

Lecture Notes, 9/16/98

Machine-Independent Virtual Memory in Mach

I. Machine-Independent VM in Mach

Goals:

- o Machine-independent VM management
- o Flexible address space manipulation
- o Multiprocessor support

Functionality requirements:

- o Large, sparse address spaces.
- o Copy-on-write and read/write memory sharing between tasks.
- o Memory-mapped files.
- o User-provided backing store objects and pagers.

Details of Mach abstractions are mostly irrelevant for this paper: need a way for kernel and pagers to communicate.

VM programming interface: (see paper for details)

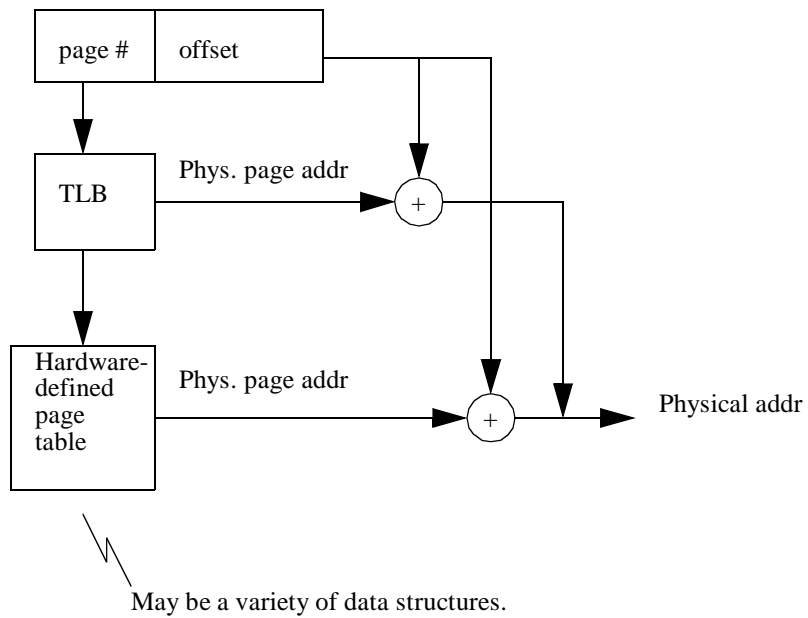
- o address space manipulation
- o memory protection
- o memory inheritance
- o misc.

Types of VM data structures:


- o resident page table: machine-independent info. about physical memory pages.
- o memory object: unit of backing storage managed by kernel or user-level pager task.
- o address map: describes mapping from a virtual address space to regions of memory objects.
- o pmap: machine-dependent memory mapping data structure (i.e. hardware-defined physical addr. map).

Lecture 6

Review of VM paging:




Resident page table: describes state of machine's physical memory




Lecture 6

Address maps: inheritance and protection of virtual memory regions


Virtual addr space 1




Virtual addr space 2



Address map 1




Address map 2




Copy-on-write: read-only case


Virtual addr space 1




Virtual addr space 2



Address map 1

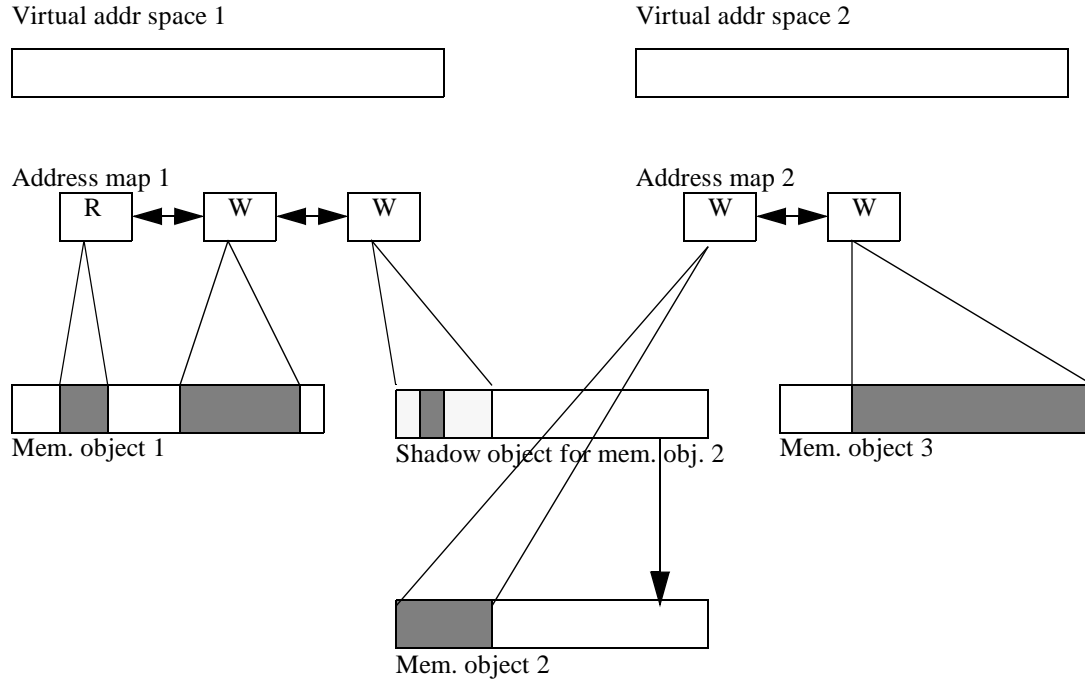


Address map 2

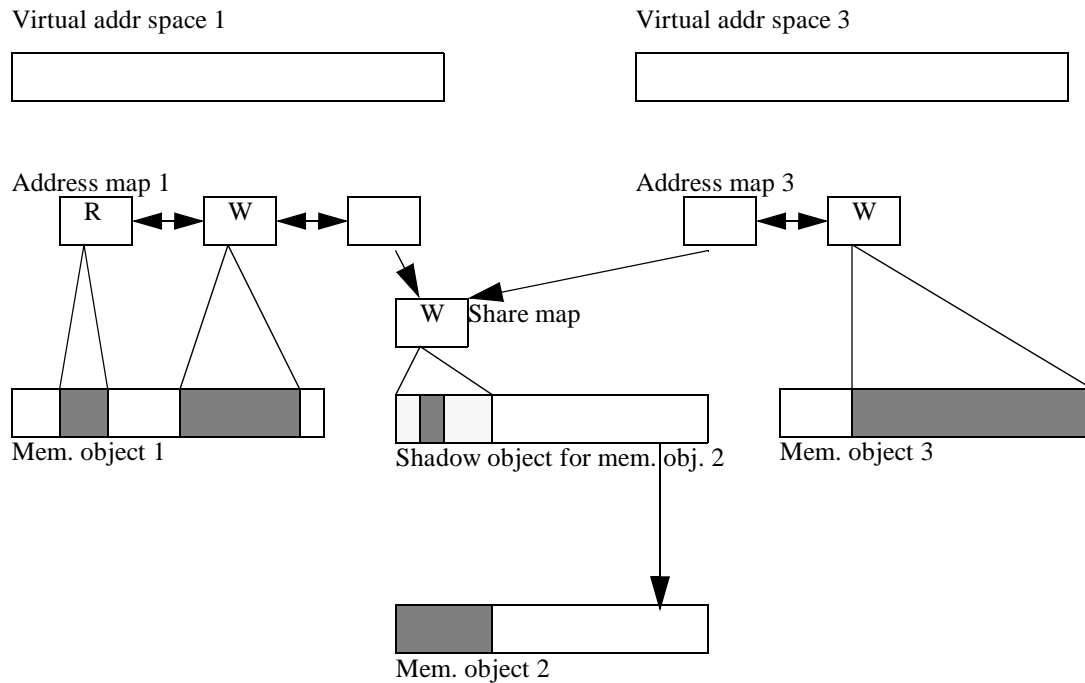


Lecture 6

Shadow objects: supporting copy-on-write with modifications



Share maps: supporting read/write sharing



Lecture 6

An example of remapping an address map entry: sending data copy-on-write to a shared VM region.

Pmap interface hides all machine-dependent aspects of VM (including hardware-defined page table facilities).

Performance:

- o Zero fill:
 - Win against uVax because of 4x page size.
 - Win against Sun & IBM because they simulate VAX VM.
- o Fork: Mach should compare itself to vfork, not fork! Worse than Unix vfork on uVax.
- o Parallel performance:
 - Demonstrates that Mach VM works in an MP environment.
 - Fault-copy case is faster because there is less cache thrashing during a fault.
- o Overall performance: paper doesn't say how much of the performance gain in Mach is due to the object cache.

MP issues:

- o TLB consistency and “shoot-down”.
- o No mention of how much trouble they had with making the kernel concurrent.

3 key features of paper:

- o First machine-independent, easily portable VM design.
- o Supports multi-address space sharing and user-definable pagers through address maps, share maps, memory objects, and shadow objects.
- o Machine-dependent part of VM described by a single abstraction: the pmap.

Some flaws:

- o Not much discussion of MP issues (none about restructuring Unix kernel to be concurrent).
- o In general, lots of things left unsaid or only alluded to.
- o Weak performance evaluation section.

A lesson: Easily portable, machine-independent designs are possible for low-level OS components.