

MEMO TO:

synchronized

SUBJECT: The use of High Speed Disks in Scientific Problems.

I. The Two Methods of using high speed disks.

If we examine the high-speed disks from the scientific user's point of view, there seems to be two basically different ways in which they can be used: (1) closely-coupled (or synchronized) with the calculation, or (2) decoupled (or buffered) from the calculation. Of course there is an in-between area where parts of a calculation may be synchronized and parts not, but we will assume these can be studied separately, continuously.

When the disk is being read it deposits a word in memory every Δt_D microseconds, where Δt_D is very rigidly fixed time interval. It must be up to the asynchronous calculation to adjust itself to some multiple of this disk interval.

II. Some I/O Properties of mesh-type scientific problems.

In mesh-type scientific problems the fundamental building block of the calculation is the processing of one general mesh point. The major complexity in the code may come from the special cases, the boundary conditions, etc. but the speed of the calculation depends almost entirely on how long it takes to process the individual mesh point. This processing depends on three parts (1) the reading in of M_1 words of data, (2) ^{(3) the} performing N_1 arithmetic and logical operations, and (3) ^{(3) the} reading out of M_2 words of results. Two and perhaps all three of these may be going on simultaneously.

a necessary

A condition for the calculation to be closely coupled to the disk in a mesh point by mesh point basis is that the calculate time be less than the time to transmit the required data ^{to and} from the disk. If the calculation takes slightly longer than the disk time it is still possible to synchronize them by leaving a few blanks on the disk between each section of data. # If the calculation takes about twice the length of the disk time, one can consider skipping every other section on one revolution of the disk and picking them up on the next revolution. Presumably one could carry this interleaving on to more and more revolutions, but there is probably a practical limit of about two or three beyond which one would go over to the completely decoupled mode of operation.

We may write the above conditions algebraically as follows:

Define t_p = time to process one mesh point.

t_c = time to calc. the quantities for one point

t_d = time to transmit the data for one point

n_c = no. of operations to calculate one point

n_d = no. of data words read in and out.

Δt_c = average time for one operation (* define)

Δt_d = time between successive words to or from disk

Then for closely coupled calculations

$$t_p = t_c (1 + \delta)$$