COMPARISON OF CONTROL DATA 3000 SERIES WITH IBM 2000 SERIES

The 3000 Series Product Line

This section of the report provides a discussion of the 3000 Series of computers. These computers divide, on the basis of machine language, into two separate groups. These are: The CONTROL DATA 3100, 3200 and 3300 Group, and the 3400, 3600 and 3800 Group. The total family relationship of this computer line is carried forward by inclusion of the following common elements in the computers.

1. A CONSISTENT FORM OF CABINET CONSTRUCTION Across the board in all of these computers is a common cabinet structuring and a consistent styling, making all of these machines recognizable as a family of machines. Economies in purchasing cabinets can be achieved, when building all of these machines of the same cabinet parts. Economies in documentation can also be achieved when this singleness of purpose is implemented.

2. A COMMON TYPE OF BUILDING BLOCK The development of a set of building blocks to be used for all of these machines (with some minor exceptions) is another method of making these machines of a common class. Once again, economies of purchase, inventory, construction and documentation can be achieved by this commonness of objectives.

3. A COMMON FORM OF LOGIC REPRESENTATION leads to an economy in training for the field force and checkout personnel.

COMPARISON OF THE CONTROL DATA 3100, 3200 AND 3300 GROUP WITH THE IBM 2040 AND 2050 GROUP

In discussing the evolution of the 3100, 3200 and 3300 Group of computers, we must accept the following rules:

1. Physical appearance of all three machines is to be of the same general form.

2. The same building blocks shall be used across the board with minor exceptions.

3. The machines are to be completely machine-language compatible with the penalty, in going from the top of the line to the bottom of the line, to be measured only in speed.

Table I
Relationship Between 3100, 3200 and 3300

	3100	3200	3300
[1]Memory Size	4k, 8k, 16k, or 32k	4k, 8k, 16k, or 32k	8k, 16k, 32k possibly more
[1]Memory Speed	1.75 usec	1.25 usec	.8 usec
[1]Memory Modularity	8k or 16k	8k or 16k	8k
[1]Normal Add	3.5 usec	2.5 usec	1.6 usec
[1]Normal Multiply	14 usec	9 usec	6 usec
[2]Floating point Mult.	240 usec	180 usec or 19 usec	11 usec
Floating point hdw.	none instructions trapped	optional	included
BCD Hardware	none instructions trapped	optional	optional
Basic Pyramid Time	.175 usec	.125 usec	.100 usec
Arithmetic Timing	synchronous	quasi-asynchronous	asynchronous
Indexing	main pyramid	main pyramid	separate pyramid or priority
I/O channels	2-3107's or 4-3106's or combination	1-3207 or 8-3206's or combination	4-3307's or 8-3306's or combination
I/O speeds	160kc transfer rate	500kc transfer rate	1,000kc transfer rate
I/O Machine Relationship	machine related	partly independent	independent
Register File	main memory	special core matrix	special core matrix

1 24-bit arithmetic

2 48-bit arithmetic

The preceding table gives a quick look at the characteristics of these three machines on a comparative basis. Increases in capability will be accomplished by reorienting of the machine organization. The machine organizations to produce these effects are outlined in the accompanying diagrams, Figures 1, 2, and 3.

No attempt will be made in this report to describe these computers in detail, since it will be assumed that the reader will be familiar with the 3200. The 3100 and 3300 are essentially different versions of the same type of machines.

Comparison between the 2040 and 2050 will be the next point to be considered. A comparative reference to the previous table on machine performances and the succeeding table on the performance of the 2040 and 2050 leads to some important conclusions.

Table II
Characteristics of the 2040 and 2050

	2040	2050
[1]Memory Size	5.4k, 10.8k, 21.6k, 43.2k and 86.4k	21.6k, 43.2k and 86.4k
[1]Memory Speed	3.75 usec	1.5 usec
[1]Memory Modularity	total memory	total memory
[1]Normal Add	11.25 usec	4.0 usec
[1]Normal Multiply	84 usec	28 usec
[2]Floating Point Multiply	170 usec	30 usec
Floating Point Hardware	stored logic	stored logic
BCD Hardware	stored logic	stored logic
Basic Pyramid Time	unknown	unknown
Arithmetic timing	synchronous	synchronous
Indexing	main pyramid	main pyramid
I/O Channels	2 selector, 1 multiplexer	3 selector, 1 multiplexer
I/O Speeds (max)	800k bytes/sec	2,000k bytes/sec
I/O Relationship	partly independent	partly independent

1 24-bit arithmetic 2 48-bit arithmetic

A quick perusal of the above comparative tables leads one to the following conclusions:

1. The arithmetic capability of the 2040 and 2050 cannot match the arithmetic speed of the 3200.

2. The I/O rates of the 3200 are in the same class with the 2050.

3. Leading to the conclusion that the throughput of the 3200 will be better than that of the 2050.

4. Memory size in excess of 32k words is a problem which we will have to face and which we will provide an answer for in the near future.

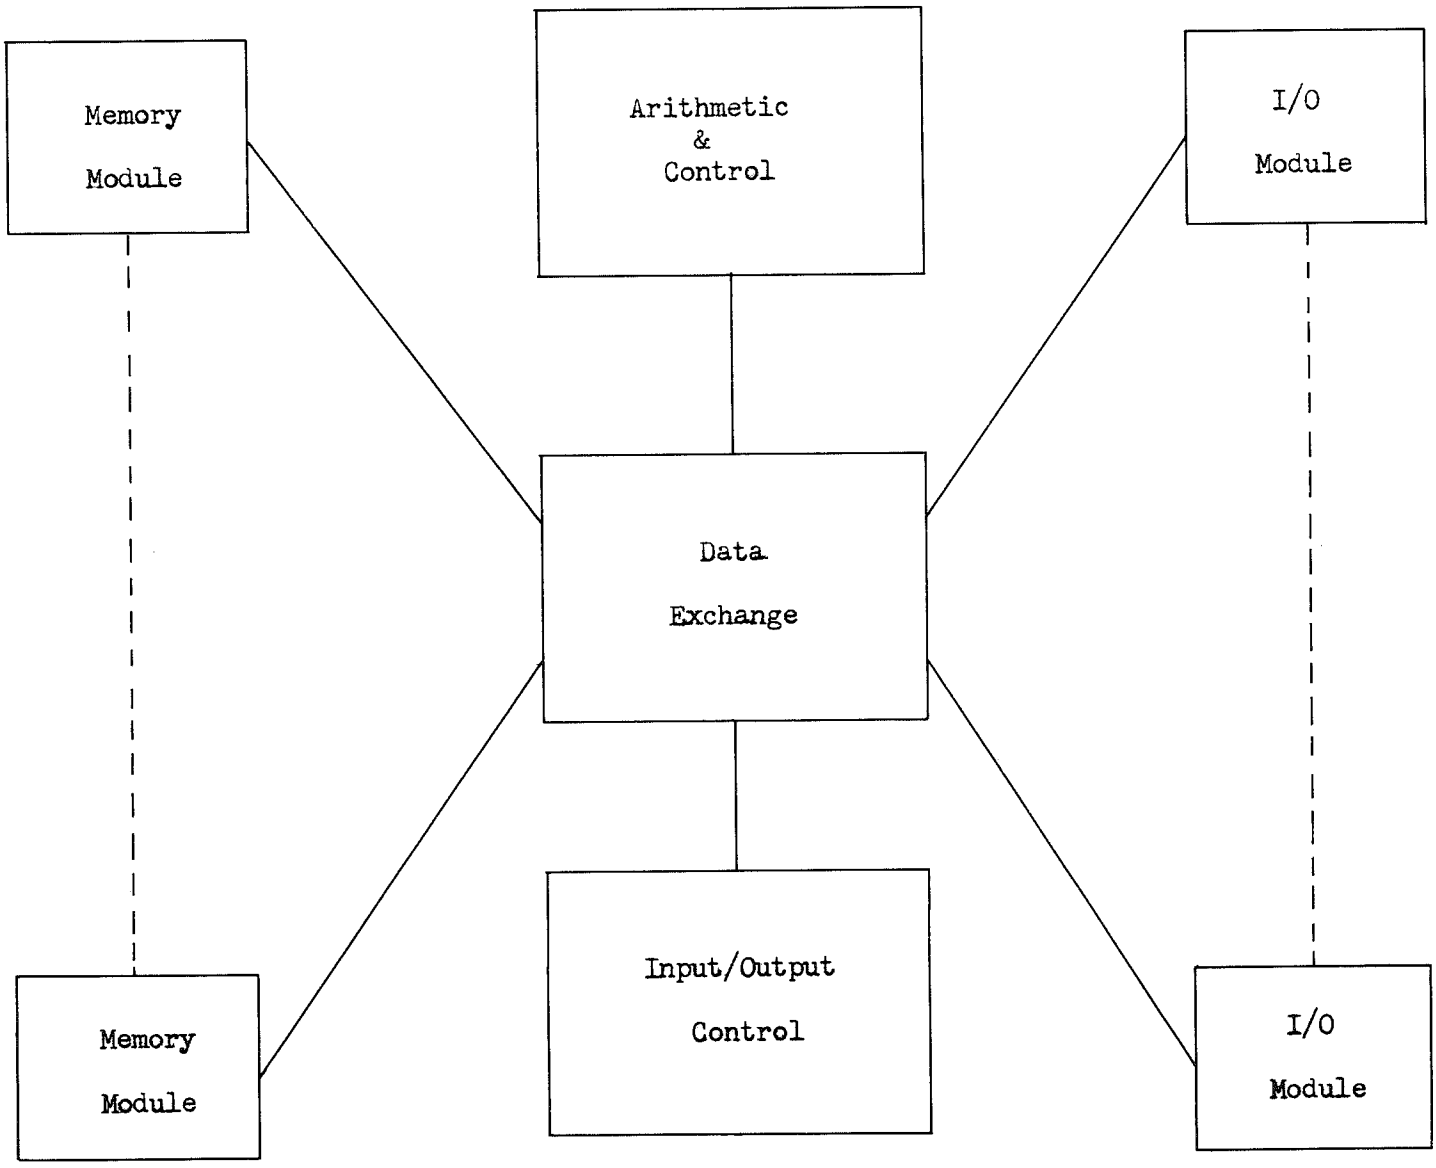

3100 MACHINE ORGANIZATION

Figure 1

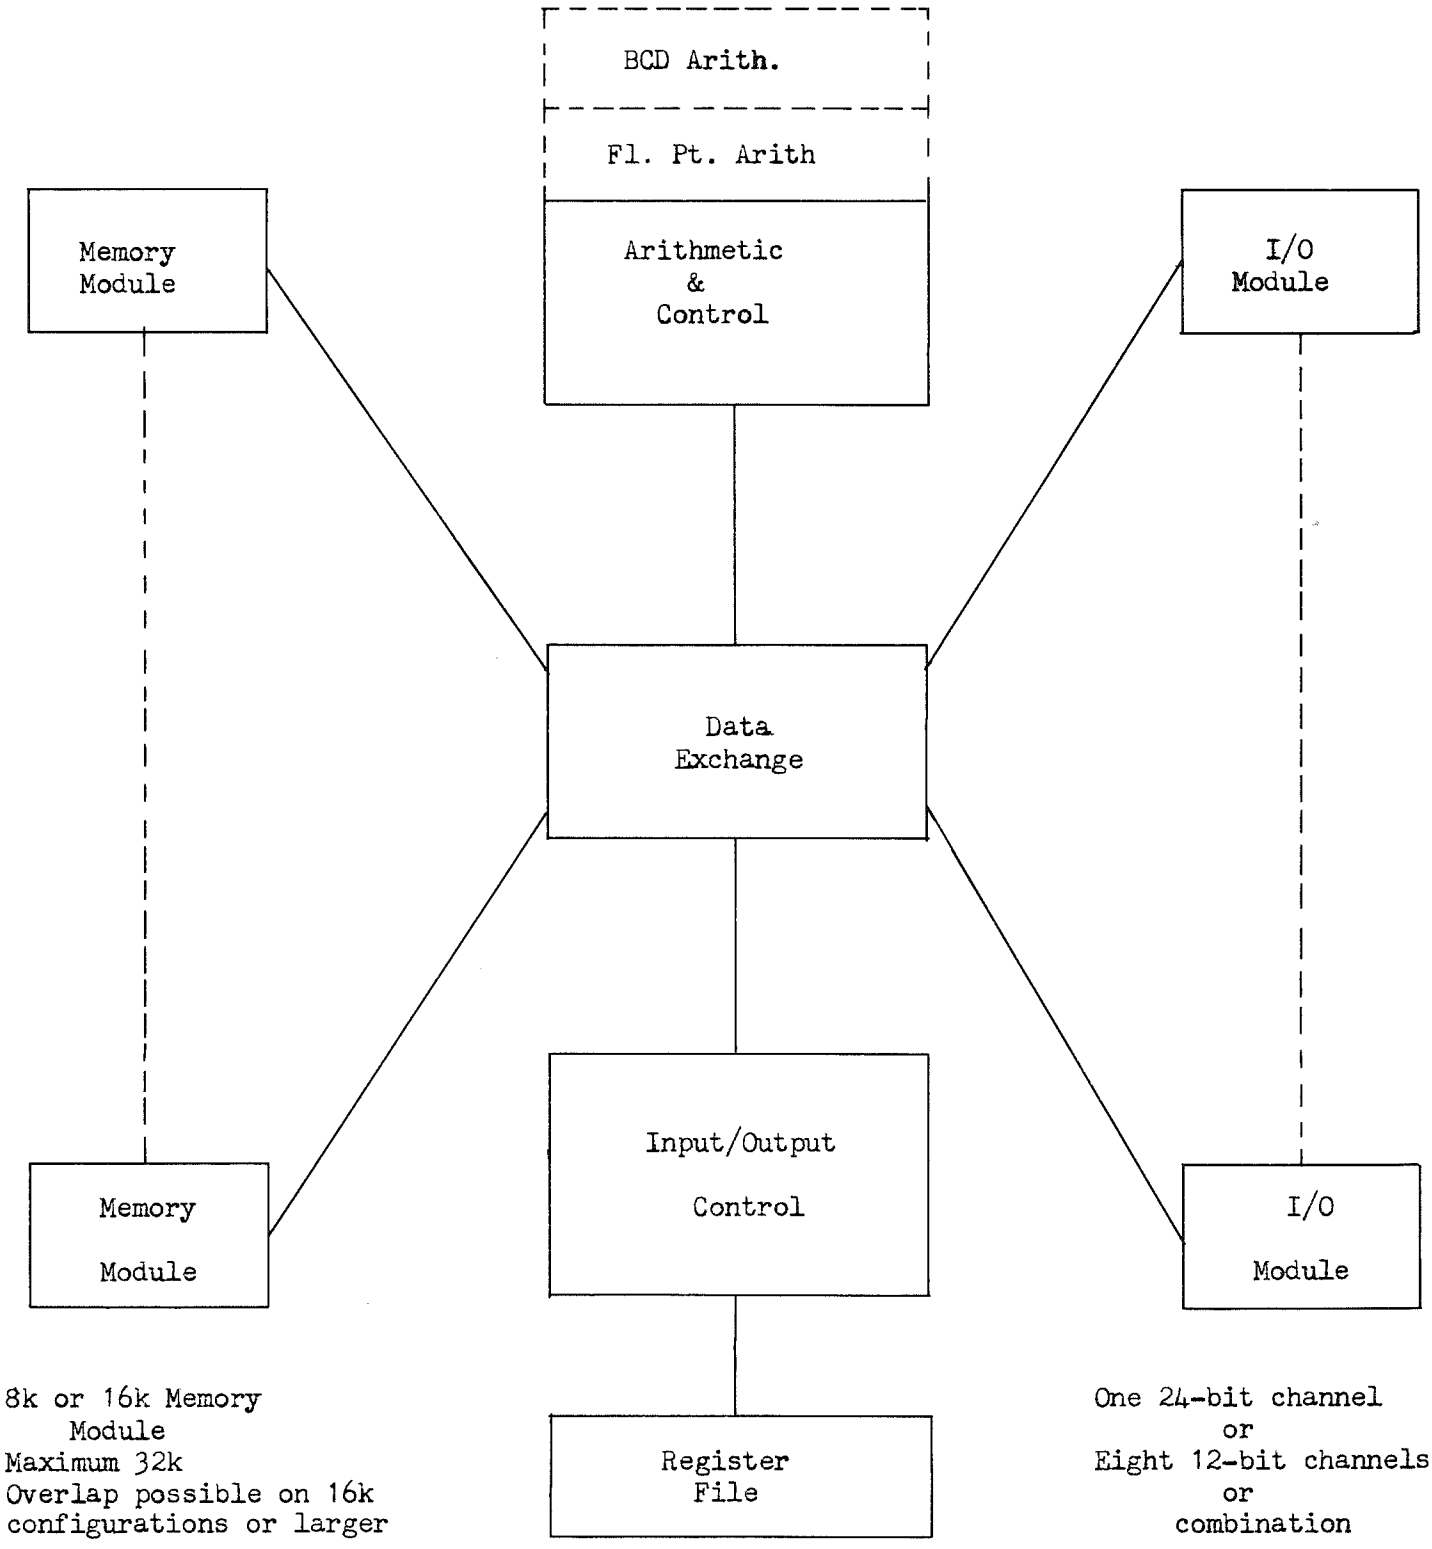

3200 MACHINE ORGANIZATION

Figure 2

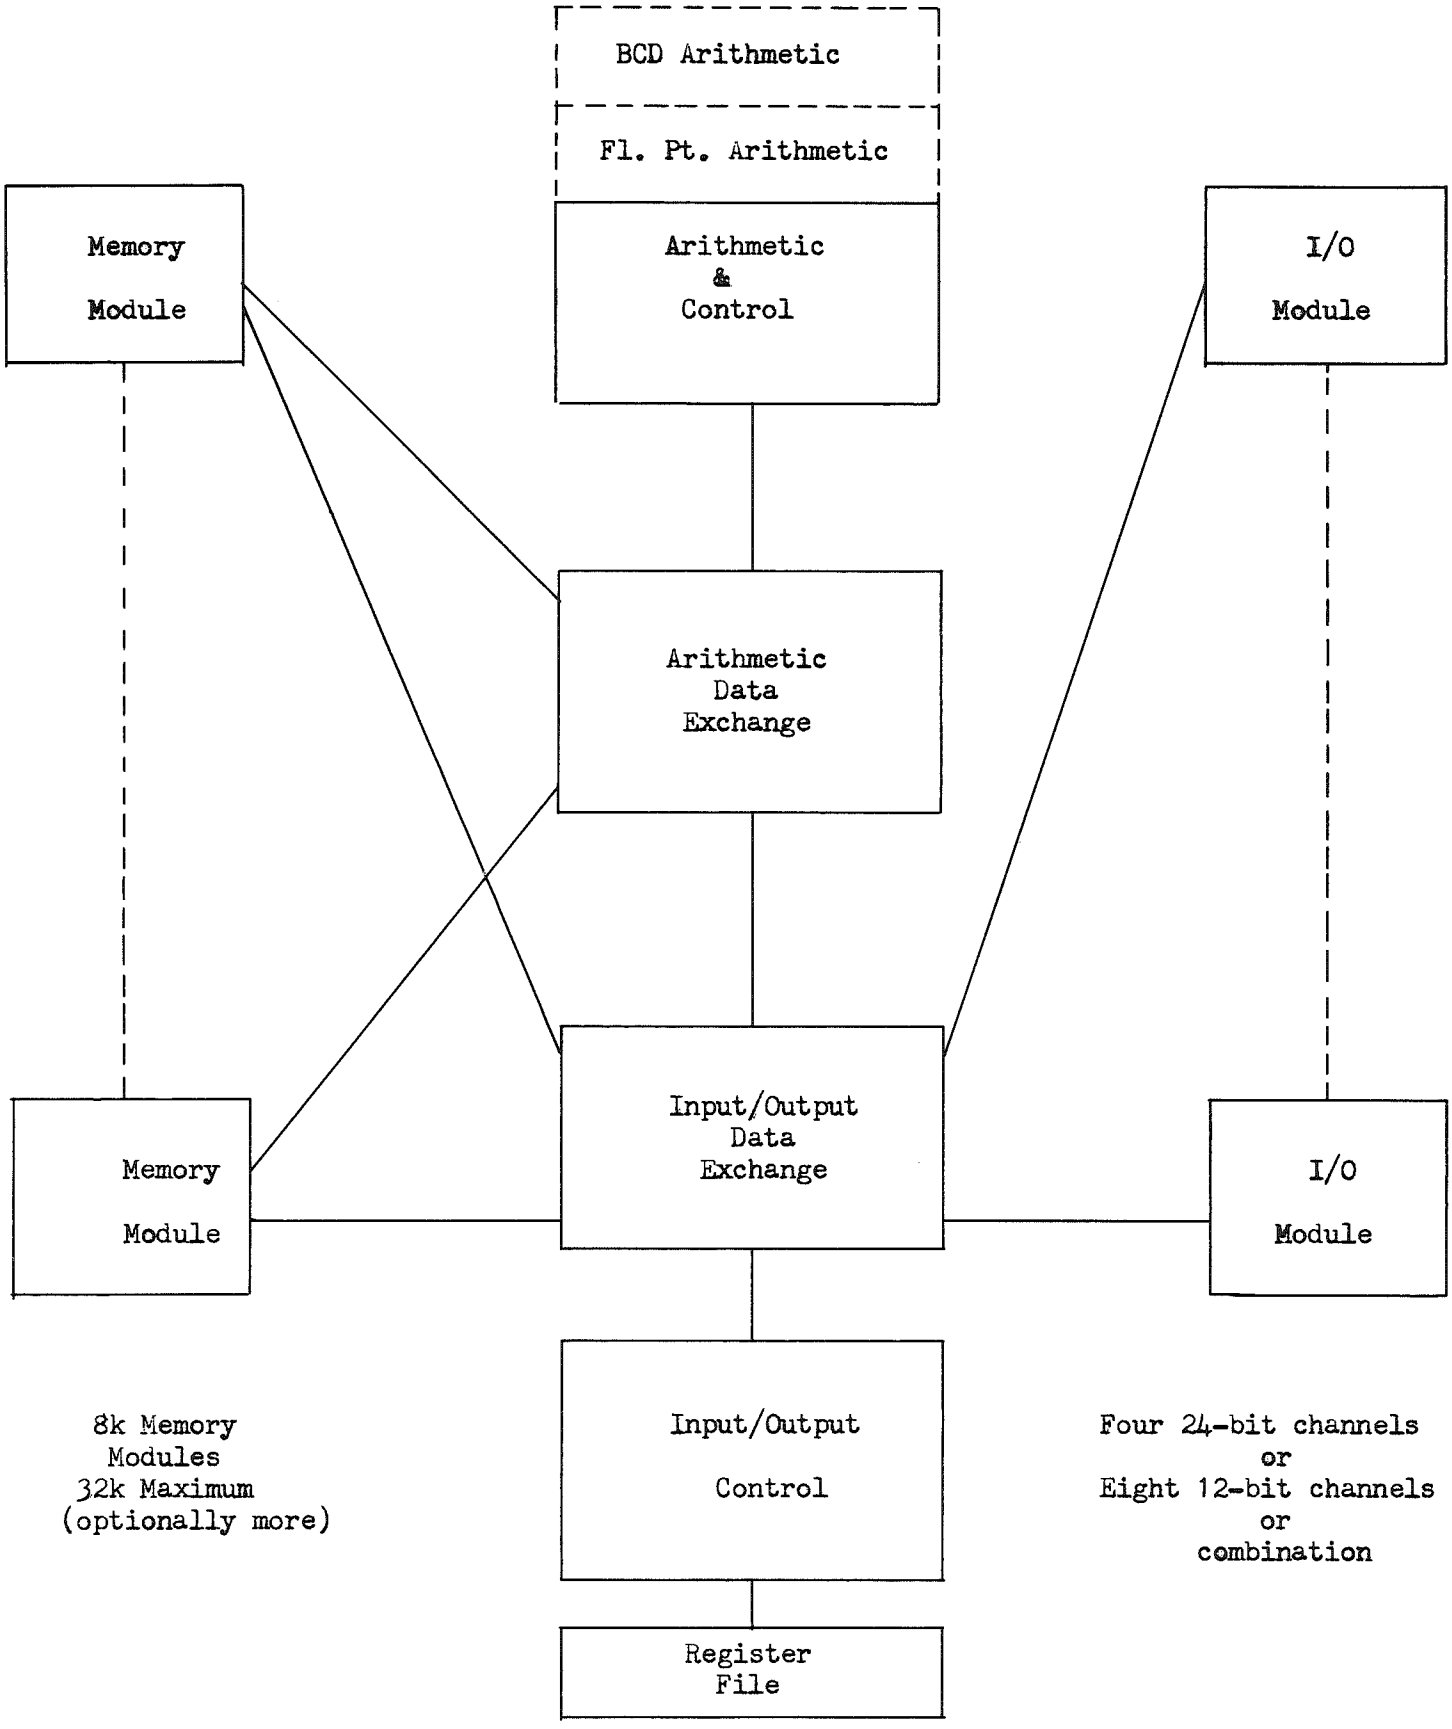

3300 MACHINE ORGANIZATION

Figure 3

COMPARISON OF CONTROL DATA 3400, 3600 AND 3800 WITH THE IBM 2050, 2060, 2062 AND 2070

This comparison of CONTROL DATA large-scale 3000 Series with IBM's large-scale 2000 Series will be divided into five parts:

1. PERFORMANCE
2. PRICE
3. COMPATIBILITY
4. PRESENT SITUATION
5. FUTURE SITUATION

Performance:

The Product Assessment Group evaluation of the performance of these large-scale systems was based entirely on internal performance of the central processor. The evaluation included mix comparisons, sample programs, and qualitative factors. Qualitative factors were included because of the radically different organization of the two series of machines.

This radical difference between the two series made it very difficult to compare the performance of the two series. The more the group studies these series, the clearer it became that the only real way or the only valid way of comparing the two would be to run a series of programs on both machines and compare results. However, since the IBM Series still does not run and will not be delivered for approximately two years, this benchmark comparison was impossible. Therefore, we used the inadequate techniques described above.

We felt that these techniques were inadequate for another more important reason. Internal performance of the CPU is only one small factor in the measurement of the performance of the system. Throughput, which involves the balance of peripheral equipment, speed and internal performance along with software efficiency, is more important to the customer than how many adds can be carried out in one second. The only method to measure throughput that we have today is to run a series of programs and compare results. Again, since neither the IBM programming systems or the hardware is available, we cannot compare results through this method.

Figure 1 gives the performance comparisons between the IBM large-scale group and the Control Data large-scale group. The IBM 7090 is used as the base machine. For instance, Control Data's evaluation of the IBM 2050 is that the 2050 is .89 times as fast as the 7090. IBM's evaluation is that it is .7 times as fast as the 7090. The chart shows that a Control Data evaluation of the IBM machines gives a higher speed rating for all IBM models. We don't know how IBM evaluated or what terms IBM used to evaluate their machines, but since the difference between our evaluations is consistent, we probably used the same relative criteria.

PERFORMANCE

	EVALUATION	
	Control Data Corp.	IBM
2050	.89	.7
3400	1.93	
2060	2.52	2.0
3600	2.94	
2062	3.21	
3800	5.3 (Est.)	
2070	6.82	5.0
	X7090	

Figure 1

Using the criteria previously discussed, the 3400, 3600 and 3800 neatly sandwich between the 2050, 2060, 2062 and 2070 as far as internal CPU performance is concerned. As far as overall system performance is concerned, Control Data systems software and hardware will be more efficient at the time the IBM systems are introduced because they will have had two more years of field use. Control Data will be introducing advanced software systems at the time IBM is introducing basic software systems. Therefore, it is our estimate that at the time of the introduction of the IBM systems, the 3400 total system will have greater throughput than the 2060. The 3600 will have greater throughput than the 2062 and the 3800 will have greater throughput than the 2070.

Price:

Since Control Data has expanded and restructured its product line, the large-scale 3000 Series is in a favorable position regarding price performance. Figure 2 compares the 3400 with the 2050 and 2060. The 3400 is priced below the 2050, and has approximately twice the performance. The 2060 is approximately 30% faster than the 3400, but costs $14,000 a month more for rental. In most cases, the 3400 will probably be able to handle the same job as the 2060 with a tremendous price advantage. If IBM tries to offer a low-price 2050 to the public, Control Data then has a tremendous performance advantage. We are in a favorable position under either circumstance.

Figure 3 shows how the 3600 and 3800 compare on price and performance with the 2060, 2062 and 2070. The 3600 is priced below the 2070 and is 1.4 times as fast and available now. The 3800 is priced below the 2062, and is about twice as fast. In fact, it is almost as fast as the 2070, although the 2070 is about twice as expensive. Thus, the same situation holds true for the 3600 and 3800 as was described in the 3400 price-performance comparison.

Compatibility:

As stated earlier, compatibility will be a big IBM selling point. However, this compatibility across the line does not exist. In IBM's literature, they point out three rules that must be followed, if program compatibility is to be maintained by the machines.

1. Machine configurations must be the same.

2. No programs can be dependent on the machine speed.

3. Certain I/O functions must be available on both machines.

PRICE

(32K.....8 Tapes)

	3400	2050	2060
COST	$24,700	$25,360	$38,800
P/L	44:1	48:1	48:1
PERFORMANCE	1.0	.46	1.3

YOUR JOB/FOR LESS MONEY

Figure 2

PRICE

(32K.....8 Tapes)

	Cost	Performance
3600	$43,800.	1.0
2060	45,600.	.85
3800	46,800.	1.8
2062	50,500.	1.05
2070	59,300.	2.25

YOUR JOB/FOR LESS MONEY

Figure 3

These rules are very difficult to meet. Take the first criteria: Same machine configuration. It is impossible for the 2060, 2062 and 2070 to have the same machine configuration as the 2030 and 2040. The reason for this is that the multiplexer channel, which connects all low speed I/O devices, such as card readers, printers, punches, etc., can only be attached to the 2030, 2040 and 2050. This means that any programs that involve these devices must be written in some form of higher-level language, and cannot be written in a machine-dependent language. This also means that the 2060, 2062 and 2070 must always have another central processor in order to handle slow-speed devices. The indication is that the software systems will require a 2040 attached to a 2060, a 2049 or a 2050 attached to a 2062 and a 2050 attached to a 2070. This obviously raises the system price, and immediately increases the complexity of the separate systems.

The second criteria involving non-dependence on machine speed is very difficult to meet, going from a very slow-speed machine like the 2030 to a very high-speed machine like the 2070. There are no statistics available on how often a program is not speed-dependent on a machine, but it seems abundantly clear that in almost no case involving I/O can an efficiently coded program run on both the 2030 and the 2070, without extensive modification. Experience shows that it is even difficult to run a program on two different models of the same machine class.

The final criteria involving same I/O functions appearing on both machines appears to be a very limited factor. There are a number of restrictions that appear in the IBM manuals. If these restrictions are indeed inconsistencies, it means that there is almost no compatibility between machines in the I/O area. Therefore, the only way you could run programs on different machines would be to operate under a monitor system, which is exactly the same way we achieve compatibility between our 3400 and 3600. This description is only a brief one, but gives a reasonable picture of what compatibility really is in IBM's line of new equipment.

There are really two groups of machines -- the 2030, 2040 and 2050 in the peripheral processing machine group and the 2050, 2060, 2062 and 2070 in the central processing group. The 2050 appears in both groups, and can attach to all peripheral equipment. It is the so-called swing machine, and it can be expected that the 2050 will be the machine that appears most often in competition with the Control Data machines. The compatibility offered the Control Data line of equipment is as extensive as that offered in IBM's line. There is complete software compatibility in the 3400 to the 3600 and there will be complete compatibility between the 3600 and the 3800. Therefore, the term "Compatibility where it counts" is certain to be a larger marketing advantage for the Control Data product line.

Present Situation:

Where does Control Data stand today? The 3600 is operating in numerous customer installations throughout the country. Its hardware and software is field-proved.

Comparisons can be run against any other operating machines. It is substantially superior to IBM's present 7000 series of equipment in both price and performance. The 3400 has been shown to the Spring Joint Computer Conference. It compared and executed FORTRAN programs at that conference. Two models of the 3400 are presently operating: One for engineering use, and one totally devoted to systems programming checkout in Los Angeles. The systems programmers will have the machine completely for their own use for over six months, before the first 3400 is delivered. We fully expect the 3400 to be the first system that will deliver both hardware and software, when the first system is placed in a customer installation.

Future Situation:

In the future, and by the future we mean when IBM delivers their first large-scale 2000 Series equipment, we will have a drum-oriented operating system, multiple-programming systems for the 3600 and 3800, a 3800 computer with new and powerful random accesses, mass memory devices and a full library of applications programs. All in all, Control Data stands very well in the large-scale market compared to IBM's 2000 Series of equipment.

AN ANALYSIS OF THE CONTROL DATA 6600 AS RELATED TO THE IBM 2000 SERIES

Shortly after our group began studying the 2000 Series of machines, it became apparent that there was no competitor to the CONTROL DATA 6600 in their new line. We would like to review with you briefly, the reasons for this conclusion. First of all, we would like to discuss two systems concepts which apply to both of these machines; namely, multi-processing and multi-programming; and then review a few of the particular hardware features of each of the machines.

Both Control Data and IBM see their large machines operating in a system where there is one central processor and one or more peripheral processors. The central processor is the large, powerful, problem-solving element in the system, and the peripheral processors provide the monitor and control function for the entire system. In the case of the 6600, there is the one central processor and the ten independent peripheral processors, giving a total of 11 logically-independent computing elements in the system. For the top of the 2000 Series, it appears that IBM's systems' programs will support a 2070-2050 combination, in which the 2070 is the central-processing unit and the 2050 will be performing the monitor and control functions. The peripheral processors will also be controlling the I/O operations. The 6600 has twelve high-speed data channels, which can be communicating with any one of the peripheral processors independent of the operation of the central processor. The 2070-2050 combination, on the other hand, has six high-speed channels available on the 2070, which must be controlled by the central processor itself if they are used, and the three intermediate-speed data channels, along with the one multiplexer channel, which are controlled by the 2050. The maximum attainable data rates are far greater in the 6600 with its twelve high-speed channels than in the 2070-2050 combination. Another more subtle but very important point is the manner of communication between the peripheral processors and the central processor. Any one of the ten peripheral processors of the 6600 can communicate directly with the large central memory, whereas in the 2000 Series, a channel connection must be established between two processors to permit communication. This can be a real bottle-neck in the flow of data through the entire system. As a multi-processing system, the 6600 with its eleven independent processors and its superior data-handling ability is clearly ahead of anything in the 2000 Series.

Looking at the central processor, the 6600 also has superior multi-programming ability. First of all, why should the additional complexity of a multi-programming system be a desirable feature? The aim is to keep the central-processing unit busy at all times. The problems in achieving this goal have actually been aggravated by the fact that the gap between the speeds of central processors and I/O devices has actually been widening rather than narrowing. The answer proposed by multi-programming, then, is to time-share the central processor between several programs, which all exist concurrently in a large, central high-speed memory. Then if an I/O wait condition occurs in one program, the CPU can be switched to another program. To implement

multi-programming, there are two hardware features which are essential; namely, storage protection and the ability to dynamically relocate programs. In the 6600 System, the relative address and field-length registers provide for a storage protect scheme which requires a minimum of attention from the software operating system. It is a flexible system, which permits a program to be located anywhere in memory and to be of any length which can be contained in memory. The 2000 System provides for memory protection on a block basis. This means that any one program must occupy an integral number of blocks and the operating system required to keep track of the block assignment is more complex.

The relative address register of the 6600 also provides a method for simple relocation of programs. To perform this operation, the operating system merely moves the program and then resets the relative address register to the new location. The relocation scheme in the 2000 Series is dependent upon the use of the base address registers. When a program is relocated, all of the values which are to be used as base addresses in the program must be readjusted to the new location. If the operating system is to be able to effect dynamic relocation in the 2000 Series, the programmer must conform to certain rules. He could not, for example, decide to use the base address register as a second index register. Nor could he decide to interchange the base address register designator with the index register designator. Concerning the multi-programmed operation of the machine, then, the hardware features in the 6600 present a much more simple and flexible scheme to implement this than does the 2000 Series.

Now, let us consider some of the more detailed features of the machine. First of all, concerning the central processor, the 6600 achieves its speed from the concurrency of its organization and from the brute speed of the individual functional units. There are two primary aspects of the concurrency in the CPU. One is the ability to overlap references to the main memory, that is, the fetching and storing of operands can be overlapped with computing. Then there are the ten independent functional units. These permit a multiply operation to be proceeding along with an add operation, along with another multiply operation, etc. In contrast to this, the 2000 machines operate sequentially. They do not permit overlap memory fetches and stores. Although there are multiple internal registers, there is only one arithmetic unit, which means that operations must proceed sequentially in an unoverlapped fashion.

Looking at the speeds of the various arithmetic operations, we see that the 6600 has an edge of $2\frac{1}{2}$ to 4 times that of the 2070. Therefore, it is clear that the 6600 CPU has an advantage in both organization and speed over the best of the 2000 Series.

Concerning the main high-speed memory, the 6600 has 131,000 60-bit words of one-microsecond memory. The 2070, on the other hand, has an apparent limitation of 65,000 words of one-microsecond high-speed memory. In addition, the 6600

memory is overlapped in 32 banks, whereas the 2070 memory is overlapped in only 4 banks. For the 6600, this may seem like "gilding the lily," but if one looks at the requirements that can be made on the high-speed memory when several I/O operations are going on simultaneously with the CPU requirements, the degree of concurrency permitted with the 32 independent banks becomes very important.

The I/O section of the machine is one of the most difficult parts to compare, because of the wide variety of equipments that can be attached on either machine. Before talking about individual I/O units, it is important to remember that the effectiveness of a unit is dependent on both its characteristic speed and the ability to logically connect it to the computer when it is desired. The 6600 permits distributing the I/O devices across 12 high-speed channels which will permit a very high degree of concurrency in using these devices. The card readers, printers, punches and tape units for both Control Data's and IBM machines are essentially equal. The 6603 Disk Unit on the 6600 System is a unique device and is very essential to the operating system of the machine. It transfers information in 12-bit parallel form which gives it a particularly high transfer rate. When we compare it to the 1302 Disk Unit, we see that its capacity is somewhat less, but the transfer rate is nearly five times that of the 1302.

The bulk-core storage is a new device from IBM which has been highly touted to their scientific customers. This is being included in the I/O section of the discussion, because we believe that is where it belongs. Data would most likely be transferred in blocks over a channel between a bulk memory and the high-speed memory. Therefore, the channel becomes the limiting factor. If we should compare the bulk storage to the 6603 Disk Unit, we see that the capacity of the disk unit is very much larger, and the transfer rate from the bulk storage is only slightly faster than that of the instantaneous transfer rate from the disk unit. The cost-per-bit of the bulk storage is about 60 times that of the cost of the 6603 disk storage. However, the fact that one has immediate access to the bulk-core storage can be of considerable importance to certain problems, and the importance of bulk core storage should not be minimized.

If we look beyond the 2070 of the 2000 Series, we know that IBM has proposed to certain customers, a 2090. Although our information on this machine is sketchy, it contained three new features which set it aside from the rest of the 2000 Series. These are the ability to overlap memory references, multiple functional units and an instruction stack. Each of these features is presently contained in the 6600 System, and it is questionable how effectively they can be implemented in a machine which still must retain program compatibility with the lower end of the 2000 Series. The claim is that the 2090 Processor would be 4 times faster than the 6600 Central Processor, with delivery possible in 1967. Based on past experience, we have every reason to believe that on that time schedule, Control Data's engineering group will be able to deliver a machine that will retain Control Data's lead in the super-computer area.

In conclusion, we have attempted herein to enumerate the points which led us to conclude that the 6600 remains in a class by itself.

I/O OPERATIONS

There are two types of I/O channels on 2000 Series Computers, multiplexer and selector channels. The selector channels operate in a manner very similar to Control Data's buffer channels, i.e., a device is selected, the operation started, and then I/O takes place while the computer is operating. The multiplexer channel allows operating many slow-speed devices (typewriter, card equipment, line printer) on one channel. This is done with a scanner that automatically interrogates activated devices. The multiplexer channel is available only on Models 2030, 2040, and 2050.

The entire I/O operation is very complex and will not be gone into at this time.

PERIPHERAL EQUIPMENT

Probably the most impressive feature of the 2000 Series is not the computers (they exhibit a reorganized approach, but offer nothing really new in computer technology), but the number of peripheral equipment.

BULK CORE STORAGE

In addition to main core storage in Models 2050, 2060, 2062 and 2070, large capacity storage modules of one or two million bytes can be expanded to eight million. Such storage is often referred to as bulk core. It is addressed contiguously with main storage. Memory transfer size is 64 bits and the memory cycle time is 8 microseconds.

IBM is pushing this bulk core for almost every 2070 System. It is not evident at this time how the core can be used efficiently in a system. There appear to be many problems yet to solve. While its access time is fast compared to electromechanical methods of storage, its transfer time is slower than Control Data's high-performance drum. At this time, it appears to us that IBM is over-selling this bulk core without giving many facts to show how it is used.

The lease price of bulk core is:

$ 9,000 per month for 1 million characters

$13,000 per month for 2 million characters

MASS RANDOM ACCESS DEVICES

IBM offers a storage control for four different devices. These devices and their characteristics are listed in Table 2. All listed items except the 2301 Drum can be controlled by the storage control unit that leases for $525 per month. However, attachments must be added to the controller for the following:

1301 attachment		$250 per month
2321	"	175 " "
7320	"	200 " "

The 2311 Disk Storage Drive does not need an attachment.

The 2301 Drum Storage requires its own control that leases for $2,300 per month.

TABLE 2

	Storage Capacity (Bytes)	Average Access Time Milliseconds	Transfer Time (Bytes/Sec.)	Lease Price ($ per Month)
2311 Disk Storage Drive	7,250,000	100	156,000	575
1302 Disk Storage				
N1	112,140,000	165	156,000	5,600
N2	224,280,000	165	156,000	7,900
2321 Data Cell Drive	400,000,000	175-600	55,000	2,800
7320 Drum Storage	830,000	8.6	135,000	2,300
2301 Drum Storage	4,000,000	8.6	1,200,000	4,400
Bulk Core	1,000,000	.008	8,000,000	9,000
" "	2,000,000	.008	8,000,000	13,000

MAGNETIC TAPE UNITS

IBM has announced new 9-channel tapes called the 2400 Series. There are three different models in the series:

	Model 1	Model 2	Model 3
Data Rate (cps)	22.5K	45K	90K
Tape Speed (ips)	28.125	56.25	112.5
Rewind Time (minutes)	3.2	2.6	1
Lease ($ per month)	335	485	785

Characteristics common to all models are: Bit density of 800 bits per inch, .6" inter-record gap, .5" tape width, and read backwards capability. A seven-track compatible feature can be installed which allows reading and writing tape under the old 7-channel format. A control unit to handle eight Model 3 tape handlers on one data channel leases for $550 per month.

In addition to the above, IBM offers a Hypertape unit. Its characteristics are:

Tape Speed	112.5	ips
Rewind Time	90	sec.
Density - low	1,511	bytes per inch
high	3,022	" " "
Inter-Record Gap	.38	inches
Tape Width	1	inch
Data Rate - low	170,000	bytes per sec.
high	340,000	" " "

Hypertape tape units lease for $1,350 per month. The lease price of a controller to handle eight tape handlers on one channel is $2,100 per month.

CARD I/O AND LINE PRINTERS

Two card read/punch units are available, the 1402 which reads 1,000 cpm

and punches 250 cpm, and the 1442 which reads 400 cpm and punches 160 cpm. These devices lease for:

1402	$660 per month
1442 with controller	425 " "

IBM has 5 models of line printers. These printers are listed with characteristics and prices:

	Speed (Lines per minute)	Lease Price ($ per month)
1403 II	600	775
1403 III	1,100	900
1404 II	600	1,550
2201 III	1,100 or 1,400 numeric	900
1443 N1	200	875*
1445 N1	190	1,425*

*included controller

The 2821 control unit is used for control of the 1403, 2201, 1404 and 1402 card printing units. It leases for $970 per month.

COMMUNICATIONS EQUIPMENT

IBM has changed its approach to communication systems. The new method follows that used in the CONTOL DATA 8050 System in which systems are made of standard communication models together with a computer and peripheral devices necessary to handle the application. In the past, IBM has had a special communication computer system called the 7740. In the 2000 Series, however, Models 2030, 2040, or 2050 can be used in communications' applications.

The communication systems are designed to handle four types of application:

The 1030 Data Collection System collects data from remote locations. This system is similar to the CONTROL DATA 8010 Data Collector.

The 1050 Data Communication System is a multi-purpose office-oriented communication system.

The 1060 Data Communication System is designed for applications such as bank and loan association teller window communication with a central computer.

The 1070 Process Communication System is designed for on-line data transmission between remote locations and the central computer. This system can have:

Multiplexers
Digital and binary display
Digital and binary input
Analog-to-digital conversion equipment

DISPLAY AND CONSOLE EQUIPMENT

Consoles are optional on all computer models. Typewriters are used in Models 2030, 2040, and 2050 for operator communication. Typewriters and visual display units are used in Models 2060, 2062, and 2070.

There are two display units, the 2250 Display Unit which displays a 12" by 12" area and the 1015 Inquiry Display Unit which displays on a small scope, a 30- by 40-character array of data.

Both of these displays can have an optional input keyboard. The 2250 Display Unit displays an array of 52 lines of 74 characters. A light pen and display buffer memory are optional features.

OPTICAL AND MCR EQUIPMENT

A series of Magnetic Character Readers, Optical Readers, and Optical Mark Page Readers can be attached to the 2000 Series of computers. The equipment is designed basically for commercial data processing applications.

The discussion above is a general description of the 2000 Series. Without going into great detail, this description is intended to familiarize you with IBM's new line of computers so that you can understand Control Data's competitive position.

SUMMARY OF THE COMPARATIVE ANALYSIS OF THE I/O IN THE 2000 SERIES AND THE 3000 SERIES EQUIPMENT

2000 SERIES

2030	One multiplexer channel. Two optional selector channels.
2040	One multiplexer channel. Two optional selector channels.
2050	One multiplexer channel. Three optional selector channels.
2060	No multiplexer channels. Six optional selector channels.
2062	No multiplexer channels. Six optional selector channels.
2070	No multiplexer channels Six optional selector channels.

The multiplexer channel is a channel capable of servicing a multiplicity of slow-speed devices, such as paper tape readers, paper tape punches, typewriters, keyboard printers and other similar devices. Higher-speed devices may be attached to this channel, but will require that the channel be used in a special manner when data transfer is to be achieved from such devices. The normal mode of operation of the multiplexer channel is one of scanning a multiplicy of low-speed inputs into, or outputs from the channel, and servicing each in turn.

The multiplexer channel is only used in the 2030, 2040 and 2050, and in every case is incorporated as part of the CPU, sharing CPU facilities. This sharing of CPU facilities leads one to the conclusion that certain maximum I/O rates are inherent in the system, even when the machine is doing nothing but servicing I/O on this single channel. Although maximum rates are not given in the system information, they do state that there are rates which they cannot achieve. These rates are:

340,000 bytes/sec. on the 2030 selector or multiplexer channel

340,000 bytes/sec. on the 2040 or 2050 multiplexer channel

By implication, this indicates that in the 2030, the channel information is stored in memory for both the multiplexer channel and the selector channels. In the 2040 and 2050, the channel information for the multiplexer channel is stored in memory while the channel information for the selector channels is stored in separate hardware.

In the 2060, 2062 and 2070, no multiplexer channel is available for use in the machine. This omission will involve a complete break in the line at this point. In the area of I/O, no general compatibility in programming can exist across this juncture. All I/O channels available in the 2060, 2062 and 2070 are hardware implemented, and operation of these channels only impinges on the machines' operation, when the actual store or retrieval of information from memory takes place. In the case of these last named models, maximum I/O rates are limited by memory rates.

Maximum memory rates for the systems are:

2030	2 usec/byte	= 500 kc byte rate
2040	2.5 usec/2 bytes	= 800 kc byte rate
2050	2 usec/4 bytes	= 2 mc byte rate
2060	2 usec/8 bytes/interleaved	= 8 mc byte rate
2062	1 usec/8 bytes	= 8 mc byte rate
2070	1 usec/8 bytes/interleaved	= 16 mc byte rate

The number of devices which can be attached to the I/O channels in the 2000 Series are as follows:

2030

Multiplexer channel - 96 subchannels - 1 control/subchannel
Selector channel - 1 subchannel - 8 controls
Selector channel - 1 subchannel - 8 controls

Inherently, there is an addressing capability of 256 devices/channel.

2040

Multiplexer channel - 128 subchannels - 1 control/subchannel
Selector channel - 1 subchannel - 8 controls
Selector channel - 1 subchannel - 8 controls

2050

Multiplexer channel - up to 256 subchannels - 1 control/subchannel
Selector channel - 1 subchannel - 8 controls
Selector channel - 8 controls
High-speed selector channel - 8 controls

2060, 2062, 2070

6 - selector channels/8 controls/channel

Performance Characteristics of the 3000 Series I/O

Number of channels which may be attached:

Model 3200	eight 3206's or six 3206's and one 3207
Model 3400	four 3406's
Model 3600	thirty-two 3606's

* The 3206's, 3207, 3406's and 3606's can be considered to be comparable in performance to the IBM selector channel.

Maximum I/O rates of information transfer:

	Word Size	Maximum Word Rate	Six-bit Character Rate	Character Addressable Rate
3200	24 Bits	600 kc	2.4 mc	600 kc
3400	48 Bits	675 kc	5.2 mc	
3600	48 Bits	675 kc	5.2 mc	

Comparison of 2000 and 3000 Series:

A. Number of I/O Channels

3000 Series		2000 Series	
3200	8 channels	2030	2 channels
3400	4 channels	2040	2 channels
3600	32 channels	2050	3 channels
		2060/62/70	6 channels

B. Incorporation of 2 multiplexer channels:

3000 Series		2000 Series	
3200	no	2030	yes
3400	no	2040	yes
3600	no	2050	yes
		2060/62/70	no

(CDC will provide an answer to this deficiency)

C. Maximum Rates for I/O

3000 Series		2000 Series	
3200	2.4 mc Char/sec*	2030	less than 340 kc byte/sec**
3400	5.2 mc Char/sec	2040	less than 800 kc byte/sec
3600	5.2 mc Char/sec	2050	2 mc bytes/sec
		2060/62	8 mc bytes/sec
		2070	16 mc bytes/sec

* Six-bit characters
** Eight-bit bytes

Summary: Although the I/O philosophy of the 2000 Series is quite complex, it does not offer the customer any material advantage over the system used in the 3000 Series, except in the presence of 2 multiplexer channels in the 2000 Series. Control Data will provide an answer to this problem. With this exception, I/O in the 2000 Series does not provide a serious challenge to our product line.

COMPARATIVE ANALYSIS OF INTERRUPT IN THE 2000 SERIES AND THE 3000 SERIES

Interrupt in the 2000 Series is a rather complex operation. In discussing this aspect of machine operation, it must be understood that there exists a basic and opposite point of view on the part of IBM and Control Data on the meaning and use of interrupt.

Control Data has always accepted the point of view that interrupt usage in a computer should be minimized. IBM, on the other hand, leaves the impression that they are incorporating the interrupt as a normal and important part of the mainstream programming in the 2000 Series. It would be rather futile to argue the relative merits of either of these approaches, for any knowledgeable person would approach such a discussion with preconceived notions. In either case, a similar degree of capability and flexibility exists in the systems.

The 2000 Series has five types of interrupts. These, listed in their order of priority, are:

MACHINE CHECK
PROGRAM OR SUPERVISOR CALL
EXTERNAL
INPUT/OUTPUT

The machine check interrupt is an interrupt which has priority over the other interrupts in the machine. It is masked by bit 13 in the Program Status Word, and its purpose is to provide a means to recover from faults in the machine.

The Program Interrupt results from the improper specification of address or the improper use of instructions. Some of the program interrupts may be masked. These are interrupts dealing with arithmetic overflows. In general, these interrupts do allow the machine to respond in a reasonable manner to programming errors and data errors.

Supervisory Call Interrupt allows the machine to go from problem state to supervisory state. It cannot be masked. Supervisory state is the condition the machine must be in to execute any of the I/O instructions. Any attempt to execute an I/O instruction, while in the problem state would result in a program interrupt, and more specifically, a privileged operation interrupt.

External interrupt provides a means by which the CPU responds to signals from the timer, from the interrupt key and from external units. These interrupts are masked by mask bit 7. A total of six signals in lines may be connected to the CPU for receiving external signals.

Interrupts from I/O devices are processed one at a time. The interruption code, signifying the interrupting device, is stored at the time the interrupt is processed. These interrupts may be masked on a channel basis. Any interrupt, not processed, remains active for later processing.

task in programming the 2000 Computers, because it is necessary to adjust and use the base designator almost 100% of the time. (The situation is analagous to use of the bank register in the 3600; however, the 3600 can operate within 32k of 48-bit word size before it is necessary to adjust the bank control whereas the 2000 operates with 4K of 8-bit word size.

It is apparent in the literature and presentations from IBM that it is intended to implement multi-programming in the 2000 Series. This means that instructions must be provided which facilitate switching between programs. It is also apparent from programming a few examples on the 2000 Series that the 16-fixed points are quickly over-taxed and much swapping is necessary. To facilitate this changing, they have a R-S instruction. Its main purpose is to load and unload the 16 fixed-point registers. In an R-S instruction, R_1 and R_3 define the registers to be transferred. For example, if registers three through eight are to be loaded, then R_1 would contain 3, R_3 would contain 8, the operation code would be a multiple load, and B and D would form the byte address of the first location of data. Subsequent data would be accessed from sequential core locations.

The S-I and S-S instructions are commercial data processing and compiler-type instructions.

The S-I instruction sets a condition that can subsequently be checked by a test instruction. For example, it is necessary to check a location in core to see if it contains the character H. The S-I instruction is used, with its literal portion containing the character H, and B and D portions containing the address of the character to be checked, and the operation code containing a compare. The execution would set a condition that can subsequently be checked by a test instruction.

The S-S instruction is similar to that used in the IBM 1400 Series. It operates from memory to memory (i.e., it does not use an accumulator). The S-S instruction operates only with characters and digits. B_1 and D_1 form one operand address and B_2 and D_2 form the second operand address. L_1 and L_2 are the length of the first and second operands.

SAMPLE PROBLEM

To point out some of the shortcomings of the 2000 Series, a sample problem is shown in Figure 4. The problem is a summation of products A_i and B_i. The first five instructions listed are needed to set the loop control and base address. The next five instructions are the main loop that is then

SAMPLE PROBLEM

$$SUM = \sum_{i=1}^{i=10} A_i B_i$$

FORTRAN CODE

```
      SUM  = 0.0
      DO1I = 1,10
    1 SUM  = SUM + A(I)  *  B(I)
```

2000 PROGRAM

```
      L      R4   ONE
      L      R5   ADDRA+10
      L      R2   ADDRB
      L      R1   ADDRA
      LD     F2   ZERO
LOOP  LD     F4   R1
      MD     F4   R2
      ADR    F2   F4
      AR     R2   R4
      BXH    R1   R4    LOOP
      STD    F2   SUM
```

Figure 4

EXECUTION TIMES

Model	Time (usec.)	
2030	13,280	
2040	3,386	11 Inst. 340 Bits of Pgm.
2050	635	4 Index Reg. 2 Float Reg.
2060	194	1 Base Reg.
2062	167	
2070	83	

Control Data	Time (usec.)	
3200	424	9 Inst.
3400	323	216 Bits of Pgm. 1 Index Reg.
3600	175	A&Q Reg.
6600	19.6	

Figure 5

repeated 10 times. The last instruction stores the results.

Figure 5 compares the timing on the 2000 Series Computers with that of Control Data's computers. Note not only the time relationships, but the difference in program storage and number of registers used. The 2000 Series Computers require 340 bits of program, five fixed-point registers, and two floating-point registers while Control Data's computers require only 216 bits of program, one index register, and the A and Q registers.

Compare the execution time of the CONTROL DATA 6600 and the fastest model in the 2000 Series, the 2070.

PROGRAM STATUS WORD

As mentioned previously, the 2000 Series was designed for implementing multiple programming. One of its features is a 64-bit program status word which contains, among other things, a program counter (P register), memory tag, the BCD code used, interrupt mark register, and generally the status of the computer. By changing the PSW, the status of the computer can be changed. This is similar to the exchange jump provision on the 6600.

However, with the PSW used on 2000 Series Computers, none of the registers are changed. If a jump between programs is done, registers must be changed by program.

The computers in the 2000 Series have two operating states; the program state and the supervisor state. Certain instructions (such as I/O and control instructions) are called privilege instructions. Privilege instructions can only be executed in the supervisory state. If a privilege instruction is attempted in the program state, an interrupt will occur, forcing the computer into the supervisory state. Because of this, all programs have to be operated with a monitor or control program.

MEMORY PROTECTION

Memory blocks of 2,048 bytes have a 4-bit tag associated with them. This tag can be set or changed with privilege instructions. When a program is operating in one of these blocks, its PSW memory tag must match the tag on the memory block. The same type of protection is used on I/O operations.

In the medium-scale class, the CONTROL DATA 3200 far surpasses the IBM 2050 on an operation-per-dollar basis, although the gross price for the 3200 is greater than the 2050. The CONTROL DATA 3100 will be available for those prospects requiring a more capable system at a lower gross price.

Advantage of Control Data Technology

IBM claims superior performance to cost ratios, because of such items as hybrid-chip construction, stored logic, multiple floating-point registers, multiplex channels, etc. Control Data will invite any prospect to witness a demonstration so that he can compare performance, based on technical excellence of equipment that is available today.

HYBRID-CHIP CIRCUITRY Even though IBM is advertising their hybrid-chip circuitry as being a major technological advance, the present Control Data circuitry can equal or exceed their performance. There is no need to make a time-consuming, expensive change. The IBM hybrid-chip circuits are no faster, at each circuit stage, than the present Control Data built-up circuit cards. Significant increases in performance can be obtained through more efficient system organization rather than by changing the basic circuit cards.

16 FIXED-POINT REGISTERS At first glance, these IBM registers seem to be a very desirable feature. IBM will allow relative addressing within any 4k (bytes), and will undoubtedly be touted as the "key to easy relocation." However, the direct, unindexed addressing of any operand from memory requires the use of at least one of these registers. In addition, registers are used for fixed-point arithmetic and I/O control. There is no capability for multiple indexing or indirect addressing. Studies of the equipment in use, indicate that the continual shuffling of information into and out of these registers, may create a bottleneck in program execution.

MEMORY PROTECTION The IBM Memory Protect operates on 2k (bytes) sections. The Control Data Memory Protect is superior, since we can designate any size. Their Program Protect is equated in our operating systems, and will be adapted into our hardware.

PROGRAM STATUS WORK IBM will undoubtedly claim that the program status word opens the door to multiple programming. Their philosophy is different from that of Control Data, but does not accomplish anything that makes it better than ours. They still must save registers, etc.

THEREFORE:

FROM ALL INDICATIONS AND INFORMATION AVAILABLE TO DATE,

IBM'S ANNOUNCEMENT THAT THE 2000 SERIES WILL BE THE

BEST AVAILABLE IS SUBSTANTIALLY OVERSHADOWED BY THE FACT

THAT PROSPECTIVE CUSTOMERS CAN GET MORE FROM CONTROL DATA NOW!

August, 1964