

For a Microkernel, a BIG Lock is fine !

Sean Peters, Adrian Danis, Kevin Elphinstone, Gernot Heiser

NICTA and UNSW, Sydney, Australia



Australian Government





Approaches for Multicore Kernels





Multicore Kernel Trade-Offs



Really?	User threa d d d d d d d d d d d d d d d d d d d	User threa d Xernel t Core Core	User threa d d d d d d d d d d d d d d d d d d d
Property	Big Lock	Fine-grained Locking	Multikernel
Data structures	shared	shared	distributed
Scalability (poor	good	excellent
Concurrency in kernel	zero	high	zero
Kernel complexity	low	high	low
Resource management	centralised	centralised	distributed

©2015 Gernot Heiser, NICTA

A Microkernel is NOT an Operating System! **NICTA** All device drivers, OS services VM are usermode processes App Strong **Isolation** Linux File NW Memory **Device** Process App Mgmt System Stack Driver Mgmt IPC seL4 microkernel (= context-switching engine) Processor Controlled Communication

Microkernel vs Linux Execution







Cost of Locking: Round-Trip Intra-Core IPC



Cycles





- Amount of locked code is small anyway, 100–200 instructions
- Corresponds to fine- to medium-grained locks in Linux
- Cost of locks is within an OoM of kernel execution time
- Kernel times are short \Rightarrow contention is low



Cache Line Migration Latencies







- Amount of locked code is small anyway, 100–200 instructions
- Corresponds to medium-grained locks in Linux
- Cost of locks is within an OoM of kernel execution time
- Kernel times are short \Rightarrow contention is low

Assertion 2: Don't share mirokernel data without shared cache

• Migrating only a few cache lines takes longer than rest of syscall

seL4 Multicore Design: Clustered Multikernel





©2015 Gernot Heiser, NICTA



- Amount of locked code is small anyway, 100–200 instructions
- Corresponds to medium-grained locks in Linux
- Cost of locks is within an OoM of kernel execution time
- Kernel times are short \Rightarrow contention is low

Assertion 2: Don't share mirokernel data without shared cache

• Migrating only a few cache lines takes longer than rest of syscall

Assertion 3: Big lock will perform for closely-coupled cores

- Shared caches presently have moderate core counts
- Big lock in a *well-designed* microkernel will scale there









Realistic Case: Syscall-Intensive Macrobenchmark

NICTA









- Amount of locked code is small anyway, 100–200 instructions
- Corresponds to medium-grained locks in Linux
- Cost of locks is within an OoM of kernel execution time
- Kernel times are short \Rightarrow contention is low

Assertion 2: Don't share mirokernel data without shared cache

• Migrating only a few cache lines takes longer than rest of syscall

Assertion 3: Big lock will perform for closely-couple cores

- Shared caches presently have moderate core counts
- Big lock in a *well-designed* microkernel will scale there
- Seems to work fine for 8 cores, should still work for 16

Summary



- 1. Big lock <u>fine-grained enough</u> for a <u>well-designed</u> microkernel
 - Short system calls \Rightarrow low contention
 - Big lock performs as well as fine-grained for closely-coupled cores
 - Complexity of fine-grained locking is not worth it!
- 2. Should not even *try* to share a kernel without shared caches

http://seL4.systems

gernot@nicta.com.au

http://microkerneldude.wordpress.com

@GernotHeiser