

Section 2

The SDS 910-9300 series, a planned family

The Scientific Data System 900-9000 series consists of the SDS 910, 920, 925, 930, 940, 945, and 9300 computers. The series includes capabilities and features found in most 24-bit machines. The design implementation is among the best for 24-bit machines, as measured by equipment utilization, the processor state, implementation technology, and ease of use.

The first delivery dates for the members of the series are 910 (August, 1962), 920 (September, 1962), 925 (February, 1965), 930 (June, 1964), 940 (April, 1966), 945 (~1968), and 9300 (December, 1964).

The 910 and 920 were designed at the same time as a planned series of compatible computers which spanned a range of performance. The 910 has instructions which facilitate defining 920 instructions by software. For example, these include the multiply and divide step¹ (see page 544) instructions in the 910 for programming the multiply and divide instruction in the 920.

The I/O facility evolved to a clean structure, with the potential for having a high degree of T and Ms data-transfer concurrency at a comparatively low cost. The IBM 7094 should be studied for a contrasting (more expensive) approach.

The instructions which help manipulate floating-point data are interesting and useful. The machine's ability to execute closed floating-point arithmetic subroutines is fairly good considering that the instructions are not hardwired.

The Programmed Operator (POP) instructions provide the ability to define an instruction set for efficient encoding. The idea appeared earlier in Atlas. However, the POP instruction calls subprograms in primary memory, instead of in fixed memory like Atlas.

A nice scheme¹ is described for increasing the memory address space from 16,384 to 32,768 words. Other schemes which switch memory banks, like those in the PDP-8 (Chap. 5)

¹We believe this appeared originally in the DEC PDP-1 introduced in November, 1960.

and in the 65,384-word 7094 II (Chap. 41), tend to be less desirable and flexible.

The SDS 930 was used at the University of California (Berkeley) as the base machine for the design of the Berkeley Time Sharing System (Chap. 24). SDS later marketed the system as the SDS 940.

The 9300 was not a member of the original 910-930 series. There is almost symbolic language program compatibility. Several registers and extra memory transfer paths were added to form the 9300 from the 930. The power of the 9300 is only a factor of 2 times the 930 for simple instructions. However, the hardwired floating-point instructions in the 9300 increases the power over the 930 by a factor of almost 10 for arithmetic problems. It is hard to believe that the incompatible 9300 was a wise choice. (We suggest a more reasonable alternative could have been a two-processor 930'. The 930' processor would be a 930 but with hardwired floating-point arithmetic instructions.) The 9300 has interesting twin-mode instructions for simultaneously operating on 12-bit data pairs. The 24-bit fixed-point word is sufficient for the real-time applications for which the computer was designed.

A flaw in the series is the sharing of K's among peripheral T's and Ms's. This problem can be seen by looking at the PMS structure (Chap. 42, Fig. 2, page 546). The connection to the peripheral K from K('Channel) requires a continuous connection during the data-transfer dialogue to Mp. This structure is especially bad in the case of a slow T, for example, a typewriter. A single character transmission requires that K('W, 'Y) be assigned to the typewriter during the complete message transmission (at a connected time of 100 milliseconds/character). The problem can be avoided by placing a character memory in each slow KT. Multiple devices could then run concurrently without requiring the elaborate K('W, 'Y) to be attached to them. The structure does not preclude such an improvement.

A complete description of the input/output and interrupt system is given and should be read carefully.

Chapter 42

The SDS¹ 910-9300 series

Introduction

The SDS 910, 920, 925, and 930 form a compatible series of computers. The 9300, though not compatible with the series, was an outgrowth of it. The 9300 uses the Ms and T devices of the 930. The 940 was designed initially at the University of California, Berkeley (see Chap. 24) for time sharing, and the 945 is a successor to the 940. The word length is 24 bits, and one single address instruction is encoded per word. The state of the machine consists of Mp(2048 ~ 32768 w) and Mps('P/Program Counter, 'A/Accumulator, 'B/Extended Accumulator, 'X/Index register).

These computers have been designed to process data originating from physical processes in real time. This design goal leads to a priority interrupt system with many (1,024) levels. The multiple interrupts facilitate programming and decrease the interrupt response time. A 24-bit word or two 12-bit words are a reasonable size for the problem types encountered. A multiple of 6 bits was chosen because of the (then) standard 6-bit magnetic-tape character. The relatively efficient storage representation and processing of floating-point data allow these computers to be used for general-purpose computation. However, only the 9300 has built-in floating-point operations. The 9300 has extensive capability for more general-purpose use. It is also used for operations on half-length data.

The data types processed by the 910-930 include words, integers, addresses, and boolean vectors. Several special instructions aid processing of types floating-point and double-length integers. The 9300 processes the additional data-types single- and double-length floating point. The 9300 has twin-mode instructions which operate on two half-length data (12 b) simultaneously. The two's complement representation is used for negative numbers.

The multiply, divide, and several other instructions are not wired into the 910, and compatibility between the 910 and 920-930 cannot be completely obtained by programming, although the 910 is a subset of the 920-930. Likewise, a smaller minimum Mp is available on the 910 (2,048 word versus 4,096 word). The 920 and 930 have identical instruction sets and differ in memory and logic performance. The 930 has a t.cycle: 1.75 μ s, and the 910-920 has t.cycle: 8 μ s. The more elaborate PMS structure of the 930 allows for greater growth, (e.g., by having more access ports to Mp).

The 9300's instruction set is different from the 930's. There are three index registers. The PMS structure is similar (and nearly compatible) with the 930. There are more (and better) working registers in the 9300 Pc to increase performance. The 9300 has two memory-access links, and the Pc can fetch instructions and data simultaneously. The instructions in the various C's appear in Table I for comparison purposes.

The SDS 925, a 1.75- μ s version of the SDS 910, was available only for a brief time and will not be discussed further.

The machines process instructions (operations to the accumulator) in the following times (microseconds):

<i>Instruction</i>	<i>910</i>	<i>920</i>	<i>930</i>	<i>9300</i>
Fixed-Point Add	16	16	3.5	1.75
Fixed-Point Multiply	248	32	7.0	7.0
Floating-Point Add	896	384	92	14.0
Floating-Point Multiply	1696	656	147	12.25

Structure

The structure of these computers is given with PMS and conventional diagrams in Figs. 1 to 4.

The SDS channel is a Kio('Channel) and not a Pio, since it has no program counter and uses Pc. However, it can be as effective as a Pio. Of course, the cost is lower since Pc is shared. If K('W, 'Y) requires memory accesses, they must wait until suitable times in the Pc instruction-interpretation process to communicate with memory (Fig. 1).

The PMS structural detail (Fig. 2) does not show the algorithm by which simultaneous Kio('W, 'Y, 'C, 'D) and Pc requests for Mp are resolved. K has the highest priority, and further resolution among K's is determined by the K with the fullest buffer memory. Thus the priority is variable.

There are three basic K types, or channels (Fig. 2), in the 930 and 9300:

- 1 K('Time Multiplexed Communications Channel/TMCC)
- 2 K('Direct Access Communications Channel/DACC)
- 3 K('Data Subchannel/DSC)

¹Scientific Data Systems merged with Xerox Corporation in 1969. The divisional name became Xerox Data Systems (XDS).

Table 1 SDS 910, 920, 930 and 9300 instruction sets¹

Mnemonic		Name	Mnemonic		Name
LOAD/STORE			REGISTER CHANGE		
+LDA	M, T	Load A	RCH	M, T	Register Change
+STA	M, T	Store A	AXB	M, T	Address to Index Base
+LDB	M, T	Load B	‡°CLA		Clear A*
+STB	M, T	Store B	‡°CLB		Clear B*
LDP	M, T	Load Double Precision	°CLR		Clear AB
STD	M, T	Store Double Precision	‡°CAB		Copy A into B*
LDS	M, T	Load Selective (Masked)	°ABC		Copy A into B, Clear A
STS	M, T	Store Selective (Masked)	‡°CBA		Copy B into A*
+LDX	M, T	Load Index X	°BAC		Copy B into A, Clear B
+STX	M, T	Store Index X	°XAB		Exchange A and B
+EAX	M, T	Copy Effective address into Index Register 1	‡°CBX		Copy B into Index*
STZ	M	Store Zero	‡°CXB		Copy Index into B*
+XMA	M, T	Exchange M and A	‡°XXB		Exchange Index and B*
XMB	M, T	Exchange M and B	‡°STE		Store Exponent*
XXM	M, T	Exchange M and Index Register	‡°LDE		Load Exponent*
			‡°XEE		Exchange Exponents*
			‡°CXA		Copy Index into A*
			‡°CAX		Copy A into Index*
			°XXA		Exchange Index and A*
			‡°CNA		Copy Negative into A*
			°CLX		Clear X
			COPY		Copy
ARITHMETIC			BRANCH		
+ADD	M, T	Add M to A	+BRU	M, T	Branch Unconditionally
DPA	M, T	Double Precision Add	+BRX	M, T	Increase Index and Branch
+SUB	M, T	Subtract M from A	+BRM	M, T	Mark Place and Branch
DPS	M, T	Double Precision Subtract	BRC	M, T	Branch and Clear Interrupt
MPO	M, T	M Plus One	BMA	M, T	Branch and Mark Place or Argument Address
°MIN	M, T	M Increment (M + 1)	+BRR	M, T	Return Branch
MPT	M, T	M Plus Two			
‡+ADM	M, T	Add to Memory			
‡+MUL	M, T	Multiply			
‡+DIV	M, T	Divide			
TMU	M, T	Twin Multiply			
DPN	M, T	Double Precision Negate			
×MUS	M, T	Multiply Step			
×DIS	M, T	Divide Step			
‡°SUC	M, T	Subtract with Carry			
‡°ADC	M, T	Add with Carry*			
×°MDE	M, T	M Decrement			
ARITHMETIC, FLOATING-POINT (OPTIONAL)			TEST/SKIP		
FLA	M, T	Floating Add	‡+SKE	M, T	Skip if A Equals M
FLS	M, T	Floating Subtract	+SKG	M, T	Skip if A Greater than M
FLM	M, T	Floating Multiply	SKL	M, T	Skip if A Less than M
FLD	M, T	Floating Divide	+SKM	M, T	Skip if A equals M on B Mask
			SKU	M, T	Skip if A Unequal M
			SKQ	M, T	Skip if Masked Quantity in A Greater than M
			SKF	M, T	Skip if Floating Exponent in B is Greater than or Equal
			+SKA	M, T	Skip if A and M do not Compare Ones Anywhere
			‡+SKB	M, T	Skip if B and M do Compare Ones Anywhere
			+SKN	M, T	Skip if M is Negative
			‡+SKR	M, T	Reduce M, Skip if Negative
LOGICAL					
+ETR	M, T	Extract			
+MRG	M, T	Merge			

Table 1 SDS 910, 920, 930 and 9300 instruction sets (Continued)

<i>Mnemonic</i>		<i>Name</i>	<i>Mnemonic</i>		<i>Name</i>
+EOR	M, T	Exclusive OR	SKP	M, T	Skip if Bit Sum Even
REGISTER SHIFT			+SKS	M, T	Skip if Signal Not Set
SHIFT	M, T	Shift	°SKD	M, T	Difference Exponents; Skip*
ARSA	N, T	Arithmetic Right Shift A	FLAG REGISTER		
ARSB	N, T	Arithmetic Right Shift B	FRTS	M	Flag Indicator Reset Test/Set
°RSH,		Arithmetic Right Shift AB	FLAG	M	Flag
ARSD	N, T;	Arithmetic Right Shift Double	FIRS	M	Flag Indicator Reset/Set
ARST	N, T	Arithmetic Right Shift Twin	FSTR	M	Flag Indicator Set Test/Reset
LRSA	N, T	Logical Right Shift A	FRST	M	Flag Indicator Reset/Set Test
LRSB	N, T	Logical Right Shift B	SWT	M	SENSE Switch Test
°LRSH	(930 only),	Logical Right Shift AB	INTERRUPTS		
LBSD	N, T;	Logical Right Shift Double	+EIR		Enable Interrupts
LRST	N, T	Logical Right Shift Twin	+DIR		Disable Interrupts
CRSA	N, T	Circular Right Shift A	+EIT		Interrupt Enabled Test
CRSB	N, T	Circular Right Shift B	+IDT		Interrupt Disabled Test
°RCY,		Circular Right Shift AB	+AIR		Arm Interrupts
CRSD	N, T;	Circular Right Shift Double	MEMORY EXTENSION (930 ONLY)		
CRST	N, T	Circular Right Shift Twin	°		Set Extension Register
°LSH;		Arithmetic Left Shift AB	°		Extension Register Test
ALSA	N, T	Arithmetic Left Shift A	BREAKPOINT TESTS (SENSE SWITCHES IN 9300)		
ALSB	N, T	Arithmetic Left Shift B	°BPT	4	Breakpoint No. 4 Test
ALSD	N, T	Arithmetic Left Shift Double	°BPT	3	Breakpoint No. 3 Test
ALST	N, T	Arithmetic Left Shift Twin	°BPT	2	Breakpoint No. 2 Test
LLSA	N, T	Logical Left Shift A	°BPT	1	Breakpoint No. 1 Test
LLSB	N, T	Logical Left Shift B	OVERFLOW (FLAG IN 9300)		
LLSD	N, T	Logical Left Shift Double	°ROV		Reset Overflow
LLST	N, T	Logical Left Shift Twin	°REO		Record Exponent Overflow
CLSA	N, T	Circular Left Shift A	°OVT		Overflow Test; Reset
CLSB	N, T	Circular Left Shift B	PROGRAMMED OPERATORS		
°LCY,		Circular Left Shift AB	°POP	M, T	Programmed Operator (64 instructions)
CLSD	N, T;	Circular Left Shift Double			
CLST	N, T	Circular Left Shift Twin			
NORA	N, T	Normalize A			
°NOD	N, T	Normalize; Decrement X			
NORD	N, T;	Normalize Double			
CONTROL					
+HLT		Halt			
+NOP	M, T	No Operation			
+EXU	M, T	Execute			
INR	M, T	Interpret			
REP	M, T	Repeat			

*M-Memory or Memory Address; N-number of shifts; T-tag field; +-also in the 910, 920 and 930; x-910 only; °-not in the 9300; †-not in the 910.

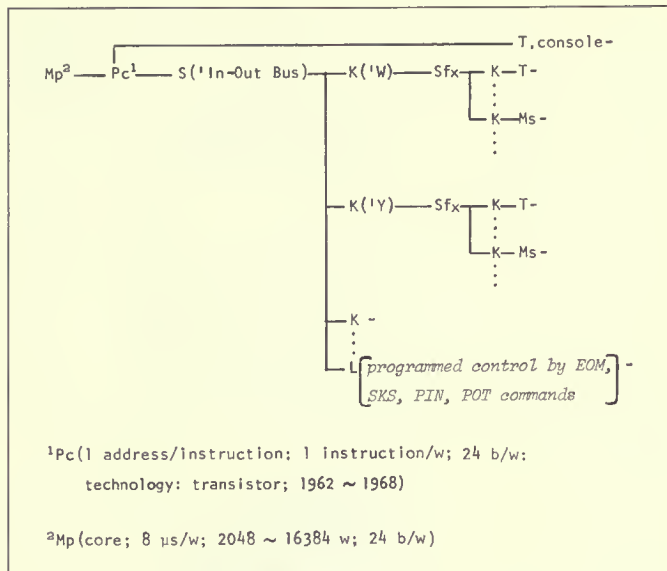


Fig. 1. SDS 910 and 920 PMS diagram.

The links between KT or KMs and any one of K('TMCC), K('DACC), and K('DSC) are identical. The KT or KMs assembles/disassembles characters into/from words and transmits/receives them to/from the Kio('Channel). The channel communicates with Mp or Pc for data transmission and finally communicates with Pc at task completion (the block of data transferred). Task alarms may cause Kio to interrupt Pc. Each Kio('Channel) can assemble data on a 6-, 12-, or 24-bit basis for Mp accesses. A K('Channel) recognizes two types of information: data being transmitted between Mp and the peripheral K, and initialization or controlling information from Pc.

In the 930 or 9300 K's the principal distinction is that the actual data-path switching routes differ. From a program operation and control viewpoint the Time Multiplexed and the Direct Access Communication Channels (TMCC and DACC) and the Data Subchannels (DSC) behave almost identically. The TMCC and DSC differ from DACC in that the block control information (number of words and location in memory) for the channel may be either in primary memory or in local hardware memory associated with the channel hardware.

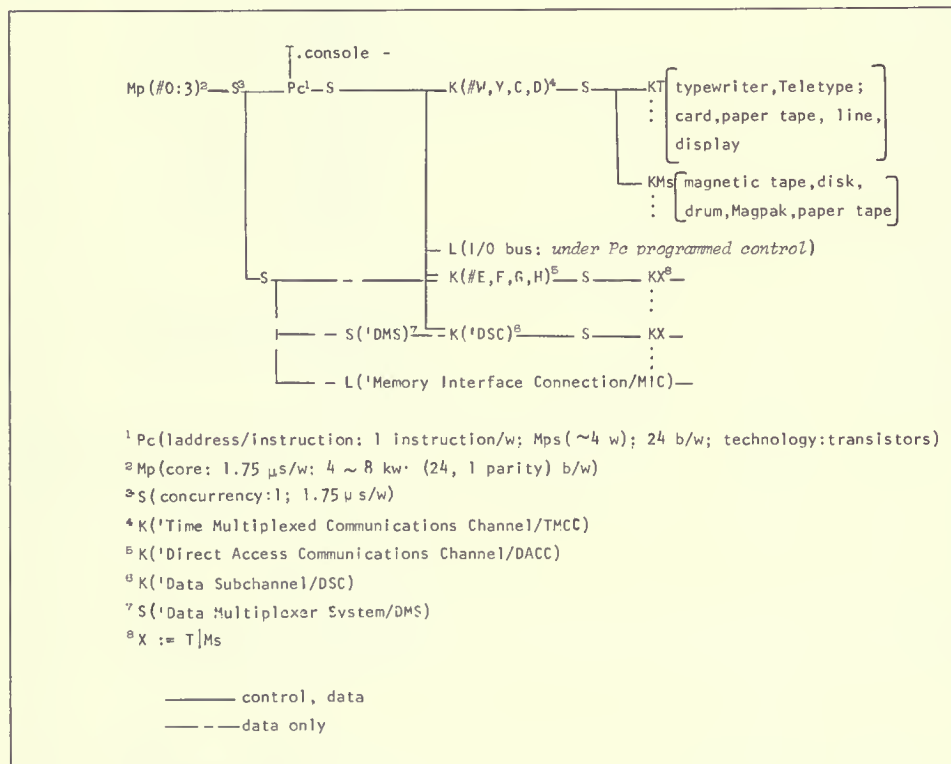


Fig. 2. SDS 930 PMS diagram.

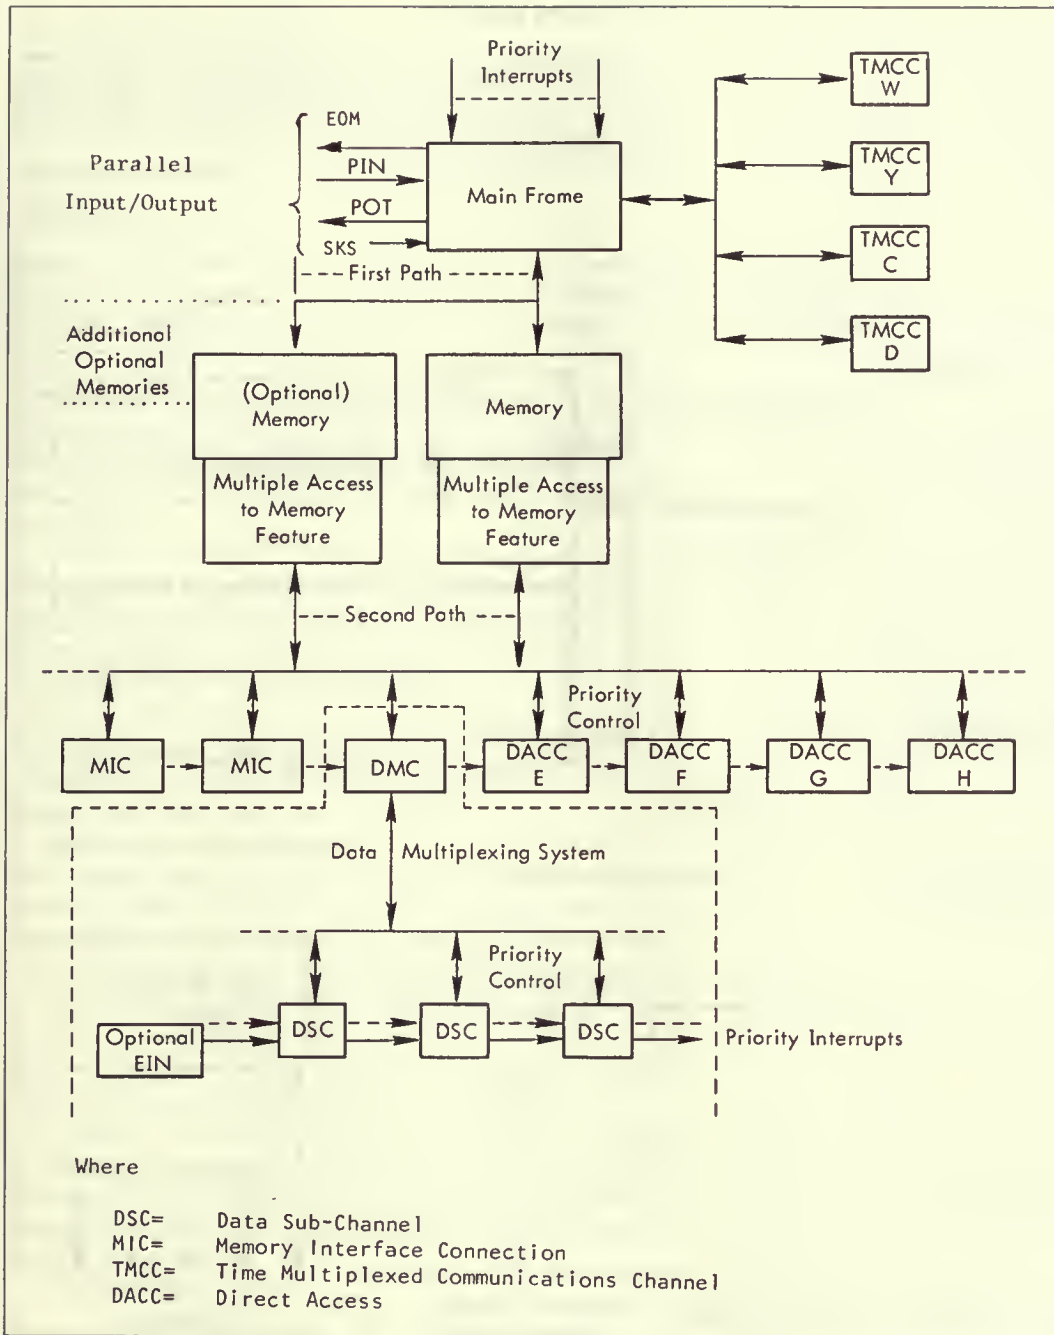


Fig. 3. SDS 390 computer-configuration diagram. (Courtesy of Scientific Data Systems.)

The 9300 structure, though not given in the PMS diagram, is essentially that of the 930 (Figs. 3 and 4). In the 9300, Mp has three access ports or a S(Memory-Processor; 8 Mp; 3 P,K). The Pc(9300) requires two of the access ports for independent access of instructions and data, leaving one for K transfer to Ms and T.

Instruction-set processor

The interesting parts of the ISP are discussed informally below. The formal ISP description given in Appendix I of this chapter should be read. The descriptions are partially taken from the SDS Programming Reference Manuals.

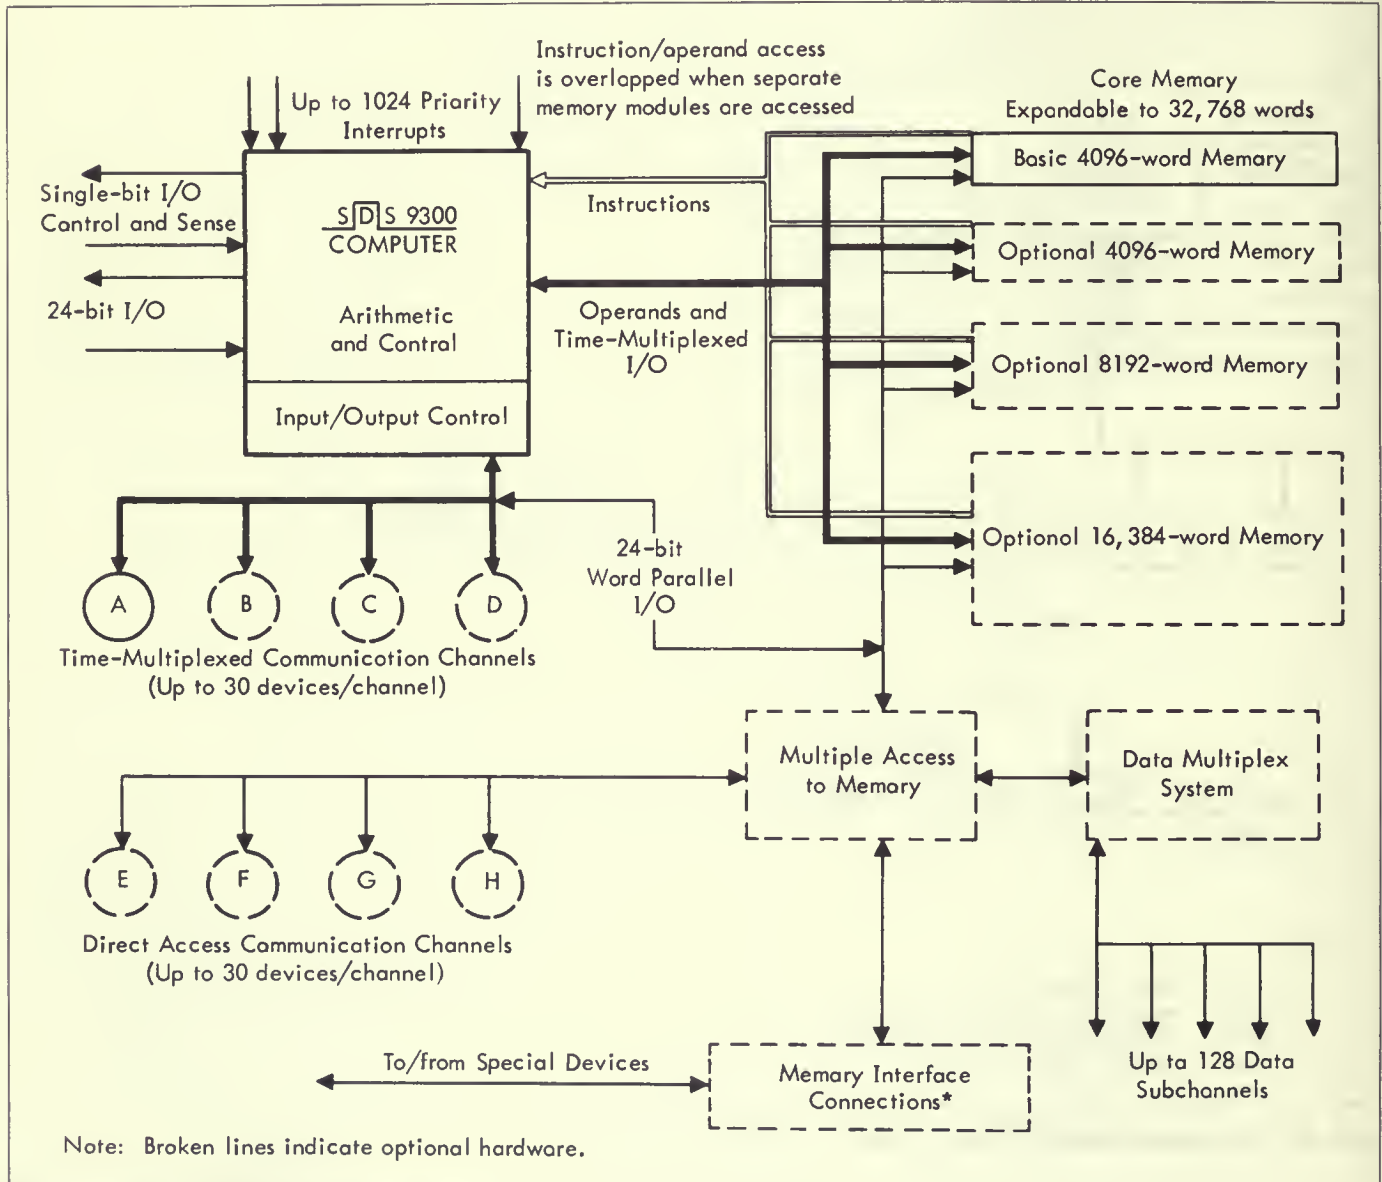


Fig. 4. SDS 9300 computer-configuration diagram. (Courtesy of Scientific Data Systems.)

Registers and memory (930)

The Pc state is declared in the ISP description. The ISP registers are A, B, X, P, M, and miscellaneous bits for overflow, carry, etc. Overflow can be turned on for arithmetic overflow in addition, subtraction, multiplication, division, and left-shift instructions.

Data formats

General. A computer word, W , is 24 binary digits (bits) or 8 octal digits. A word is numbered $W\langle 0:23 \rangle$ from left to right or alternatively $W\langle 0:7 \rangle_8$.

Fixed-point data format. Fixed-point numbers are represented in two's complement form with the sign at $W\langle 0 \rangle$. A 23-bit fraction $W\langle 1:23 \rangle$ can be assumed. The binary point is to the left of bit position 1 ($W\langle 1 \rangle$). For integers, the binary point is to the right of $W\langle 23 \rangle$.

Floating-point data format. Subroutines perform double- and single-precision floating-point arithmetic. A floating-point word is defined as $f\langle 0:47 \rangle := W[n:(n+1)]\langle 0:23 \rangle$. Of course, single-precision floating point requires less processing time.

The fractional portion (mantissa), $f\langle 0:38 \rangle$, of a double-precision floating-point number is a 39-bit proper fraction with the leading bit being the sign bit and the binary point located to the left of the most significant magnitude bit, $f\langle 1 \rangle$.

The floating-point exponent is a 9-bit integer, $f\langle 39:47 \rangle$, with the leading bit being the sign, $f\langle 39 \rangle$. The standard routines operate on both fraction and exponent in two's complement form. If F represents the contents of the fractional field and E represents the contents of the exponent field, the number has the form $F \times 2^E$.

Standard subroutines assume that the more significant word is in the A register and that the less significant word is in the B register. Correspondingly for M_p , the more significant word is in $M_p[x]$ and the least significant word in $M_p[x+1]$.

The single-precision floating-point representation is identical to that of double-precision floating point; i.e., it takes two words. However, the least significant bits of the mantissa, $f\langle 24:38 \rangle$, are not processed; thus there is a saving in time but not in space for using single precision.

Instruction word format (930)

The computer instruction word format is given in Fig. 5.

$W\langle 0 \rangle$ is the Relative Address bit, R. Standard software loading programs use this bit; central processor decoding logic does not use or sense this bit. A 1 in $W\langle 0 \rangle$ causes some loading programs




Fig. 5. SDS 930 instruction-format diagram.

to add the assigned location of the instruction to the address field contents prior to actual storage into the assigned location.

$W\langle 1 \rangle$ is the Index Register bit, X. It determines whether or not the index register will be added to calculate the effective address.

$W\langle 2:8 \rangle$ is the Instruction Code field and determines the operation to be performed. The Programmed Operator facility is selected by $W\langle 2 \rangle$; it is part of the Tag field $W\langle 0:2 \rangle$.

$W\langle 9 \rangle$ is the Indirect Address bit, I. It determines whether or not e or $M[e]$ is to be used as the effective address (see below).

$W\langle 10:23 \rangle$ is the Address field and for most instructions represents the location of the operand called for by the instruction code.

Address modification. Index and indirect addressing, used singly or in combination, perform address modification after bringing the instruction from memory but before executing it. The instruction remains in memory in its original form. The results of indexing and/or indirect addressing form the "effective address," e.

INDEXING If the content of the index bit in an instruction is a 1, prior to execution the computer adds the contents $X\langle 10:23 \rangle$, of the index register to the contents of the address field of the instruction. This addition does not keep any overflow or carry beyond the fourteenth address bit. This addition occurs prior to any indirect action.

INDIRECT ADDRESSING A 1 in the indirect address bit causes the computer to decode the contents of the effective address, accessed as described above, as if it were an instruction without an instruction code; that is, the address logic reinitiates address decoding, using the word in the effective location (the memory cell whose address is the effective address). This is an iterative process and provides multilevel indirect and indexed addressing. Each level of indirect addressing adds an additional cycle time to the instruction execution time.

930 memory extension control registers. Core memory in the 930 is expandable to 32,768 words. However, the address field in the

instruction format is 14 bits long, allowing direct access of only up to 16,384 words. Memory extension in the 930 contains two 3-bit memory extension registers, EM2 and EM3, and allows addressing of memories of 32,768 words. The program loads either or both of the registers and activates them as desired. Each register can become the most significant digit (fifth octal) of any operand address.

The program uses the first extension register, EM3, by calling for an address with an 11_2 in the most and next most significant address bits, respectively (a 3 for the most significant octal digit). The program calls for EM2, the second extension register, by setting the same two address bits to 10_2 (a 2 for the most significant octal digit). In this way, normal addressing compatible with the 910 and 920 occurs by setting a 3 in EM3, and a 2 in EM2.

910-930 instructions

Programmed Operators (POP's) enable subroutines to be called with a single instruction. This provides definable instructions of the same form as built-in machine instructions. The computer decodes the operation codes $100_8 \sim 177_8$ as special instructions and transfers to a subroutine whose address is uniquely determined by the code. The computer records the address of the POP instruction at location 0 together with an indirect address bit so that the program continuity may be maintained. By indirect addressing which refers to location 0, which in turn refers to the POP instruction, the subroutine can gain access to the effective address of the operand associated with the POP instruction.

The instruction set for the computers in this series is listed in Table 1. The table should be used to compare the machines.

There are two instructions in the 910 which are not in the 920 or 930: Multiply Step and Divide Step. These instructions facilitate writing subroutines for multiplication and division. The Multiply Step (MUS) instruction is defined:

$$\text{MUS} \rightarrow (\text{B}\langle 23 \rangle \rightarrow \text{A} \leftarrow \text{A} + \text{M}\langle e \rangle; \text{next AB} \leftarrow \text{AB}/2);$$

9300 instructions

The instruction word format in the central processor is shown in Fig. 6.




Fig. 6. SDS 9300 instruction-format diagram.

$W\langle 0 \rangle$ contains the Indirect Address bit I.

$W\langle 1:2 \rangle$ contains the Index Register bits $X\langle 0:1 \rangle$.

$W\langle 0:2 \rangle$ is called the Tag field.

$W\langle 3:8 \rangle$ contains the Instruction code; the contents of this field determine the operation to be performed.

$W\langle 9:23 \rangle$ contains the Address; for most instructions, the contents of this field represent the memory location of the operand called for by the instruction code.

Address modification. Each index register contains an unsigned base address of 15 magnitude bits and a signed increment of 9 bits. The increment contains 8 magnitude bits and a sign bit and is held in two's complement form.

Index registers are modified by adding the signed-increment value to the base address using two's complement arithmetic. Since the increment and base address fields are of unequal lengths, the sign bit (bit 0) of the increment field is extended six positions to the left prior to the addition. This 15-bit sum is then stored in the base address field of the index register. The index register may be incremented by any value from -256_{10} to 255_{10} using a single instruction. Incrementing and testing for a "terminal condition" is done by the instruction Increase Index And Branch (BRX), as follows:

If the index register has been negatively incremented, a terminal condition exists when the base address has been reduced below the zero value.

If the index register has been positively incremented, a terminal condition exists when the resultant base address has been increased beyond the maximum address value (077777_8).

If the terminal condition exists, the next instruction is taken in sequence. If the terminal condition does not exist, program control is transferred to the location specified.

The instruction set for the 9300 is given in Table 1.

Pc implementation

All the processors of the series have basically similar register configurations because of the common Instruction-set Processor. However, the increasing complexities of the machines can be seen by comparing the register structures of the 910-930 (Fig. 7) with the 9300 (Fig. 8). The figures show both the registers accessible to the program or defined by the ISP (denoted by *) and the temporary registers which are necessary for the implementation.

910, 920, 930 registers (Fig. 7)

ISP registers (*). The A register is the main accumulator of the computer. The B register is an extension of the A register. The




Fig. 7. SDS 910, 920, and 930 registers diagram.

B register contains the less significant portion of double-length numbers. Overflow and carry bits are used with A and B operations.

The index register X, used in address modification, is a full-word register. Index-register operations use the least significant 14 bits.

The P register is a 14-bit register that contains the memory address of the current instruction. Unless modified by the program, the contents of P increase by I at the completion of each instruction.

The memory extension registers, EM3 and EM2, are 3-bit registers that specify the portion of extended memory being used. They exist only in the 930.

Hardware registers not in the ISP. The S register is a 14-bit register that contains the address of the memory location to be accessed for instructions or data. The 15-bit address is formed by S and one of the memory extension registers.

The 24-bit C register communicates with memory. Instructions are temporarily held in C before instruction decoding. It is used as an arithmetic and control register in multiply, divide, and other operations. Address modification and parity generation/detection use the C register.

The O register is a 6-bit register that contains the instruction or operation code of the instruction being executed.

The M' register is a 24-bit register that holds each word as it comes from memory. Recopying of a word into memory takes place from the M' register.

9300 registers (Fig. 8)

ISP registers (*). The A and B registers of the 9300 are the same as in the 900 series computers; however, the P register is P<9:23>.

There are three 24-bit index registers, X[1:3]. Each index register is composed of a base address of 15 bits and a signed increment of 9 bits.

The Flag register, F, is a 6-bit register that may be set and/or sensed by the program. The first bit position of this register is the overflow indicator.

Hardware registers not in the ISP. The C register holds the 24-bit operand word as it is transmitted to, or received from, memory.

The D register holds the next 24-bit instruction word as it is received from memory.

The 15-bit S register contains the address of the memory location to be accessed for either instruction or operand.

The 6-bit O register contains the instruction code of the instruction being executed.

The A' register is an optional 15-bit register used for the floating-point option. It temporarily extends the A register during the execution of floating-point instructions.

The B' register is an optional 15-bit register which temporarily extends the B register during the execution of floating-point instructions.

Instruction interpretation in the 900 series

The instruction-interpretation process can be explained in terms of the processor's registers (Fig. 7). The ADD instruction execution (not including memory mapping) defined in ISP as $A \leftarrow A + M[e]$ is interpreted as

$S \leftarrow P$; $P \leftarrow P + 1$; next *fetch the instruction*

$M' \leftarrow \text{Memory}[S]$; next

$C \leftarrow M'$; next




Fig. 8. SDS 9300 registers diagram. (Courtesy of Scientific Data Systems.)

$O \leftarrow C\langle 0:5 \rangle$; next

$(O = 05) \rightarrow$ (

ADD execution

operand effective-address-calculation process (including indexing and indirect addressing)

$S \leftarrow C\langle 10:23 \rangle$; next

final operand fetch

$M' \leftarrow \text{Memory}[S]$; next

$C \leftarrow M'$; next

$A \leftarrow A + C$)

add operation

Input/output processing

Introduction

There are several methods of transferring data between M_p and the K 's. These methods will be described independently, and in

order of increasing complexity. They are:

- 1a Single bit sent to a selected K (EOM instruction).
- 1b Single bit sense (or bit detection) from a K (SKS instruction).
- 2 Word parallel to/from a K (POT/PIN instruction).
- 3 Interrupt from one of 1,024 K 's on a priority basis to P_c . K can signal P_c to execute a particular program.
- 4a Time Multiplexed Communication Channel/TMCC (Internal Interlace¹ feature).
- 4b Time Multiplexed Communication Channel (External Interlace²).
- 5 Direct Access Communication Channel/DACC (External Interlace).

¹The control information for the location of the next word transferred and the number of words to transfer are kept in M_p .

²The control information is taken from registers within K .

- 6a Data Subchannel/DSC (Internal Interlace).
- 6b Data Subchannel (External Interlace).
- 7 Memory Interface Connection/MIC link. A component has a link to Mp.

Methods 1 to 3 above are completely under control of a program and are simple time-independent instructions (or methods) of transferring data to K's (and onto KT or KMs). The ISP description (Appendix I of this chapter) has a detailed description of the I/O devices and these I/O instructions.

Single-bit control and sense

Two instructions provide for single-bit ON/OFF control signals. The first, EOM, transmits a control signal and a 14-bit address to an external device or a function within the computer. The second, SKS, selects an external device or computer function and skips in response to a false (0) signal. Up to 16,384 control signals can be sent and 16,384 input signals tested theoretically. (A more reasonable number of physical destinations would be 50.) Execution of an EOM causes a signal of approximately 1.4 microseconds duration to be transmitted.

EOM instruction format. EOM is used to select a specific I/O device by placing a 1 in its select register. EOM requires one cycle.

W<2> = 0.

W<0:1> is reserved for special system address bits.

W<3:8> contains the EOM instructions code, 02.

W<10:11> contains the system mode specifier.

W<12:23> contains the 12-bit address field that specifies the special system destinations.

SKS format. The SKS instruction format has each corresponding bit field identical to the system EOM format. Execution of an SKS causes a 14-bit address to be presented to all K's; the K being addressed responds and is tested. If the addressed external K supplies a "set" signal to the central processor, the computer executes the next instruction in sequence from the SKS. If no signal is set, the computer skips the next instruction in sequence and executes the following instruction. No registers are affected except the P register. SKS requires two or three Mp cycles if no skip or skip, respectively, is executed.

Word parallel instructions

Two instructions, Parallel Output (POT) and Parallel Input (PIN), permit any word in Mp to be presented in parallel on a physical connector to a K or, inversely, permit signals sent from a K to

be stored in Mp. The execution of a POT or PIN instruction sends a signal to the external device involved in the input/output operation, which notifies the device to send its data word as soon as it is operational. When the device becomes operational during a Read or PIN operation, it transmits a Ready signal to the central processor while at the same time presenting a data word to Pc.

During the execution of a POT instruction, the central processor transmits a signal to the external device, alerting it to receive a data word. When the device becomes operational, it transmits a Ready signal to the central processor, which releases the data word to the external device.

Selective input/output with these devices is accomplished by preceding POT or PIN with an EOM to alert (select) the desired device by a specific address. By preceding the POT or PIN with an SKS, the Ready signal of the special device can be tested after the execution of the EOM but prior to execution of the parallel transfer instruction; a possible Pc "hangup" can thus be avoided. The Ready signal can also set one of the priority interrupts.

PIN stores the contents of 24 input lines in parallel in the effective-memory location. PIN or POT requires four cycles plus any waiting time for Ready.

Interrupt

The interrupt provides program control of input/output operations, aids in programming simultaneous input/output and compute operations, and allows immediate recognition of special external conditions by causing Pc to execute an instruction in a selected Mp location at the end of the execution cycle of the current instruction. Without disturbing the program register, the processor executes an instruction in one of a selected set of memory locations. A Mark Place and Branch (BRM) instruction in this location saves the contents of the program register, EM3, EM2, and overflow indicator and transfers to the particular interrupt servicing routine required. To exit from the interrupt service routine, a Branch Unconditionally (BRU) instruction using indirect addressing returns control to the next instruction in proper sequence in the main program; it also clears the interrupt. Processor state (that is, A, B, Overflow, and X) must be preserved and restored by the program if the registers are used by the program.

The priority interrupt system has up to 1,024 interrupts arranged in levels. The levels have priority according to a priority number; the higher priority levels have a smaller number. Interrupt channels are installed in Pc in groups of 16. The assignment of physical memory locations to interrupt levels is shown in Appendix I of this chapter; the assignment is in order of decreasing priority from location 200₈ (highest) to 1477₈ (lowest). Interrupt requests can also be programmed. The power fail-safe (for power

supply off) interrupts and out-of-order interrupts have the highest priority.

Besides the interrupt mechanism just discussed, there is also a single instruction interrupt. This permits the execution of only one instruction before automatically being cleared and returning to the program that was interrupted. For example, if an external clock source is connected to the computer so that it pulses an interrupt line at set intervals, the program can maintain a programmed real-time clock. Each time the external pulse causes an interrupt, the program executes the single instruction, Memory Increment (MIN), to add 1 to the memory word selection for use as a programmed real-time clock. (The main program can examine this memory location whenever necessary to determine how many time increments have elapsed since the clock was started.)

Interrupts can be single or normal-instruction interrupts in any combination desired.

An interrupt has three operational states: inactive, waiting, and active states.

In the inactive state, no interrupt signal has been received into the level and none is currently being processed by its interrupt servicing subroutine.

In the waiting state, an interrupt has been received but is not being processed. This situation may arise when an interrupt of higher priority is being processed. When all higher waiting interrupts have been processed, this level goes to the active state.

In the active state, the interrupt has caused the main program to recognize its presence and has transferred to its assigned interrupt location where it is being processed.

Two program control features are Arm/Disarm and Enable/Disable. Arm/Disarm controls whether an interrupt can proceed from the inactive state to the waiting state. When armed, an interrupt signal sets the interrupt to the waiting state. Enable/




Fig. 9. SDS 930 direct-access communication-channel register diagram. (Courtesy of Scientific Data Systems.)

Disable operates on the entire interrupt system. (When the interrupt system is enabled, interrupts can occur.)

Communications channels—Kio('Channel)'s

Kio('Communication Channels) provide buffering, input/output control, and data transmission simultaneously with computation. There can be up to eight independent communication channels and a large number of subchannels in a single system. Figure 9 shows the registers in a K('Channel).

Each channel can control up to 30 KT's or KM's's. The channel handles character, word assembly and disassembly, input/output parity detection and generation, data transmission to and from memory, and end-of-transmission detection.

All channels are bidirectional and can communicate with 6-bit character devices or word devices in 6, 12, and 24 bits. The main program that initializes a K specifies the number of characters to be contained in each word during the transmission.

The channel interlace controls the transfer of the data words going through the associated channel buffer, supplies the memory address of data coming from or going to memory, and maintains the word count determining the number of words transferred. This interlace information can be either in K hardware (external interlace) or in Mp (internal interlace). The terminal interrupts, End of Record and Zero Word Count, come from the interlace and are under its control.

The time-multiplexed channels use the memory-access logic of Pc to transmit input and output of data words and require two memory cycles (see Fig. 2). Each direct-access channel has independent memory-access logic and requires one memory cycle (see Fig. 2).

Communication-channel description. Up to 30 peripheral devices (K's for T or Ms) may be connected to one K('Channel) (Fig. 9). Each device has a unique, 2-digit, octal address by which it is selected for an input/output operation. To select the peripheral device, the program loads the proper unit address into the 6-bit Unit Address Register (UAR) in the channel. This address selects both the device and, if appropriate, the function to be performed. Placing a nonzero unit address in the unit address register connects the peripheral unit addressed to the channel, and the unit becomes active. When the UAR contains a zero address, or any time that a terminal or initial condition clears the contents of UAR, the channel becomes inactive.

The 24-bit data Word Assembly Register (WAR) contains the data word actively being received or transmitted during an input or output operation. During input, 6-bit characters (plus parity)

enter the Single-Character Register (SCR) where the channel buffer assembles them, one at a time, into the WAR.

The channel interlace contains two working registers: the Word Count Register (WCR) and the Memory Address Register (MAR). A channel may have these registers either in K or in Mp. In the setup sequence for an interlaced input/output operation, the POT instruction transmits to the interlace a data word made up of the word count (that is, length) and the starting address of the data block. The 15-bit Word Count Register (WCR) contains the data word count during a data transfer. The number of data words is decremented by 1, and the new count replaces the old one in the WCR for each word transmitted.

The Memory Address Register (MAR) contains the starting destination or source address in memory of the transmitted data. The memory locations to or from which data words are to be transmitted enter the MAR at the same time the word count does. During transmission of data, the interlace increments the MAR after each word as it decrements the contents of the WCR. These two registers provide the interlace control of block transmissions. Obviously, if the interlace control registers are in Mp, then two extra accesses are required for each word transferred.

Memory interface connection link

Once a computer is equipped with a multiple-access-to-memory feature, one or more Memory Interface Connections (MIC) can be attached. The MIC is a general interface to the computer that allows special devices to access Mp. It preserves the integrity of the memory by generating the parity of incoming data words and checking the parity of words read from memory to indicate memory failures. The device that is connected to the MIC must hold both the data and the address until the transmission to/from memory is completed (that is, MIC does not have registers).

Conclusions

The SDS computers appear to be the first attempt to design several computers at the same time with a common ISP. Over a longer time span other compatible computers were added to the original 910 and 920 as technology (and marketing) dictated. The series is characteristic of well-designed typical 24-bit computers. By increasing the arithmetic capability, the series could also be used more generally.

References

Scientific Data Systems Reference Manuals for the 930 and 9300 computers

APPENDIX 1 SDS 930 ISP DESCRIPTION

Appendix 1

SDS 930 ISP Description

The description defines the Instruction Set without exact assignment of operation codes to instruction names. Input-output instruction actions are given for the simple controls, but do not include the action of the channels or the devices.

Pc State

A<0:23>	Accumulator; main arithmetic register
B<0:23>	secondary arithmetic register for multiplier, quotient, etc.
AB<0:47> := ACB	combined 48 bit arithmetic register
X<0:23>	Index Register
P<10:23>	Program or instruction location counter for 16 kw
Overflow/0v	set on integer operations
Carry := X<0>	used in multiple precision operations to link words
Run	

Mp State

Memory [0:7777 ₈] ₈ <0:23>	32 kw primary memory
Two 3 bit map (or extension) registers extend the address space of Mp to 32 kw. EM2 holds a 4 kw block number when addresses 20000 ₈ -27777 ₈ are used. EM3 holds the 4 kw block number for addresses 30000-37777 ₈ .	
EM2<0:2>	Extension Memory registers
EM3<0:2>	

Memory Mapping Process

This process maps the 16 kw address space into the 32 kw physical memory.

M<0:23>[a] := (
(a < 20000 ₈) → Memory [a] ₈ <0:23>	
(20000 ₈ ≤ a ≤ 27777 ₈) → Memory [EM2<0:2>□a<2:23>] ₈ <0:23>	
(30000 ₈ ≤ a) → Memory [EM3<0:2>□a<2:23>] ₈ <0:23>	

Pc Console State

Individual registers in Pc can be read and written from the console.

BPT<1:4>	Breakpoint or sense switches
----------	------------------------------

Instruction Format

instruction/i<0:23>	
relative := i<0>	unused by ISP; software relocation bit
index_bit/xb := i<1>	
op_code/op<2:8> := i<2:8>	
pop_code<0:5> := i<3:8>	programmed operation code value
indirect_bit/ib := i<9>	
y<10:23> := i<10:23>	address field for 16 kw
μ microcoded instruction bits within an instruction	

Effective Address Calculation Process

e<10:23> := (¬ ib → (iterative process of indefinite indirect addressing until no indirect bit, ib, is found
¬ xb → y;	
xb → y + X);	
ib → (
¬ xb → (i<0:23> ← M[y] ₈ <0:23>	
xb → (i<0:23> ← M[y + X] ₈ <0:23>); next e)	
e<18:23> := e<18:23>	shift count

Instruction Interpretation Process

<pre> Interrupt_interpretation → (Instruction ← M[P]; P ← P + 1; next Instruction_execution); Interrupt_interpretation → (Instruction ← M[200₈ + 20₈ × K₁address + I₁address]; next Instruction_execution) </pre>	<pre> normal interpretation fetch execute interrupt interpretation </pre>
---	---

Instruction Set and Instruction Execution Process

Instruction_execution := (

Load and Store Group

LDA → (A ← M[e]);	load A
STA → (M[e] ← A);	store A
LDB → (B ← M[e]);	load B
STB → (M[e] ← B);	store B
LDX → (X ← M[e]);	load index
STX → (M[e] ← X);	store index
EAX → (X ← e);	load index from e
XMA → (M[e] ← A; A ← M[e]);	exchange A and M

Arithmetic Group

SUB → (Ov, Carry ← A - M[e]);	subtract
ADD → (Ov, Carry ← A + M[e]);	add
SUC → (Ov, Carry ← A - M[e] - Carry);	subtract with Carry
ADC → (Ov, Carry ← A + M[e] + Carry);	add with Carry
MIN → (Ov, M[e] ← M[e] + 1);	memory increment
ADM → (Ov, M[e] ← M[e] + A);	add to memory
MUL → (Ov, AB ← A × M[e]);	multiply
DIV → (Ov, B ← AB/M[e]; A ← AB mod M[e]);	divide

Logical Group

ETR → (A ← A ∧ M[e]);	extract
MRG → (A ← A ∨ M[e]);	merge
EOR → (A ← A ⊕ M[e]);	exclusive or

Microcoded Register Exchange Instruction

Each instruction can be formed from a series of microprogrammed operations. Compound microcoded instructions are shown below without a μ .

CLA → (A ← 0);	μ , clear A
CLB → (B ← 0);	μ , clear B
CLR → (AB ← 0);	clear A and B
CLX → (X ← 0);	μ , clear X
CAB → (B ← A);	μ , copy A into B
CBA → (A ← B);	μ , copy B into A
XAB → (A ← B; B ← A);	exchange A and B
CXB → (B ← X);	μ , copy X into B
CBX → (X ← B);	μ , copy B into X
XXB → (X ← B; B ← X);	exchange X and B
CAX → (X ← A);	μ , copy A into X

CXA $\rightarrow (A \leftarrow X);$
 XXA $\rightarrow (A \leftarrow X; X \leftarrow A);$
 CNA $\rightarrow (A \leftarrow \neg A);$
 BAC $\rightarrow (A \leftarrow B; B \leftarrow 0);$
 ABC $\rightarrow (B \leftarrow A; A \leftarrow 0);$
 STE $\rightarrow (X<15:23> \leftarrow B<15:23>; X<0:14> \leftarrow \text{sign_extend}(B<15>));$
 $B<15:23> \leftarrow 0);$
 LDE $\rightarrow (B<15:23> \leftarrow X<15:23>);$
 XEE $\rightarrow (B<15:23> \leftarrow X<15:23>; X<15:23> \leftarrow B<15:23>;$
 $X<0:14> \leftarrow \text{sign_extend}(B<15>));$

End of microcoded instruction group

Shift Group

LRSH $\rightarrow (AB \leftarrow AB / 2^{e1} \{\text{logical}\});$
 RSH $\rightarrow (AB \leftarrow AB / 2^{e1});$
 RCY $\rightarrow (AB \leftarrow AB / 2^{e1} \{\text{rotate}\});$
 LSH $\rightarrow (0v, AB \leftarrow AB \times 2^{e1});$
 LCY $\rightarrow (AB \leftarrow AB \times 2^{e1} \{\text{rotate}\});$
 NOD $\rightarrow (X \leftarrow X - \text{normalize_exponent}(AB);$
 $AB \leftarrow \text{normalize}(AB));$

Skip Test Group

SKE $\rightarrow ((A = M[e]) \rightarrow (P \leftarrow P + 1));$
 SKB $\rightarrow ((M[e] \wedge B) = 0) \rightarrow (P \leftarrow P + 1);$
 SKN $\rightarrow (M[e]<0> \rightarrow (P \leftarrow P + 1));$
 SKR $\rightarrow (0v, M[e] \leftarrow M[e] - 1; \text{next } M[e]<0> \rightarrow (P \leftarrow P + 1));$
 SKM $\rightarrow ((M[e] \wedge B) = (A \wedge B)) \rightarrow (P \leftarrow P + 1);$
 SKG $\rightarrow (A > M[e]) \rightarrow (P \leftarrow P + 1);$
 SKO $\rightarrow (XR<0:23> \leftarrow \text{abs}(B<15:23> - M[e]<15:23>);$
 $(M[e]<15:23> > B<15:23>) \rightarrow (P \leftarrow P + 1));$
 SKA $\rightarrow ((M[e] \wedge A) = 0) \rightarrow (P \leftarrow P + 1);$

Branch Group

BRU $\rightarrow (P \leftarrow e);$
 BRX $\rightarrow (X \leftarrow X + 1; X<9> \rightarrow P \leftarrow e);$
 BRM $\rightarrow (M[e]<0> \leftarrow 0v; M[e]<3:5> \leftarrow EM3; M[e]<1,2,9> \leftarrow 0;$
 $M[e]<6:8> \leftarrow EM2; M[e]<10:23> \leftarrow P; \text{next}$
 $P \leftarrow e + 1);$
 BRR $\rightarrow (P \leftarrow M[e] + 1; 0v \leftarrow 0v \vee M[e]<0>);$

Control Group

HLT $\rightarrow (\text{Run} \leftarrow 0);$
 NOP $\rightarrow :$
 EXU $\rightarrow (\text{instruction} \leftarrow M[e];$
 $\text{instruction_execution});$

Overflow Test Group

OVT $\rightarrow (0v \rightarrow (P \leftarrow P + 1); (0v \leftarrow 0));$
 ROV $\rightarrow (0v \leftarrow 0);$
 REO $\rightarrow (X<14> \oplus X<15>) \rightarrow (0v \leftarrow 1);$

μ , copy X into A
 exchange X and A
 μ , not A
 copy B into A, clear B
 copy A into B, clear A
 μ , store exponent; exponent control bit

load exponent
 exchange exponent

logical right shift
 right shift
 right cycle
 left shift
 left cycle
 normalize, decrease X

skip if A = M
 skip if B and M don't compare 1's
 skip if M negative
 reduce M, skip < 0
 skip on masked M
 skip if greater than M
 difference exponents and skip
 skip if A and M don't compare 1's

branch unconditionally
 increment Index, Branch
 mark place and branch
 used to call subroutines

branch return; used in terminating subroutines

halt
 no operation
 execute

overflow test
 reset overflow
 record exponent

Breakpoint Test Group

((BPT 1 \wedge BPT<1>) \vee (BPT 2 \wedge BPT<2>) \vee (BPT 3 \wedge BPT<3>) \vee (BPT 4 \wedge BPT<4>)) \rightarrow (P \leftarrow P + 1);

Memory Extension Register Control Group

SET \rightarrow (instruction<17> \rightarrow (EM2 \leftarrow instruction<21:23>);
instruction<16> \rightarrow (EM3 \leftarrow instruction<18:20>));

EXT \rightarrow condition \rightarrow (P \leftarrow P + 1);

condition := ((instruction<22> \wedge (EM2 = 2)) \wedge (instruction<23> \wedge (EM3 = 3)))

POP \rightarrow (M[0]<0,9:23> \leftarrow 0v01CP; P \leftarrow 100₈ + pop_code);

programmed operator; 64 user defined instructions called via subroutine link in M[0]

EOM \rightarrow IO_instruction_execution;

see the definition of the IO instruction set below

POT \rightarrow IO_instruction_execution;

PIN \rightarrow IO_instruction_execution;

SKS \rightarrow IO_instruction_execution;

)

end Instruction_execution; not including Input Output instructions

*Input-Output Control from the Pc**KT and KMs State*

Devices consist of the following parts:

IO_Device[0:7777₈]

name (or address) of a specific IO device; the EOM command is first given to select the specific device; subsequent commands are implicitly to the selected device

IO_Output[0:7777₈]<0:23>

Input and Output Data buffers associated with specific devices

IO_input[0:7777₈]<0:23>

bit for each device to denote when device is ready to transmit data

IO_Ready[0:7777₈]

a bit within each device denoting it has been selected for an operation

IO_Select[0:7777₈]

the particular io device selected by the EOM command;

io_unit<0:14>

IO Instruction Set

EOM \rightarrow (io_unit \leftarrow e);

command to select or address the device; energize output M output data command

POT \rightarrow (IO_Select[io_unit] \wedge IO_Ready[io_unit] \rightarrow (
IO_Output[io_unit] \leftarrow M[e]; io_unit \leftarrow 0);
IO_Select[io_unit] \wedge \neg IO_Ready[io_unit] \rightarrow (POT));

wait until ready

PIN \rightarrow (IO_Select[io_unit] \wedge IO_Ready[io_unit] \rightarrow (
M[e] \leftarrow IO_input[io_unit]; io_unit \leftarrow 0);
IO_Select[io_unit] \wedge \neg IO_Ready[io_unit] \rightarrow (PIN));

input data command

SKS \rightarrow (io_unit \leftarrow e: next

skip if signal is not set

(IO_Select[io_unit] \wedge IO_Ready[io_unit] \rightarrow (
P \leftarrow P + 1);
io_unit \leftarrow 0);

Interrupt System States

Interrupt

controls whether interrupts will be processed

I_RQ[0:63]<0:15>

array of 1024 interrupt requests

I_ON[0:63]<0:15>

array of interrupt enable to enable or inhibit interrupt requests

I_Signal[0:63]<0:15> := I_RQ[0:63]<0:15> \wedge I_ON[0:63]<0:15>

K_address<0:5>

group number

I_address<0:3>

level number within a group of the active interrupt

The $I_address$ and $K_address$ combine ($200_8 + 20_8 \times K_address + I_address$) to establish an interrupt address, 200_8 is the highest priority and $200_8 + 1477_8$ the lowest priority.

$Interrupt_level_state[0:63] < 0:15 >_3$

There are three states associated with each interrupt, Inactive, Waiting, and Active:

Inactive means no I_signal is present.

Waiting means the I_signal has been received but is waiting to be processed.

Active means the interrupt has caused the main program to recognize its presence.

The instruction in $M[200_8 + 20_8 \times K_address + I_address]$ is executed upon interrupt. There are two kinds of interrupts: single instruction allows one instruction to be executed and the interrupt level state is changed from active to inactive; and normal requires that a mark place and branch, BRM, instruction to be executed to save P. At the completion of the interrupt program, a branch unconditional (BRU) indirectly via the BRM instruction restores the interrupt level. (That is, the $Interrupt_level_state$ is changed from Active to Inactive, and another I_Signal can be processed.)

Interrupt Interpretation

A state denoting that an interrupt is to be processed or the interrupt level state to be changed from Waiting to Active for normal interrupts and Waiting to Active to Inactive for single interrupts. The interrupt processed is the highest of those waiting provided there are no interrupts of highest level in the Active state.

Interrupt Control Instructions

$EIR \rightarrow (Interrupt \leftarrow 1);$

enable interrupt; turn on mode

$DIR \rightarrow (Interrupt \leftarrow 0);$

disable interrupt; turn off

$IET \rightarrow (Interrupt \rightarrow P \leftarrow P + 1);$

interrupt test; skip if on

$IDT \rightarrow (\neg Interrupt \rightarrow P \leftarrow P + 1);$

interrupt disable test; skip if off

POT instruction to control the Interrupt System. $EOM[20020]$ is first given to select the Interrupt System.

$(POT \wedge IO_Ready[20020]) \rightarrow ($

interrupt control instructions

$(c = 1) \rightarrow I_DN[a] < 0:15 > \leftarrow I_DN[a] < 0:15 > \vee B < 0:15 >;$

arm a channel level group

$(c = 2) \rightarrow I_DN[a] < 0:15 > \leftarrow I_DN[a] < 0:15 > \vee \neg B < 0:15 >;$

disarm a channel level group

$(c = 3) \rightarrow I_DN[a] < 0:15 > \leftarrow b < 0:15 >;$

set a channel level group

$a < 0:5 > := M[e] < 0:5 >$

group select or $K_address$

$b < 0:15 > := M[e] < 8:23 >$

data for $I_address$

$c < 0:1 > := M[e] < 6:7 >$

command control bits