

## Controlling Extended Core Storage Access Limits

### The Problem

The Extended Core Storage (ECS) address space available to the 6400 CPU is controlled by a basing register (ECSRA) and a limit register (ECSFL) in the same way that the central memory address space is defined by the Reference Address (RA) and Field Length (FL). The 18 bit values of RA and FL are maintained in words 2 and 3 of the exchange jump package and are transferred to the appropriate central processor registers by the Exchange Jump instruction(s). The ECS counterparts of the central memory RA and FL are allocated 24 bits each in words 5 and 6 of the exchange jump package. However, only 18 bits are maintained in the ECSRA and ECSFL registers of the CPU. In particular, the low order 6 bits of ECSRA and ECSFL are wired down to zero, and any information to the contrary in the exchange jump package is ignored by the hardware. The restriction implies that access to ECS can only be controlled to the nearest 100B words in ECS.

### Usage of Extended Core Storage

Unlike other random-access mass storage devices, the access time to any word in ECS is not hindered by rotational latency or positioning delays. In fact the difference between the access time for ECS and for rotating mass storage devices is so great that one may consider new strategies in utilizing this fast access mass storage. Access delays associated with disc and drum necessitate the transfer of a large number of bits in order to offset the expense of the long access time. Large records on disc and drum are reflected in the methods used to allocate these devices. Usually, the hardware itself has a fixed length, addressable block comprising the minimum data transfer.

On the other hand, the small access penalty (approx. 4  $\mu$ sec) associated with ECS makes it feasible to transfer small records of arbitrary length from ECS to CM. The system design no longer need tolerate the wasteful "breakage" currently associated with fixed length block or record

sizes. Small blocks could now be used to hold infrequently accessed but logically independent objects. While some hardware advantage could be gained by allocations which initiate transfers on 8 word boundaries, the large fixed length units of allocation used with disc and drum could be avoided.

Furthermore, it would also be desirable to give the user direct access to a contiguous portion of ECS, thereby allowing him to access shared objects or I/O buffer areas. If files (or portions of files) are stored in ECS, direct user access to frequently used or shared data would enhance system performance.

However, in these respects, the hardware is restrictive; since the ECS access limits (ECSRA and ECSFL) can only be set at 100B word intervals, the control of user access is limited to 100B word blocks. This restriction, dictated by the "missing" bits in ECSRA and ECSFL, not only shatters the hope of avoiding the wasted space associated with the "breakage" of block allocation, but also renders impractical the use of small areas of shared data and small system objects. The high cost of ECS makes this restriction seem unreasonable.

#### Summary

Although the fast access speed of ECS seems to encourage variable length block allocation and the use of small blocks, and although system performance and versatility would profit from controlled user access to ECS, the limitations in the ECS access control registers make these two goals incompatible.