

Virtual Memory: Working Sets

Questions answered in this lecture:

How to allocate memory across competing processes?

What is thrashing? What is a working set?

How to ensure working set of all processes fit?

Allocating Memory across Processes

Scenario:

- Several physical pages allocated to processes A, B, and C. Process B page faults.
- Which page should be replaced?
- Three options...

Per-process replacement

- **Each process has separate pool of pages**
 - Fixed number of pages (e.g., Digital VMS)
 - Fixed fraction of physical memory ($1/P$)
 - Proportional to size of allocated address space
- **Page fault in one process only replaces pages of that process**
 - Perform replacement (e.g., LRU) over only those pages
- **Advantage: No interference across processes**
- **Disadvantage: Potentially inefficient allocation of memory**
 - How to handle sharing of pages?

Allocating Memory across Processes

Per-user replacement

- Each user has separate pool of pages
- Advantage: Fair across different users
- Disadvantage: Inefficient allocation

Global replacement

- Pages from all processes lumped into single replacement pool
 - Example: Run clock over all page frames
- Each process competes with other processes for frames
- Advantages:
 - Flexibility of allocation
 - Minimize total number of page faults
- Disadvantages
 - One memory-intensive process can hog memory, hurt all processes