

# Wrap-up

## Advanced Operating Systems Lecture 19

# Lecture Outline

- Review of Rationale, Aims and Objectives, and the Intended Learning Outcomes
- Review of Lectures
- Summary and Discussion
- Assessment and Examination

# Rationale

- Radical changes to computing landscape;
  - Desktop PC becoming irrelevant
  - Heterogeneous, multicore, mobile, and real-time systems – smart phones, tablets – now ubiquitous
- Not reflected by corresponding change in operating system design practice
- This course will...
  - review research on systems programming techniques and operating systems design;
  - discuss the limitations of deployed systems; and
  - show how the operating system infrastructure might evolve to address the challenges of supporting modern computing systems.

# Aims and Objectives

- To explore programming language and operating system facilities essential to implement real-time, reactive, and embedded systems
- To discuss limitations of widely-used operating systems, introduce new design approaches to address challenges of security, robustness, and concurrency
- To give an understanding of practical engineering issues in real-time and concurrent systems; and suggest appropriate implementation techniques

# Intended Learning Outcomes (1)

- At the end of this course, you should be able to:
  - Clearly differentiate the different issues that arise in designing real-time systems;
  - Analyse a variety of real-time scheduling techniques, and prove the correctness of the resulting schedule; implement basic scheduling algorithms;
  - Understand how to apply real-time scheduling theory to the design and implementation of a real-world system using the POSIX real-time extensions, and be able to demonstrate how to manage resource access in such a system;
  - Describe how embedded systems are constructed, and discuss the limitations and advantages of C as a systems programming language, understand how managed code and advanced type systems might be used in the design and implementation of future operating systems;
  - ...

# Intended Learning Outcomes (2)

...

- Discuss the advantages and disadvantages of integrating garbage collection with the operating system/runtime, understand the operation of popular garbage collection algorithms, and alternative techniques for memory management, and know when it might be appropriate to apply such techniques and managed runtimes to real-time systems and/or operating systems;
- Understand the impact of heterogeneous multicore systems on operating systems, compare and evaluate different programming models for concurrent systems, their implementation, and their impact on operating systems;
- Construct and/or analyse simple programming to demonstrate understanding of novel techniques for memory management and/or concurrent programming, to understand the trade-offs and implementation decisions.

# Course Outline

- Key topics in systems programming and operating systems:
  - Systems programming
  - Hardware trends
  - Memory management
  - Implications of concurrency
  - High performance networking
  - Real-time systems
  - Virtualisation
- Focus is on ideas and advanced concepts
  - Research topics and new approaches

# Systems Programming

- What do we mean by systems programming?
  - Systems programs interact with hardware
  - Systems programs have memory and data layout constraints
  - Systems programs strongly driven by bulk I/O performance
  - Systems programs maintain long-lived, concurrently accessed, state
  - This is writing operating system kernels, low-level services, embedded systems, etc.
- Why is systems programming challenging?
- How are systems programs implemented? How should they be implemented?



# Hardware Trends

- How is hardware evolving? How does this affect systems programming?
  - Implications of Moore's law and Denard scaling
  - Implications of concurrency
  - Implications of solid-state storage
  - Implications of high performance networking

# Memory Management

- Implications of NUMA and the caching hierarchy
- Memory management
  - Layout of a processes address space
  - Manual memory management
  - Region based memory management; ownership and borrowing; Rust
  - Garbage collection algorithms

# Implications of Concurrency

- Memory models, threads, and locks
- Why threads and shared state concurrency are not sufficient
- Alternative concurrency models
  - Transactions
  - Message passing: Akka, Erlang, etc.

# High Performance Networking

- Interactions between high-performance networking and Moore's law
  - Why high-performance networking is becoming more challenging
- APIs and programming models
  - Message passing
  - netmap and StackMap
  - User-space protocol stacks

# Real-time Systems

- Definition of a real-time system
  - Scheduling theory
  - Need for formalisation
- Scheduling periodic systems
  - Rate monotonic algorithm
  - Earliest deadline first algorithm
- Scheduling aperiodic and sporadic tasks
  - Sporadic servers

# Virtualisation

- What is virtualisation?
  - Hypervisors
  - Cloud computing
- Full system virtualisation
  - Xen
- Containers
  - FreeBSD jails, Docker, etc.
  - Unikernels

# Summary and Discussion

- Computing systems are changing – hardware and usage changes
- Operating systems need to change too:
  - Mobile, Real-time, Embedded, Concurrent, Reliable, Secure
  - Is Unix the right basis for the operating system of the future? Unclear –many interesting new systems, approaches, languages, runtimes
  - Systems programming should move away from C
    - Tremendous potential in languages such as Rust for low-level systems and software – security and robustness benefits due to memory and type safety
    - Higher-level approaches for networked services/applications: message passing, “your server as a function” – significant productivity and safety benefits
- Significant ongoing changes – make sure you look beyond traditional approaches

# Assessment

- Coursework (20%)
  - Exercises that have been completed
- Examination (80%)
  - Two hours duration; sample and past papers are available on Moodle or the course website
  - All material in the lectures, tutorials, and cited papers is examinable
  - Aim is to test your understanding of the material, not to test your memory of all the details; explain why – don't just recite what
  - The contents of the highlighted “Further Reading” research papers is **examinable** – but exam questions covering that material will focus on concepts, rather than on details