****
   Paging continued
****

== Ways of allocating physical frames to processes

* How many number of pages should be allocated to each process initially:
  The boundaries for number of pages that can be allocated to a process:
       Min: Minimum number of pages a single instruction of the process requires.
       Max: Frame available in physical memory.

* How to divide the available physical frames among processes:
 -> Equal division :
       Not practical, as each new process will cause the page distribution to
       change.
 -> On process demand:
       The process can request the initial number of pages as per its need.
 -> Total virtual memory requested by process:
       The total virtual memory request by the process can be used to estimate
       the amount of pages should be allocated to it initially.
 -> Priority of the process:
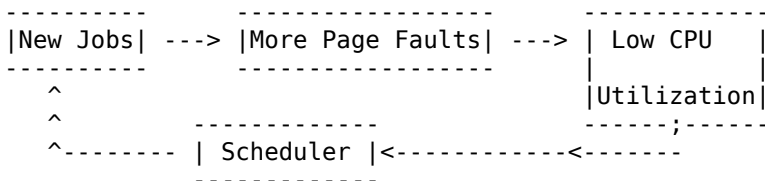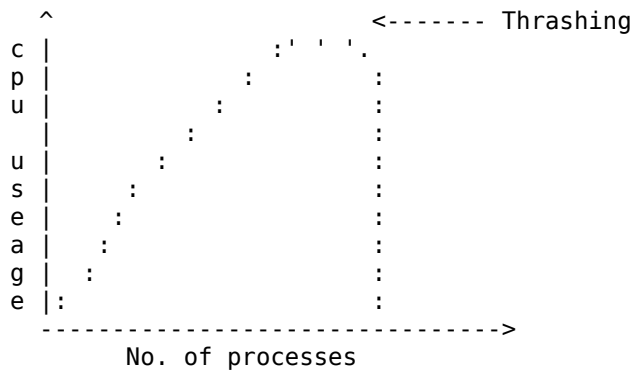   High priority process can be allocated higher number of pages.

== Page Replacement

 -> Local Page replacement:
       In cases of page fault a page from the same process(which caused
       page-fault) is removed to place the new page.
 -> Global Page replacement:
       Any of the total pages in memory is chosen to be paged out, for the new
       page. Less predictable performance, as performance for same process
       at each run can vary depending on other processes in system.

 -> Location of page in physical memory matters in case of multiple
    processors. A page on the same board must be chosen for a process
    running on one processor.
    Also applicable on NUMA (Non Uniform Memory Access) architectures.

== Disadvantages of Paging

 -> Makes the process execution slower
 -> If number of frames allocated < Locality ----> Thrashing
    * Locality: Set of pages that a process calls at a time.
    * Thrashing: is a situation in which the system keeps on dealing with
      only page faults with no actual progress.

                    === Thrashing ===

       ----------   ------------------   ------------
       |New Jobs| ---> |More Page Faults| ---> | Low CPU |
       ----------   ------------------ | |
          ^ |Utilization|
          ^   -------------   ------;------
       ^-------- | Scheduler |<------------<-------

```
       ^                        <------- Thrashing
   c |                  :'  '  '.
   p |              :            :
   u |            :              :
     |          :                :
   u |        :                  :
   s |      :                    :
   e |     :                     :
   a |   :                       :
   g |  :                        :
   e |:                          :
     -------------------------------->
              No. of processes
```

    Processes -> { P1, P2 ----- Pn}

                   _ __
    To avoid thrashing: _>__   locality (Pi) > Physical Memory

     -> Computing Locality: Locality can be expressed as working set. Working
        set is the number of different pages accessed in delta time.

       e.g. P1 :   1 2 {7 4 3} 3 7 4
                         ^

                         |
                       | dt | <---- working set

== Other Paging Issues

 -> Pre-paging: Putting next possibly pages to be used into the memory
                 initially. Opposite of demand paging.
 -> Page size:
     = Small Page size -> Low internal fragmentation
     = Large Page size -> Smaller page table
                       -> fewer page faults
                       -> Better usage of TLB, increased hit rate and better
                          coverage of process virtual memory
 -> Page Locking: Locking mechanism (using locks)to avoid replacement of a
                   page. Can be used to save the process form dropping the
                   critical task.
                   e.g. while performing I/O operation
```