

To: Jerry Saltzer  
Bob Daley  
Corby  
Ted Glaser  
Carla Marceau

From: Art Evans

Date: June 8, 1967

A basic concept in the Mark II scheduler is the idea of an interaction. As it happens, it is not ~~immediately~~ immediately obvious just what an interaction is in Multics, or how to tell when one has taken place. The attached document attempts to raise some of the relevant issues and to suggest a ~~possible~~ possible solution. Comments are requested.

The second Multics scheduler <sup>uses</sup> a multi-queue, CTSS-like algorithm. Basic to ~~this~~ <sup>the idea of</sup> this algorithm is an interaction. Intuitively, it seems that an interaction takes place when ~~the~~ a user, seated at a console, has completed a "think time" and has ~~not~~ given the system something to do. A basic idea in Multics is that such a user should be favored, at least to the extent of giving him ~~extra priority the next time he is scheduled.~~ <sup>extra priority the next time he is scheduled.</sup>

~~The~~ The mechanism of favoring him is not relevant here -- instead, we are concerned with just what ~~an~~ an interaction is and how we can tell one has taken place.

This problem has no real counterpart in CTSS. There, ~~an~~ an interaction is defined to take place whenever a program goes into "input wait" status. (Output wait also produces an interaction -- ~~in aspect of CTSS that many find curious.~~ <sup>in aspect of CTSS that many find curious.</sup>) In the nature of things, a CTSS program can be in input wait only if it is waiting for console input, and it is ~~reasonable~~ reasonable to regard the arrival of that input as motivation for high priority scheduling. (One can "beat" the system with judicious use of interconsole messages.)

Because of the read-ahead feature in Multics, it is harder to tell just ~~when~~ when an interaction has taken place. We want to be sure that an interaction is reported only when ~~the input actually arrives from the device.~~

(a) a process must block itself for lack of console input, ~~and~~ and

(b) the input ~~actually~~ actually arrives from the device.

The problem has to do with the diffuseness of the code that processes console input. Let us consider the ~~case~~ case of input from a console with type-ahead. At some stage the working process asks for input. In due course, <sup>the</sup> a Device Strategy Module (DSM) in <sup>the process'</sup> Ring 1 considers the request. If there is available text that has <sup>already</sup> been <sup>typed,</sup> it is merely returned to the working process. If not, however, the DSM must wait for text, by calling the Wait Coordinator. It is at this time that ~~the~~ the DSM knows that there is the possibility of an interaction. (As we will see, events may transpire that result in no interaction's taking place.) ~~\_\_\_\_\_~~  
~~\_\_\_\_\_~~

The DSM actually gets input from a Device <sup>different process</sup> Control Module (DCM) ~~in a~~ with which it communicates via suitable Event Channels and shared data bases. To simplify a ~~complicated~~ complicated situation, the DSM leaves its request in coded form in a place available to the <sup>DCM</sup> and then informs the <sup>DCM</sup> of the existence of the request. Now we return to our case. ~~The~~ The DSM in its request for the next line includes a bit indicating that the line may well produce an interaction. Normally, when the <sup>DCM</sup> gets ~~input for~~ data for a process, it makes it available in <sup>the</sup> shared data base and calls wakeup on behalf of the ~~working~~ working process. In the case when the interaction bit ~~is~~ is on, however, the <sup>DCM</sup> calls a special entry of wakeup (or, perhaps, supplies a ~~special~~ special parameter) to indicate that the working process has <sup>apparently</sup> experienced an interaction. Wakeup will <sup>then</sup> set an interaction bit in the working process' Active Process Table (APT) entry before calling ready\_him. The working process' scheduler, observing this bit, will take appropriate

scheduling action and then reset the bit. The cycle is complete. <sup>To recoup</sup> The DSM, asked for input when there is none, sets a switch indicating that an interaction is possible. The <sup>DSM</sup> observes the <sup>switch</sup> and sets a bit in the APT indicating that the interaction has taken place. The scheduler then gives the user special action as desired and resets the bit so that the user will not be given special consideration twice for one interaction.

Although it should be clear that the scheme described will work, the perceptive reader should have the feeling that it is overly complicated. However, all of the complexity is necessary, as we now show. ~~Some of the~~ We first show why both the DSM and the <sup>DSM</sup> must <sup>take</sup> part <sup>in</sup> the decision, and we then present some more subtle problems.

It would not be possible for the <sup>DSM</sup> alone to detect an interaction. Consider the situation in which the <sup>DSM</sup> indicated an interaction whenever it called wakeup on behalf of a process. The clever user will, in his working process, start the I/O system in read-ahead input mode, and then go about a long computation. The man at the console can then type carriage return every few seconds, secure in the knowledge that he is thereby moving his working process to the top of the queue <sup>each time</sup>. (The working process never asks for input.) What this solution misses is the ability to know when the working process ~~cannot~~ cannot proceed without input.

~~There are also~~ There are also problems in trying to detect interaction <sup>in</sup> the DSM alone. One might propose the following solution: On realizing that input is requested and not available, the DSM before going blocked would set an interaction bit in the process' APT entry. This would then entitle the process to high priority on its next scheduling. Unfortunately, this

solution also can be beaten. Setting the interaction bit this way has the effect that the process gets priority on its next scheduling, no matter why the scheduler is called. The clever user arranges to have some other friendly process send his working process ~~an event~~ periodically ~~over an Event Call Channel~~. Then the working process would ask for input every two minutes. The person would carefully keep his hands off the keyboard so that ~~these~~ these input requests would all produce calls to block with the interaction<sup>bit</sup> set, and then the friendly event would result in scheduling with priority.

It should now be clear that both the DSM and the ~~DCN~~<sup>DCN</sup> must contribute to the decision that an interaction has transpired. The DSM knows that work cannot proceed without input, ~~the~~<sup>DCN</sup> the ~~DCN~~ knows that input has arrived, and both ~~are~~ are needed. There is ~~one~~ one more problem: Consider the interactive user who, in typing, "gets ahead" of his process. That is, he supplies data faster than the process, considering the share of the processor available to it, can eat it up. Stated differently, in the quantum available to it the process does not use up all ~~the~~ available input. Then on succeeding executions things are worse, since each is at lower <sup>(This problem exists in CTSS)</sup> priority. In some sense it seems intuitive ~~that~~ that this user is interacting and ~~is~~<sup>is</sup> entitled to ~~some~~ preferential treatment. Unfortunately, however, there ~~is~~<sup>seems to be</sup> no way to give him ~~some~~ priority without opening an immense loophole to beat the system. We must stick to our decision that an interaction has taken place only if the process cannot proceed without input and if the input then becomes available.

Firm ~~the~~ adherence to this principle produces one more change to the algorithm as described. We said that wakeup, when called at its priority entry, would set the interaction bit in the working process' APT entry before calling ready\_him. We now add the proviso that it do so only if the working process is currently blocked. ~~the~~ Consider <sup>a</sup> ~~the~~ process<sub>a</sub> <sup>which</sup> blocks waiting for input, and suppose that while waiting an event arrives for it on an Event Call Channel, producing a wakeup. The <sup>input</sup> ~~input~~ comes from the console while the process is either ready or running as a result of this wakeup. We do not in this case give priority because ~~the~~ one critical requirement for an interaction is missing: that the process be unable to proceed without input. (Clearly, the process was proceeding.) This ~~part of~~ part of the algorithm does not close any loopholes, but it ~~is~~ is consistent with announced principles.

A few problems ~~remain~~ remain:

1. For there to be an interaction, must the input come from the command console ~~or~~ or is any attached console good enough?
2. Presumably the Multics equivalents of QUIT and RSTART should produce an interaction. Also, QUIT should be processed with high priority. We have yet to see how.
3. We do not yet ~~know~~ know how to call wakeup so as to indicate that an ~~interaction~~ interaction has taken place.