

Concurrency

Advanced Operating Systems (M)
Tutorial 7

Tutorial Outline

- Review of lectures
- Key learning outcomes
- Discussion

Review of Lectures

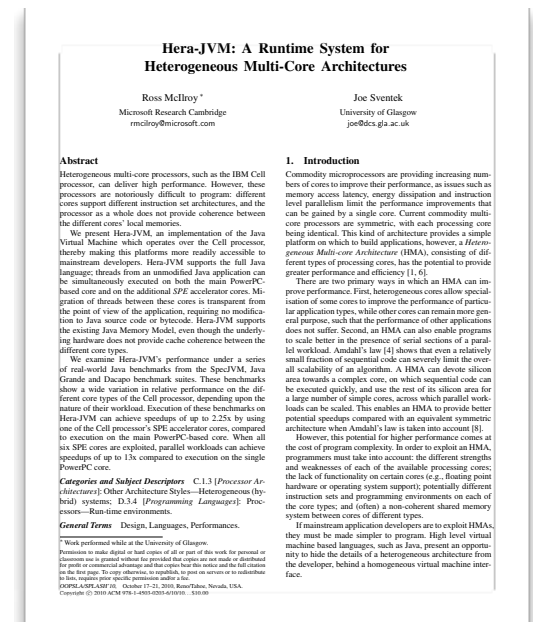
- Hardware trends towards multicore
 - Moore's law, power consumption, cores, and interconnects
 - NUMA optimisations: locality-aware memory allocation; replication of kernel data structures; locality-aware scheduling
 - Difficulty of maintaining cache coherence; benefits of message passing
- Multi-kernel model
 - Distributed system model; explicit communication; replicated state
 - Example: Barrelfish
- Heterogenous instruction set systems
 - Programming models: slave core, multi-kernel model, VM with transparent offload, hybrid VM with device-specific features
 - Examples: Helios, OpenCL, Hera-JVM, Accelerator

Key Learning Outcomes

- Understanding of the hardware trends towards multicore systems, heterogeneity
- Benefits and limitations of the multi-kernel model for structuring operating systems
- Understanding the trade-off between the different approaches to handling heterogeneity
- Familiarity with concepts underlying the example systems: Helios, OpenCL, Hera-JVM, Accelerator

Discussion: Hera-JVM

- Compilation to the SPE core – relatively standard
- Software caching of heap objects
 - DMA access to main memory only, used to cache entire objects
 - Java memory model maintained through cache invalidation (§5.3) – expensive
 - Software-managed cache
- Software caching of methods
 - DMA used to copy methods to SPE when they're invoked
 - Each object may be JIT compiled twice, once for PPE, once for SPE
 - Problems with returning from a method, if caller has been evicted (§5.2)
- How to trigger migration to SPE core?
 - Explicit core annotations on a method
 - Annotating methods with the sort of operations performed, with the runtime deciding which core is appropriate – transparent but difficult
- Performance
- Programming model – is this a good approach?



Discussion: Accelerator

- GPGPU offload using data parallel arrays in C#
- Architecture of GPGPU device
- Programming model
 - Data parallel arrays
 - Lazy evaluation, JIT compilation and execution when result assigned to a standard array
- Comparison with OpenCL
 - Programming model is simpler but more restrictive
 - Is it sufficient?
- Trade-offs compared to Hera-JVM?

