

ETH zürich spci.inf.ethz.ch  
@spci\_eth

ADRIAN PERRIG & TORSTEN HOEFLER

## Networks and Operating Systems (252-0062-00)

### Chapter 5: Memory Management

<http://support.apple.com/kb/HT5642>

“Description: The iOS kernel has checks to validate that the user-mode pointer and length passed to the copyin and copyout functions would not result in a user-mode process being able to directly access kernel memory. The checks were not being used if the length was smaller than one page. This issue was addressed through additional validation of the arguments to copyin and copyout.”



ETH zürich spci.inf.ethz.ch  
@spci\_eth

## Oldskool: `signal()`

```
void (*signal(int sig, void (*handler)(int))) (int);
```

- **Unpacking this:**
  - A handler looks like  
`void my_handler(int);`
  - Signal takes two arguments...  
*An integer (the signal type, e.g., SIGPIPE)*  
*A pointer to a handler function*
  - ... and returns a pointer to a handler function  
*The previous handler,*
- **“Special” handler arguments:**
  - SIG\_IGN (ignore), SIG\_DFL (default), SIG\_ERR (error code)

2



ETH zürich   [spcl.inf.ethz.ch](http://spcl.inf.ethz.ch)  
[@spcl\\_eth](https://twitter.com/spcl_eth)

## Unix signal handlers

- Signal handler can be called at **any time!**
- Executes on the current user stack
  - If process is in kernel, may need to retry current system call
  - Can also be set to run on a different (alternate) stack

⇒ User process is in *undefined* state when signal delivered

3

ETH zürich   [spcl.inf.ethz.ch](http://spcl.inf.ethz.ch)  
[@spcl\\_eth](https://twitter.com/spcl_eth)

## Implications

- **There is very little you can safely do in a signal handler!**
  - Can't safely access program global or static variables
  - Some system calls are *re-entrant*, and can be called
  - Many C library calls cannot (including `_r` variants!)
  - Can sometimes execute a `longjmp` if you are careful
  - With `signal`, cannot safely change signal handlers...
- **What happens if another signal arrives?**

4

ETH zürich spcl.inf.ethz.ch  
@spcl\_eth

## Multiple signals

- If multiple signals of the *same* type are to be delivered, Unix will **discard all but one**.
- If signals of *different* types are to be delivered, Unix will deliver them **in any order**.
- **Serious concurrency problem:**  
How to make sense of this?

5

ETH zürich spcl.inf.ethz.ch  
@spcl\_eth

## A better signal () POSIX sigaction ()

```

#include <signal.h>

int sigaction(int signo,
              const struct sigaction *act,
              struct sigaction *oldact);

struct sigaction {
    void (*sa_handler) (int);
    sigset_t sa_mask;
    int sa_flags;
    void (*sa_sigaction) (int, siginfo_t *, void *);
};

```

New action for signal signo




Previous action is returned

Signal handler

Signals to be blocked in this handler (cf., fd\_set)

More sophisticated signal handler (depending on flags)




6

ETH zürich   [sycl.inf.ethz.ch](http://sycl.inf.ethz.ch)  
 @sycl\_eth

## Signals as upcalls

- **Particularly specialized (and complex) form of *Upcall***
  - Kernel RPC to user process
- **Other OSES use upcalls much more heavily**
  - Including Barrelfish
  - “Scheduler Activations”: dispatch every process using an upcall instead of return
- **Very important structuring concept for systems!**

7

ETH zürich   [sycl.inf.ethz.ch](http://sycl.inf.ethz.ch)  
 @sycl\_eth

## Our Small Quiz

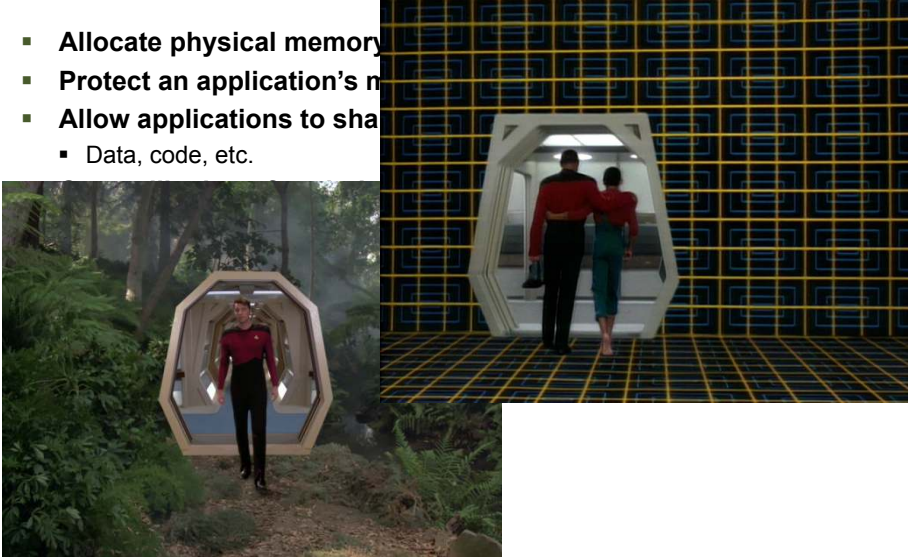
- **True or false (raise hand)**
  - Mutual exclusion on a multicore can be achieved by disabling interrupts
  - Test and set can be used to achieve mutual exclusion
  - Test and set is more powerful than compare and swap
  - The CPU retries load-linked/store conditional instructions after a conflict
  - The best spinning time is 2x the context switch time
  - Priority inheritance can prevent priority inversion
  - The receiver never blocks in asynchronous IPC
  - The sender blocks in synchronous IPC if the receiver is not ready
  - A pipe file descriptor can be sent to a different process
  - Pipes do not guarantee ordering
  - Named pipes in Unix behave like files
  - A process can catch all signals with handlers
  - Signals always trigger actions at the signaled process
  - One can implement a user-level tasking library using signals
  - Signals of the same type are buffered in the kernel

8

ETH zürich spci.inf.ethz.ch  
@spci\_eth

## Goals of Memory Management

- Allocate physical memory
- Protect an application's memory
- Allow applications to share memory
  - Data, code, etc.





9

ETH zürich spci.inf.ethz.ch  
@spci\_eth

## In CASP last semester we saw:

- Assorted uses for virtual memory
- **x86 paging**
  - Page table format
  - Translation process
  - Translation lookaside buffers (TLBs)
  - Interaction with caches
- **Performance implications**
  - For application code, e.g., matrix multiply



10

ETH zürich   [spci.inf.ethz.ch](http://spci.inf.ethz.ch)  
[@spci\\_eth](https://twitter.com/spci_eth)

## What's new this semester?

- **Wider range of memory management hardware**
  - Base/limit, segmentation
  - Inverted page tables, etc.
- **How the OS uses the hardware**
  - Demand paging and swapping
  - Page replacement algorithms
  - Frame allocation policies



11

ETH zürich   [spci.inf.ethz.ch](http://spci.inf.ethz.ch)  
[@spci\\_eth](https://twitter.com/spci_eth)

## Terminology

- **Physical** address: address as seen by the memory unit
- **Virtual** or **Logical** address: address issued by the processor
  - Loads
  - Stores
  - Instruction fetches
  - Possible others (e.g., TLB fills)...

12



ETH zürich   [sycl.inf.ethz.ch](http://sycl.inf.ethz.ch)  
[@sycl\\_eth](https://twitter.com/sycl_eth)

## Memory management

1. Allocating physical addresses to applications
2. Managing the name translation of virtual addresses to physical addresses
3. Performing access control on memory access

- Functions 2 & 3 usually involve the hardware Memory Management Unit (MMU)

13

ETH zürich   [sycl.inf.ethz.ch](http://sycl.inf.ethz.ch)  
[@sycl\\_eth](https://twitter.com/sycl_eth)

## Simple scheme: partitioned memory

14

ETH zürich spci.inf.ethz.ch  
@spci\_eth

## Base and Limit Registers

- A pair of **base** and **limit** registers define the logical address space

The diagram illustrates memory layout with the following address ranges and components:

- 0x0000000: Operating System
- 0x1000000: Process
- 0x5600ba0: Process (Base register points here)
- 0x8ff0010: Process (Limit register points here)
- 0xB000000: (Empty space)
- 0xffffffff: (End of memory)

Labels on the right indicate the **base** register at 0x5600ba0 and the **limit** register at 0x39ef470.

15

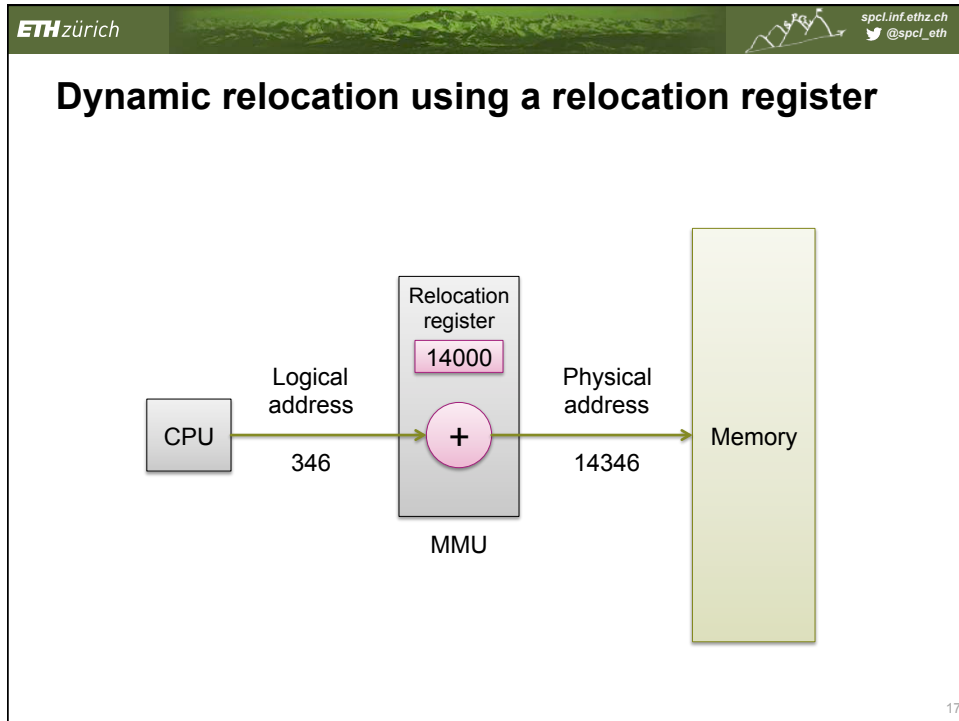
ETH zürich spci.inf.ethz.ch  
@spci\_eth

## Issue: address binding

- Base address isn't known until load time
- Options:
  - Compiled code must be *completely position-independent*, or
  - Relocation Register** maps compiled addresses dynamically to physical addresses

16



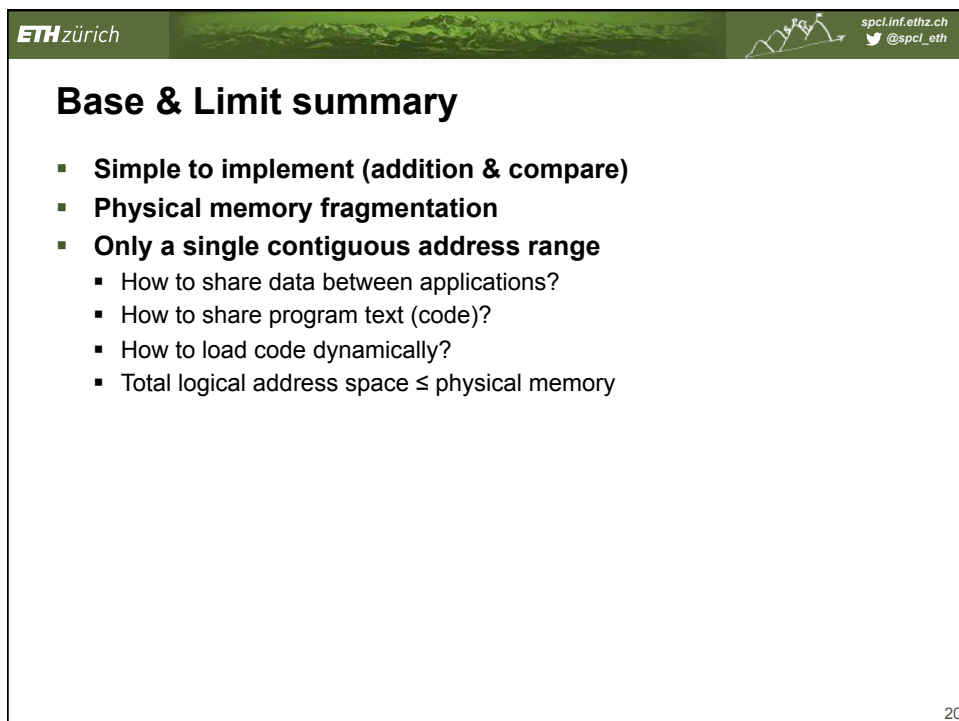
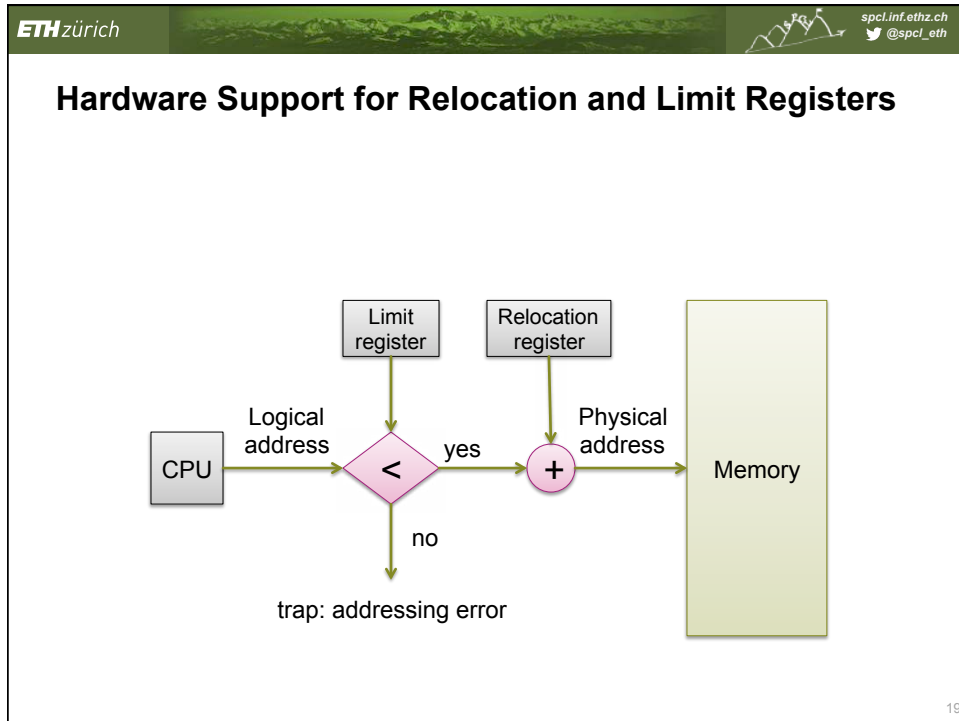




ETH zürich spci.inf.ethz.ch  
@spci\_eth

## Contiguous Allocation

- **Main memory usually into two partitions:**
  - Resident OS, usually in low memory with interrupt vector
  - User processes in high memory
- **Relocation registers protect user processes from**
  1. each other
  2. changing operating-system code and data
- **Registers:**
  - **Base register** contains value of smallest physical address
  - **Limit register** contains range of logical addresses  
*each logical address must be less than the limit register*
  - MMU maps logical address dynamically



18



ETH zürich   spci.inf.ethz.ch  
@spci\_eth

## Segmentation

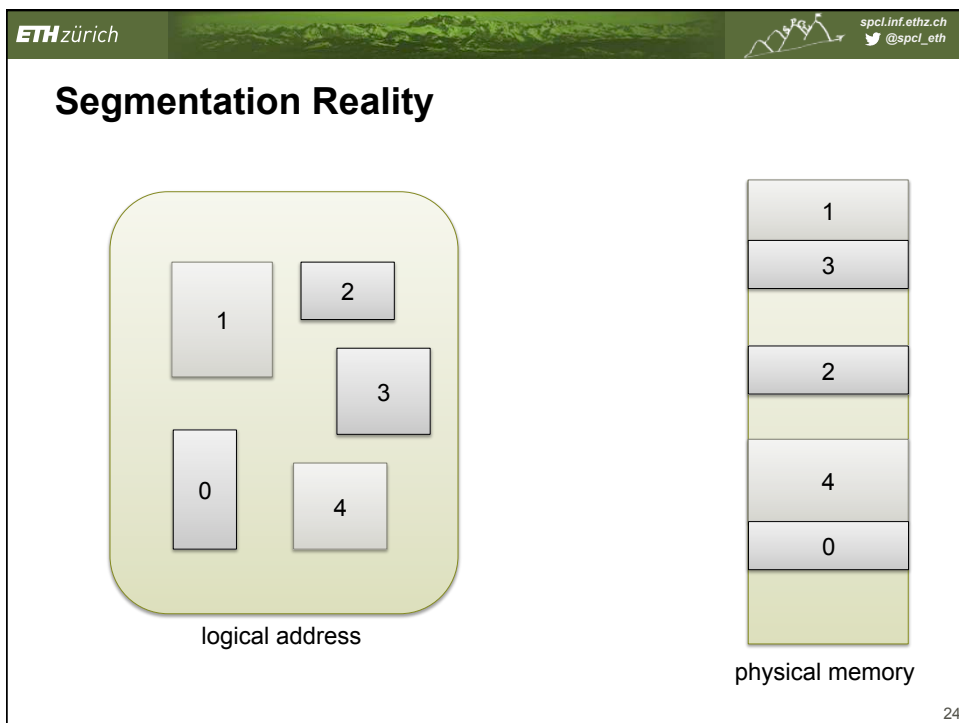
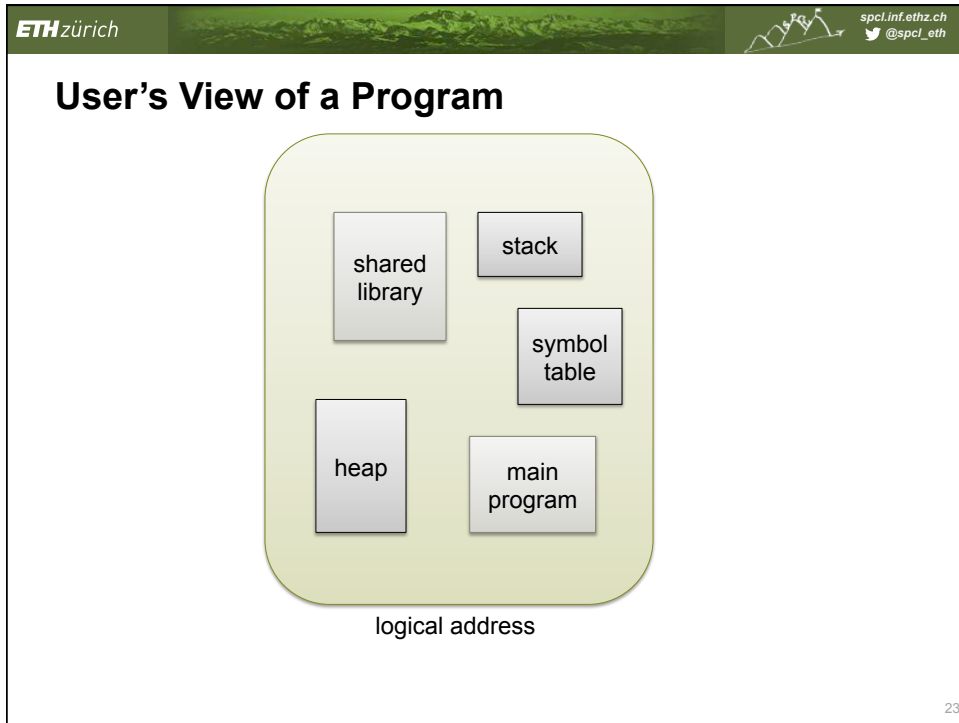
21

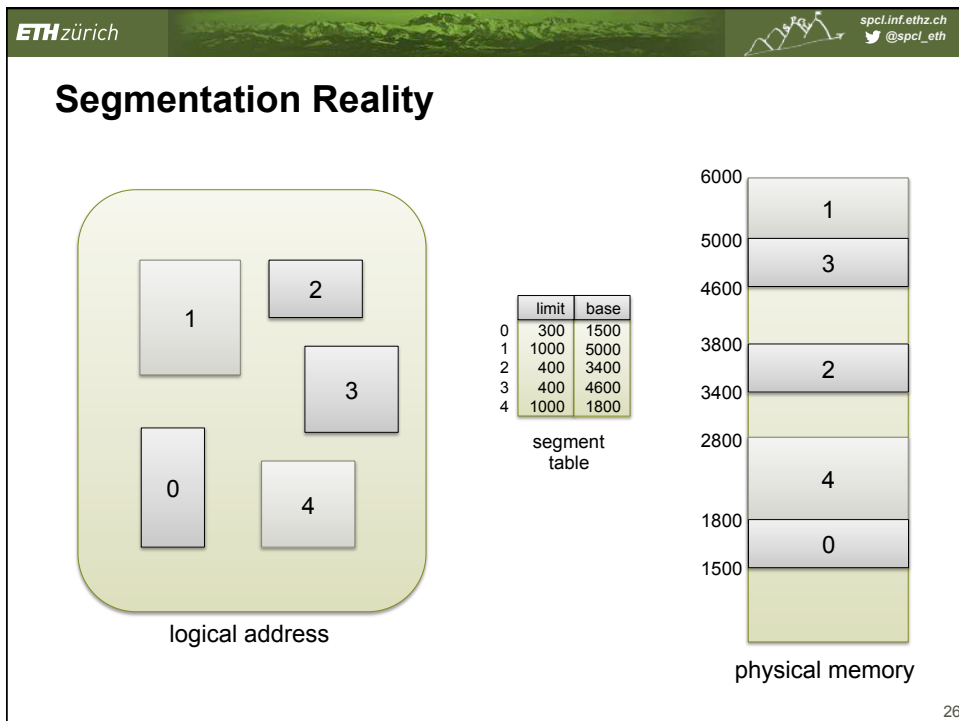
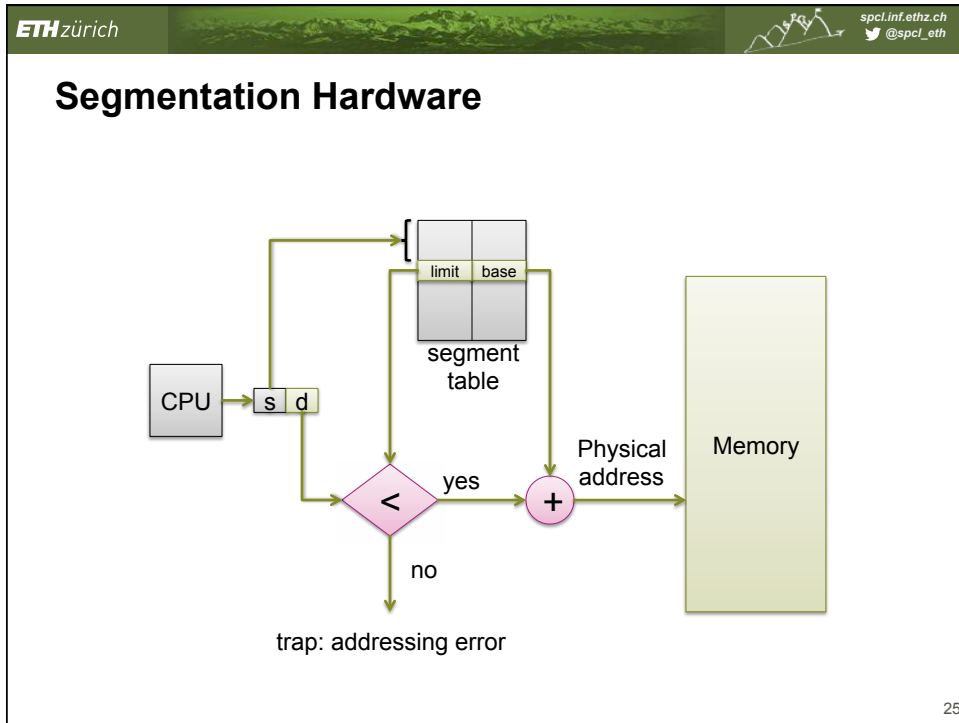
ETH zürich   spci.inf.ethz.ch  
@spci\_eth




## Segmentation

- **Generalize base + limit:**
  - Physical memory divided into *segments*
  - Logical address = (segment id, offset)
- **Segment identifier supplied by:**
  - Explicit instruction reference
  - Explicit processor segment register
  - Implicit instruction or process state

22








ETH zürich   [spci.inf.ethz.ch](http://spci.inf.ethz.ch)  
 @spci\_eth

## Segmentation Architecture

- **Segment table** – each entry has:
  - **base** – starting physical address of segment
  - **limit** – length of the segment
- **Segment-table base register (STBR)**
  - Current segment table location in memory
- **Segment-table length register (STLR)**
  - Current size of segment table

**segment number  $s$  is legal if  $s < \text{STLR}$**


27

ETH zürich   [spci.inf.ethz.ch](http://spci.inf.ethz.ch)  
 @spci\_eth

## Segmentation Summary


- **Fast context switch**
  - Simply reload STBR/STLR
- **Fast translation**
  - 2 loads, 2 compares
  - Segment table can be cached
- **Segments can easily be shared**
  - Segments can appear in multiple segment tables
- **Physical layout must still be contiguous**
  - (External) fragmentation still a problem

28

ETH zürich  [spcl.inf.ethz.ch](http://spcl.inf.ethz.ch)  
[@spcl\\_eth](https://twitter.com/spcl_eth)

## Paging

29

ETH zürich  [spcl.inf.ethz.ch](http://spcl.inf.ethz.ch)  
[@spcl\\_eth](https://twitter.com/spcl_eth)

## Paging

- **Solves contiguous physical memory problem**
  - Process can always fit if there is available free memory
- **Divide physical memory into *frames***
  - Size is power of two, e.g., 4096 bytes
- **Divide logical memory into *pages* of the same size**
- **For a program of  $n$  pages in size:**
  - Find and allocate  $n$  frames
  - Load program
  - Set up ***page table*** to translate logical pages to physical frames

30

ETH zürich spl.inf.ethz.ch  
@spl\_eth

## Page table jargon

- **Page tables maps VPNs to PFNs**
  - Page table entry = PTE
- **VPN = Virtual Page Number**
  - Upper bits of virtual or logical address
- **PFN = Page Frame Number**
  - Upper bits of physical or logical address
- **Same number of bits (usually).**

31

ETH zürich spl.inf.ethz.ch  
@spl\_eth

## Recall: P6 Page tables (32bit)

Logical address: 20  
VPN 12  
VPO

VPN1 VPN2

20  
PPN 12  
PPO

PDE p=1 PTE p=1

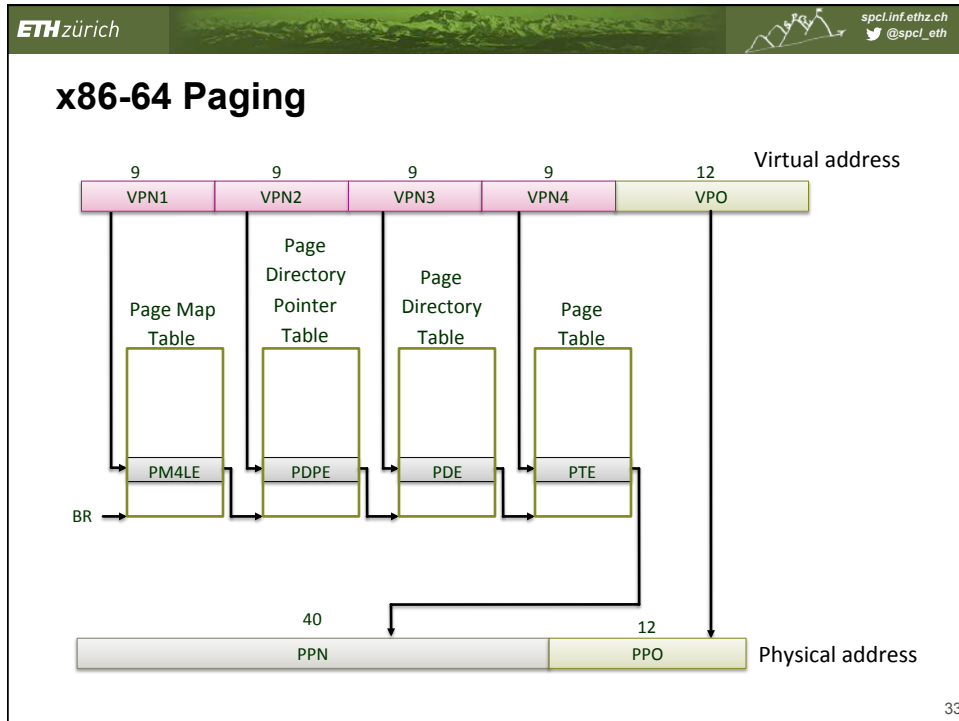
data

PDBR → Page directory → Page table → Data page

- **Pages, page directories, page tables all 4kB**

32





ETH zürich spci.inf.ethz.ch  
@spci\_eth

## Problem: performance

- **Every logical memory access needs more than two physical memory accesses**
  - Load page table entry → PFN
  - Load desired location
- **Performance ⇒ half as fast as with no translation**
  - Solution: cache page table entries

34

ETH zürich spcl.inf.ethz.ch  
@spcl\_eth

## Translating with the P6 TLB

1. Partition VPN into TLBT and TLBI.
2. Is the PTE for VPN cached in set TLBI?
3. **Yes:** Check permissions, build physical address
4. **No:** Read PTE (and PDE if not cached) from memory and build physical address

35

ETH zürich spcl.inf.ethz.ch  
@spcl\_eth

## In fact, x86 combines segmentation and paging

- Segments do (still) have uses
  - Thread-local state
  - Sandboxing (Google NativeClient, etc.)
  - Virtual machine monitors (Xen, etc.)

36

ETH zürich spci.inf.ethz.ch  
@spci\_eth

## Effective Access Time

- **Associative Lookup =  $\epsilon$  time units**
- **Assume memory cycle time is 1 time unit**
- **Hit ratio  $\alpha$  =**
  - % time that a page number is found in the TLB;
  - Depends on locality and TLB entries (**coverage**)

Then **Effective Access Time:**

$$\begin{aligned} \text{EAT} &= (1 + \epsilon) \alpha + (2 + \epsilon)(1 - \alpha) \\ &= 2 + \epsilon - \alpha \end{aligned}$$

Assuming single-level page table.

*Exercise:* work this out for the P6 2-level table

37

ETH zürich spci.inf.ethz.ch  
@spci\_eth

## Page Protection

38

ETH zürich spci.inf.ethz.ch  
@spci\_eth

## Memory Protection

- **Associate protection info with each frame**
  - Actually no - with the PTE.
- **Valid-invalid bit**
  - “valid” ⇒ page mapping is “legal”
  - “invalid” ⇒ page is not part of address space, i.e., entry does not exist
- **Requesting an “invalid” address ⇒ “fault”**
  - A “page fault”, or....

39

ETH zürich spci.inf.ethz.ch  
@spci\_eth

## Remember the P6 Page Table Entry (PTE)

<b>31</b>	<b>12</b>	<b>11</b>	<b>9</b>	<b>8</b>	<b>7</b>	<b>6</b>	<b>5</b>	<b>4</b>	<b>3</b>	<b>2</b>	<b>1</b>	<b>0</b>
Page physical base address			Avail	G	0	D	A	CD	WT	U/S	R/W	P=1

**Page base address:** 20 most significant bits of physical page address (forces pages to be 4 KB aligned)

**Avail:** available for system programmers

**G:** global page (don't evict from TLB on task switch)

**D:** dirty (set by MMU on writes)

**A:** accessed (set by MMU on reads and writes)

**CD:** cache disabled or enabled

**WT:** write-through or write-back cache policy for this page

**U/S:** user/supervisor

**R/W:** read/write

**P:** page is present in physical memory (1) or not (0)

40



ETH zürich spci.inf.ethz.ch  
@spci\_eth

## Page sharing

43

ETH zürich spci.inf.ethz.ch  
@spci\_eth

## Shared Pages Example

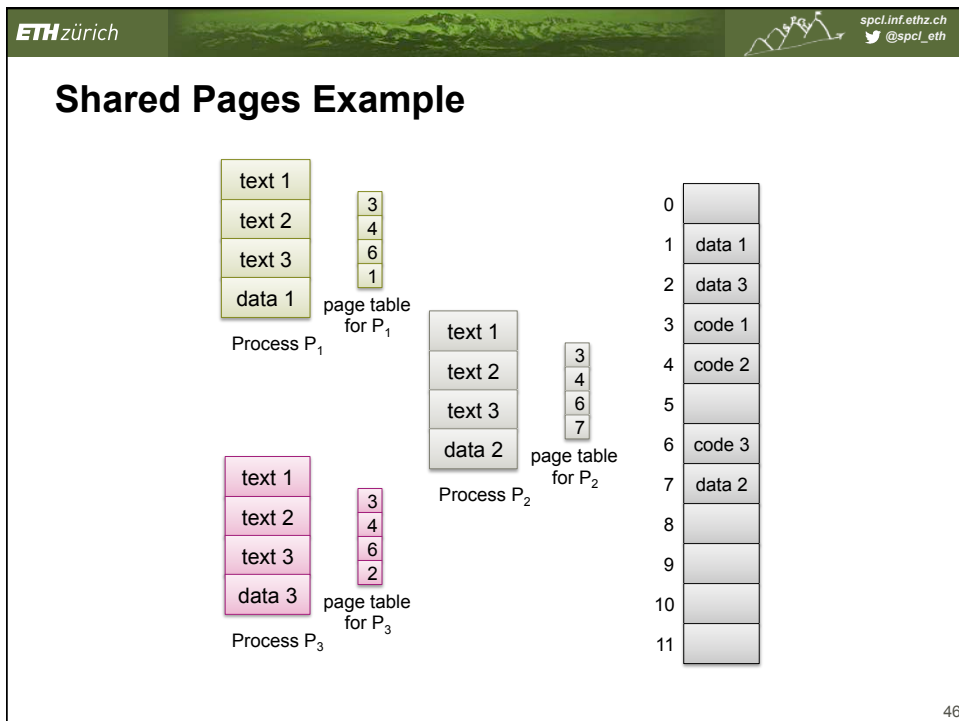
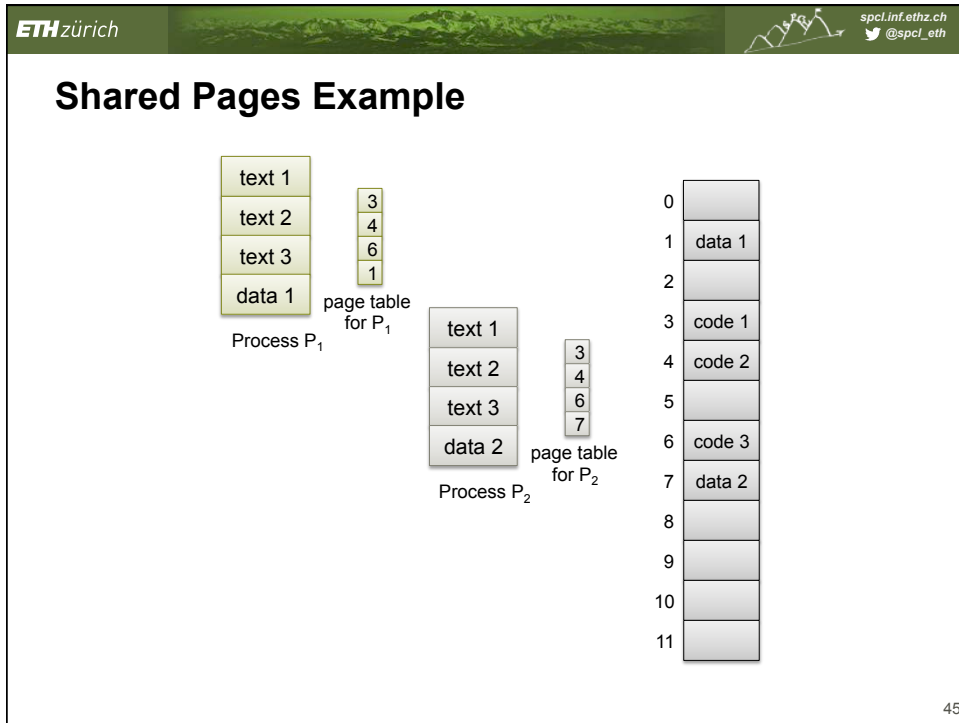
text 1
text 2
text 3
data 1




3
4
6
1

page table  
for  $P_1$

0	
1	data 1
2	
3	code 1
4	code 2
5	
6	code 3
7	
8	
9	
10	
11	

44



ETH zürich   [spci.inf.ethz.ch](http://spci.inf.ethz.ch)  
 @spci\_eth

## Shared Pages

- **Shared code**
  - One copy of read-only code shared among processes
  - Shared code appears in same location in the logical address space of all processes
  - Data segment is not shared, different for each process
  - But still mapped at same address (so code can find it)
- **Private code and data**
  - Allows code to be relocated anywhere in address space

47

ETH zürich   [spci.inf.ethz.ch](http://spci.inf.ethz.ch)  
 @spci\_eth

## Per-process protection

- **Protection bits are stored in page table**
  - Plenty of bits available in PTEs
- ⇒ **independent of frames themselves**
  - Different processes can share pages
  - Each page can have different protection to different processes
  - Many uses! E.g., debugging, communication, copy-on-write, etc.




48



ETH zürich    [spcl.inf.ethz.ch](http://spcl.inf.ethz.ch)  
[@spcl\\_eth](https://twitter.com/spcl_eth)

## Page Table Structures



49

ETH zürich    [spcl.inf.ethz.ch](http://spcl.inf.ethz.ch)  
[@spcl\\_eth](https://twitter.com/spcl_eth)

## Page table structures

- **Problem: simple linear page table is too big**
- **Solutions:**
  1. Hierarchical page tables
  2. Virtual memory page tables
  3. Hashed page tables
  4. Inverted page tables

50



ETH zürich   [sycl.inf.ethz.ch](http://sycl.inf.ethz.ch)  
[@sycl\\_eth](https://twitter.com/sycl_eth)

## Page table structures

- **Problem: simple linear page table is too big**
- **Solutions:**
  1. Hierarchical page tables
  2. Virtual memory page tables (VAX)
  3. Hashed page tables
  4. Inverted page tables

} Saw these last Semester.

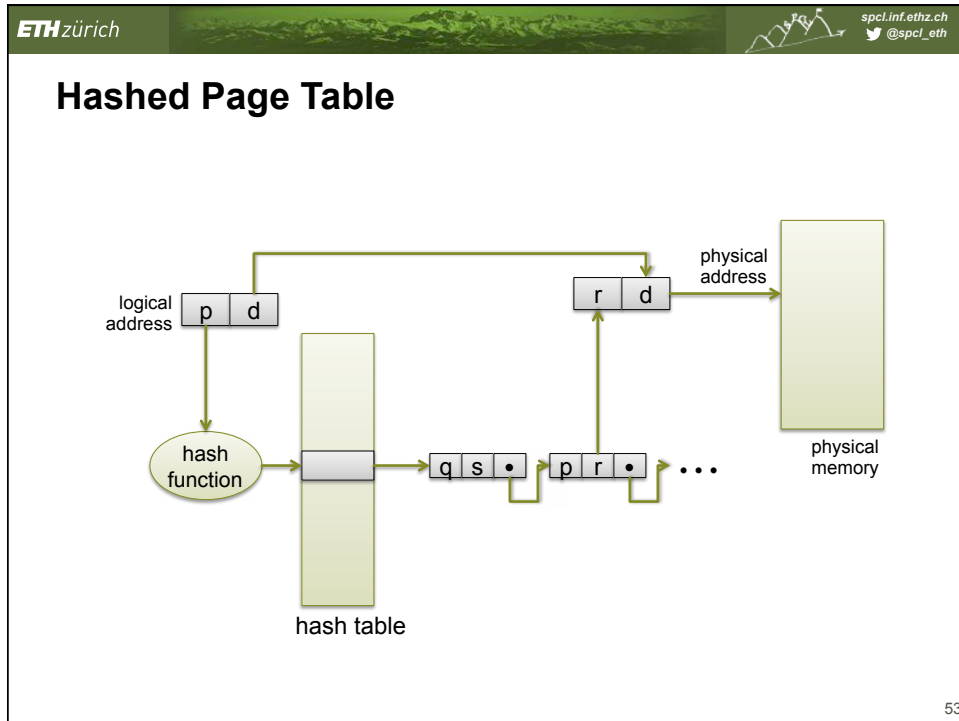
51

ETH zürich   [sycl.inf.ethz.ch](http://sycl.inf.ethz.ch)  
[@sycl\\_eth](https://twitter.com/sycl_eth)

## #3 Hashed Page Tables

- **VPN is hashed into table**
  - Hash bucket has chain of logical->physical page mappings
- **Hash chain is traversed to find match.**
- **Can be fast, but can be unpredictable**
- **Often used for**
  - Portability
  - Software-loaded TLBs (e.g., MIPS)

52

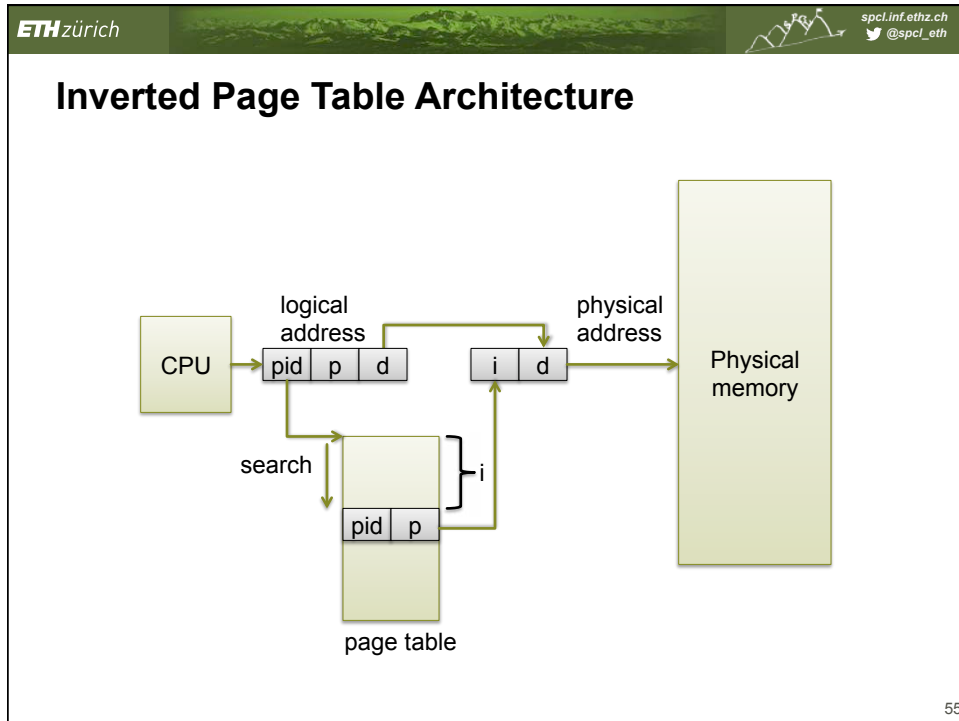


ETH zürich spci.inf.ethz.ch  
@spci\_eth

## #4 Inverted Page Table

- **One system-wide table now maps PFN -> VPN**
  - One entry for each real page of memory
  - Contains VPN, and which process owns the page
- **Bounds total size of all page information on machine**
  - Hashing used to locate an entry efficiently
- **Examples: PowerPC, ia64, UltraSPARC**

54





ETH zürich spci.inf.ethz.ch  
@spci\_eth

## The need for more bookkeeping



- **Most OSes keep their own translation info**
  - Per-process hierarchical page table (Linux)
  - System wide inverted page table (Mach, MacOS)
- **Why?**
  - Portability
  - Tracking memory objects
  - Software virtual → physical translation
  - Physical → virtual translation

56

ETH zürich   [spci.inf.ethz.ch](http://spci.inf.ethz.ch)  
[@spci\\_eth](https://twitter.com/spci_eth)

## TLB shutdown

57

ETH zürich   [spci.inf.ethz.ch](http://spci.inf.ethz.ch)  
[@spci\\_eth](https://twitter.com/spci_eth)

## TLB management

- Recall: the TLB is a **cache**.
- Machines have many MMUs on many cores  
⇒ many TLBs
- Problem: TLBs should be coherent. Why?
  - Security problem if mappings change
  - E.g., when memory is reused

58

ETH zürich spci.inf.ethz.ch  
@spci\_eth

## TLB management

	Process ID	VPN	PPN	access
Core 1	0	0x0053	0x03	r/w
	TLB: 1	0x20f8	0x12	r/w
Core 2	0	0x0053	0x03	r/w
	TLB: 1	0x0001	0x05	read
Core 3	0	0x20f8	0x12	r/w
	TLB: 1	0x0001	0x05	read

59

ETH zürich spci.inf.ethz.ch  
@spci\_eth

## TLB management

	Process ID	VPN	PPN	access
Core 1	0	0x0053	0x03	r/w
	TLB: 1	0x20f8	0x12	r/w
Core 2	0	0x0053	0x03	r/w
	TLB: 1	0x0001	0x05	read
Core 3	0	0x20f8	0x12	r/w
	TLB: 1	0x0001	0x05	read

60

Change to read only

ETH zürich spci.inf.ethz.ch  
@spci\_eth

## TLB management

	Process ID	VPN	PPN	access
Core 1	0	0x0053	0x03	r/w
	TLB: 1	0x20f8	0x12	r/w
Core 2	0	<del>0x0053</del>	0x03	r/w
	TLB: 1	0x0001	0x05	read
Core 3	0	0x20f8	0x12	r/w
	TLB: 1	0x0001	0x05	read

Change to read only

61

ETH zürich spci.inf.ethz.ch  
@spci\_eth


## TLB management

	Process ID	VPN	PPN	access
Core 1	0	0x0053	0x03	r/w
	TLB: 1	0x20f8	0x12	r/w
Core 2	0	<del>0x0053</del>	0x03	r/w
	TLB: 1	0x0001	0x05	read
Core 3	0	0x20f8	0x12	r/w
	TLB: 1	0x0001	0x05	read

Change to read only

Process 0 on core 1 can only continue once shutdown is complete!


62

ETH zürich  [spcl.inf.ethz.ch](http://spcl.inf.ethz.ch)  
[@spcl\\_eth](https://twitter.com/spcl_eth)

## Keeping TLBs consistent

- 1. Hardware TLB coherence**
  - Integrate TLB mgmt with cache coherence
  - Invalidate TLB entry when PTE memory changes
  - Rarely implemented
- 2. Virtual caches**
  - Required cache flush / invalidate will take care of the TLB
  - High context switch cost!  
⇒ Most processors use physical caches
- 5. Software TLB shutdown**
  - Most common
  - OS on one core notifies all other cores - Typically an IPI
  - Each core provides local invalidation
- 6. Hardware shutdown instructions**
  - Broadcast special address access on the bus
  - Interpreted as TLB shutdown rather than cache coherence message
  - E.g., PowerPC architecture

63

ETH zürich  [spcl.inf.ethz.ch](http://spcl.inf.ethz.ch)  
[@spcl\\_eth](https://twitter.com/spcl_eth)

## Tomorrow: demand paging

64