

MEMORANDUM

TO: Former Members of Ad Hoc Group to Evaluate HARVEST

FROM: J. M. Willard, ANEQ-1

SUBJECT: Changes in the Logical Design of HARVEST

ENCLOSURES: Logical Design of HARVEST (Streaming Mode)
10 sheets.

1. I.B.M. has informally submitted a new logical design for the Streaming Mode of HARVEST. This new design reflects the recommendations of the Ad Hoc Group to Evaluate HARVEST.

The Table Address Assembler now contains an adder so that bytes may be added to the base address. A large number of inputs have been specified for the counter and the Statistical Accumulator. More conditions for reset of these units are available. The data gates have been redefined so as to be made more flexible. Operations with groups and records may now be defined in connection with the LU. There are now nine possible inputs to the adjuster and many more adjustment functions. Several new instructions have been defined, including one to read out the contents of the counters. Your comments on these changes will be appreciated.

2. The Streaming Mode has become a distinct mode of operation of HARVEST with its own set of instructions and its own Indicator Register. A "branch to Y and Enter Streaming Instruction" is used to enter the Streaming Mode from the Arithmetic Mode. The return to the Arithmetic Mode is accomplished by a bit (38) in the Stream Instruction format (table g sheet 8). The Arithmetic Instruction is in the process of change (line 1 of sheet 8). The most important part of the change is the reduction of the Operand Address to 24 bits (18 bit Word Address) with a sign bit. The main reason for this reduction is to include a shift in the arithmetic operations. The OP Code has been increased to 8 bits. The sign bit of the Operand Address is designed to completely generalize relative coding. There are sixty one instructions in the stream mode. The first 16 instructions (sheet 7) are the 16 logical connectives of the former Stream Logic Zero Instruction. The next 16 instructions are the 16 logical connectives of the former Stream Logic Parity instruction. The following 16 instructions are 16 options of the former Stream Select Instruction. Two add modular and two subtract modular instructions are included. The additional instructions are housekeeping instructions, except for the Collate Instruction which include all of the former Merge and Search Instructions. The housekeeping instructions are designed to keep the machine in the Streaming Mode.

3. The Stream Instruction format has 16 data gating bits including one spare (sheets 1 and 8). These gates have been renumbered, and, in the case of the output of the LU, redefined. The Statistical counter has been eliminated and its function transferred to a single CTR. The Supplementary Adjust field is one of 9 adjust fields. These will be discussed below. The U field is an enable on the three one-bit outputs of the LU. The F flip flop may receive its input from U or from Match Unit W. The F modes have been expanded in the same direction as previously defined. The GS field defines the group size as a function of a particular end of level indicator. The modulus field is 8 bits. SKIP means, I believe, the number of instructions skipped over after completion of this instruction. The C field defines whether or not: (1) the words in memory are to be brought into C and selectively replaced (VFL storage); or (2) complete replacement of the word (some speed advantage may be gained in the latter case). The other bit of the C field determines whether or not C is stored at the end of the instruction. The Indirect Load and Store Instruction is an expansion of the former Stream Indirect Store Instruction. The Store and Clear Instruction is used to store the contents of the SACC and CTR. Either or both may be cleared. The reading out of the counters has not been made a function of adjust operations as has been suggested. The Branch Instruction includes the ability to set a special bit in order to remember a branch out of the Stream Mode (in case of Interrupt). The Conditional Branch Instruction has the ability to interrogate and set any bit in memory. (This is a powerful housekeeping tool for multi-programmed operations. It makes possible the expansion of the Indicator Register). The Conditional Adjust Instruction also addresses any bit in memory and contains 2 of the 9 adjust fields. (This makes possible special adjustments during an interruption dependent on pre-set indicators). The Clear Memory Instruction now clears a complete table in 4 memories, instead of requiring four separate instructions. The Collate Instruction includes the option of automatically running out the last records in a file.

4. The Statistical Accumulator is now 24 bits plus sign to correspond to the shortened Operand Address. The Counter is 16 bits. The SACC and CTR are set up by word 18. The Threshold for SACC is set up by word 17. Also in word 18 are the count and reset controls for SACC and CTR. The SACC STEP field allows two to be added to the SACC under any one of 15 possible conditions (sheet 5). This count is in addition to any bytes that may be entered. SACC may be reset under any one of 7 conditions. The CTR may be stepped on any one of 60 conditions (63 are possible). The CTR may be reset on any one of 7 conditions.

5. The TAA indexing no longer counts the number of bytes. The length is determined by the End of First Level signals (\mathcal{L} , A, \mathcal{L} , B, * \mathcal{L} , E) from the Unit specified by the Data Gating. The bytes are added in to TAA in the following order:

* NOTE: The script \mathcal{L} is shorthand for End of Level signal.

- (1) Bytes from Source 1 are added to the base address.
- (2) Bytes from Source 2 are added to the base address.
- (3) MD shift (if any).
- (4) Count Displacement is Or-ed to the result.

6. The Stream Indexing has been changed (tentatively) (Sheets 3 and 6). Instead of the Length ($L = NI$), the number of times (N) is specified. Since the accumulated increment (P) must be stored in order to reset S (start address), this change has increased the size of the Stream Unit by about 12% or 2000 transistors. The Reset suppress bit has been replaced by a Mode (M) bit. This determines if the address is to be reset or not. If $M = 1$, the address is reset (left side of sheet 6). If $M = 0$, there is no reset. In this case P may be used as an increment between levels since it is not used to reset S . The V bit indicates that this is the highest level. In the higher level setup words, an E bit is used to enable the End of Higher Level Indicator bit in the Stream Indicator Register. There is no Read Enable bit.

7. There are now three 1 bit outputs from the LU. The meaning of these three bits (k , l and m) varies according to the instruction specified (sheet 7). For instance, for the comparison instructions, 32-47, a 1 from k means that the byte from A is greater than the byte from B ; a 1 from l means that the byte from A is equal to the byte from B ; and a 1 from m means that the byte from A is less than the byte from B . k , l or m may be counted by CTR . These three outputs also go to U and G . In U , each of the bits is gated by the corresponding bit in the U field of the Stream Instruction Format (line 2 of sheet 8). The outputs of the "and" gate are "or-ed" together to form a one bit output from U . Thus, if it is desired to signal when the byte from A is less than or equal to the byte from B , the bits in the Stream Instruction field U corresponding to l and n are set to 1. The one bit output of U may go to Gate 8 and/or F . The output of Gate 8 may go to C or $SACC$ thru Gate 9. The output of U may be counted in CTR and $SACC$. U sets indicators in the Stream Indicator Register.

8. The F Flip Flop is operated according to the F mode field in the Stream Instruction format (line 2 of sheet 8). The input may be from U or Match Unit W according to the "Fin" bit in the Stream Instruction format. The output of F may be counted in $SACC$ or CTR . It sets indicators in the Stream Indicator Register. It may be gated to C or $SACC$ thru gates 8 and 9. The F flip flop is reset to \emptyset at the end of a group.

9. The G unit has an output describing the characteristics of a group (sheet 10). The Group Size (GS) is specified by a 4 bit field in the Stream Instruction format. The group size is indicated by specifying a unit (ABC or E) and the end of level signal that specifies the end of the group. Hence, a group may be a byte (L_0), a group of characters or bytes (L_1), a record (L_2), or a file (L_3). The inputs to G are

the k, l, and m bits from IU. The output from G is an 8 bit byte which gives the result of group analysis. This byte may be gated to C thru gate 10. Since the information in the byte is redundant, it will probably not be gated to SACC. The Byte mask and addressing mechanism of C can remove the unwanted information from the byte. The first two bits of the byte may have state 0, 1, 2, 3. The 0 state is a waiting or reset state. The meaning of states 1, 2 and 3 depend upon the meaning of k, l and m, in the instruction specified. For instance, for the comparison instructions (32-47):

1. Means that the group from A is greater than the group from B
2. Means that the group from A is equal to the group from B
3. Means that the group from A is less than the group from B

It should be noted that 1 is the output for groups, that corresponds to the k output from IU for bytes. Likewise 2 corresponds to l and 3 corresponds to m. This carries over the meanings of k, l, & m for the Logical Connective Instructions (0-15 and 16-31). The transition matrices on sheet 10 are a formal method of presenting the technique of arriving at the 1, 2 or 3 output. Thus #3 is a matrix to describe the conditions for the Comparison Instructions and Mod. Subtract.

#3 Next Input

		k=1	l=1	m=1	
Present Status	0	1	2	3	New Status
	1	1	1	1	
	2	1	2	3	
	3	3	3	3	

The rows indicate the status of the output before the new bytes are entered into IU. The columns indicate the next input coming from IU. The result of the present status combining with the new input is located at the intersection of the row and column. The new status then becomes the present status for the next pair of bytes into IU. For matrix 3 (above) the present status at the beginning of a group is always 0. Then, if the first byte from A is greater than the first byte from B, k=1 and the new status is 1. If k=1 for the next pair of bytes, then Gp A is greater than Gp B (for 2 bytes) and the intersection of row 1 & k=1 is 1. Since the first byte from A was greater than the byte from B, it does not matter whether the remaining pairs of bytes are greater, equal or less than, so the rest of row 1 is filled with 1. Starting a new group, with G=D, if the first pair of bytes are equal (l=1), then G=2. With the new present status, G=2, then if the next pair of bytes are equal, the new

status of G is the same. However if $A > B$ ($k=1$) or $A < B$ ($m=1$), then $G=1$ or $G=3$ and group A is then either greater or less than group B and $G=1$ or $G=3$ at the end of the group. Again, at the beginning of a group, $G=0$, then if the first input is $m=1$ ($A < B$), the new status is $G=3$. The status of G will then remain at $G=3$ for the remainder of the group. Transition matrices #1 and #2 are evaluated in a similar manner. A transition matrix for the Mod. Add Instructions has not been made yet (I believe that #3 would be satisfactory). Bits 2 through 7 of the output byte contain redundant information about the status. For the Comparison Instructions, they are defined as follows:

Bit 2=1 if $G=1$ ($G_p A > G_p B$)
 Bit 3=1 if $G=1$ or $G=2$ ($G_p A \geq G_p B$)
 Bit 4=1 if $G=2$ ($G_p A = G_p B$)
 Bit 5=1 if $G=2$ or $G=3$ ($G_p A \leq G_p B$)
 Bit 6=1 if $G=3$ ($G_p A < G_p B$)
 Bit 7 =1 if $G=1$ or $G=3$ ($G_p A \neq G_p B$)

If one of these 6 bits is desired as an output, it may be selected by the C controls. The output $G=1, 2$ or 3 may be used to step the CTR either at the end of a group or at all intermediate states. They may step SACC at the end of the group. They set indicators in the Stream Indicator Register. G is reset after the end of a group.

10. There are nine possible inputs to the Adjuster. Six of these are from the Match Units (one from each Match Unit and two from combinations of Match Units.) In set-up words 16 and 17 are the fields for the 4 MU's. The "model MU" under word 18 shows the fields within the MU's. The W_4, X_4, Y_4 , or Z_4 , field of 6 bits describes the adjustment function desired on a match (Table on page 10). Each field (W_4, X_4, Y_4, Z_4) may be set independently. If it is desired to make a single adjustment when two MU's match simultaneously, there are two fields in word 17 (W-X & Y-Z) that may be set with the single adjustment desired for that pair. This adjustment replaces the two individual adjustments unless the adjustment function specified is equal to 1. In this case both of the single adjustments are carried out. The seventh adjustment field is located in the Stream Instruction format (line 2 of sheet 8). Here is a 10 bit field divided into a 4 bit field describing the stimulus and a 6 bit field describing the adjustment. The stimulus may be U, F, G, $awMU$ (note that in this way two independent adjustments may be made from a particular MU), End of Group, Stop and the end of first level signal from A, B or C (Table (b) sheet 8). The following two inputs to the adjuster are not used during streaming, these are two 6 Bit Adjust fields in the Conditional Adjust Instruction. Contingent upon the status of any specified bit in the memory one or two adjustment functions may be accomplished. (This may possibly be useful in resetting the SU's after an interruption as well as special conditions). There are 63 adjustments that can be defined by the 6 bit adjustment function field. These are described in the table on sheet 10. 50 of these options that have been defined (including "no OP") are single function adjustments. 11 of the defined adjustments (53-63) are multiple function adjustments, combining two or more of the adjustments as specified by the crosses in the columns opposite the single valued adjustment functions. The adjustments are divided into

groups according to the unit adjusted.

11. The CTR is stored and cleared through the use of the stream instruction Store and Clear in conjunction with an adjustment function and the skip feature. Assume that through adjustment function 2 the machine executes the instruction Store and Clear. ~~Should~~ the latter be located in the memory location just after the stream instruction which caused the adjustment, a skip of -1 would return control to the original streaming process. Obviously, the original stream instruction must employ a skip to avoid an endless loop after streaming is completed.

J. M. WILLARD
ANEQ-1