POUGHKEEPSIE
Dept. 539, Bldng. 965
May 16, 1958

FILE MEMO

SUBJECT: 709-7000A Performance Comparison on a Sparse Matrix Multiplication Problem

By: F. P. Brooks, Jr.

The manipulation of matrices whose elements are mostly zeros is often required in linear programming, scheduling, and transient analysis of electrical networks. This is an example of a technical problem in which logical power is important, for it is almost always desirable to store and handle only the non-zero elements of the operand matrices.

The 7000A operates from 30 to 80 times as fast as the 709 on typical problems of this type. Programs were written to perform multiplication of sparse matrices on the 7000A and 709. Both programs use a similar technique. They represent the operand matrices with lists of non-zero elements and bit matrices whose zero or non-zero bits designate zero or non-zero elements. The 709 program takes advantage of the fast zero multiplication and the indirect addressing features. The 7000A program takes advantage of variable field length connectives, bit testing, and automatic index refilling.

The programs were written to calculate in single precision. Then the necessary modifications were made to calculate in double precision, which is almost always required for 709 programs operating on matrices of fortieth order and above. These modifications considerably lengthen both program times, but the 709 program is lengthened about twice as much as the 7000A program.

Both programs were written to operate on matrices of order up to the word length. The modifications required to operate on larger matrices were then written. In neither program did these represent a significant change in execution time.

The program statistics are given below for four cases, all for matrices with 10% of their elements non-zero. The first, multiplication of 36th order matrices is most favorable to the 709. The second, multiplication of 64th order matrices, shows the increase in 7000A advantage as matrix size increases. In the third case, it is assumed that the 64th order multiplication requires from 28 to 48 bits of mantissa precision. The 709 uses double precision routines, but the 7000A has the required precision normally. The fourth case shows the same matrix multiplication when from 48 to 54 bits of mantissa precision are required, so that both programs must use double precision. Above 54 bits of required precision, the 709 must go to triple precision.

The program statistics follow:

Case I: 36th order, 10% non-zero elements.

	709	7000A	Ratio
Time	6.23 sec	.206 sec	30.2 *
Instructions executed	242,600	23,160	10.5
Time per instruction	25.7 usec	8.9 usec	2.9
Instructions written	73	41	1.8
Words of program storage	73	29.5	2.5
Bits of program storage	2628	1888	1.4

* Ratio greater than 35 in case of 704.

Case II: 64th order, 10% non-zero elements.

Time	33.5 sec	.853 sec	39.3*
Instructions executed	1,297,000	91,314	14.2

Other categories essentially unchanged.

* Greater than 45 in case of 704.

Case III: 64th order, 10% non-zero elements, 28-48 bits precision. 709 only, double precision.

Time	69.3 sec	.853 sec	81.3
Instructions executed	2,425,000	91,314	26.6
Time per instruction	28.6 usec	8.9 usec	3.2
Instructions written	134	41	3.3
Words of program storage	134	29.5	4.5
Bits of program storage	4852	1888	2.6

Case IV: 64th order, 10% non-zero elements, 48-54 bits precision. Both double precision.

Time	69.3	1.384 sec	50.0
Instructions executed	2,425,000	121,542	20.0
Time per instruction	28.6 usec	11.4 usec	2.5
Instructions written	134	52	2.6
Words of program storage	134	37	3.6
Bits of program storage	4852	2368	2.0

Conclusions:

1. The 7000A performs many times more rapidly than the 709 on this problem. The ratio increases rapidly with N, and decreases as the operand matrices become more dense.

2. The high performance ratio is due mostly to the power of the instruction set, since the powerful variable field length instructions sharply reduce the number of instructions executed.

3. The longer 7000A floating point word permits double precision operation to be deferred until very big or ill-conditioned matrices are encountered. The floating point instruction set provides economical and convenient double precision operations when they are required.

4. The 7000A requires the writing of about half as many instructions to accomplish the task. The directness of the instructions written contrast with the circumlocutions necessary on the 709.

FPB/pkb