal check characters. Each bit of the check characters represents one of the parity bits of the erroneous word, being a one if the corresponding parity bit is in error and a zero otherwise.

Although there is but one machine instruction for the print function, ARGUS recognizes three separate print instructions having the following mnemonic operation codes:

PRA	alphanumeric print
PRD	hexadecimal print
PRO	octal print

The operation code may be followed in the command code field by a comma and the letter M (denoting more information to follow before carriage return) or the letters MR (denoting more information to follow after carriage return). If neither appears, the carriage is returned after printing, and the typewriter is released to print from another program.

SECTION XII. MISCELLANEOUS INSTRUCTIONS

The A and C address fields may contain any valid address format. The B address field, which is used only to specify the active typewriter, contains a C, an S, or a 2-digit number specifying the typewriter which is to print. Either C or 00 indicates the console typewriter; S or 01 indicates the slave.

The time required to select and execute a print instruction with a direct memory location A address and an inactive C address is four memory cycles. One additional cycle is required if C is a direct memory location address. When the contents of A have been transferred to the appropriate buffer register and the sequence change, if any, has been effected, the central processor proceeds with the next instruction from the same program. When a character is printed (approximately once every 100 milliseconds), seven memory cycles are required to transfer the buffer contents to the accumulator, shift the next character into printing position, and return the result to the buffer.

Control Program, MPC

To facilitate parallel processing, the Honeywell 800 central processor contains a non-addressable 48-bit register which is used to store information required for efficient control of a multi-program operation. Access to this register, called the program control register, is provided by the machine instruction control program. A description of the information stored in the program control register is shown below:

- Bits 1-3: Peripheral Fixed-Start Bits. The values of these bits are generated manually by setting a switch on an on-line peripheral control unit. They are written into the program control register during the execution of the first peripheral read or write instruction addressed to the peripheral unit after the above mentioned manual control has been activated. The same instruction creates an unprogrammed transfer of control to $U + 6$ or $U + 7$ (see Figure V-5, page 61). The bits in the program control register retain their values until the beginning of another peripheral instruction, at which time they are cleared to zeros.
- Bits 4-8: Unassigned.
- Bits 9-16: Input Buffer Interlock Bits (one for each input channel). Each such bit is a one if the corresponding input buffer is interlocked and a zero otherwise.
- Bit 17: Unassigned.
- Bits 18-20: Control Group Indicator. These bits store an octal value 0 through 7 denoting the special register group under whose control the current instruction was selected.

- Bits 21-28: Bisequence Bits (one per special register group). Each such bit is a one if the next instruction to be executed under control of the respective special register group is to be selected by the cosequence counter or a zero if the next such instruction is to be selected by the sequence counter.
- Bits 29-36: Program Demand Bits (one per special register group). Each such bit is a one if the corresponding sequence counter or cosequence counter is in demand and a zero otherwise. By "in demand" is meant program running, including the case in which the instruction is stalled to await availability of a peripheral device.
- Bits 37-40: Console Fixed-Start Bits. The values of these four bits are determined by striking a hexadecimal key at the time a console fixed-start command is entered. They retain the same values until another console fixed-start command is entered.
- Bits 41-48: Output Buffer Interlock Bits (one for each output channel). Each such bit is a one if the corresponding output buffer is interlocked and a zero otherwise.

For any group of eight bits which represent the eight special register groups (input buffer interlock bits, bisequence bits, program demand bits, output buffer interlock bits), the highest-order bit corresponds to special register group 0, and the succeeding bits correspond to group 1, 2, 3, etc. For example, bit 29 is the program demand bit for group 0, bit 30 represents group 1 and so forth up to bit 36, which represents group 7. It should be noted particularly that this assignment differs from the order of priority for traffic control. Figure XII-1 summarizes the relationship of each group to priority in traffic control and to the program control register.

The instruction, control program, is defined as follows. Ignore the A address group. Place the contents of the program control register in the location specified by the C address group. Then alter the bits of the program control register specified by bits 5 through 12 of the B address group of this instruction. Bits 1 through 4 of the B address define the manner in which the bits of the program control register are altered. Hunt for the next program in demand if the memory designator for the B address (bit 5 of the command code) is one; otherwise do not hunt. Figure XII-2 shows the manner in which B_{1-4} (the high-order four bits of B) determine the interpretation of B_{5-12} (the low-order eight bits of B).

The command code for the control program instruction allows the programmer to specify whether the next instruction shall be selected by the sequence or cosequence

Special Register Group Indicator (PCR Bits 18-20)	Program Demand Bits PCR Position	Bisequence Bits PCR Position	Output Buffer Interlock Bits PCR Position	Input Buffer Interlock Bits PCR Position	Peripheral Channel Address (Command Code Bits 1-3)	Traffic Control Priority
						Output
						Input
0	29	21	41	9	0	8
1	30	22	42	10	1	9
2	31	23	43	11	2	10
3	32	24	44	12	3	11
4	33	25	45	13	4	12
5	34	26	46	14	5	13
6	35	27	47	15	6	14
7	36	28	48	16	7	15

Figure XII-1. Honeywell 800 Special Register Group Indicator Relationships

SECTION XII. MISCELLANEOUS INSTRUCTIONS

		B ₁₋₄	Interpretation of B ₅₋₁₂ *
These four combinations do not affect the bisequence bits		0000	Turn off the program which initiated this instruction.
		0001	Turn off all program demand bits corresponding to zeros in B ₅₋₁₂ ; do not alter program demand bits corresponding to ones.
		0010	Turn on all program demand bits corresponding to ones in B ₅₋₁₂ ; do not alter program demand bits corresponding to zeros.
		0011	Turn on all program demand bits corresponding to ones in B ₅₋₁₂ ; turn off all program demand bits corresponding to zeros.
These four combinations may affect the bisequence bits.		0100	Do not alter the contents of the program control register.
		0101	Turn on the program demand bits and set to SC the bisequence bits corresponding to zeros in B ₅₋₁₂ ; do not alter the program demand bits or the bisequence bits corresponding to ones in B ₅₋₁₂ .
		0110	Turn on the program demand bits and set to CSC the bisequence bits corresponding to ones in B ₅₋₁₂ ; do not alter the program demand bits or the bisequence bits corresponding to zeros in B ₅₋₁₂ .
		0111	Do not alter any program demand bits; set to SC the bisequence bits corresponding to zeros in B ₅₋₁₂ ; set to CSC the bisequence bits corresponding to ones in B ₅₋₁₂ .
<p>* Bits 5 through 12 of the B address bear a one-to-one correspondence with the bisequence bits (21 through 28) and the program demand bits (29 through 36) of the program control register. In other words, B₅ controls special register group 0, B₆ controls special register group 1, etc.</p>			

Figure XII-2. B Address Group Function in Control Program Instruction

counter. However, if a conflict exists between the value assigned to a bisequence bit by the B address of a control program instruction and the value assigned to the same bit by command code bit 1 of the same instruction, the value corresponding to bit 1 takes precedence. The affected bisequence bit in the PCR is set according to the value of bit 1, but not until the PCR contents have been stored in the memory.

SECTION XII. MISCELLANEOUS INSTRUCTIONS

Since the A address of the instruction is ignored, it may be used to store 12 bits of data if the programmer desires. The C address of the instruction may be a direct, indexed, or indirect memory location address. If the C address is inactive, the contents of the program control register are not stored. From the standpoint of hunting for the next program in demand, it is irrelevant for this one instruction whether the C address is active or inactive; hunting is unconditionally controlled by the B address memory designator bit in the command code. If a control program instruction which does not permit hunting turns off the demand of the program which selected it, this program remains in control until it encounters an instruction which allows hunting.

Six of the eight operations which can be performed by the control program instruction are represented in ARGUS notation by a group of program control instructions, each designated by a unique mnemonic operation code. These six instructions, which perform most of the control operations normally required, are explained in detail in Figure 9, Section VII of the ARGUS Manual of Assembly Language. In order to perform the other two functions, ARGUS also recognizes a control program instruction written in machine language with the mnemonic operation code MPC. In this case, the 12 bits of the B address group are written as three hexadecimal digits.

The time required to execute the control program instruction is always four memory cycles, without exception.

Proceed, PR

This instruction, whose 8-bit operation code consists entirely of zeros, results in no operation other than the normal incrementation of the sequencing counter which selected it. Bits 1 and 4 through 6 of the command code are irrelevant; thus, proceed cannot specify the source of the next instruction and is always followed by an instruction selected by the same sequencing counter. Since the A, B, and C address groups are irrelevant, they may be used to store information. However, if the information stored in the C address consists of 12 binary ones (inactive address), hunting for the next program in demand is inhibited.

The time required to execute the proceed instruction is two memory cycles.

Simulator, S

Simulator instructions permit the programmer to represent a subroutine with a single instruction in his program. For each simulator instruction used, he codes a subroutine which is stored elsewhere in memory, beginning with the next memory location higher than the address specified by the command code. The simulator instruction sets the cosequence counter to the starting address of the subroutine and then gives control to this counter.

A simulator instruction is specified by a machine command code in which the low-order three bits are ones. If the high-order bit of the command code is a zero, the low-order 11 bits are interpreted as a main memory subaddress ending in octal seven. If the high-order bit is a one, the low-order 11 bits are interpreted as a 3-bit index register number and an 8-bit augments ending in octal seven. Since the high-order bit of the command code is used to indicate the type of addressing, the programmer does not have the option of specifying the source of the next instruction. In fact, the next instruction is automatically taken from the cosequence counter. This is the only instance in which the machine makes a functional distinction between the sequence counter and the cosequence counter.

When a simulator instruction is executed, the instruction itself is stored in the location specified by the command code. If this address refers directly to a memory location, then the instruction is stored in the same bank of memory from which it was executed; if the address is indexed, the instruction may be stored in any bank, according to the value of the bank indicator bits in the referenced index register. The cosequence counter is set to the next higher address and the next instruction is taken from the cosequence counter. The contents of the source counter, after normal incrementing, are stored in the cosequence history register to provide a return to the main program.

Since the address portions of a simulator instruction have no assigned functions, they may be used to store subroutine parameters such as the sources of operands, the location in which a result is to be stored, the number of words to be manipulated, and so forth. Since the machine automatically stores in special register AU1 an address generated from the A address group and in AU2 an address generated from the C address group, it is advantageous to use these address groups for operand or re-

SECTION XII. MISCELLANEOUS INSTRUCTIONS

sult locations. The A and C address groups may contain direct memory location addresses (bit 1 = zero) or indexed memory location addresses (bit 1 = one). If bit 1 is zero, the low-order 11 bits of the address group and the 4-bit bank indicator from the sequencing counter which selected the instruction are placed in the AU-CU counter to form a 15-bit address. If bit 1 is one, the augmented contents of the referenced index register are placed in the counter. The contents of the two counters may then be referenced easily through the technique of indirect memory location addressing.

The command code for an ARGUS simulator instruction is S, followed by a comma and an address designated by a symbolic tag or by an index register number with augments equal to seven. The A and C address fields may contain symbolic tags, with or without address modifiers, or index register numbers with augments.

The time required to execute the simulator instruction is seven or nine memory cycles, as indicated in Appendix C.

Compute Orthocount, CC

The technique of orthotronic control incorporated in the Honeywell 800 assumes that every record written on tape will include two orthowords, one associated with the odd-numbered data words in the record, the other associated with the even-numbered words. If the record contains an even number of data words, then orthoword 1 represents the orthocount of the odd words (words 1, 3, 5 etc.) and orthoword 2 represents the orthocount of the even words. If an odd number of words is orthocounted, then orthoword 1 contains the orthocount of the even words (words 2, 4, 6 etc.) and orthoword 2 contains the count of the odd words. Each orthoword is the complement of the binary half add of the words associated with it.

The machine instruction which performs the orthocount of a record and generates the two orthowords is called compute orthocount. When this instruction is executed, the machine first generates an end-of-record word which is placed in the location specified by the C address group. It then orthocounts the record, starting with the location specified by the A address group and ending when an end-of-record word is reached. The first end-of-record word sensed will terminate the operation, regardless of whether the location specified in C has been reached. Orthoword 1 is stored in the location specified by C; orthoword 2 is stored in C + 1; and an end-of-record

word is placed in $C + 2$. If address B is inactive, control is not changed for distributed item handling. If the B address is active, distributed item control is exercised, using a memory location table stored in consecutive words beginning with the location specified by B.

The behavior of the system is unspecified if either the A or C address group of the compute orthocount instruction contains a direct or indexed special register address. If C is an indirect memory location address with a non-zero increment, the behavior of the system is also unspecified.

At the beginning of the instruction, special register AU1 is set to the location designated by the A address group. As each word is orthocounted, the counter is automatically incremented by one to specify the location of the next word to be orthocounted. At the conclusion of the instruction, AU1 contains the address of the end-of-record word which terminated the instruction, plus one. If the B address of the instruction is active, special register AU2 is initially set to the location specified by the B address group. This counter performs the same function for the distributed orthocount as the distributed read address and distributed write address counters perform for the distributed read and write instructions (see Section XI, pages 96 and 102). As each item is handled, the counter is automatically incremented by one to reference the entry in the address table where the starting address of the next item is found. At the conclusion of the instruction, special register AU2 always contains the address of the location specified in the C address of the instruction plus three, regardless of whether the instruction was normal or distributed.

The provision for generating an end-of-record word and storing it in the location designated by the C address guarantees the programmer that the record being orthocounted has a valid end-of-record word. This is a necessary precaution, since the compute orthocount instruction is terminated only when an end-of-record word is sensed. During execution of the instruction, every other word is half-added, in binary, to form orthoword 1, and the alternate words are similarly added to form orthoword 2. The accumulator is set to all binary ones before performing the first half-add for each orthoword. The inclusion of a word of binary ones in the half add assures that each orthobit is an "odd" parity check on the identical bit positions of all the associated words. The accumulator and mask register are used to store partial results as

SECTION XII. MISCELLANEOUS INSTRUCTIONS

the successive alternate words are half-added. When the orthocount is terminated by an end-of-record word, which is not included in the orthocount, the contents of the accumulator are transferred as orthoword 1 to the location specified by the C address, and the contents of the mask register are transferred as orthoword 2 to C + 1. If distributed item control is executed, the end-of-item words are included in the orthocount.

The compute orthocount instruction is used in conjunction with the check parity instruction (see below) to reconstruct data in which an error has been detected. The channel in which the parity failure occurred may be identified by computing new orthowords for the record in error and comparing them with the orthowords accompanying the record. The check parity instruction may then serve to identify the erroneous word or frame. The formation and use of orthowords is described in greater detail in Appendix B.

The time required to execute a compute orthocount instruction with inactive B address and direct memory location A and C addresses is $11 + n$ memory cycles, where n equals the number of words orthocounted. Variations in timing for distributed item handling are shown in Appendix C.

Check Parity

When an error is detected during a read from tape, the check parity instruction is used to pinpoint the incorrect word(s) or frame(s). The instruction may be performed without a mask to check the parity of the entire word, or it may be used with a frame mask to test parity on one or more frames. The checking is accomplished by transferring the word to be checked (stored in the location specified by the A address group) to the location specified by the B address group. During this transfer, parity is automatically checked by the parity-check circuits, and an indicator is set if parity is incorrect. The correct parity bits from the parity-check circuits are delivered to B, together with the information bits contained in A, thus guaranteeing that the word stored in B has proper parity so that it can be manipulated in the machine without causing a control error. The parity-error indicator is then sampled, and if the parity of the word at address A was incorrect, the sequencing counter specified as the source of the next instruction is changed to select the next instruction from the memory address designated by C. The C address group may contain a direct memory

SECTION XII. MISCELLANEOUS INSTRUCTIONS

location address, an indexed memory location address, an indirect memory location address, or an indexed indirect memory location address, as described under the instruction "transfer A to B, go to C" (see Section VIII, pages 78 and 79).

When the check parity instruction is used with a mask, only those information bits corresponding to binary ones in the mask are checked for parity. If the mask defines one or more complete frames, the sequencing counter will be changed if any frame in the masked operand has incorrect parity. The instruction is undefined for any other type of mask.

In order to test parity on an individual frame basis, the programmer must be aware of the relationship between the information bits of a word on tape and their positions in memory. This relationship is shown in Figure XII-3.

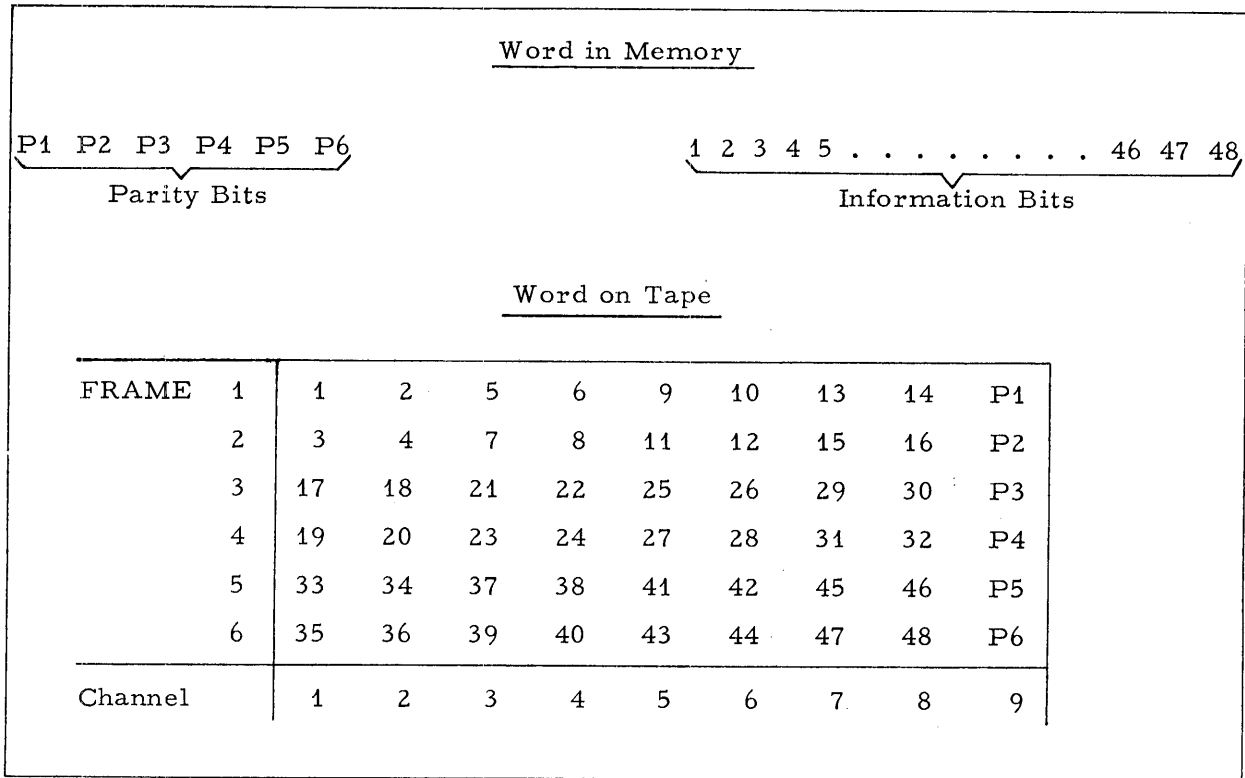


Figure XII-3. Bit Layout -- Memory and Tape

From this figure it will be seen that the mask used to check the parity of frame 1 must have the following bit configuration:

SECTION XII. MISCELLANEOUS INSTRUCTIONS

1100 1100 1100 1100 0000 0000 0000 0000 0000 0000 0000 0000,

while the mask required to check the parity of frame 2 would have the configuration:

0011 0011 0011 0011 0000 0000 0000 0000 0000 0000 0000 0000.

The time required to execute an unmasked check parity instruction with direct memory location addresses is four memory cycles.

SECTION XIII
SUMMARY OF INSTRUCTIONS

General Instructions - Unmasked or Masked

MNEMONIC OPERATION CODE	DESCRIPTION	TIME IN ¹ MEMORY CYCLES
BA ⁴	<p>Binary Add algebraically (A) to (B). Store sum in C. If overflow occurs, take next instruction from U + 8 if the sequence counter selected this instruction, or from U + 9 if the cosequence counter selected this instruction.</p> <p>The sign of either operand is positive if any of its four sign bits is one. The sign of the sum is 0000 if negative, 1111 if positive.</p>	4
DA ⁴	<p>Decimal Add algebraically (A) to (B). Store sum in C. Otherwise same as Binary Add.</p>	4
WA	<p>Word Add. Binary add (A) to (B), considered as unsigned 48-bit numbers. Store 48-bit result in C. If overflow occurs, observe same conventions as in Binary Add.</p>	4
BS ⁴	<p>Binary Subtract algebraically (B) from (A). Store result in C. Observe same overflow and sign conventions as in Binary Add.</p>	4
DS ⁴	<p>Decimal Subtract algebraically (B) from (A). Store result in C. Observe same overflow and sign conventions as in Binary Add.</p>	4
WD	<p>Word Difference. Binary subtract (B) from (A). Otherwise identical to Word Add.</p>	4
NA	<p>Inequality Comparison, Alphabetic. Compare (A) and (B) including sign positions. If (A) ≠ (B), change sequencing counter to select next instruction from location specified by C. Plus zero is not equal to minus zero.</p>	4
NN	<p>Inequality Comparison, Numeric. Compare algebraically (A) and (B). If (A) ≠ (B), follow procedure for NA. Plus zero equals minus zero.</p>	4

SECTION XIII. SUMMARY OF INSTRUCTIONS

General Instructions - Unmasked or Masked (cont.)

MNEMONIC OPERATION CODE	DESCRIPTION	TIME IN ¹ MEMORY CYCLES
LA	Less Than Or Equal Comparison, Alphabetic. Compare (A) and (B) including sign positions. If $(A) \leq (B)$, change sequencing counter to select next instruction from location specified by C. Plus zero is not equal to minus zero.	4
LN	Less Than Or Equal Comparison, Numeric. Compare algebraically (A) and (B). If $(A) \leq (B)$, follow procedure for LA. Plus zero equals minus zero.	4
TX	Transfer (A) to C. Ignore B.	3
TS	Transfer (A) to B. If C is active, change sequencing counter to select next instruction from location specified by C.	4
HA	Half Add. Binary add without carry (A) to (B), considered as unsigned 48-bit numbers. Store 48-bit result in C. Result is zero wherever corresponding bits of (A) and (B) are identical, one wherever corresponding bits of (A) and (B) are different. This is "logical exclusive OR."	4
SM	Superimpose (A) and (B). Store result in C. Result is zero wherever bits of (A) and (B) are both zero, one everywhere else. This is "logical inclusive OR."	4
CP	Check Parity. Test (A) for correct parity. Place (A) with correct check bits in B. If (B) differs from (A), change sequencing counter to select next instruction from location specified by C.	4

General Instructions - Unmasked

BM	Binary Multiply (A) by (B). Store high-order product with proper sign in C and accumulator, low-order product with proper sign in low-order product register. Product signs are 0000 if negative or 1111 if positive.	33
DM	Decimal Multiply (A) by (B). Store high-order and low-order products as in Binary Multiply with same sign conventions.	27

SECTION XIII. SUMMARY OF INSTRUCTIONS

General Instructions - Unmasked (cont.)

MNEMONIC OPERATION CODE	DESCRIPTION	TIME IN ¹ MEMORY CYCLES
BT ²	Binary Accumulate. Place low-order 44 bits of (A) in accumulator. Perform this transfer B' times, adding in binary (with positive sign implied) the successive 44-bit words transferred. B' (high-order 6 bits of B) = 0 to 63. Store 44-bit result, with sign of first word transferred, in C. Observe same overflow conventions as in Binary Add. Note that if A is an indirect address with non-zero increment, B different numbers are accumulated.	3 + n
DT ²	Decimal Accumulate. Same as Binary Accumulate except that transferred words are added as 11-digit decimal numbers.	3 + n
MT ²	Multiple Transfer. Transfer (A) to C. Perform this instruction B' times. B' (high-order 6 bits of B) = 0 to 63. Note that if A and C are indirect addresses with non-zero increments, B' different transfers are performed.	1 + 2n
TN ²	N-Word Transfer. Transfer B' words from consecutive locations starting at A to consecutive locations starting at C. B' = 0 to 63.	5 + 2n
CC ^{2,5}	Compute Orthocount. Write a generated end-of-record word in the location specified by C. Orthocount the record starting at A to the end-of-record word. Store orthoword 1 in C and orthoword 2 in C + 1. Place end-of-record word in C + 2. If B is inactive, control is <u>not</u> changed for distributed item handling. If B is active, end-of-item words are sensed and control is changed for distributed item handling.	11 + n
IT ²	Item Transfer. Substitute an end-of-item symbol for the high-order 32 bits of (B), protecting the low-order 16 bits of (B). Transfer words from consecutive memory locations starting with A to consecutive memory locations starting with C until an end-of-item word is transferred.	7 + 2n
RT ²	Record Transfer. Store an end-of-record word in B. Transfer words from consecutive memory locations starting with A to consecutive memory locations starting with C until an end-of-record word is transferred.	7 + 2n

SECTION XIII. SUMMARY OF INSTRUCTIONS

General Instructions - Unmasked (cont.)

MNEMONIC OPERATION CODE	DESCRIPTION	TIME IN ¹ MEMORY CYCLES
MPC	Control Program. Ignore A. Place (PCR) in the location specified by C. Then alter the bits of PCR specified by bits 5-12 of B, using bits 1-4 of B to define how the bits are altered. If B address memory designator bit is 1, hunt for next program in demand. Otherwise, do not hunt.	4
PR	Proceed. If C is inactive, do not hunt for next program in demand.	2

Inherent Mask Instructions

SWS ³	Shift Word and Substitute. Shift right end around including sign (A) as directed by B'. Mask result and store in C (protected). B' (high-order 6 bits of B) specifies the number of 1-bit shifts.	5 + k
SPS ³	Shift Preserving Sign and Substitute. Shift right end around excluding sign (A) as directed by B'. Otherwise same as SWS.	5 + k
SWE ³	Shift Word and Extract. Same as SWS except that the unmasked portions of (C) are unprotected.	5 + k
SPE ³	Shift Preserving Sign and Extract. Same as SPS except that the unmasked portions of (C) are unprotected.	5 + k
SSL ³	Shift and Select. Shift right end around including sign (A) as directed by B'. Binary add to C under mask control no more than 11 low-order bits of the shifted word, with positive sign implied. Change the sequencing counter to select the next instruction from the location specified by the modified C address. B' is interpreted as in the SWS instruction.	6 + k
SS	Substitute. Using (B) as a mask, transfer (A) to C and protect unmasked portions of (C).	5
EX	Extract or Logical AND. Using (B) as a mask, transfer (A) to C without protecting unmasked portions of (C). Result is one wherever bits of (A) and (B) are both one, zero everywhere else.	5

SECTION XIII. SUMMARY OF INSTRUCTIONS

Peripheral and Print Instructions

MNEMONIC OPERATION CODE	DESCRIPTION	TIME IN ¹ MEMORY CYCLES
RF	<p>Read Forward from peripheral device XX into consecutive memory locations starting with A. XX represents command code bits 1-6. Set the RAC to +A. If B is inactive, do not change control for distributed item handling. If B is active, set the DRAC to +B and sense for end-of-item words. Change sequencing counter to select next instruction from location specified by C.</p> <p>If end of tape is sensed, take next instruction from U + 4 if the sequence counter selected this instruction or from U + 5 if the cosequence counter selected this instruction. If a parity error was detected during the last previous read from this device, reset the parity error flip-flop, do not perform the read. Instead, take next instruction from U + 6 or U + 7. This instruction is interlocked against device XX and the associated buffer.</p>	5
RB	<p>Read Backward from magnetic tape unit XX into consecutive memory locations starting with A. This instruction is otherwise identical with RF except that the RAC is set to -A and if B is active the DRAC is set to -B.</p>	5
WF	<p>Write Forward on peripheral device XX the contents of consecutive memory locations from A through the end-of-record word. Set the WAC to +A. If B is inactive, do not change control for distributed item handling. If B is active, set the DWAC to +B and sense for end-of-item words. Change sequencing counter to select next instruction from location specified by C.</p> <p>If an error was detected during the last previous write to peripheral device XX, reset the parity error flip-flop, do not perform the write. Instead, take next instruction from U + 6 or U + 7. This instruction is interlocked against device XX and the associated buffer. End-of-tape convention is identical with Read Forward.</p>	5
RW	<p>Rewind tape on magnetic tape unit XX to end. If tape is already rewound, proceed. If A is active, set manually removable interlock after tape is rewound. B and C are ignored. If an error was detected during the last previous read or write for this tape, reset the parity error flip-flop, do not perform the rewind. Instead, take next instruction from U + 6 or U + 7.</p>	2

SECTION XIII. SUMMARY OF INSTRUCTIONS

Peripheral and Print Instructions (cont.)

MNEMONIC OPERATION CODE	DESCRIPTION	TIME IN ¹ MEMORY CYCLES
PRA, PRD, PRO	Print (A) on the typewriter and in the format specified by B. If C is active, change the sequencing counter to select the next instruction from the location specified by C. In an alphabetic print, eight 6-bit characters are printed. In a decimal print, 12 hexadecimal digits are printed. In an octal print, 16 octal numbers are printed.	5

Simulator Instructions

S	Simulator. Form a memory location address (direct or indexed) from the low-order 11 bits of the command code and store this instruction in the location thus specified. Change the sequence counter to select the next instruction from the next higher address.	7
---	--	---

NOTES:

1. One memory cycle equals six microseconds. All addresses considered active. For variations in time, see Appendix C.
2. n = number of words accumulated, transferred, or ortho-counted.
3. Values of k are based on number of 16-, 4-, and 1-bit shifts required. See Appendix C for table of values.
4. Under certain conditions, as explained in Appendix A, an additional one or two memory cycles may be required.
5. B address considered inactive (sole exception to note 1 above).

APPENDIX A

ADDITION IN THE HONEYWELL 800

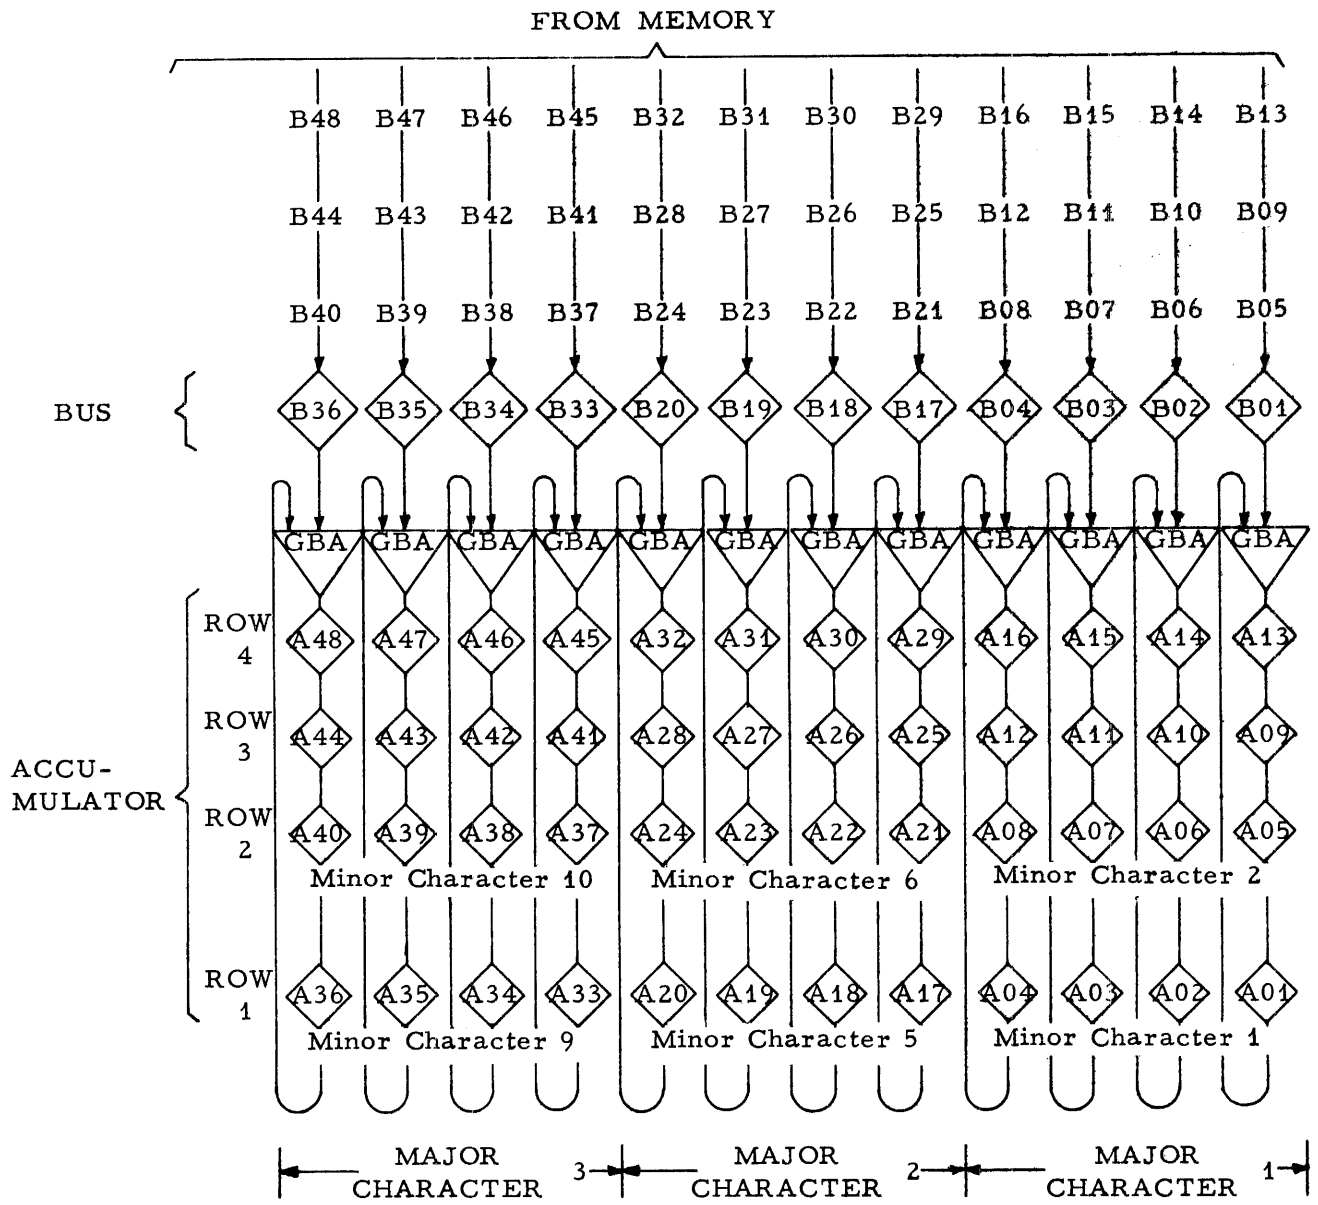
Two registers are involved in the addition function of the Honeywell 800: the accumulator and the bus. A symbolic representation of these registers is shown in Figure A-1. The accumulator is used in the performance of addition, subtraction, multiplication, shifting, comparison, transfer, and print instructions.

Since the Honeywell 800 word contains 48 information bits, the accumulator requires 48 flip-flops, each capable of storing a single binary digit. These 48 flip-flops are arranged in three 4 x 4 arrays, called "major characters." Each major character contains four "minor characters." For example, minor character No. 1 contains bits 1, 2, 3, 4¹; minor character No. 2 contains bits 5, 6, 7, 8; minor character No. 5 contains bits 17, 18, 19, 20; etc. (see Figure A-1).

The accumulator has 24 additional flip-flops and amplifiers called "carry functions" which permit carries or borrows generated during arithmetic operations to propagate through the arrays without interfering with the information bits. These carry functions are used to implement the anticipating techniques discussed below.

All the operations and manipulations accomplished by the accumulator are based on its ability to recirculate bits between rows. In other words, the information stored in the 48 flip-flops is not stationary. For example, the information stored in the 12 flip-flops of the first row remains static for less than one microsecond and is then transferred to the flip-flops in the fourth row. The information in the fourth row is fed in turn to the third row, the third to the second row, and the second to the first. By this recirculation process the accumulator stores information.

1. In this Appendix only, bits are numbered from right to left, conforming to standard engineering usage. This is in contrast to the programmers' left-to-right numbering convention which is followed elsewhere in the manual.



Note: Carry Logic not shown.

Figure A-1. Schematic Representation of Bus and Accumulator

Each operand is delivered from the memory to the accumulator via the bus. In the accumulator, the least significant binary digit is stored in flip-flop 1, the next higher-order digit is stored in 2, and so on. The most significant binary digit of a signed number is stored in flip-flop 44; the most significant binary digit of an unsigned number (such as an operand of the word add instruction) is stored in 48. Since four binary digits are required to represent a decimal digit in the Honeywell 800, the least significant decimal digit is stored in the first row of major character No. 1, the next higher-order decimal digit in the second row of major character No. 1, etc. The most significant decimal digit of a signed number is represented by the four bits in the third row of major character No. 3. In the accumulator, bits 45 through 48 are used to store either the high-order four bits of an unsigned number, or an arbitrary binary configuration when signed numbers are being operated upon.

The sign of the A operand is stored in a single flip-flop designated sign control (SGA). The content of SGA, the sign of the B operand, and the operation code determine the sign of the result, and also the nature of the operation to be performed. For example, if SGA contains a plus sign, the B operand is negative, and the operation code specifies subtraction, the operation to be performed logically is addition. When addition is performed on signed numbers, bits 45 through 48 of the accumulator are filled with binary ones so that a carry out of bit 44 is reflected by a carry out of bit 48. Thus overflow may be determined by the presence of a carry out of bit 48 in addition of either signed 44-bit operands or unsigned 48-bit operands. In like manner, bits 45 through 48 are filled with binary zeros when subtracting signed numbers in order that borrows may be sensed at the same point for both signed and unsigned numbers. For addition of numbers with unlike signs, subtraction of numbers with like signs, or multiplication, the contents of SGA may change to correspond with the sign of the result generated in the accumulator.

As information is transferred from the bus to the accumulator, the four bits in each minor character receive the information simultaneously (in parallel) and the minor characters within a major character receive information serially. However, a group of three minor characters in the same row (1, 5, and 9, or 2, 6, and 10 -- see Figure A-1) receive their information in parallel. This method of receiving and processing data characterizes the Honeywell 800 as a parallel-serial-parallel machine.

The following detail steps are performed in executing an addition instruction. First, the A operand is transferred from the memory to the bus and then is received by the accumulator in parallel-serial-parallel fashion. The B operand is then transferred from memory to the bus. As the first row of information is transmitted via the bus to the accumulator, each corresponding pair of bits, one from the bus and one from row 1 of the accumulator, is input to a gate buffer amplifier (GBA), a logical element in which the addition of the two binary digits takes place. The sum of the two digits is then transmitted to the fourth row. If a carry will occur within a minor character from this addition, it is anticipated and introduced as an input into the next higher-order GBA to be added to the bit from the accumulator (A operand) and the bit from the bus (B operand). If a carry occurs from the high-order binary digit of a minor character during the arrival time of rows 1, 2 or 3 at the GBA's, the carry is brought to the low-order GBA associated with the same major character at the time of arrival of the following row. Therefore, carries within a minor character or within a major character require no extra processing time.

If a carry occurs from the high-order bit position of a minor character in row 4 (in other words, from one major character to another), a second cycling of the accumulator is required. The carry is brought as input to the low-order GBA associated with the next major character as row 1 is again transmitted to row 4. (On this cycle, of course, there is no input to the GBA from the bus.)

A second cycle of the accumulator does not necessarily mean an extra memory cycle each time a carry occurs from a major character. The content of the accumulator is read from the output of the flip-flops of row 2 to storage and thus a "half cycle," to place the original contents of row 1 in row 2, is required before the contents of the accumulator are transferred to the bus and thence to memory. Therefore, if the carry from a major character propagates through no more than two minor characters, no additional memory cycle is needed. If the carry propagates further (including the worst case of a carry from major character No. 1 which propagates through major character Nos. 2 and 3 in turn), a maximum of one extra cycle is needed. Because the "borrow" logic includes the possibility of a carry from major character 3 end-around to major character 1, the addition of two words with unlike signs could require two extra memory cycles.

The above discussion is applicable to both binary and decimal arithmetic. In the case of decimal operations, however, an additional operation takes place. As each minor character in the sum proceeds from row 4 to row 3, the 4-bit configuration is tested to see if it represents a decimal digit (0 through 9). If so, it proceeds as described above. If not, the 4-bit configuration is reduced by 10 (binary 1010) and a carry of one is introduced into the low-order bit of the next higher-order minor character. This discussion assumes that operands in decimal operations are made up of decimal numbers.

APPENDIX B

ORTHOTRONIC CONTROL

Orthotronic control is the technique incorporated in the Honeywell 800 for the automatic detection and correction of errors that may occur with the use of magnetic tape. This file protection feature accomplishes error correction without the manual intervention and lost computer time so frequently associated with ordinary restart or rerun procedures. In other systems, the problem of restart and rerun becomes most serious when the record in error was properly recorded during a previous run but is now incapable of being read. It is to this type of error, the most costly in time and effort, that orthotronic control is applied most advantageously in the Honeywell 800 system.

The basic principle of orthotronics may be compared to the accountant's practice of crossfooting. In crossfooting, a zero balance of rows and columns is obtained to insure accuracy. If a zero balance is not obtained, the crossfooting technique has accomplished its singular function of error detection. Orthotronic control not only performs this function of error detection but also provides a unique means of error correction.

The orthotronic technique involves three interdependent elements: frame parity, orthotronic control words, and channel parity. Frame parity is recorded in the ninth channel on tape. An "odd" parity system is used in the generation of the parity bit. The parity bit is the complement of the binary half adds of the bits in each frame. Thus, the parity bit is a one if there is an even number of ones in a frame (0, 2, 4, 6, 8 ones), and the parity bit is zero if there is an odd number of ones in a frame (1, 3, 5, 7 ones).

Before a new or altered record is written on tape, the machine instruction compute orthocount is used to compute two orthowords: one associated with the odd-numbered data words in the record, the other associated with the even-numbered words. These orthowords become part of the record written on tape. Orthoword 1

may be the orthocount of either the odd words or even words of a record. This is determined by the number of words to be orthocounted in the record (see Figure B-1). If the record contains an even number of data words, then orthoword 1 represents the orthocount of the odd words. If an odd number of data words is orthocounted, then orthoword 1 contains the orthocount of the even words. However, it should be noted that by working back from the orthoword, the relationship is invariant, regardless of the number of words. Each orthobit is an "odd" parity bit checking the corresponding bit positions of all the associated words.

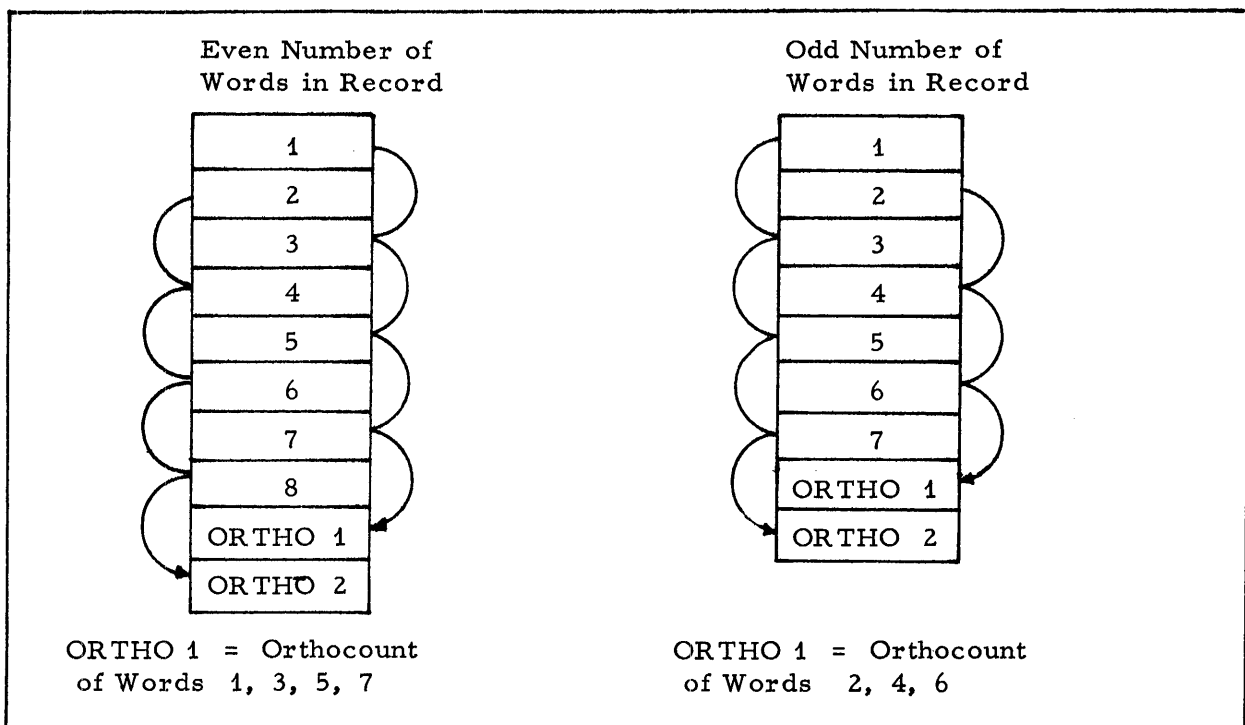


Figure B-1. Relation of Data Words to Orthowords

Consider a record in memory (for the sake of simplicity, 16-bit words are shown). Figure B-2 shows the method used to compute the orthowords. Each orthoword is the complement of the binary half add of the words associated with it.

Channel parity checking is automatically performed in the buffer to provide a longitudinal check of information being read or written. The generated orthotronic words guarantee an even channel parity up to the end-of-record word. The end-of-


		<u>Odd Words</u>				<u>Even Words</u>				
		(1111	1111	1111	1111	1111	1111	1111	1111)	Complement
Word										Bits
1	1647	0001	0110	0100	0111					
2	3018					0011	0000	0001	1000	
3	8390	1000	0011	1001	0000					
4	3764					0011	0111	0110	0100	
5	0519	0000	0101	0001	1001					
6	A B 4					0100	0101	0010	0100	
7	2613	0010	0110	0001	0011					
8	7 C D	<hr/>				0111	0100	1101	0100	
ORTHO 1	4922	0100	1001	0010	0010	<hr/>				
ORTHO 2	12973					1100	1001	0111	0011	
	END OF RECORD									

Figure B-2. Computation of Orthowords

record word is such that its inclusion guarantees an odd channel parity. Thus, double errors in a frame, which will escape detection by frame parity, will be picked up by the channel parity check.

Using the example shown in Figure B-2, assume that a bit was altered in the high-order decimal digit of word 7 to make it 1010 instead of 0010. Figure B-3 shows the words in memory and a simplified, hypothetical representation of the same words on tape, with frame parity bits.

When this record is read, the error is automatically detected by the frame parity check. The automatic channel parity check also shows up the error, since a binary half add of the bits in the left-most channel produces a one bit instead of a zero bit. If, in addition, a bit in the second high-order digit of word 7 were altered so that it became 0111 instead of 0110, this frame would now appear as follows on tape:



This combination of errors would escape the frame parity check only to be caught by the channel parity computation.

Regardless of the means of error detection, the result is the same: an unprogrammed transfer of control is initiated when the next read or write instruction to the same device is executed.¹ The procedure to be followed when the error occurs on

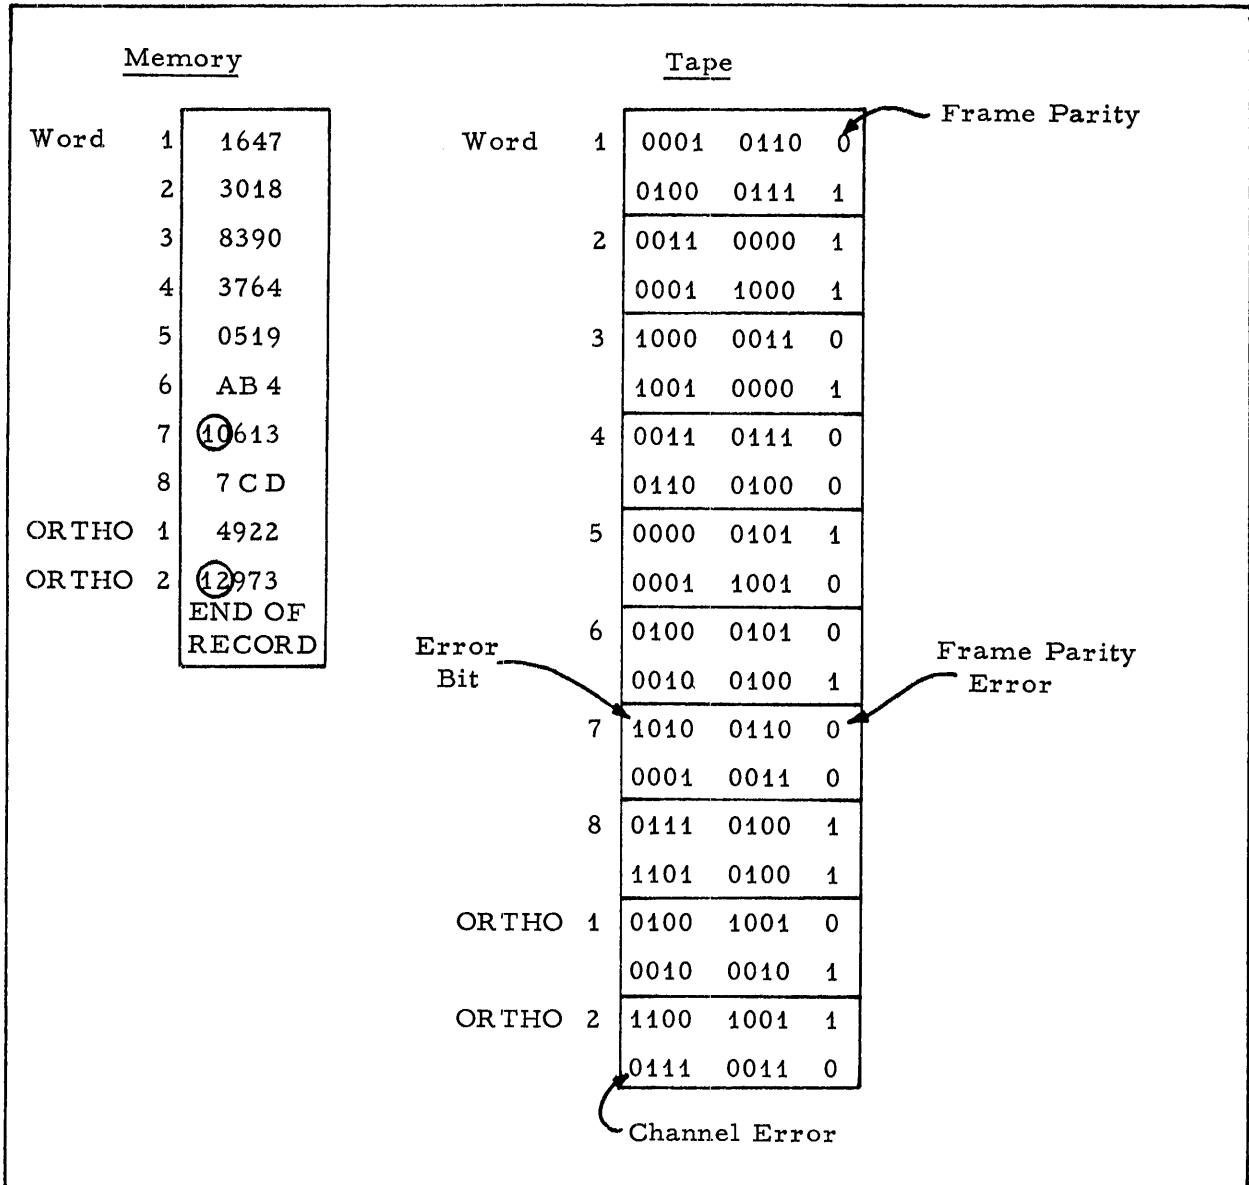


Figure B-3. Orthotronic Check of Tape Error

1. As noted in Section XI, a read instruction checks the previous read while a write instruction checks the previous write.

APPENDIX B. ORTHOTRONIC CONTROL

writing is to read backward over the erroneous record and then rewrite it. The procedure to be followed when the error occurs on reading is to orthocount the entire record again (including the bad word and the orthowords), using the compute orthocount instruction. The result of this procedure is illustrated in Figure B-4, using the example in which a bit was altered in the high-order digit of word 7.

		<u>Odd Words</u>				<u>Even Words</u>					
		(1111	1111	1111	1111	1111	1111	1111	1111)	Complement	
Word										Bits	
1	1647	0001	0110	0100	0111						
2	3018					0011	0000	0001	1000		
3	8390	1000	0011	1001	0000						
4	3764					0011	0111	0110	0100		
5	0519	0000	0101	0001	1001						
6	AB 4					0100	0101	0010	0100		
7	ⓐ613	1010	0110	0001	0011						
8	7 C D					0111	0100	1101	0100		
ORTHO	1	4922	0100	1001	0010	0010					
ORTHO	2	ⓑ2973	_____				1100	1001	0111	0011	
NEW ORTHO	1	8000	1000	0000	0000	0000	_____				
NEW ORTHO	2	0000					0000	0000	0000	0000	

Figure B-4. Orthocount of Error Record

This computation shows that the error lies in one of the odd words (1, 3, 5, or 7). The check parity instruction is then used to isolate the bad word (and/or the bad frame within the word, by masking the instruction). The check parity instruction also assigns correct parity to the word. If correction is to be made on a word basis, a binary half add is used to add the new orthoword and the bad word (with adjusted parity) to restore the correct word. If correction is to be made on a frame basis, a properly masked half add is used to restore the frame. If correction is made on a word basis, the binary half add produces the result shown in Figure B-5.

Bad Word	7	1010	0110	0001	0011	
New Ortho	1	<u>1000</u>	<u>0000</u>	<u>0000</u>	<u>0000</u>	
Restored Word	7	0010	0110	0001	0011	(cf. Figure B-2)

Figure B-5. Correction of Tape Error -- First Method

A different procedure for reconstructing the bad word would be first to locate it by the check parity instruction and then to cancel it by substituting a word of all binary zeros. Next, orthocount the entire record including the cancelled word and the orthowords, as shown in Figure B-6.

		<u>Odd Words</u>				
		(1111	1111	1111	1111)	Complement
Word						Bits
	1	0001	0110	0100	0111	
	3	1000	0011	1001	0000	
	5	0000	0101	0001	1001	
(Bad Word)	7	0000	0000	0000	0000	(Binary Zeros
Old ORTHO		<u>0100</u>	<u>1001</u>	<u>0010</u>	<u>0010</u>	Substituted)
New ORTHO		0010	0110	0001	0011	

Figure B-6. Correction of Tape Error -- Second Method

The new orthoword equals the reconstructed word (cf. Figure B-2) and may now be substituted for the word of binary zeros. This second technique can be used only when errors are detected in a single word associated with each orthoword. For example, this method can be used to correct an error in one even word and one odd word in the record, but not to reconstruct two odd or two even words. The technique illustrated by Figure B-5, on the other hand, may be used to reconstruct errors in several words or frames, as long as there is not more than one frame parity or ortho error associated with each of the twelve frames comprising the two orthowords.

APPENDIX C
TIMING SUMMARY

Instruction	Basic Time in Memory Cycles ¹	Modifications of Basic Time
<p><u>BA⁵, DA⁵, WA, BS⁵, DS⁵, WD, HA, SM</u></p> <p>A. Unmasked</p> <ol style="list-style-type: none"> 1. If B is inactive, and <ol style="list-style-type: none"> a. A is active, C inactive b. A is inactive, C active c. A and C are both active 2. If B is active <p>B. Masked</p> <ol style="list-style-type: none"> 1. If B is inactive, and <ol style="list-style-type: none"> a. A is active, C inactive b. A is inactive, C active c. A and C are both active 2. If B is active 	<p>3</p> <p>4</p> <p>5</p> <p>4</p> <p>4</p> <p>6</p> <p>7</p> <p>6</p>	<p>Add: 1 mc if A is indexed</p> <p>Add: 1 mc if A is indexed</p> <p>Add: 1 mc if A is inactive</p> <p>1 mc if A or B is indexed</p> <p>*1 mc if A is a direct special register or indirect memory location address <u>and</u> B is indexed</p> <p>*1 mc if B is a direct special register or indirect memory location address <u>and</u> C is indexed</p> <p>*2 mc if A and B are direct special register or indirect memory location addresses <u>and</u> C is indexed</p> <p>Add: 1 mc if A is indexed</p> <p>Add: 1 mc if A is indexed</p> <p>Add: 1 mc if A is inactive</p> <p>1 mc if A or B is indexed</p> <p>Sub: 1 mc if C is inactive</p>
<p><u>BT², DT²</u></p> <ol style="list-style-type: none"> A. If B address (n) = 0 B. If B address (n) > 0 	<p>2</p> <p>3 + n</p>	<p>Add: 1 mc if A is indexed</p>
<p><u>BM</u></p>	<p>33</p>	<p>Add: 1 mc for each indexed address</p>

APPENDIX C. TIMING SUMMARY

Instruction	Basic Time in Memory Cycles ¹	Modifications of Basic Time
<u>DM</u>	27	Add: 1 mc for each indexed address
<u>NA, NN, LA, LN</u> A. Unmasked B. Masked	4 5	Add: 1 mc if A or B (or both) is indexed 1 mc if A is direct special register or indirect memory location address <u>and</u> B is indexed 1 mc if B is direct special register or indirect memory location address Add: 1 mc if A or B (or both) is indexed
<u>TX</u> A. Unmasked B. Masked	3 5	Add:*1 mc if A is indexed *1 mc if A is direct special register or indirect memory location address <u>and</u> C is indexed *2 mc if A is indexed special register or indirect memory location address <u>and</u> C is indexed Add: 1 mc if A or C (or both) is indexed
<u>TS</u> A. Unmasked 1. If B is inactive, and a. A is inactive b. A is active 2. If B is active	4 5 4	Add: 1 mc if A is indexed Sub: 2 mc if C is inactive Add:*1 mc if A is inactive *1 mc if A is inactive and B is indexed *1 mc if A or B is indexed *1 mc if A is direct special register or indirect memory location address <u>and</u> B is indexed *2 mc if A is indexed special register or indirect memory

APPENDIX C. TIMING SUMMARY

Instruction	Basic Time in Memory Cycles ¹	Modifications of Basic Time
<p><u>TS (cont.)</u></p> <p>B. Masked</p> <p>1. If B is inactive, and</p> <p> a. A is inactive</p> <p> b. A is active</p> <p>2. If B is active</p>	<p>5 } 6 } 6 }</p>	<p>location address <u>and</u> B is indexed</p> <p>1 mc if B is direct special register or indirect memory location address</p> <p>1 mc if C is direct special register or indirect memory location address</p> <p>1 mc if C is indexed</p> <p>{ Add: 1 mc if A is indexed 1 mc if C is indexed Sub: 2 mc if C is inactive Add:*1 mc if A is inactive *1 mc if A is inactive and B is indexed *1 mc if A or B is indexed 1 mc if C is indexed</p>
<p><u>MT</u>²</p> <p>A. If B address (n) = 0</p> <p>B. If B address (n) > 0</p>	<p>2</p> <p>1 + 2n</p>	<p>Add:*n if A is indexed</p> <p>*n if A and C are both indexed</p> <p>*1 mc if C is indexed</p> <p>*n + 1 if A is a direct special register or indirect memory location address <u>and</u> C is indexed</p> <p>*2n if A is indexed special register or indirect memory location address <u>and</u> C is indexed</p>
<p><u>TN</u>²</p> <p>A. If B address (n) = 0</p> <p>B. If B address (n) > 0</p>	<p>2</p> <p>5 + 2n</p>	<p>Add: 1 mc if A is indexed indirect memory location address</p> <p>n if A is direct or indexed special register address</p> <p>Add: 1 mc if C is indexed indirect memory location address</p> <p>n if C is direct or indexed special register address</p>

APPENDIX C. TIMING SUMMARY

Instruction	Basic Time in Memory Cycles ¹	Modifications of Basic Time																						
<u>IT², RT²</u>	7 + 2n	Add: 1 mc if A is indexed indirect memory location address 1 mc if C is indexed indirect memory location address																						
<u>SWS, SPS, SWE, SPE</u>	5 + k	<table border="1" data-bbox="933 556 1445 903"> <thead> <tr> <th data-bbox="933 556 1388 588"><u>Total no. 16-, 4-, 1-bit shifts</u></th> <th data-bbox="1388 556 1445 588">k</th> </tr> </thead> <tbody> <tr><td data-bbox="1112 588 1128 619">0</td><td data-bbox="1388 588 1421 619">0</td></tr> <tr><td data-bbox="1112 619 1128 651">1</td><td data-bbox="1388 619 1421 651">0</td></tr> <tr><td data-bbox="1112 651 1128 682">2</td><td data-bbox="1388 651 1421 682">0</td></tr> <tr><td data-bbox="1112 682 1128 714">3</td><td data-bbox="1388 682 1421 714">1</td></tr> <tr><td data-bbox="1112 714 1128 745">4</td><td data-bbox="1388 714 1421 745">1</td></tr> <tr><td data-bbox="1112 745 1128 777">5</td><td data-bbox="1388 745 1421 777">2</td></tr> <tr><td data-bbox="1112 777 1128 808">6</td><td data-bbox="1388 777 1421 808">2</td></tr> <tr><td data-bbox="1112 808 1128 840">7</td><td data-bbox="1388 808 1421 840">3</td></tr> <tr><td data-bbox="1112 840 1128 871">8</td><td data-bbox="1388 840 1421 871">3</td></tr> <tr><td data-bbox="1112 871 1128 903">9</td><td data-bbox="1388 871 1421 903">4</td></tr> </tbody> </table> <p data-bbox="933 913 1445 1029">Add: 1 mc if A is indexed 1 mc if C is indexed 2 mc if C is direct special register address</p>	<u>Total no. 16-, 4-, 1-bit shifts</u>	k	0	0	1	0	2	0	3	1	4	1	5	2	6	2	7	3	8	3	9	4
<u>Total no. 16-, 4-, 1-bit shifts</u>	k																							
0	0																							
1	0																							
2	0																							
3	1																							
4	1																							
5	2																							
6	2																							
7	3																							
8	3																							
9	4																							
<u>SSL</u>	6 + k	For values of k, see table for SWS, etc. Add: 1 mc if A is indexed 1 mc if C is indexed 1 mc if C is direct special register or indirect memory location address																						
<u>SS</u>	5	Add: 1 mc if A or B is indexed 1 mc if A is direct special register or indirect memory location address <u>and</u> B is indexed 1 mc if B is direct special register or indirect memory location address <u>and</u> C is indexed 2 mc if C is direct special register address																						
<u>EX</u> A. If C is inactive, and B is active	4	Add: 1 mc if A or B is indexed 1 mc if A is direct special register or indirect memory location address <u>and</u> B is indexed																						

APPENDIX C. TIMING SUMMARY

Instruction	Basic Time in Memory Cycles ¹	Modifications of Basic Time
<u>EX (cont.)</u> B. If C is active	5	Add: 1 mc if A or B is indexed 1 mc if A is direct special register or indirect memory location address <u>and</u> B is indexed 1 mc if B is direct special register or indirect memory location address <u>and</u> C is indexed 2 mc if C is direct special register address
<u>RF³, RB³, WF³</u> A. If C is inactive B. If C is active	3 5	Add: 1 mc for each indexed address Add: 1 mc for each indexed address
<u>RW³</u>	2	
<u>PRA, PRD, PRO</u> A. If C is inactive B. If C is active	4 5	Add: 1 mc if A or B (or both) is indexed Add: 1 mc if A or B (or both) is indexed 1 mc if C is indexed 1 mc if C is indirect memory location address In addition, 7 mc are required to print each character, at the rate of 1 character every 100ms (approximately)
<u>CC</u> A. If B is inactive ² B. If B is active ⁴	$11 + n$ $9 + 2j + n_1 + n_2 + \dots + n_j$	Add: 1 mc if C is indexed Add: 1 mc if A is direct special register or indirect memory location address <u>and</u> B is indexed 1 mc if B is indexed special register or indirect memory location address 1 mc if C is indexed

APPENDIX C. TIMING SUMMARY

Instruction	Basic Time in Memory Cycles	Modifications of Basic Time
<u>CP</u> A. Unmasked B. Masked	4 6	Add: *1 mc if A or B (or both) is indexed *1 mc if A is direct special register or indirect memory location address <u>and</u> B is indexed *2 mc if A is indexed special register or indirect memory location address <u>and</u> B is indexed 1 mc if B is direct special register or indirect memory location address 1 mc if C is direct special register or indirect memory location address 1 mc if C is indexed Add: 1 mc if A or B (or both) is indexed 1 mc if C is indexed
<u>MPC</u>	4	
<u>PR</u>	2	
<u>S</u>	7	Add: 1 mc if D/I bit = 1 1 mc if C is indexed

APPENDIX C. TIMING SUMMARY

NOTES:

1. One memory cycle equals 6 microseconds.
All addresses assumed active except as otherwise specified.
Address configurations for which the behavior of the system is unspecified have not been considered.
 2. n = number of words accumulated, transferred, or orthocounted.
 3. If the instruction is delayed by an interlock, 3 (if A is not indexed) or 4 (if A is indexed) memory cycles are required for the first attempt to perform it, and an additional 3 (or 4) cycles are required each time that any buffer completes a read or write instruction before this instruction proceeds.
 4. j = number of items in record.
 n_1, n_2, \dots, n_j = number words in item 1, number words in item 2, etc.,
up to n_j , number words in last item.
 5. Under certain conditions, as explained in Appendix A, one or two additional memory cycles will be needed in excess of those computed from this summary.
- * Conditions so marked are mutually exclusive.

Key Punch	Card Code	Honeywell 800 Code	Octal	Standard Printer	High Speed Printer	Console	Key Punch	Card Code	Honeywell 800 Code	Octal	Standard Printer	High Speed Printer	Console
0	0	000000	00	0 (zero)	0	0	-	X	100000	40	- (minus)	-	-
1	1	000001	01	1	1	1	J	X,1	100001	41	J	J	J
2	2	000010	02	2	2	2	K	X,2	100010	42	K	K	K
3	3	000011	03	3	3	3	L	X,3	100011	43	L	L	L
4	4	000100	04	4	4	4	M	X,4	100100	44	M	M	M
5	5	000101	05	5	5	5	N	X,5	100101	45	N	N	N
6	6	000110	06	6	6	6	O	X,6	100110	46	O	O	O
7	7	000111	07	7	7	7	P	X,7	100111	47	P	P	P
8	8	001000	10	8	8	8	Q	X,8	101000	50	Q	Q	Q
9	9	001001	11	9	9	9	R	X,9	101001	51	R	R	R
#	8,2*	001010	12	9*	9	9	R*	X,8,2*	101010	52	R*	R*	R*
⊙	8,3	001011	13	=	⊙	=	\$	X,8,3	101011	53	\$	\$	\$
Space	8,4	001100	14	- (minus)	⊙	=	*	X,8,4	101100	54	*	*	*
&	Blank	001101	15	Blank	Blank	Blank		X,8,5*	101101	55	**	Blank*	Blank*
A	8,6*	001110	16	=*	&	&		X,8,6*	101110	56	**	Blank*	Blank*
B	8,7*	001111	17	- (minus)*	+ &	+ &		X,0	101111	57	**	Blank*	Blank*
C	R	010000	20	+	+ +	+ +	/	8,5*	110000	60	/	/	CR
D	R,1	010001	21	A	+ A	+ A	S	0,1	110001	61	S	S	S
E	R,2	010010	22	B	A B	A B	T	0,2	110010	62	T	T	T
F	R,3	010011	23	C	B C	B C	U	0,3	110011	63	U	U	U
G	R,4	010100	24	D	C D	C D	V	0,4	110100	64	V	V	V
H	R,5	010101	25	E	D E	D E	W	0,5	110101	65	W	W	W
I	R,6	010110	26	F	E F	E F	X	0,6	110110	66	X	X	X
⊠	R,7	010111	27	G	F G	F G	Y	0,7	110111	67	Y	Y	Y
	R,8	011000	30	H	G H	G H	Z	0,8	111000	70	Z	Z	Z
	R,9	011001	31	I	H I	H I	,	0,9	111001	71	Z*	Z*	Z*
	R,8,2*	011010	32	I*	I ;	I ;	,	0,8,2*	111010	72	,	,	,
	R,8,3	011011	33	.	. ;	. ;	%	0,8,3	111011	73	(((
	R,8,4	011100	34)) ;) ;	%	0,8,4	111100	74	(*	(*	(*
	R,8,5*	011101	35)*)* ;)* ;	%	0,8,5*	111101	75	,*	,*	,*
	R,8,6*	011110	36)*)* ;)* ;	%	0,8,6*	111110	76	1/2	1/2	1/2
	R,0	011111	37)*)* ;)* ;	%	0,8,7*	111111	77	Blank*	Blank*	Blank*

Notes: Key Punch: Use MLT PCH key to overpunch omitted characters.
 Card Code: * legal in illegal punch check Mode 2 only for card readers.
 Printer: * indicates symbol which will be printed by otherwise non-standard printer bit configuration.

Table I. Honeywell 800 Coding and Punched or Printed Equivalents

APPENDIX D. TABLES

		1	2	3	4	5	6	7	8	9	10	11	12	
GENERAL UNMASKED OR MASKED	BA	BINARY ADD								0	1	0	0	1
	DA	DECIMAL ADD								0	0	0	0	1
	WA	WORD ADD								0	1	1	0	1
	BS	BINARY SUBTRACT								1	1	0	0	1
	DS	DECIMAL SUBTRACT								1	0	0	0	1
	WD	WORD DIFFERENCE								1	0	0	0	1
	NA	INEQUALITY COMPARISON, ALPHABETIC								1	1	1	0	1
	NN	INEQUALITY COMPARISON, NUMERIC								0	1	1	0	0
	LA	LESS THAN OR EQUAL COMPARISON, ALPHABETIC								0	1	0	0	0
	LN	LESS THAN OR EQUAL COMPARISON, NUMERIC								1	1	1	0	0
	TX	TRANSFER (A) TO C								1	1	0	0	0
	TS	TRANSFER (A) TO B AND GO TO C								1	0	0	0	0
	HA	HALF ADD (MOD. 2)								0	0	1	0	0
	SM	SUPERIMPOSE								1	0	1	0	1
	CP	CHECK PARITY								0	0	1	0	1
	GENERAL UNMASKED	BM	BINARY MULTIPLY								0	1	0	1
DM		DECIMAL MULTIPLY	S	0	0	A	B	C	0	0	0	0	1	1
BT		BINARY ACCUMULATE								0	1	0	1	1
DT		DECIMAL ACCUMULATE	S	1	0	A	B	C	0	0	0	0	1	1
MT		MULTIPLE TRANSFER								1	0	0	0	0
TN		N-WORD TRANSFER								1	0	0	0	0
CC		COMPUTE ORTHOCOUNT								0	1	0	0	0
IT		ITEM TRANSFER								1	1	0	0	0
RT		RECORD TRANSFER								1	1	0	0	0
MPC		CONTROL PROGRAM	S	1	1	A	B	C	0	0	0	0	0	0
PR		PROCEED	S	1	0	A	B	C	0	0	0	0	0	0
INHERENT MASK		SWS	SHIFT WORD AND SUBSTITUTE (PROTECTED)								0	0	1	1
	SPS	SHIFT PRESERVING SIGN AND SUBSTITUTE (PROTECTED)								0	0	0	1	0
	SWE	SHIFT WORD AND EXTRACT (UNPROTECTED)	S	1	0	A	B	C	0	0	1	1	0	0
	SPE	SHIFT PRESERVING SIGN AND EXTRACT (UNPROTECTED)								0	1	0	1	0
	SSL	SHIFT AND SELECT (PROTECTED)								1	0	1	1	0
	SS	SUBSTITUTE (PROTECTED)								0	0	1	1	0
PERIPHERAL AND PRINT	EX	EXTRACT (UNPROTECTED)	S	0	0	A	B	C	0	0	1	1	0	0
	RF	READ FORWARD								1	1	0	1	0
	RB	READ BACKWARD								0	1	0	1	0
	WF	WRITE FORWARD								0	1	1	1	0
	RW	REWIND								0	0	0	1	0
	PRA	PRINT ALPHA								0	0	0	1	0
	PRD	PRINT DECIMAL								0	0	0	1	0
PRO	PRINT OCTAL	S	X	X	A	B	C	1	0	0	1	1	0	
SIMU- LATOR	S	SIMULATOR												
		Direct								0	X	X	X	X
	Indexed									1	X	X	X	X
	Notes:	PPPPPP = PERIPHERAL ADDRESS												
		S = SEQUENCE, COSEQUENCE CODE												
		A, B, C, = MEMORY DESIGNATOR												
		MMMMM = MASK ADDRESS												
		X = IRRELEVANT												

Table II. Honeywell 800 Command Codes

INDEX

Accumulator	46, 49, 62, 73, 83, 106, 115, 125
Addressing	22, 27, 29, 33
Direct Memory Location Address	34, 44, 53, 77, 79, 90, 113, 116
Direct Special Register Address	35, 44, 77, 136
Inactive Address	46, 62, 64, 74, 78, 83, 93, 96, 97, 103, 106, 112, 115, 121, 123, 136
Indexed Indirect Memory Location Address	42, 44, 77, 79, 90, 117, 136
Indexed Memory Location Address	37, 44, 77, 79, 90, 113, 117, 136
Indexed Special Register Address	39, 44, 77, 136
Indirect Memory Location Address	40, 44, 77, 79, 90, 117, 136
Stopper Address	45
Algebraic Compiler	1
ARGUS (Automatic Routine Generating and Updating System)	1, 26, 30, 35, 39, 42 47, 50, 87, 92, 107, 112, 114
Arithmetic Unit	5
AU-CU Counters Numbers 1, 2 (see Special Registers)	
Binary Accumulate, BT	68, 121, 136
Binary Add, BA	64, 119, 136
Binary-Coded Decimal	3
Binary Multiply, BM	70, 120, 136
Binary Subtract, BS	66, 119, 136
Binary System	3
Bisequence Bits	109
Bus	125
Card Conversion Modes	7
Card Punches (824-1) (824-2)	10
Card Readers (823-1) (823-2)	7, 9, 23
Card Reader Control Word	9
Card Reading Error Options	9
Central Processor (801)	4, 16, 18, 24, 43
Check Parity, CP	116, 120, 134, 141
Codes (Internal, Punch, Printer, Console).	143
Command Code.	27, 144
Compute Orthocount, CC	114, 121, 130, 140
Console	12, 20
Console Fixed - Start Bits	109
Continuous Initial Segment	100
Control Group Indicator Bits	108
Control Memory (see Special Registers)	
Control Program, MPC	122, 141
Control Unit	4, 22
Cosequence Counter (see Special Registers)	
Cosequence History Register (see Special Registers)	

INDEX (cont.)

Data Words	24
Decimal Accumulate, DT	70, 121, 136
Decimal Add, DA	66, 119, 136
Decimal Multiply, DM	71, 120, 137
Decimal Subtract, DS	67, 119, 136
Decimal System	2
Direct Memory Location Address (see Addressing)	
Direct Special Register Address (see Addressing)	
Distributed Read Address Counter (see Special Registers)	
Distributed Write Address Counter (see Special Registers)	
End-of-Item Word	31, 82, 96, 116, 121
End-of-Record Word	31, 81, 99, 114, 121, 131
End of Tape	95, 123
Executive Routine	1
Extract, EX	72, 122, 139
FACT (Fully Automatic Compiling Technique)	1
Floating-Point Option	5
Gate	63, 73, 87, 128
General Purpose Registers (see Special Registers)	
Half Add, HA	75, 120, 136
Hexadecimal System	3
Hex Digits (In Decimal Numbers)	63
High-Speed Printer (822-3)	11
Hunting for Program Demands (see Multiprogram Control)	
Inactive Address (see Addressing)	
Indexed Indirect Memory Location Address (see Addressing)	
Indexed Memory Location Address (see Addressing)	
Indexed Special Register Address (see Addressing)	
Index Registers (see Special Registers)	
Indirect Memory Location Address (see Addressing)	
Inequality Comparison, Alphabetic, NA	84, 119, 137
Inequality Comparison, Numeric, NN	85, 119, 137
Input Buffer Interlock Bits	94, 108
Instruction Words	26
Item Transfer, IT	82, 121, 139
Less Than or Equal Comparison, Alphabetic, LA	85, 120, 137
Less Than or Equal Comparison, Numeric, LN	85, 120, 137
"Logical AND" Function (see Extract)	
"Logical Exclusive OR" Function (see Half Add)	
"Logical Inclusive OR" Function (see Superimpose)	
Low-Order Product Register	46, 62, 64, 73, 87

INDEX (cont.)

Main Memory Address	33
Mask Index Register (see Special Registers)	
Mask Register	46, 73, 115
Modulo-3 Check	21, 45, 70
Multiple Transfer, MT	80, 121, 138
Multiprogram Control (Hunting)	13, 17, 64, 70, 83, 93, 107, 109, 112, 122
N-Word Transfer, TN	79, 121, 138
Octal System	3
Off-Line Auxiliary Control (815, 816, 817)	12, 21
Operation Code	27, 119, 144
Orthotronic Control	20, 114, 130
Orthowords	21, 101, 114, 130
Output Buffer Interlock Bits	99, 109
Overflow	63, 65, 68, 119, 127
Parity Bit	9, 21, 107, 116, 120, 130
Parity Error Flip-Flop	95, 101, 123
Partial-Product Locations (Multiplication).	43
Peripheral Control Units	17, 94, 100
Peripheral Fixed-Start Bits	108
Positional Notation	2
Print, PRA, PRD, PRO	105, 122, 140
Proceed, PR	112, 122, 141
Program Control Register	94, 98, 108
Program Demand Bits	109
Punch Control Unit	10
Punch Control Word	10
Read Address Counter (see Special Registers)	
Read Backward, RB	98, 122, 140
Read Forward, RF	94, 122, 140
Record Transfer, RT	81, 121, 139
Rewind, RW	103, 123, 140
Sequence Counter (see Special Registers)	
Sequence History Register (see Special Registers)	
Shift Preserving Sign and Extract, SPE	89, 122, 139
Shift Preserving Sign and Substitute, SPS	88, 122, 139
Shift Word and Extract, SWE	89, 122, 139
Shift Word and Select, SSL	90, 122, 139
Shift Word and Substitute, SWS	89, 122, 139
Sign Flip-Flop (SGA)	62, 127
Simulator, S	27, 30, 113, 122, 141

INDEX (cont.)

Special Registers	5, 16, 23, 33, 49
AU-CU Counters, Numbers 1, 2	50, 70, 77, 113, 115
Cosequence Counter	50, 53, 61, 77, 109, 113, 124
Cosequence History Register	50, 55, 113
Distributed Read Address Counter	50, 58, 96, 123
Distributed Write Address Counter	50, 58, 102, 123
General Purpose Registers	50, 58, 94, 96
Index Registers	37, 49, 55
Mask Index Register	29, 50, 56
Read Address Counter	50, 58, 94, 96, 98, 123
Sequence Counter	50, 53, 61, 77, 83, 93, 105, 109, 116
Sequence History Register	50, 55
Unprogrammed Transfer Register	50, 60, 95
Write Address Counter	50, 58, 99, 102, 123
Special Register Address	33
Special Register Words	26
Standard-Speed Printers (822-1) (822-2)	10
Stopper Address (see Addressing)	
Subroutine Library	1
Substitute, SS	74, 122, 139
Superimpose, SM	75, 120, 136
Tape Control Unit (803)	5, 16, 94, 100
Tape Protection	6, 101
Tape Unit (804)	5, 92
Traffic Control	13, 18, 94, 100
Transfer A to B and Go to C, TS	78, 93, 106, 117, 120, 137
Transfer A to C, TX	78, 120, 137
Typewriter Buffers	43, 106
Typewriter, Console	12, 105
Typewriter, Inquiry Station	12, 105
Typewriter, Slave	12, 105
Unprogrammed Transfer Register (see Special Registers)	
Unprogrammed Transfers	60, 63, 65, 95, 98, 104, 103, 108, 119, 123
Variable-Length Recording	6
Vertical-Format Word	10
Word Add, WA	67, 119, 136
Word Difference, WD	68, 119, 136
Word Structure	24, 31
Write Address Counter (see Special Registers)	
Write Forward, WF	99, 122, 140

Honeywell



Electronic Data Processing