Graduate Operating Systems

Spring 2023

1

Today's Paper

- **[3]** Dawson R. Engler, M. Frans Kaashoek, and James O'Toole Jr., "Exokernel: An Operating System Architecture for Application-Level Resource Management", Proc. of the 15th Symposium on Operating Systems Principles, December 1996.

2

Traditional Operating Systems

- Traditionally, operating systems were built as a monolithic kernel:

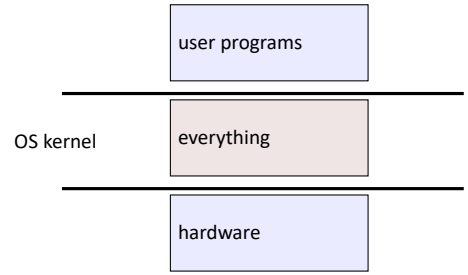

3

Traditional Operating Systems

- Traditional operating systems use a "one-size-fits-all" interface and implementation of OS abstractions

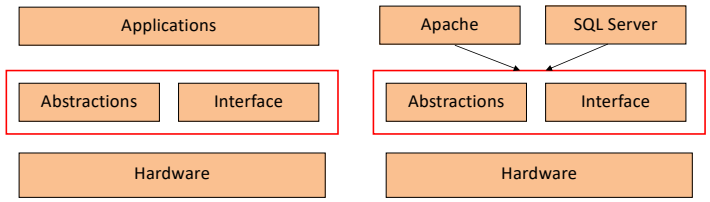

4

Problems with Traditional OS

- Performance
  - Denies applications the advantages of **domain-specific optimizations**
- Flexibility
  - Restricts the flexibility of application builders
  - Concept: **"with more information exposed, resources can be utilized 'better'"**!
- Functionality
  - Discourages changes to the implementations of existing abstractions

5

Microkernels

- Popular in the late 80s, early 90s
  - recent resurgence of popularity for small devices
- Goal:
  - minimize what goes in kernel
  - organize rest of OS as user-level processes
- This results in:
  - better reliability (isolation between components)
  - ease of extension and customization
  - poor performance (user/kernel boundary crossings)
- First microkernel system was Hydra (CMU, 1970)
  - follow-ons: Mach (CMU), Chorus (French UNIX-like OS), early Windows NT (Microsoft), OS X (Apple)

6

Example: Exokernel

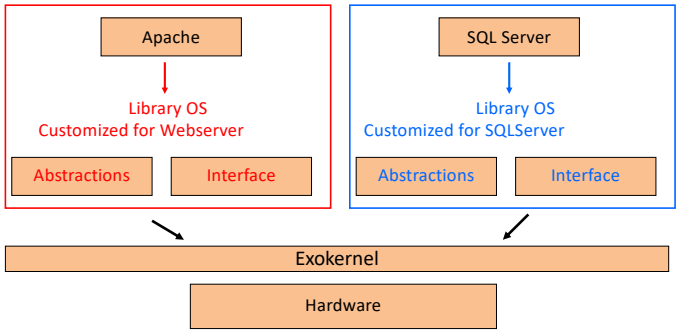

7

Solution: Exokernel

- **Separate protection from management!**
  - Allows user level to manage resources
    - Application libraries implement OS abstractions
  - Exokernel exports (not emulate) resources
    - Low level interface
    - Protects, does not manage
    - Exposes hardware
- End-to-end argument; "applications know better"

8

Exokernel + Library OS

- **Exokernel**'s resource management:
  - Allocate, revoke, share, track ownership
- **Library OS**:
  - Uses low-level exokernel interface, provides higher-level abstractions; provides special purpose implementations

An application can choose the library which best suits its needs, or even build its own.

9

Exokernel

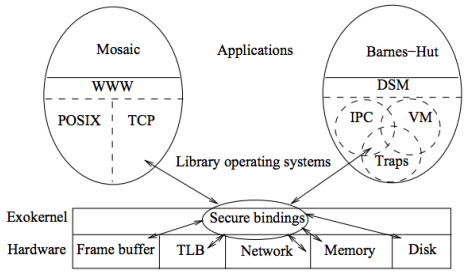

10

Exokernel

- Hypotheses:
  - Exokernels can be very **efficient**
  - Low-level, secure multiplexing of hardware resources can be implemented **efficiently**
  - Traditional operating system abstractions can be implemented **efficiently** at application level
  - Applications can create **special-purpose implementations** of these abstractions

11

Library Operating Systems

- Simpler
- Specialized
- Multiple can exist
- Few kernel crossings

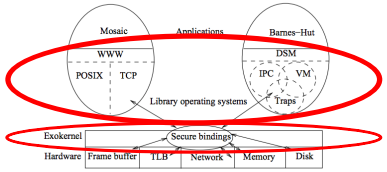

12

Design Challenge

- How can an Exokernel allow libOSes to freely manage physical resources while protecting them from each other?
  - Track ownership of resources
    - Secure bindings – libOS can securely bind to machine resources
  - Guard all resource usage
    - Invisible/visible resource revocation
  - Revoke access to resources
    - Abort protocol

13

Design Principles

- Securely expose hardware
  - Kernel should provide secure low-level primitives that allow all hardware resources to be accessed as directly as possible.
- Expose allocation
  - Allow to request specific physical resources
- Expose names
  - Export physical names.
  - Remove a level of indirection: Translation
- Expose revocation
  - Utilize a visible resource revocation protocol

14

Secure Bindings

- Exokernel allows LibOSes to bind resources using secure bindings
- Decouples authorization from the actual use of a resource
- Multiplex resources securely
- Performs authorization only at bind time
  - Allows the kernel to protect resources without having to understand them

15

Some Terminology

- Packet filters
- TLB
- Physical memory: capabilities for page
- Downloadable code (ASH)
- RPC
- DMA

16

Kernel Comparisons

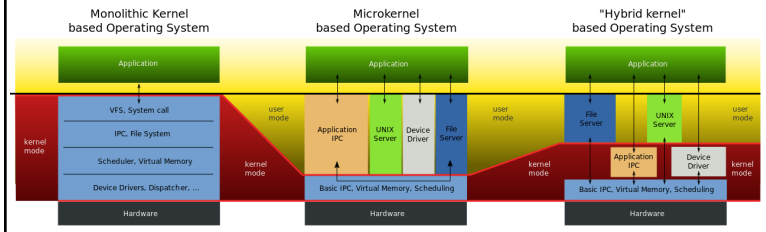

17

Microkernels

- A good idea in the 1970s and 80s
- Not up to demands of modern processors
  - Virtual memory
  - Heavy caching
- Not up to demand of modern operating systems
- "Resurrection":
  - Compare to concepts of virtual machines
  - Mobile and wearable devices:
    - Fixed or limited functionality
    - No general purpose files
    - No dynamic virtual memory
    - Simple context switches
    - All code already in memory
    - Easy IPC

18