

~~Saltzer~~
file - Ushin
proposal

TO: F. J. Corbató
R. C. Daley
J. W. Gintell

FROM: R. J. Feiertag, V. L. Voydock

DATE: May 17, 1972

SUBJECT: High Speed Interprocess Communication

This document proposes some new primitives which will provide high speed paths through interprocess communication. The primitives are carefully chosen to provide the high speed paths for the most commonly occurring cases of interprocess communication, i.e. a process waiting for an I/O device.

The basis of the high speed path is the special event channel. The special event channel is identical to an ordinary event channel with the exceptions that no message may be passed on a special channel and wakeups do not queue in the channel. In other words the special channel is a binary switch indicating whether or not a wakeup is pending on the channel. (It is possible to add queueing without loss of speed but this will not be done initially.) A primitive will be provided to create special channels and the special channels will be identified by a fixed bin(71) number as are ordinary channels. Each process will have only a finite number of special channels (initially 6) and they will be dispensed on a first come basis. Initially special channels may not be converted to event call channels.

why?
Not
Kathleen
event
bin 71
special
channel

why
etc?

1.6 + C1
make the 1.6
done on
process IPC
interface

The user ring block primitive will be rewritten. This new primitive will be functionally identical to the current block primitive. It will, however, be much faster and take fewer page faults than the current block primitive in cases where it must wait on special channels only and when no event call wakeups are pending. All of the cumbersome ECT and ITT mechanisms now always invoked by block will be completely bypassed when the above conditions are met. However, no loss of generality results because when the conditions are not met, the regular IPC mechanism will be invoked. The typewriter and possibly some other IOSIMS will be modified to use special channels. Note that the conditions for high speed are those usually encountered in most processes.

For closed subsystems that are willing to sacrifice generality for even further increase in speed, an entry directly into the supervisor will be provided. This entry will be passed a single special channel and will wait for a wakeup on that channel. Note that using this mechanism means that most of the normal IPC mechanisms, e.g. event call channels, will not work. Users that are willing to sacrifice the normal IPC mechanism to use the direct supervisor entry will have eliminated the user ring call. This entry will also return if any IPS (e.g. quit) has occurred. In this case a special code will be returned and the caller should immediately call again.

*Want something close to this or
only new mechanism.*