

II. THE SWAPPER

The primary goal of a time-sharing system, as opposed to other methods of computer utilization, is to provide each of a number of users with immediate access to data and to processing facilities. Each user is given the illusion that he has complete control over the functions of the computer system. Time-sharing systems achieve this illusion by dividing the total available computing time into small discrete time units (time quanta) and distributing these time quanta among the system's users according to some predetermined queueing algorithm. Because of the extremely high speeds of computers, each user still feels that his job is being performed at computer speed. In the case of an interactive system, in which the user program is delayed by some slow I-O device (e.g., Teletype) or by the time spent by the user in thinking about his problem, a given user program requests time quanta for computation relatively infrequently. Thus, a program engaged in computation is normally in competition for system resources not with all other users but with only a small percentage of the user population.

Some time-sharing systems utilize a "paging" concept which allows some portion of many user programs to be present in core memory at the same time. The system decides when a new portion (page) of a user program is required in core, allocates a space for the new page, and retrieves it from some bulk-storage device. The Hospital Computer System permits only one user program to be in core memory (for computation) at any one time; the space available is 4K words and the user may perform his own program segmentation. When the time quantum allocated to a user program has elapsed, or when the program becomes "hung" waiting for an external action