The Pebble Component-Based Operating System

Gabber, Small, Bruno, Brustoloni and Silberschatz

Bell Laboratories

4. August 2006

Goals

- flexibility
- safety
- performance

Philosophy and Architecture

- minimal privileged mode nucleus
- system services by replacable user-level components with minimal privileges
- reducing the cost of transferring a thread from one protection domain and another

Related Work

Title

- Mach, Windows NT - client-server OS
- L4 Microkernel - very similar
- Fluke - layered system structure
- Exokernel - examinate OS abstractions
- Plan 9 - similar isolation mechanism (resources are files)

Nucleus primitives

Abstractions

- protection domain = address space
- thread = execution context
- portal = communication endpoint

System services

- fork - create new thread
- fork_domain - create new domain
- portal_create
- invoke_portal

Privileged user-level servers

- portal manager - instantiates and manages portals
- name server - enforces also security policy
- memory manager - maps and revokes pages
- scheduler - implements user-level scheduling policy
- interrupt dispatcher - invokes interrupt handler
- device driver
- filesystem - backing store

Portals

Definition

- consist of IDL specification and portal code
- mapped into portal table
- portal manager generates portal code from IDL
- portal code transforms arguments, copies data, changes stack and maps pages

Portal Instantiation

- 1.step: register the portal with the portal manager
- 2.step: install the portal in the client's portal table
- client requests the portal with given name
- name server approves the access

Sandboxing

- transparent to application
- intercepting portal calls of controlled application
- claim: controlls all interaction of application
- controlling domain gets notification on portal creation

Finally

Implementation

- MIPS 1200 (software managed and tagged TLBs)
- only microbenchmarks
- comparison to OpenBSD

Questions

- handling of timeouts
- real-life scenario (legacy OS)
- size of portal table and portal code
- multiprocessor support