September 24, 1957

Memorandum to: H. K. Wild
W. Buchholz
S. W. Dunwell

Subject: Project 7000 - High Speed Disks

This memo contends that the following four points are significant enough that a redefinition of the subject disks be made and adequate funds be provided to realize the redefinition.

1. A word rate of 16 usec per word is too slow.

2. Total capacity of 1,000,000 words is too little.

3. Having two logical files per disk unit when read-while-writing is too restrictive.

4. Access to the block (after the track has been found) is too long.

This memo proposes that a redefinition can cope with these points respectively as follows:

1. Double or quadrouple the linear and track densities and hence the word rate (if revolution speed is held constant); or process the entire word in parallel by bit rather than in four bytes.

2. Double or quadrouple the linear and track densities.

3. Put the write heads on a separate access mechanism, or process the entire word in parallel.

4. Increase the number of blocks per track and provide the "roll" operations.[1]

---

1. Gibson, J. C.: Project 7000, Ultra Fast Disks - New Instructions; File Memo; July 8, 1957

Several items point to the inadequacy of the 16 usec per word rate. One is that the contract with Los Alamos clearly calls for a rate of 4 usec per word. Since the contract specifies only two other items (capacity and continuity from track to track), failure to meet the word rate specification implies 33 1/3% malcompliance on our part.

Competition both within and from outside IBM also points to this inadequacy. LARK drums operate at 33 usec per word with simultaneous reading and writing, and the Westinghouse people report that they have been told that with three read-write synchronizers the LARK drum can operate at 16 usec per word; and, of course, the LARK I/O Processor handles all details of memory arrangement, indexing, control, and interrupt without interrupting the main computer. Further, a current Business Machines Analysis report reveals that LARK has a 6 million character drum; this is equivalent to 1 million LARK words.

Endicott reports that their present state of the art would permit definition of a 100,000 word drum with average access of 5.5 ms and a rate of 2 usec per word.[2] The 704 drum is now capable of operating at 12 usec per word. The 660 drum operates at 4 usec per word. If drums can match or surpass the speed and capacity of disks, may not all our disk programs become suspect, particularly in view of the proven reliability of drums?

The 16 usec per word rate does not satisfy basic design goals in the B + S system. This system is intended first for those problems that call for high speed and large memory, second for all others. A goal of "balanced design" is to minimize the time the arithmetic units must wait for input - output. Among the typical applications for B + S uncovered to date are the reactor and the hydrodynamic problems and many of these average as few as 6 operations per input word. Assuming an average of 1 usec per floating point operation, this means 6 usec of computing per input word. Westinghouse reports an application demanding only 2 usec of computing per input word. Obviously, then, if the input rate is 16 usec per word, the arithmetic units will wait 14 out of every 16 usec for input in the Westinghouse job - - hardly an example of balanced design. Were such problems not so large, others could be run with them and thereby use some of this idle compute time, but such is not the case. Further, one does not know how much more demanding the typical problem of the 1960's will be, but if one extrapolates from past experience it will certainly be more demanding.

There are many applications (e.g. the supervisory program, random relaxation techniques, simulation, linear programming) whose nature prevents neat overlapping of input-output with computing. For these, any word rate less than the maximum possible on the busses and in the memories is intolerable.

---

2. Hemphill, L. E.: Endicott Drum Development; Memo to D. W. Pendery; June 11, 1957

Programming costs are so high and schedules are so short for users today that "one-shot" problems are not usually optimized. In particular, for these situations the programming effort to overlap computing with input-output will not be made. Faster word transmission from disks will reduce the time and dollar penalties for this type of operation.

The B + S is to be a prestige machine. To be this its goals must be fastest effective computing speed and largest main memory possible. This goal, if achieved, is pointless unless input-output speeds are high enough to support the computing speed. Also, when considering the high cost of memory, it is obvious that the faster the disk unit the less time memory will be occupied with input-output and the lower the cost per application for memory.

Many applications for which B + S is intended require more than 1 million words of disk storage. These jobs require multiple accesses to the same data with a level of random access that is low, yet high enough to make rewind time intolerable if tapes were used. Among these are those employing three dimensionsal networks in which a small problem involves 100^3 points, each point represented by two or more words. These problems require 2 million words up. Fewer points can be postulated but the resultant coarseness of the grid may render any results meaningless. Factor analysis and linear programming matrices of order 1200 x 2000 are also problems in the over 1 million word class.

If we halve disk memory when simultaneouly reading and writing (read from one half and write to the other), then we have imposed two basic restrictions. One of these is capacity, e.g. to use read-while-writing on a network problem, two matrices must be stored, a new and the old, each of .5 million words or less. The other is disk memory layout where every simultaneous operation will have to be predicted so that the input record is not in the same file as the output. On the other hand, without dividing the disk unit into two logical files the matrix can fill memory without awkward layouts.

Access time to date on the disks is so slow when compared to average floating point operation time (100,000* usec vs. 1 usec) that any improvement in effective access time can have tremendous effect; e.g. saving 1 ms of access that cannot overlap computing can mean the saving of enough main frame time to execute 1000 instructions. Multiply this by the number of accesses required per problem or per day to observe the full effect of a 1 ms improvement. The proposed "roll" type of operation can save as many as 28 ms in a large number of cases. This type of operation become even more efficient (i.e. the number of cases in which a saving is made becomes larger) as the number of blocks per track increases.

*maximum access.

One should observe that the current definition of the disk is an output of the Stretch Engineering Planning Group and does not necessarily represent mechanical hardware restraints as would, say, a definition proposed by the San Jose people. Rather it represents a desire to economize on electronics by sharing disk circuits with tape circuits - - a laudable goal if it can be achieved without compromising performance. This memo contends, however, that performance is severely compromized by the current definition and that this goal of circuit sharing is only achievable if performance is regained by improving density.

The San Jose people have already stated that they believe they can increase this density if given funds to undertake the necessary research and development.[3]

If the density is not increased, then the goal of circuit sharing should be abandoned. Performance can then be restored by a new definition that stipulates processing 72 bits in parallel rather than 18, i.e. by providing for 1 bit per track per word rather than 4. This was the original definition proposed by San Jose.

This memo contends that a satisfactory and reasonable definition for the disk units is to return the density to about 400 bits per inch which is the density San Jose reported was possible for commercial disks. This quadroupling of density will permit a word rate of 4 usec per word. Since this satisfies the A.E.C. contract, the simultaneous reading and writing feature may be dropped to meet memory and bus problems. The density improvement should also make it possible to return to San Jose's original track density of 256 tracks per face (double that of the current definition). Thus the capacity of a disk unit will be 8 million words. This capacity plus the elimination of simultaneity of reading and writing solves the problems currently imposed by having two logical files per disk unit so nothing more is required here. Finally, even though rental may be increased somewhat by the density improvements, the cost per word of storage is dramatically reduced, the speed is four times higher, and programming (for read-write simultaneity) has been simplified. Lack of simultaneity here may also reduce main memory requirements in many applications. These changes plus use of "roll" type of operation should be imposed on the current definition; i.e. the other features of the current definition need not be discarded. This includes circuit sharing.

---

3. Haanstra, J. W.; Minutes on Stretch Ramac Meetings, August 7, 8, & 9, 1957; August 9, 1957.

Dr. H. G. Kolsky has been consulted on the question of word rate and has expressed serious doubt that 16 usec per word is adequate.

D. W. Sweeney

J. C. Gibson

DWS:JCG:pw

cc: Mr. G. A. Blaauw
Mr. F. P. Brooks
Mr. J. D. Calvert
Mr. E. F. Codd
Mr. E. W. Coffin
Mr. C. F. Earley, San Jose
Mr. R. A. Gregory
Mr. P. H. Howard
Mr. H. G. Jones
Mr. H. G. Kolsky
Mr. J. Lyons
Mr. H. A. Mussell
Mr. D. W. Pendery
Mr. B. L. Sarahan, WHQ