COMPANY PRIVATE
RETURN TO RICK ELLINGER
AUTO./MEAS./DIV.

RSX-11D SPEC

TO: RSX-11D Distribution

FROM: H. Krejci

DATE: 17 Jun 72

SUBJ: MEMORY MANAGMENT

DOC: 130-101-039-00

The material included in this functional specification, including but not limited to, instruction times and operating speeds is for information purposes only. All such material is subject to change without notice. Concequently DEC makes no claim and shall not be liable for its accuracy.

Unless specified otherwise, the terms "RSX" and "RSX-11" imply "RSX-11D".

Under RSX-11D, memory is divided into (1) an area for the executive's code, (2) an area for the system's tables, lists, and subroutines, (3) an I/O Handler Library, (4) a Utility Library, (5) a Buffer Pool, (6) Global Common Blocks, and (7) Task Partitions. These memory divisions and allocations are fixed at System Configuration (SGEN).

The I/O Handler Library is a mechanism for sharing code (reentrant) that is common to more than one (large) Handler Task, and is used primarily for file structure code. The Handler Library is accessable only by Privileged Tasks. This library may not be required in all systems, and is optionally created at System Configuration.

The Utility Library is a mechanism for sharing code (re-entrant & position-independent) that is common to more than one Task. It is available to all Tasks[1] and is used primarily for the FORTRAN Object Time System (OTS). This library may not be required in all systems, and is optionally created at System Configuration.

The Buffer Pool is an area of memory to be allocated to privileged Tasks for their internal usage. The primary use of the Buffer Pool is by Handler Tasks that provide a file structure. A Buffer Pool may not be required in all systems, and is optionally allocated at System Configuration.

Global Common Blocks are COMMON areas that are intended for Inter-Task communication and extra-Task data storage. As many Global Common Blocks, including none, may be defined at System Configuration.

Task Partitions, are areas of memory dedicated for the execution of Tasks. A Partition name is specified, or implied, whenever a Task image is loaded from disk into memory. At least one Task Partition must be defined. Each Task Partition in a system is identified as being either a "User Controlled Partition" or a "System Controlled Partition".

A User Controlled Partition can accommodate only one Task at a time, regardless of its size. If all Task Partitions are User Controlled Partitions, then all dynamic memory allocation is under the control of the user.[2]

A System Controlled Partition will accommodate as many Task as can fit in the Partition ("first fit" algorithm). If there is only one Task Partition in a system, and it is a System Controlled Partition, then all dynamic memory allocation is under the control of the system.

With either type partition, a checkpointable Task will be written out if space within its partition is required for a higher priority Task. A checkpointed Task is reloaded when sufficient space within its partition is available and not required for higher priority Tasks. A checkpointed Task is always brought back into the same partition it was initially loaded, but not necessarily at the same location within that partition.

---

[1] Privileged Tasks are not mapped into the Utility Library, but because they can set ASRs, a Privileged Task can map itself into the Utility Library.

---

[2] The EXECUTE, FIX, and UNFIX Directives are particularly useful in this case.

With either type partition, a non-checkpointable Task will remain resident thru completion (until it issues an EXIT Directive), even if its memory is required for a higher priority Task.

System controlled memory allocation and checkpointing are powerful features, but their use does increase system overhead and can cause a significant increase in the time required to complete a Task's execution (due to excessive checkpointing).

The number and type of Task Partitions, and the use of checkpointing, is dependent upon the predictability of an application and the judgment of the person in charge.