



Gelato

UNSW

Performance and
Scalability on Itanium

www.gelato.unsw.edu.au



Filesystem-related Work at UNSW

Peter Chubb

Gelato Project

peterc@gelato.unsw.edu.au




Computer Science and Engineering

UNSW

Thursday 15th October



- Not working on Cluster filesystems
yet
- Some filesystem-related work:
 - Allow large block devices
 - Improve I/O performance
 - User-mode drivers for disk and network





FS (blkosz)	file limit	filesystem	notes
ext[23] (1k)	16Gb	2Tb	Slow to check, wastes space, needs regular fsck
ext[23] (4k)	2Tb	16Tb	
JFS (4k)	9Eb	4Pb	Not for IA64 (yet)
XFS (4k)	9Eb	9Eb	VM-friendly fsck
ReiserFS 3.5	1Eb	16Tb	Uses lots of physical memory
ReiserFS 4.0	?	?	Not yet mainstream

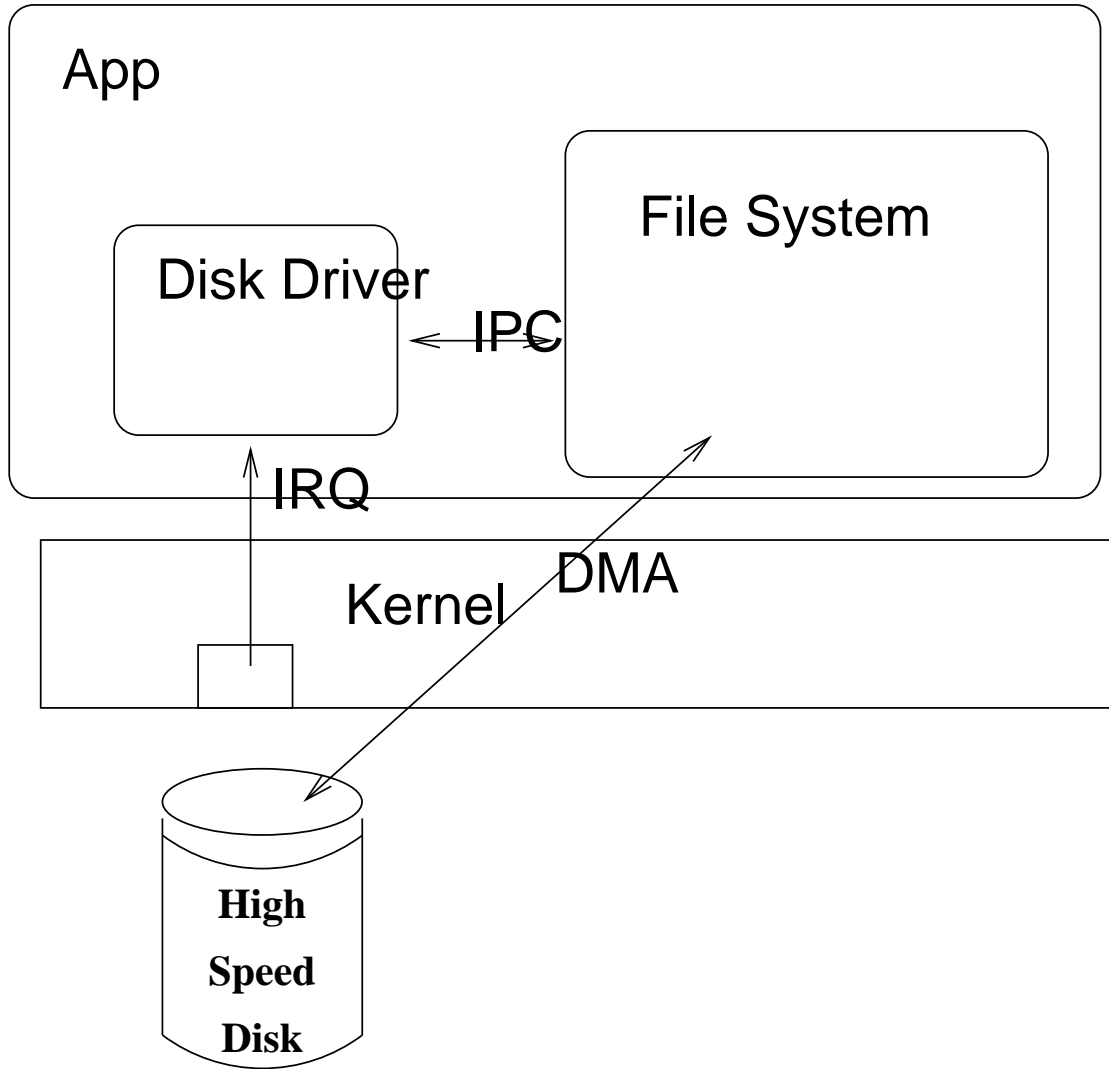


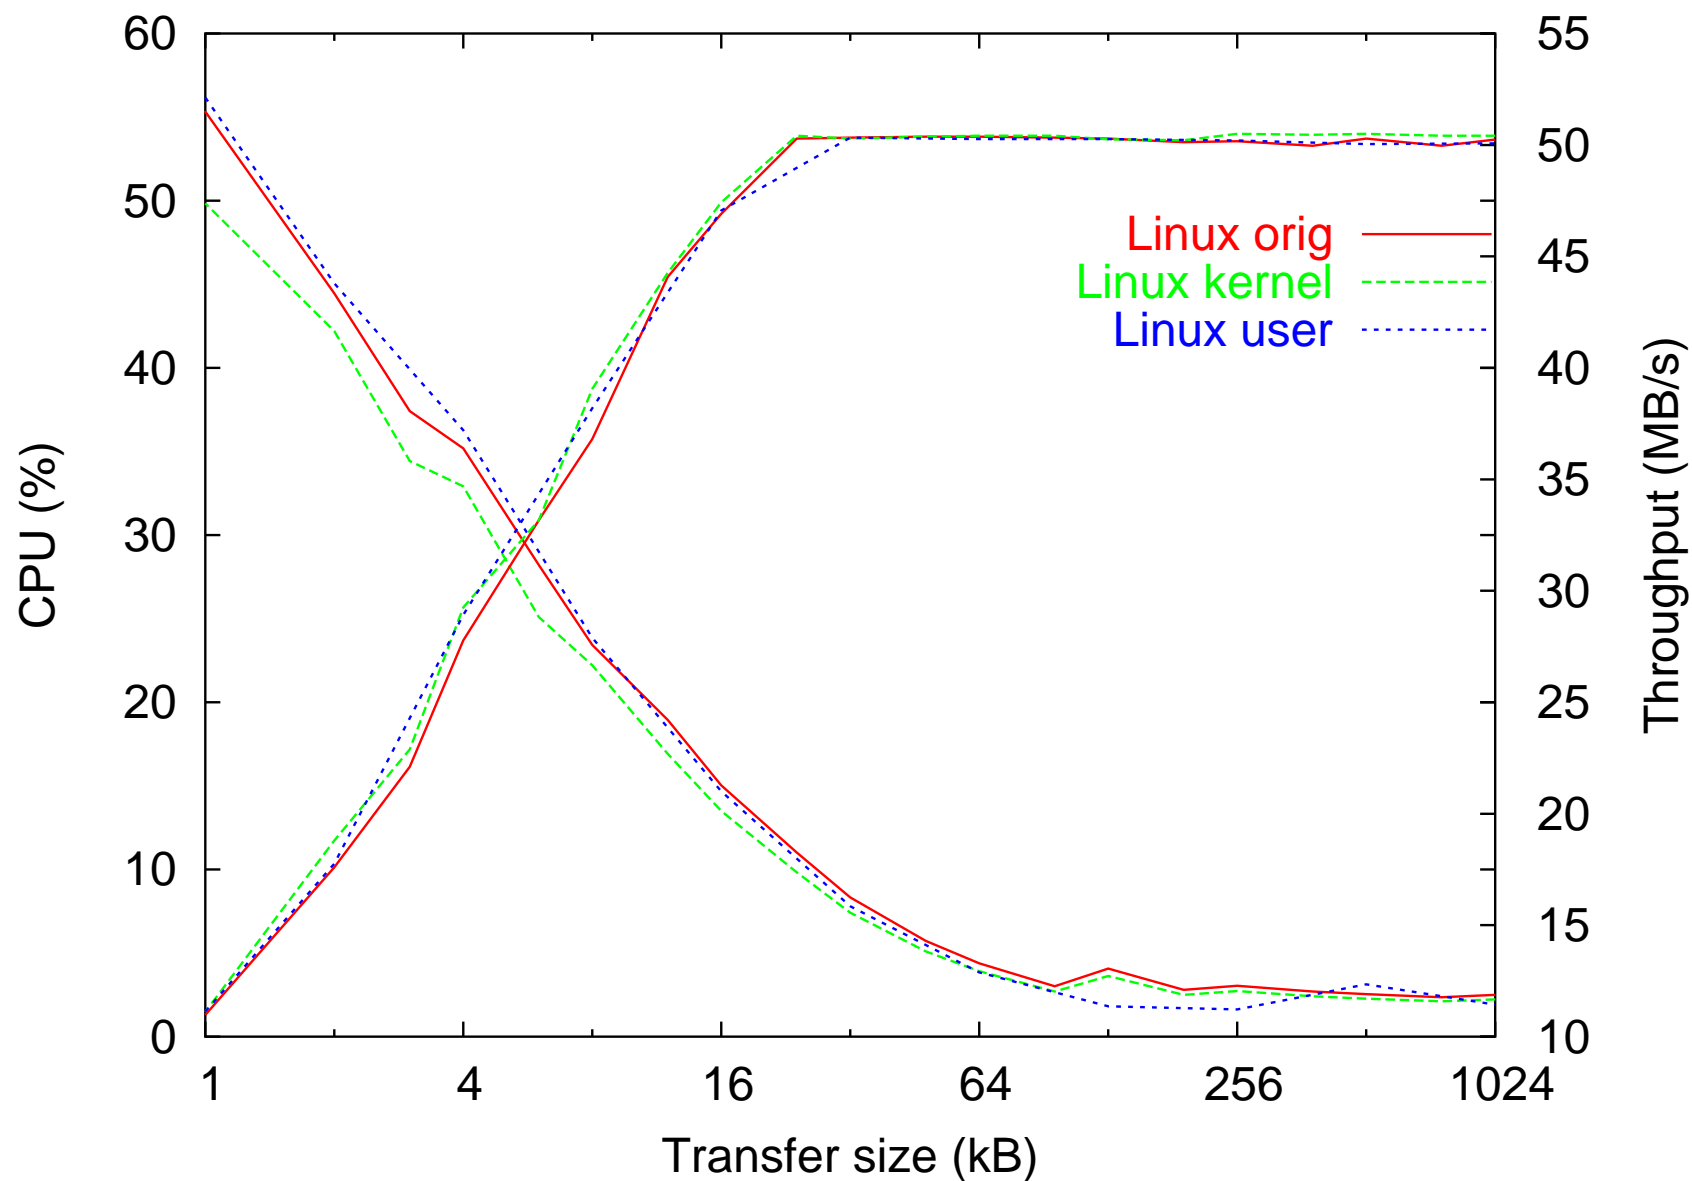
I/O Performance

- Micro-state accounting patches allow measurement
- IOMMU
- Faster system call support
- User-mode device drivers avoid system call overhead



User-mode drivers







Current Work

- Export VFS and Block interfaces to userland
- Optimize driver interface
- Cleanup and release... soon?



Network performance

- Unaligned accesses in drivers a big problem
- Beginning extensive benchmarking and profiling
- There's something funny in there... still working out what.