



First INFN International School on Architectures, tools and methodologies for
developing efficient large scale scientific computing applications

Ce.U.B. – Bertinoro – Italy, 12 – 17 October 2009



Andrew Hanushevsky: Exercises



The Setup

- First, create a 53687111 byte file
 - `dd of=~testfile if=/dev/zero bs=536871112 count=1`
- Second, locate the cache clearing command
 - **clearcache**
 - `/home/nfsexp/track2_furano/Track2progs/clearcache`
- Make sure to issue this *before* each test
 - This will avoid FS memory caching effects

Excercise 1: Simple Copy Program

- Create two copy programs
 - 1st reads, say 1MB, into a buffer and writes it
 - 2nd that does the same using segmented **mmap()**
- Use simple command line arguments
 - E.g., *mycopy infn outfn*
- Use the **time** command to measure performance differences between the two

Excercise 1: Optional Part

- Create a copy that uses **O_DIRECT** open flag
 - Just do it for the **mmap()** version
 - There are many caveats here...
 - **O_DIRECT** requires you write whole pages
 - Some file systems will allow sector sized writes
 - Requires deft programming for the last file segment
 - Not trivial
- Use **time** command to measure performance differences

Excercise 2: lio_listio

- Augment the previous **mmap()** copy to. . .
 - Use **lio_listio()** to create two output files
 - Use **LIO_WAIT** for simplification
 - Simple command line api
 - `liocopy infile outfile1 outfile2`
 - Make sure to include **-lrt** with **g++**
- Measure difference between two successive copies and a single one creating two files