
Scientific Computing 1
Handout 8
December 3, 2014

Memory Architecture and Memory Management

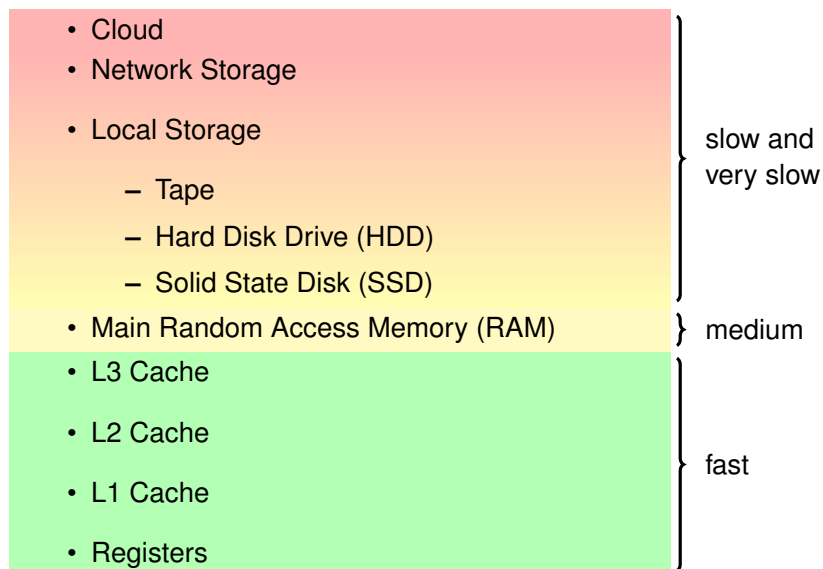


Figure 1: Memory Classes in Scientific Computing

- Operating system provides one abstract **virtual memory** to disguise local data storage (RAM, HDD, SSD) and fragmentation from users.
- virtual memory arranged in **separate virtual address spaces** for all process.
- smallest memory unit is a **page** of at least 4k Bytes.
- by default data can be anywhere in a page. **Page aligned** memory guarantees data to start at page start. **Page locked** memory can not be swapped.
- Processes trying to access data in foreign address spaces are aborted by a `SIGSEGV`. Processes that access symbols from a shared object library that was replace during their operation inaccurately cause abortion by `SIGBUS`, the bus error signal.

Cache very fast small memory portion that is especially close to processing units.

L1 closest to processing units. Separated for instructions and data. Arranged per processor core

L2 Larger, but shared by data and instructions. Also available per processor core.

L3 Largest of the three. Connects to main memory. Shared by all cores of the processor.

Main Memory The general purpose volatile storage system. Operates at average speed.

Local Storage Rather slow and should be used for **double buffering**, i.e., caching of unused data to slower memory to free up main memory for temporary data.

Network Storage and Cloud Space Mainly useful for storing final data. The exception are Infiniband Server Network storage systems that operate at similar speed as local disks.

Secret of a fast method: program and data locality