from Computer Measurement and Evaluation, Selected Papers from the SHARE project, ed. Janet Wixson, volume III, December 1973-March, 1975.

SOME CAUSES OF VARIABILITY IN CPU TIME

1. Bit-spin loops - used in MFT when pseudo-disabled during RLDL/FETCH I/O using PANIC workarea - number of loops will be affected by a change in DASD arm position and by a change in location of module being fetched.

   Such loops are also used in MP65 supervisor.

2. Timer resolution - 16 2/3 ms. on 360 - 3 1/3 ms. on 370. Instructions are executed in Us and sometimes ns - the variability caused by this is a function of the length of the interval during which a task has a burst of CPU activity - partially dependent on task characteristics (CPU/IO relationship) - probably more significantly affected by the number and type of concurrently running tasks - i.e. number of task switches.

3. Improper allocation of CPU time for I/O interrupt handling - it is 'charged' to the interrupted task and consists of these elements:

   a. basic IOS time for the initial I/O interrupt

   b. more of the same for any pending interrupts which arise during the disabled time of handling the original request.

   c. IOS time for channel restart operation to process any queued I/O requests which may now be started on the channel(s) and devices freed by the interrupts in (a) and (b).

   d. time spent executing IOS appendages for PCI, EOE, CE, ABE for all of the interrupts handled in (a) and (b) and for SIO for any new operations scheduled via (c).

   NOTE: For MFT, all of these times are now excluded from step CPU time as of OS release 21 (APAR 40380).

4. CPU's using a high speed buffer (360/85 and 195, 370/155 and 165) are subject to variable CPU time due to change in buffer hits from run to run. This will usually occur due to buffer interference caused by concurrent tasks but may also be due to partial or full disabling of the buffer because of storage errors.

5. Cycle stealing on systems with integrated channels (360/50 down and 370/155 down) - CPU instruction execution is temporarily suspended when channels require the use of hardware resources shared with the CPU.

6. Lockout time due to CPU being prevented from access to memory while a channel is using it - this factor will depend on CPU architecture such as interleaving, data widths and paths.

7. MFT - transient area refresh - the dispatcher enqueues a task's TQE before checking if it requires a transient SVC routine which has been overlaid. If a refresh is required, certain preliminary processing is performed using FINCH under the requestor's TCB. The task then

loses control (TQE is dequeued) to the transient area fetch task for the actual fetch operation and is subsequently redispatched (TQE enqueued again) to continue execution in the interrupted transient SVC. Variability is introduced by the additional ENQ and DEQ of the TQE and by the additional FINCH processing executed under the user's task.

8. Number of task switches - a certain amount of time spent in the dispatcher and timer routines before dequeuing or after enqueuing a task's timer are chargeable to that task. As the task switching rate increases, this additional time may become a more significant portion of the CPU time attributable to a task since his bursts of CPU activity tend to shorten.

9. Availability of serially reusable resources - system ENQ routine time will vary depending on whether the requested resource is available. If not available, additional time will be taken to queue up the current request and additional task switches will be necessary. DEQ time will increase if other tasks have subsequently requested the resource which the current task is releasing.

10. DASD space allocation - if the number of extents is not exactly the same as before, additional end of extent processing will be required.

11. When a task requires the use of transient (type 3 or 4 and not in LPA) SVC routines, execution time may vary depending on whether or not the requested routine is already in a transient area.

12. Temporary I/O errors will have an effect on time - additional SVC's (15 - ERREXCP, 16 - PURGE, 55 - EOV) may be required. In MFT without subtasking, execution of error routines (IGExxxxx) is performed under the user's TCB which further increases the variability due to temporary errors.

13. The time to process a getmain request for SQA will vary with the length of the FQE chain existing at the time. SQA freemain time will vary depending on the status (free or allocated) of the adjacent areas.

14. The time to process an EXCP request (SVC 0) will vary depending on the availability status of the requested channel and device. Some possibilities are:

    a. The device and primary channel are available - the SIO preparation and execution can be done immediately.

    b. The device is free, the primary channel is busy, but one of the alternate channels is available. Time to test alternate channels is added to time required in (a).

    c. All paths are busy and the request must be queued.

    The times for (a), (b) and (c) are individually different and

further, for (c) if priority or ordered seek queuing is used, the enqueue time itself will vary somewhat depending on the number of other I/O requests queued and the processing necessary to find the correct location in the chain for insertion of the current request.

15. Timer enqueue routine processing time for task and jobstep timing requests will vary with the number of elements on the timer queue which must be passed over to find the proper slot for the current request - this depends on the TOX of other elements. Variability could be introduced by something as basic as the time of day at which your job is running.

16. If extensive use is made of BPAM, a time difference for processing BLDL/FIND requests will be experienced if there has been a change in the PDS directory. A change in the location of the entry for the required member will be reflected by a change in the time to scan the block containing it.

17. The occurrence of timer/external interrupts may cause variation from run to run since all processing handled directly off the associated FLIH is charged to the interrupted task. This variation might be particularly noticeable if your job is subsequently run with, say, a time interrupt driven software monitor. The incorrectly allocated time would include that required for stage 2 exit effector processing or pseudo-clock updating and for handling the next TQE to be placed in the interval timer.

18. Recovery management processing - CCH/MCH processing will be done under control of the interrupted task.

19. If GTF (release 21 on) is active, a user task will experience an increase in CPU time dependent on what system functions (SVC, SIO, IO, PCI, DSP) have been selected for current GTF recording. If USR functions are to be recorded and the application contains GTRACE macros a more pronounced difference may be noted.

20. STOW - processing time may vary because of a difference in the PDS directory which would cause reordering or bumping to be done.

21. ENQ service time can vary with a different environment of resources enqueued due to a difference in the number of queue control blocks to be examined and chained and whether or not storage must be acquired for new QCB's.

22. WTO, WTOR and WTL processing time can vary depending on the processing necessary to find a free WQE and/or RQE and possibly on whether a getmain has to be done to build a new element. If the WTO or reply elements are at their limit, additional time will be required for enqueuing.

23. The transient area handler (MVT) processing necessary to find a TA to overlay (if necessary) increases with the number of tasks currently

using a transient routine. The variable processing includes that required to find all current users and determine their relative priorities to each other and to the current request.

24. Transient area refresh processing increases with more concurrent tasks and more entries on the TA user and request queues. With longer queues more time is required to determine the highest priority user of each transient area.

25. Transient area retry (subsequent attempt to find an available TA after first attempt found all to be in use) processing time will increase as the number of tasks using transient routines increases. There will be an increased probability of building up a queue of deferred requests; retries for all pending requests are scheduled whenever any SVC3-EXIT is issued and at least one area is free. If transient area contention is a bottleneck in the system, this factor may cause a time increase whose relationship to the number of queued requests is possibly worse than exponential.

26. A change in the LPA module contents or the resident BLDL list(s) will cause a time variation. If the required module is now included in the LPA (or BLDL list) a time reduction will be experienced due to reduction in processing for SVC, LINK, LOAD, XCTL, ATTACH and BLDL requests. If the LPA is expanded but not for the requested module, more time will be taken due to more CDE or BLDL chain entries to scan.

27. Allocate and scratch time will be somewhat affected by a change in the VTOC. A difference in the format 5 (free space) DSCB(s) will require a change in the amount of time required to locate an available slot or the time to merge the space being freed with adjacent free areas.

28. The dispatcher processing which is done prior to dequeuing the old task's timer is increased if asynchronous exits are to be scheduled. This includes time for examining AEQ's, queuing IQE's and RQE's to IRB's and IRB/SIRB initialization.

--Brian Currah, London Life
Insurance Company