


Assignment 3 Intro

Pointer Recap

Memory

- 4-bit addresses, i.e. address range 0 – 15



```
char *c;
int *i;
```


Examples

```
c = 5; *c = 5 42
i = 4; *i = 42
```


Converting Page/Frame numbers of addresses

Examples

```
frame_num = paddr >> 12;
paddr = frame_num << 12; ✓
```




Frame table



Indexing off Pointers

Memory

- 4-bit addresses, i.e. address range 0 – 15



```
char *c;
int *i;
```


Examples

```
c = 5; c[0] = 'h'; c[1] = 'i';
i = 4; i[0] = 42; i[2] = 7;
```

Bump pointer allocator

Memory

- 4-bit addresses, i.e. address range 0 – 15

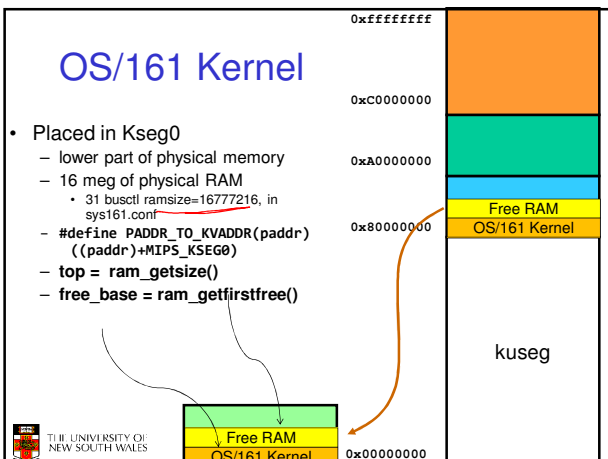
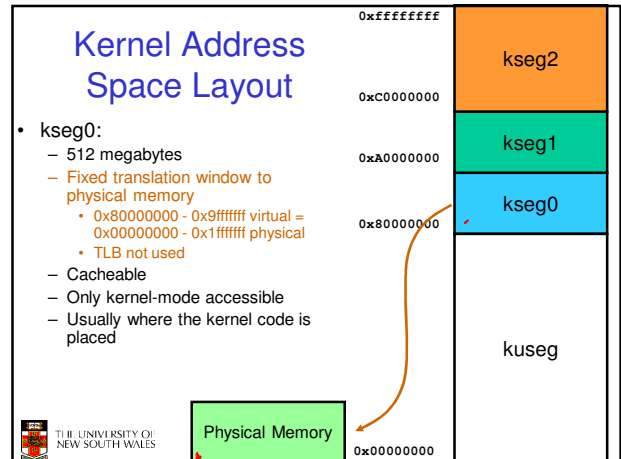
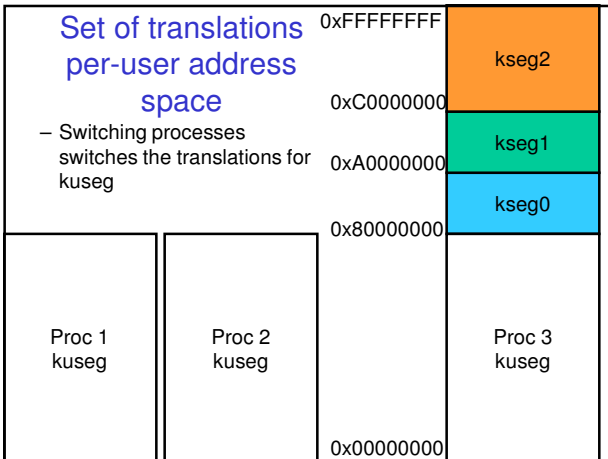
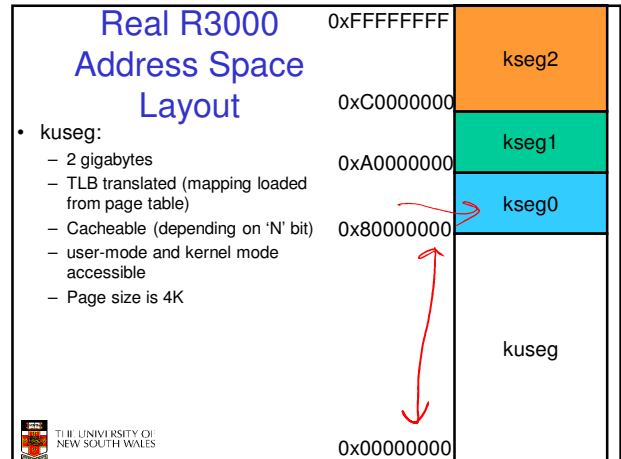
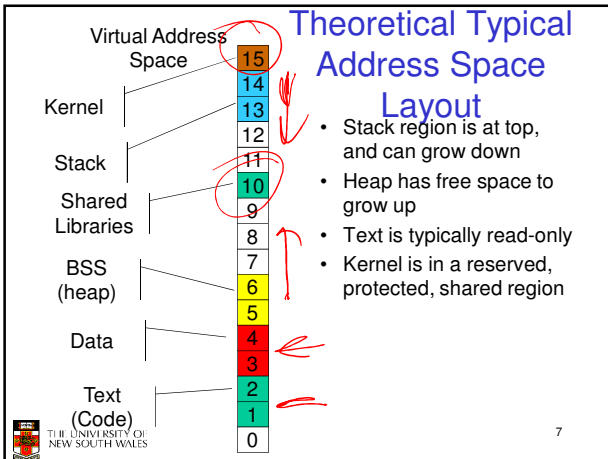


```
char *c1, *c2, *c3;
unsigned int next_free;
```

Note: Bump pointers only
allocate, not free.

Assignment 3

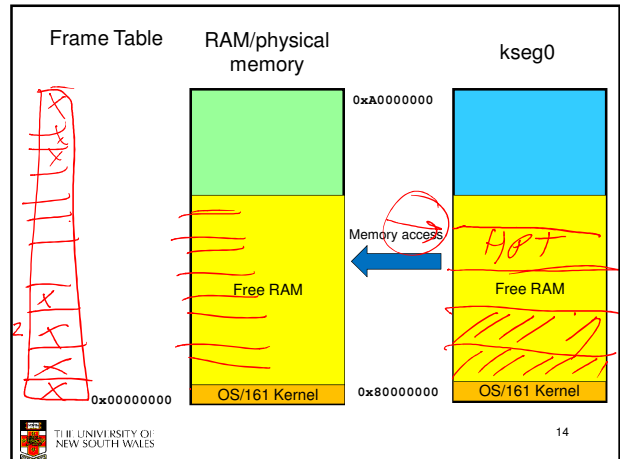

- Two parts (taking a data-structure view)
 - Frame table
 - An allocator for physical memory
 - Page table and 'region' support
 - Virtual memory for applications



Implementing alloc_kpage()/free_kpage()

- The low-level functions that kmalloc()/kfree() use to allocate/free memory in its memory pool.
 - You can assume 1 page at a time
 - But should return NULL (no memory) if you get a call for more than a page.
 - Consider using an assert() for debugging in case you get such a request
- Addresses must be in the address range of kseg0

THE UNIVERSITY OF NEW SOUTH WALES



How to 'place' my frame table?

At top:

```
struct frame_table_entry *frame_table;
location = top_of_ram - (top_of_ram / PAGE_SIZE *
    sizeof(page_table_entry));
frame_table = (struct frame_table_entry *) location;
```

How to 'place' my frame table?

At bottom:


Use existing bump allocator to allocate table, prior to switching to your allocator.

Note: You can only call `ram_getfirstfree()` once!

How can my my allocator work before and after it is initialised?

```
struct frame_table_entry *frame_table = 0;
alloc_kpages()
{
    if (ft == 0) {
        /* use ram_stealmem */
    } else {
        /* use my allocator as frame table is now initialised */
    }
}
```

Virtual Address
Space



KUseg layout

- Stack region is at top, and can grow down
- Other regions determined by ELF file
 - see `load_elf()`
 - number can vary
 - permissions specified also
 - `cs161-objdump -p testbin/huge`

```

thresher% cs161-objdump-h ./bin/true

./bin/true: file format elf32-tradbigmips

Sections:
Idx Name          Size      VMA           LMA     File off  Algn
0 .reginfo        00000018 00400094 00400094 00000094 2**2
CONTENTS, ALLOC, LOAD, READONLY, DATA, LINK_ONCE_SAME_SIZE
1 .text           000001d0 004000b0 004000b0 000000b0 2**4
CONTENTS, ALLOC, LOAD, READONLY, CODE
2 .data           00000000 10000000 10000000 00001000 2**4
CONTENTS, ALLOC, LOAD, DATA
3 .sbss           00000008 10000000 10000000 00001000 2**2
ALLOC
4 .bss            00000000 10000010 10000010 00001008 2**4
ALLOC
5 .comment        00000036 00000000 00000000 00001008 2**0
CONTENTS, READONLY
6 .pdr            000004d0 00000000 00000000 00001040 2**2
CONTENTS, READONLY
7 .mdebug.abi32  00000000 00000000 00000000 000014e0 2**0
CONTENTS, READONLY
thresher%

```

```


thresher% cs161-objdump-p ./bin/true

./bin/true: file format elf32-tradbigmips

Program Header:
0x70000000 off 0x00000094 vaddr 0x00400094 paddr 0x00400094 align 2**2
filesz 0x00000018 memsz 0x00000018 flags r-
LOAD off 0x00000000 vaddr 0x00400000 paddr 0x00400000 align 2**12
filesz 0x00000280 memsz 0x00000280 flags r-x
LOAD off 0x00001000 vaddr 0x10000000 paddr 0x10000000 align 2**12
filesz 0x00000000 memsz 0x00000010 flags rw-
private flags = 1001:[abi=O32][mips1][not 32bitmode]

thresher%

```



Zero fill fresh pages prior to mapping

Process Layout

- Process layout in KUseg
 - regions specified by calls to
 - as_define_stack()
 - as_define_region()
 - usually implemented as a linked list of region specifications
 - as_prepare_load()
 - make READONLY regions READWRITE for loading purposes
 - as_complete_load()
 - enforce READONLY again

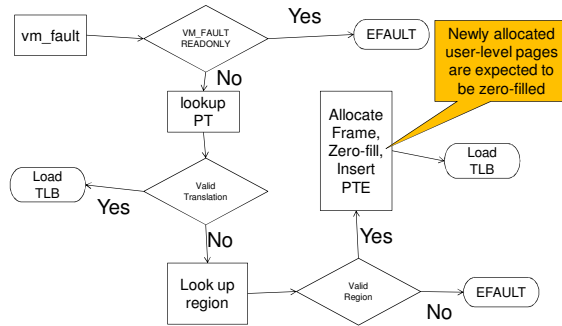
Process Layout

- Need to keep translation table for KUSEG

- as_create()
 - allocate a data structure used to keep track of an address space
 - i.e. regions
 - proc_getas() used to get access to current address space struct
 - You can use the current address space pointer as process it.
 - struct addrspace *as;
 - pid = (uint32_t) as
- as_destroy()
 - deallocate book keeping and page tables.
 - deallocate frames used

- as_copy()
 - allocates a new (destination) address space
 - adds all the same regions as source
 - roughly, for each mapped page in source
 - allocate a frame in dest
 - copy contents from source frame to dest frame
 - add PT entry for dest
- as_activate()
 - flush TLB
 - (or set the hardware asid)
- as_deactivate()
 - flush TLB
 - (or flush an asid)

VM Fault Approximate Flow Chart



kprintf()

• Do not use it in vm_fault()

- kprintf() blocks current process while printing
 - Switches to another process
 - Context switch flushes TLB
 - Flushes what you just inserted
 - Endless loop

trace161 can help with debugging

<http://cgi.cse.unsw.edu.au/~cs3231/06s1/os161/man/sys161/index.html>

- The following additional options control trace161's tracing and are ignored by sys161:
 - **-f tracefile**
 - Set the file trace information is logged to. By default, stderr is used. Specifying -f sends output to stdout instead of stderr.
 - **-t traceflags**
 - Tell System/161 what to trace. The following flags are available:
 - d Trace disk I/O
 - e Trace emul'd I/O
 - j Trace jumps and branches
 - k Trace instructions in kernel mode
 - n Trace network I/O
 - t Trace TLB/MMU activity
 - u Trace instructions in user mode
 - x Trace exceptions
- Caution: tracing instructions generates huge amounts of output that may overwhelm smaller host systems.

```

wagner% trace161 -tt kernel
sys161: System/161 release 2.0.8, compiled Feb 19 2017 14:31:56
sys161: Tracing enabled: tlb
trace: 00 tlbp: 81000/000 -> 00000 ----: [0]
trace: 00 tlbp: 81001/000 -> 00000 ----: [1]
trace: 00 tlbp: 81002/000 -> 00000 ----: [2]
trace: 00 tlbp: 81003/000 -> 00000 ----: [3]
trace: 00 tlbp: 81004/000 -> 00000 ----: [4]
trace: 00 tlbp: 81005/000 -> 00000 ----: [5]
trace: 00 tlbp: 81006/000 -> 00000 ----: [6]
trace: 00 tlbp: 81007/000 -> 00000 ----: [7]
trace: 00 tlbp: 81008/000 -> 00000 ----: [8]
trace: 00 tlbp: 81009/000 -> 00000 ----: [9]
trace: 00 tlbp: 8100a/000 -> 00000 ----: [10]
trace: 00 tlbp: 8100b/000 -> 00000 ----: [11]
trace: 00 tlbp: 8100c/000 -> 00000 ----: [12]
trace: 00 tlbp: 8100d/000 -> 00000 ----: [13]
trace: 00 tlbp: 8100e/000 -> 00000 ----: [14]
trace: 00 tlbp: 8100f/000 -> 00000 ----: [15]
trace: 00 tlbp: 81010/000 -> 00000 ----: [16]
trace: 00 tlbp: 81011/000 -> 00000 ----: [17]
trace: 00 tlbp: 81012/000 -> 00000 ----: [18]
trace: 00 tlbp: 81013/000 -> 00000 ----: [19]
trace: 00 tlbp: 81014/000 -> 00000 ----: [20]
  
```

```

.....
trace: 00 tlbp: 8103f/000 -> 00000 ----: [63]
trace: 00 tlbp: 81040/000 -> NOT FOUND
trace: 00 tlbwi: [ 0] 81000/000 -> 00000 ---- ==> 81040/000 -> 00000 ----
trace: 00 tlbp: 81041/000 -> NOT FOUND
trace: 00 tlbwi: [ 1] 81001/000 -> 00000 ---- ==> 81041/000 -> 00000 ----
trace: 00 tlbp: 81042/000 -> NOT FOUND
trace: 00 tlbwi: [ 2] 81002/000 -> 00000 ---- ==> 81042/000 -> 00000 ----
trace: 00 tlbp: 81043/000 -> NOT FOUND
trace: 00 tlbwi: [ 3] 81003/000 -> 00000 ---- ==> 81043/000 -> 00000 ----
trace: 00 tlbp: 81044/000 -> NOT FOUND
trace: 00 tlbwi: [ 4] 81004/000 -> 00000 ---- ==> 81044/000 -> 00000 ----
trace: 00 tlbp: 81045/000 -> NOT FOUND
trace: 00 tlbwi: [ 5] 81005/000 -> 00000 ---- ==> 81045/000 -> 00000 ----
trace: 00 tlbp: 81046/000 -> NOT FOUND
trace: 00 tlbwi: [ 6] 81006/000 -> 00000 ---- ==> 81046/000 -> 00000 ----
trace: 00 tlbp: 81047/000 -> NOT FOUND
trace: 00 tlbwi: [ 7] 81007/000 -> 00000 ---- ==> 81047/000 -> 00000 ----
trace: 00 tlbp: 81048/000 -> NOT FOUND
trace: 00 tlbwi: [ 8] 81008/000 -> 00000 ---- ==> 81048/000 -> 00000 ----
  
```

```

.....
trace: 00 tlbwi: [60] 8103c/000 -> 00000 ---- ==> 8107c/000 -> 00000 ----
trace: 00 tlbp: 8107d/000 -> NOT FOUND
trace: 00 tlbwi: [61] 8103d/000 -> 00000 ---- ==> 8107d/000 -> 00000 ----
trace: 00 tlbp: 8107e/000 -> NOT FOUND
trace: 00 tlbwi: [62] 8103e/000 -> 00000 ---- ==> 8107e/000 -> 00000 ----
trace: 00 tlbp: 8107f/000 -> NOT FOUND
trace: 00 tlbwi: [63] 8103f/000 -> 00000 ---- ==> 8107f/000 -> 00000 ----
  
```

OS/161 base system version 2.0.3
(with locks/CVs, system calls solutions)
Copyright (c) 2000, 2001-2005, 2008-2011, 2013, 2014
President and Fellows of Harvard College. All rights reserved.

Put-your-group-name-here's system version 0 (ASST3 #29)


16208k physical memory available
Device probe...
lamebus0 (system main bus)
emu0 at lamebus0

```

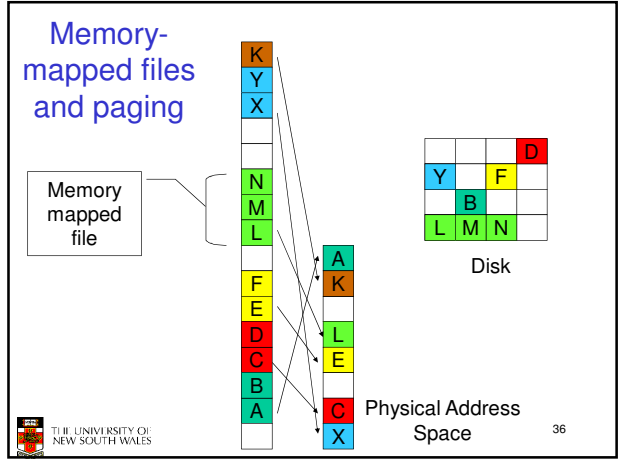
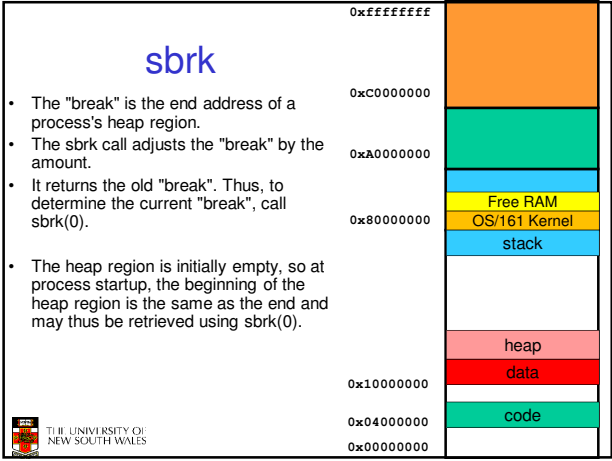
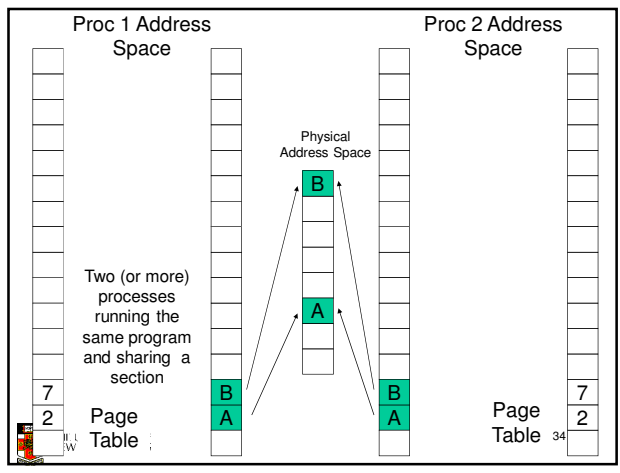
End of trace from bin/true

trace: 00 tlblookup: 00400/000 -> no match
trace: 00 tlbwr: [58] 8003a/000 -> 00000 ---- ==> 00400/000 -> 00034 -V--
trace: 00 tlblookup: 00400/000 -> 00034 -V--: [58] - OK
trace: 00 tlblookup: 00400/000 -> 00034 -V--: [58] - OK
trace: 00 tlblookup: 00410/000 -> no match
trace: 00 tlbwr: [34] 80022/000 -> 00000 ---- ==> 00410/000 -> 00036 -VD-
trace: 00 tlblookup: 00400/000 -> 00034 -V--: [58] - OK
trace: 00 tlblookup: 00400/000 -> 00034 -V--: [58] - OK
trace: 00 tlblookup: 00410/000 -> 00036 -VD--: [34] - OK
trace: 00 tlblookup: 00410/000 -> 00036 -VD--: [34] - OK
trace: 00 tlblookup: 00400/000 -> 00034 -V--: [58] - OK
trace: 00 tlblookup: 7ffff/000 -> 00035 -VD--: [25] - OK
trace: 00 tlblookup: 00400/000 -> 00034 -V--: [58] - OK
trace: 00 tlblookup: 00400/000 -> 00034 -V--: [58] - OK
trace: 00 tlblookup: 00400/000 -> 00034 -V--: [58] - OK

```



- ## Advance Assignment
- Shared pages and copy-on-write
 - Sbrk()
 - Demand loading and mmap
 - Paging



mmap semantics

```
void *mmap(size_t length, int prot, int fd, off_t offset);  
int munmap(void *addr);
```