

NICTA

UNSW
THE UNIVERSITY OF NEW SOUTH WALES

COMP9242
Advanced Operating Systems

S2/2011 Week 7:
Multiprocessors – Part 2

Australian Government
Department of Broadband, Communications
and the Digital Economy
Australian Research Council

NICTA Funding and Supporting Members and Partners

UNSW

STANFORD

NSW

Victoria

UNIVERSITY OF TORONTO

COMP9242 S2/2011 W07 3

Multiprocessor OS

Memory

- Key design challenges:
 - Correctness of (shared) data structures
 - Scalability

COMP9242 S2/2011 W07 2

NICTA

UNSW

Scalability of Multiprocessor OS

Remember Amdahl's law

- Serialisation prevents scalability
- Whenever application not running on core, scalability reduced

Sources of Serialisation:

- Locking
 - Waiting for a lock → stalls self
 - Lock implementation:
 - Atomic operations lock bus → stalls everyone
 - Cache coherence traffic loads bus → slows down others
- Memory access
 - Relatively high latency to memory → stalls self
- Cache
 - Processor stalled while cache line is fetched or invalidated
 - Limited by latency of interconnect round-trips
 - Performance depends on data size (cache lines) and contention (number of cores)

COMP9242 S2/2011 W07 3

NICTA

UNSW

More Cache Issues

- False sharing
 - Unrelated data structs share the same cache line
 - Accessed from different processors
 - Cache coherence traffic and delay
- Cache line bouncing
 - Shared R/W on many processors
 - E.g: bouncing due to locks: each processor spinning on a lock brings it into its own cache
 - Cache coherence traffic and delay
- Cache misses
 - Potentially direct memory access
 - When does cache miss occur?
 - Application runs on new core
 - Cached memory has been evicted

COMP9242 S2/2011 W07 4

NICTA

UNSW

Optimisation for Scalability

- Reduce amount of code in critical sections
 - Increases concurrency
 - Fine grained locking
 - Lock data not code
 - Tradeoff: more concurrency but more locking (and locking causes serialisation)
 - Lock free data structures
- Reduce false sharing
 - Pad data structures to cache lines
- Reduce cache line bouncing
 - Reduce sharing
 - E.g: MCS locks use local data
- Reduce cache misses
 - Affinity scheduling: run process on the core where it last ran.
 - Avoid cache pollution

COMP9242 S2/2011 W07 5

NICTA

UNSW

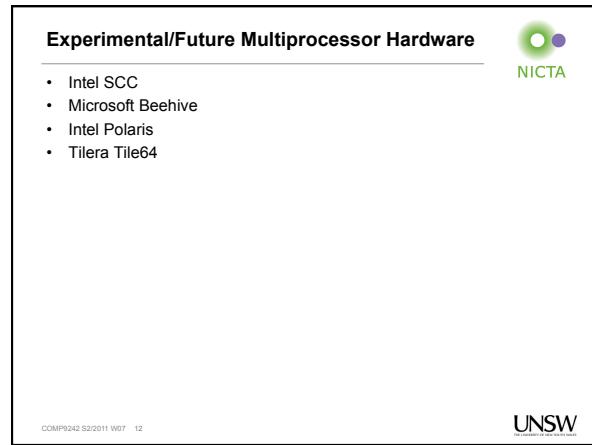
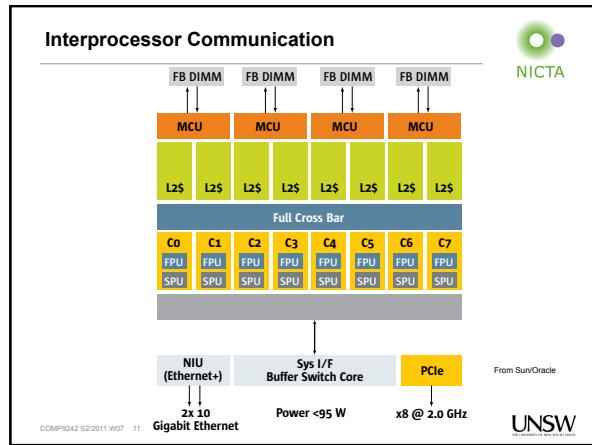
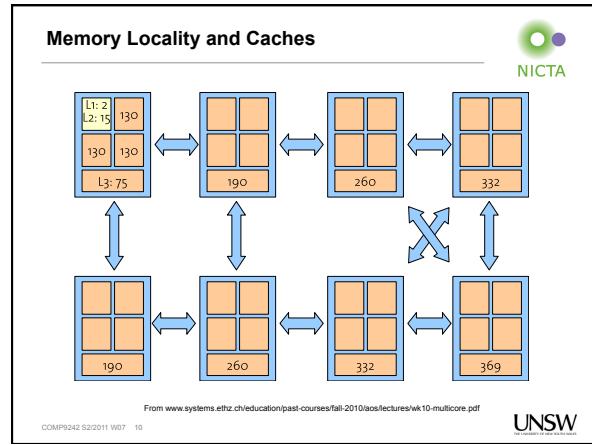
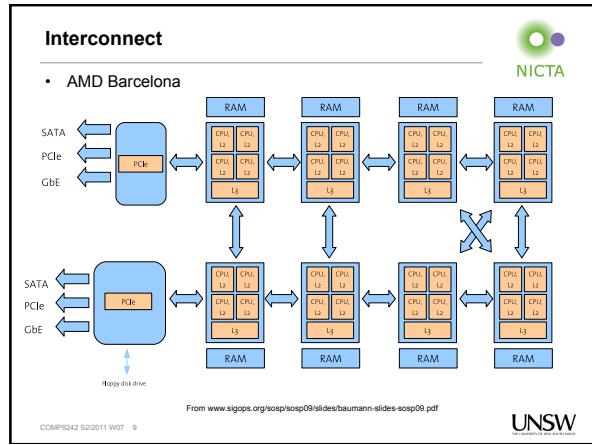
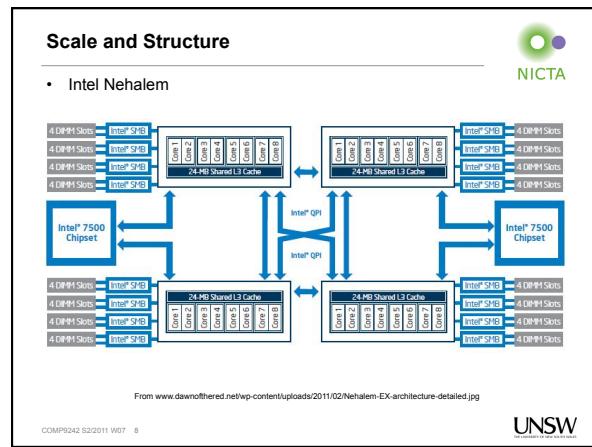
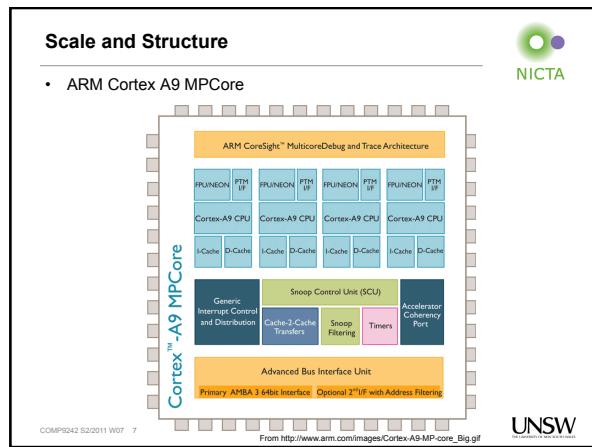
Contemporary Multiprocessor Hardware

- Intel Nehalem: Beckton, Westmere
- AMD Opteron: Barcelona, Magny Cours
- ARM Cortex A9, A15 MPCore
- Oracle (Sun) UltraSparc T1,T2,T3,T4 (Niagara)

COMP9242 S2/2011 W07 6

NICTA

UNSW



Scale and Structure

- Tilera Tile64, Intel Polaris**

From www.tilera.com/products/processors/TILE64

COMP9242 S2/2011 W07 13

UNSW

Cache and Memory

- Intel SCC**

From techresearch.intel.com/spew2/uploads/files/SCC_Platform_Overview.pdf

COMP9242 S2/2011 W07 14

UNSW

Interprocessor Communication

- Beehive**

From projects.csail.mit.edu/beehive/BeehiveV5.pdf

COMP9242 S2/2011 W07 15

UNSW

Summary

- Scalability
 - ~ 100+ cores
 - ~ Amdahl's law really kicks in
- NUMA
 - ~ Also variable latencies due to topology and cache coherence
- Cache coherence may not be possible
 - ~ Can't use it for locking
 - ~ Shared data structures require explicit work
- Computer is a distributed system
 - ~ Message passing
 - ~ Consistency and Synchronisation
 - ~ Fault tolerance
- Heterogeneity
 - ~ Heterogeneous cores, memory, etc.
 - ~ Properties of similar systems may vary wildly (e.g. interconnect topology and latencies between different AMD platforms)

OS Design for Modern (and future) Multiprocessors

- Avoid shared data
 - Performance issues arise less from lock contention than from data locality
- Explicit communication
 - Regain control over communication costs
 - Sometimes it's the only option
- Tradeoff: parallelism vs synchronisation
 - Synchronisation introduces serialisation
 - Make concurrent threads independent
- Allocate for locality
 - E.g. provide memory local to a core
- Schedule for locality
 - With cached data
 - With local memory
- Tradeoff: uniprocessor performance vs scalability

Design approaches

- Divide and conquer
 - Using virtualisation
 - Using exokernel
- Reduced sharing
 - By design
 - Brute force
- No sharing
 - Computer is a distributed system

Divide and Conquer

Disco Running commodity OSes on scalable multiprocessors [Bugnion et al., 1997] <http://www-flash.stanford.edu/Disco/>

- Context:
 - ca. 1995, large ccNUMA multiprocessors appearing
 - Scaling OSes requires extensive modifications
- Idea:
 - Implement a scalable VMM
 - Run multiple OS instances
- VMM has most of the features of a scalable OS:
 - NUMA aware allocator
 - Page replication, remapping, etc.
- VMM substantially simpler/cheaper to implement
- Modern incarnations of this
 - Virtual servers (Amazon, etc.)
 - Research (Cerebrus)

COMP9042 S2/2011 W07 19

UNSW

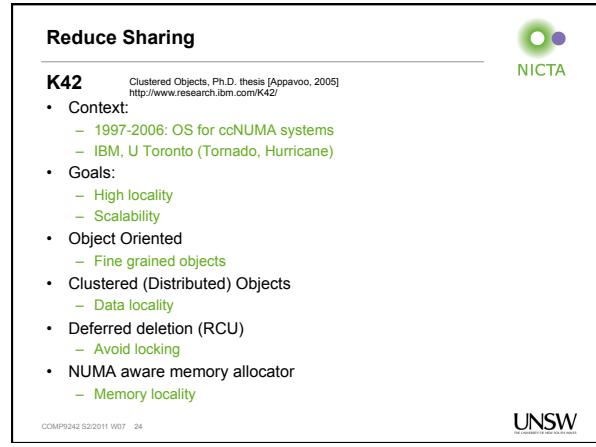
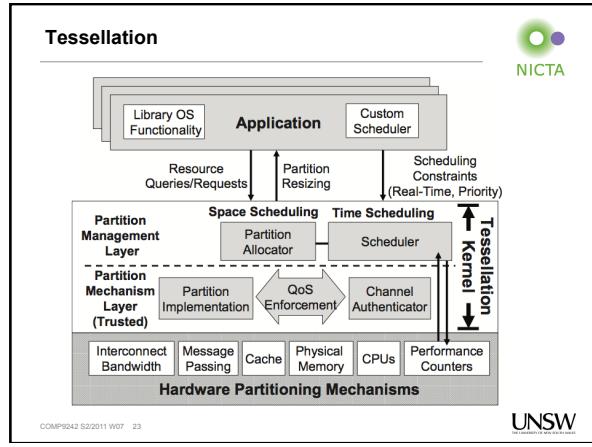
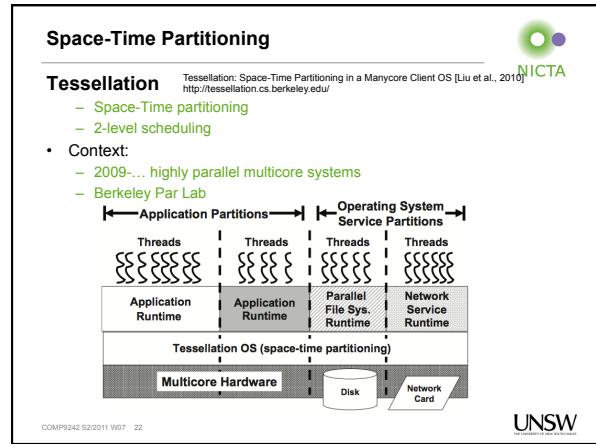
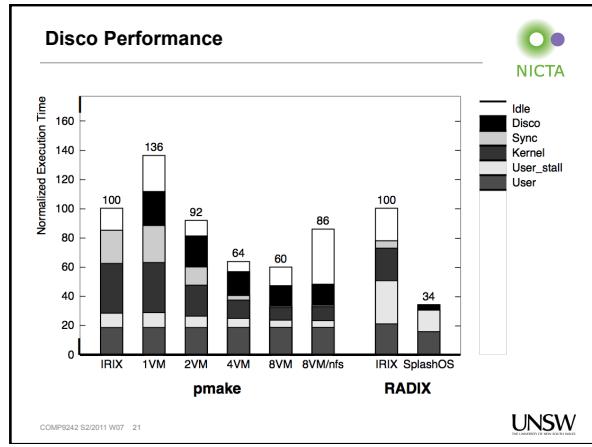
Disco Architecture

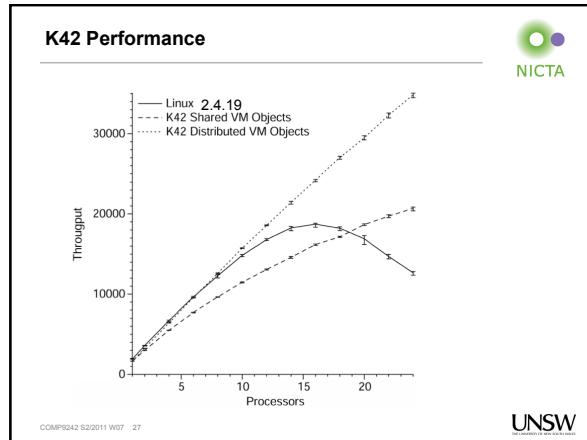
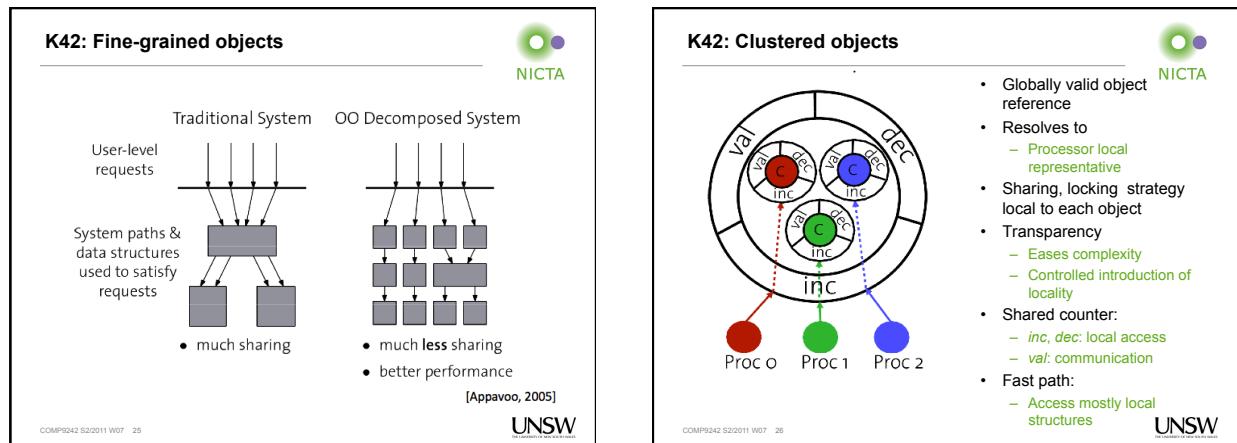
The diagram illustrates the Disco architecture. At the top, there is a row of boxes representing an SMP-OS layer, which includes Pmake, DB, NFS, and a Scientific App. Below this is a horizontal bar labeled "Disco". Underneath the Disco bar is a row of boxes labeled "PE" (Processor Element), representing the ccNUMA Multiprocessor. These PEs are interconnected via an "Interconnect". The entire system is labeled "ccNUMA Multiprocessor" at the bottom.

[Bugnion et al., 1997]

COMP9042 S2/2011 W07 20

UNSW

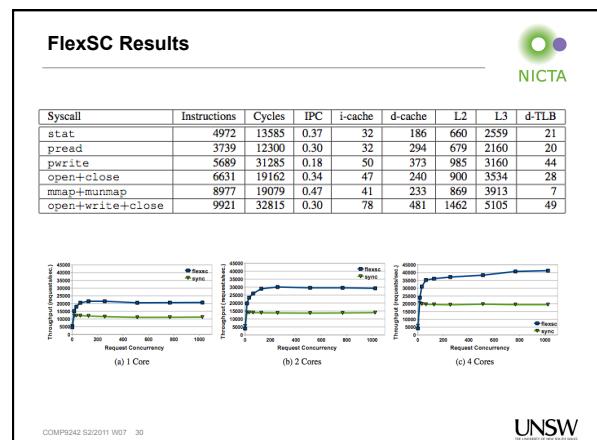
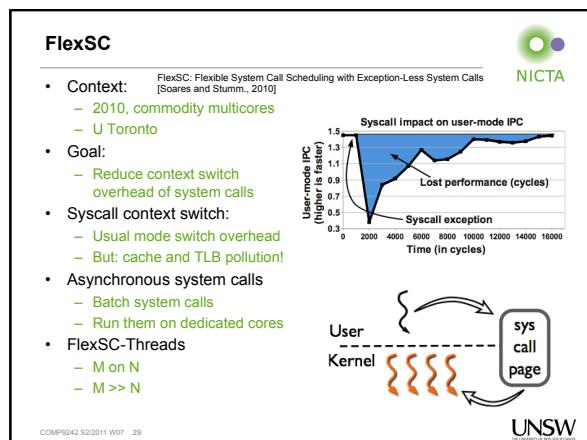




Corey

- Context
 - 2008, high-end multicore servers, MIT
- Goals:
 - Application control of OS sharing
- Address Ranges
 - Control private per core and shared address spaces
- Kernel Cores
 - Dedicate cores to run specific kernel functions
- Shares
 - Lookup tables for kernel objects allow control over which object identifiers are visible to other cores.
- Linux scalability (2010 – scale Linux to 48 cores)
 - sloppy counters, per-core data structs, fine-grained lock, lock free, cache lines : 3002 lines of code changed
 - no scalability reason to give up on traditional operating system organizations just yet.

An Analysis of Linux Scalability to Many Cores [Boyd-Wickizer et al., 2010]



No sharing

- Multikernel
 - Barreelfish
 - fos: factored operating system

The Multikernel: A new OS architecture for scalable multicore systems [Baumann et al., 2009]
http://www.barreelfish.org/

COMP9042 S2/2011 W07 31

Barreelfish

- Context:** The Multikernel: A new OS architecture for scalable multicore systems [Baumann et al., 2009] <http://www.barreelfish.org/>
 - 2007 large multicore machines appearing
 - 100s of cores on the horizon
 - NUMA (cc and non-cc)
 - ETH Zurich and Microsoft
- Goals:**
 - Scale to many cores
 - Support and manage heterogeneous hardware
- Approach:**
 - Structure OS as *distributed system*
- Design principles:**
 - Interprocessor communication is explicit
 - OS structure hardware neutral
 - State is replicated
- Microkernel**
 - Similar to sel4: capabilities

COMP9042 S2/2011 W07 32

Barreelfish

COMP9042 S2/2011 W07 33

Barreelfish: Replication

- Kernel + Monitor:
 - Only memory shared for message channels
- Monitor:
 - Collectively coordinate system-wide state
- System-wide state:
 - Memory allocation tables
 - Address space mappings
 - Capability lists
- What state is replicated in Barreelfish
 - Capability lists
- Consistency and Coordination
 - Retype: two-phase commit to globally execute operation in order
 - Page (re/un)mapping: one-phase commit to synchronise TLBs

COMP9042 S2/2011 W07 34

Barreelfish: Communication

- Different mechanisms:
 - Intra-core
 - Kernel endpoints
 - Inter-core
 - URPC
- URPC
 - Uses cache coherence + polling
 - Shared buffer
 - Sender writes a cache line
 - Receiver polls on cache line
 - (last word so no part message)
 - Polling?
 - Cache only changes when sender writes, so poll is cheap
 - Switch to block and IPI if wait is too long.

COMP9042 S2/2011 W07 35

Barreelfish: Results

- Message passing vs caching

Cores	SHM8	SHM4	SHM2	SHM1	MSG8	Server
2	2	1	1	1	1	1
4	4	2	2	2	2	2
6	6	3	3	3	3	3
8	8	4	4	4	4	4
10	10	5	5	5	5	5
12	11	6	6	6	6	6
14	11.5	7	7	7	7	7
16	11.5	7.5	7.5	7.5	7.5	7.5

COMP9042 S2/2011 W07 36

