March 3, 1968

TO: John Gintell

FROM~~P~~: J. H. Saltzer

SUBJ: Supervisor Performance Measurement

Your suggested supervisor measurement facility is quite interesting and I think that with a few minor modifications it could be made into very powerful tool for analyzing goings on within the supervisor. Consider the following minor extension of your scheme:

Define a four bit quantity called the fault state. The first bit is one whenever a page fault has occurred, and remains one until the page fault is completely processed and return has been made. Similarly, the second bit is one during the entire processing of a segment fault, including the time when a page fault may occcur. The third bit is set to one during handling of wall crossing faults and the fourth during linkage faults.

We observe that there are sixteen possible fault states, although because of supervisor restrictions as well as the fact that only one fault can occur at any instant, there are not 256 possible state transitions but only the twnety indicated in the figure, and twenty ~~the/twenty~~ which are the reverse of the paths indicated.

Since the fault interceptor has control for an instant during the begin and the end of each fault processing, it can maintain a record of the current fault state in a wired down metering segment. It can also observe state transitions, and count them in a 256-word array in another area of the metering segment. Finally, it can record the time spent in each state/~~16~~ in a 16-word array, using either real time when testing with one process, or process time when using a multiprocess system.

Finally, the fault state can be made accessible to the segment usage metering procedure, and it can be instructed to examine the bits of the fault state to decide whether or not ~~to count update~~ the metering interupt should be counted. By allowing specification of "must be zero", "don't care", or "must be one" for each of the four bits separately, it will be possible to meter segment usage while in any combination of fault states desired.

The information resulting from such a facility will give precise indication of how the supervisor is recursively using its various features. ~~In particular~~ For example, it will be possible to distinguish ~~see how many~~ page faults which occur as a result of wall crossing from all others by comparing the number of times state "1101" and "1111" occur as compared with all other states with a "1" in the ~~fourth~~ first bit. Similarly, it will be pssilbe to see if faults resulting from some sources take different average times than faults resulting from others.

I think that this scheme avoids problems of stacking fault information, and possible missing-page faults while metering state transitions.

If the scheme is implemented in machine language with an eye toward efficiency, it would probably be possible to make this particular set of meters a permanent fixture of the supervisor at very low cost. At some later time print or display programs could format the data in more appropriate forms, but for the moment the usual octal dumps will at least get some information out.

~~PSW~~ = $\overset{1111}{0000}$
LWSQ

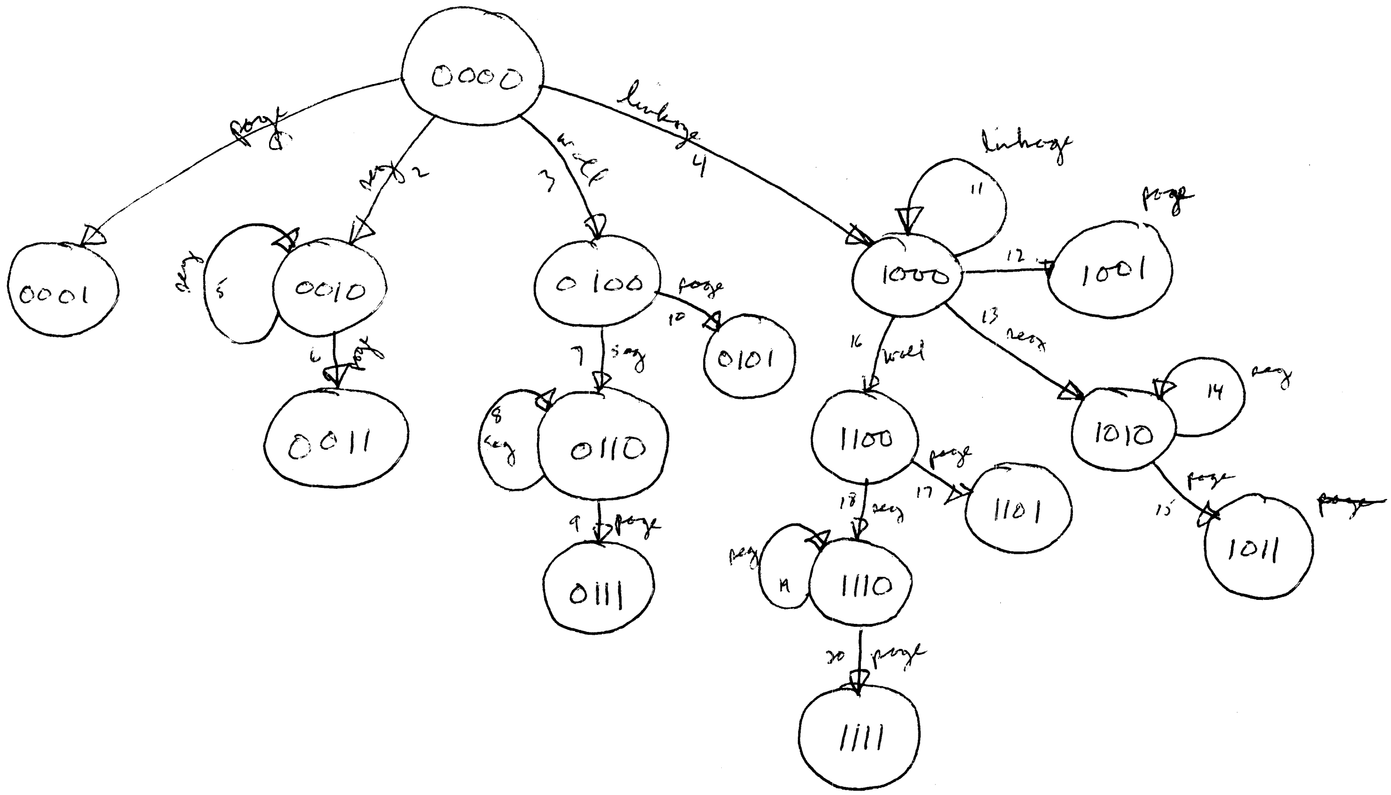

# of states 16

# of paths 20

1. Measure 1. # of times each path is exercised
   2. total time spent in each state

2. Use state switches to determine whether or not to enter seg usage.
   (for each bit, specify MBZ
   DC
   MBØ)