

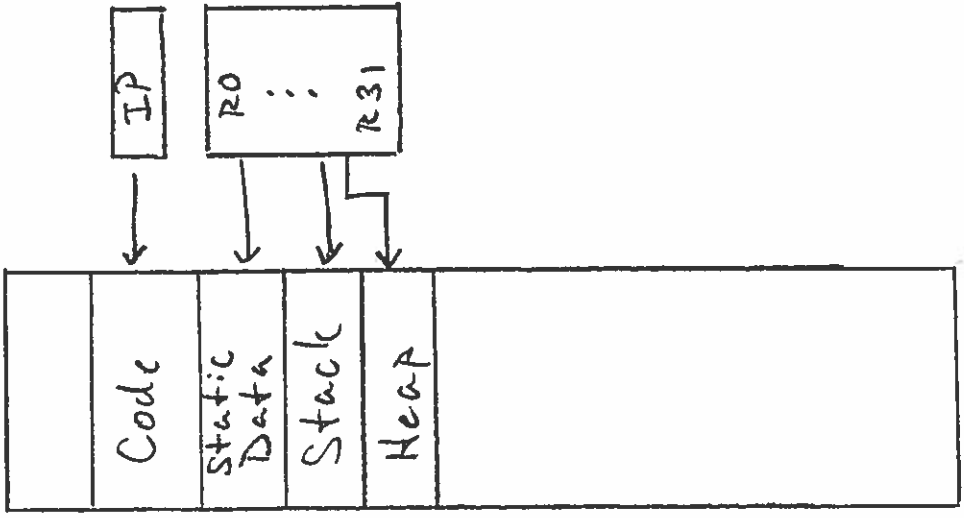
How Do Programs Use Names?

- Registers: $i, t, *t$
Register addressing mode
- Code: spaces In String
Absolute or PC-relative
addressing mode
- Static data: s
Absolute addressing mode
- Stack: i, t , return address
Relative addressing mode
- Heap:
Indirect addressing mode
- Dynamically-linked libraries:
Indirect addressing mode


```
static char s[LEN];  
  
int spacesInString(char *t) {  
    int i=0;  
    while (*t) {  
        if (*t++ == ' ')  
            i++;  
    }  
    return i;  
}
```

Relocation and Protection

Single Program



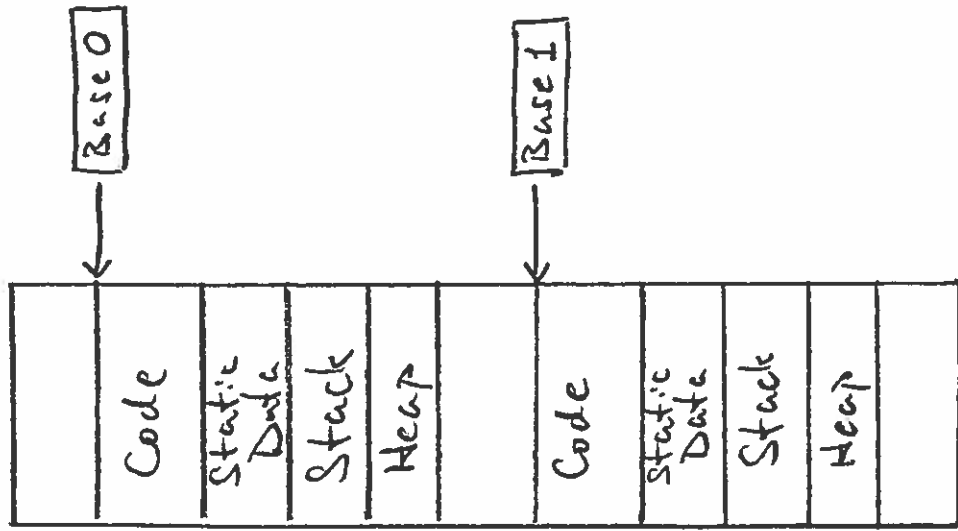
Multiple Programs




- Relocation
- No Protection

Relocation and Protection

Base Register Addressing




- Relocation
- No Protection



- Relocation
- Protection

Sharing





Logical
Name-
Space 1

Logical
Name-
Space 2

Segmented Addressing

Translation Lookaside Buffer:
 Small, Fully associative cache
 of segment
 descriptors
 (~64
 entries)



Resource Management


Given a name, where should the name reside in memory?
or, "What parts of a name space should I keep in memory?"

Problems:

- ① Physical memory is finite (i.e. small)
- ② Multiple name spaces may not fit in memory
- ③ Each name space may not fit in memory

Solution: Virtual Memory

Keep most of name spaces in secondary storage (disk) and move "important" portions into physical memory automatically.




Solve 2 problems:

- Mapping (relocation)
- Management

(very similar to hardware caches)




Memory Fragmentation

using variable-sized (segments) units




As segments come and go, the storage is "fragmented"; therefore, at some point segments must be moved around to compact storage \Rightarrow "burping the memory"

Page Addressing




Paged Addressing



Translation Lookaside Buffer:

Small, fully associative cache
of PTEs (~64 entries)



Caching vs. Paging

Caching

Block Size:
~ 16-128 bytes

Miss Rate:
1-20%

Hit Time:
1 cycle

Miss Penalty:
~100 cycles

Organization:
Direct-Mapped, Set-Associative, Fully Associative

Write Hit Policy:
Write-through or Write back

Write Miss Policy:
Write allocate or Write no-allocate

Paging

Page
~4K-64K bytes

Page Miss Rate
0.00001-0.001%

Page Hit Time
~100 cycles

Page Fault
~10⁶ cycles

Organization:
Fully Associative

Write Hit Policy:
Write back

Write Miss Policy:
Write allocate

Page Frame Management

- Allocate some number of page frames to each process.
- Maintain free page list.
- Evict pages when free pages falls below "low watermark".
- Evict pages from a process using LRU replacement policy.
- If dirty bit is clean, don't copy page back to disk.

How many page frames should each process get?

Monitor page fault frequency for each process

- Try to keep working set of each process in memory
 - Page fault frequency above some upper limit
 - ⇒ Increase page frame allocation.
 - Page fault frequency below some lower limit
 - ⇒ Decrease page frame allocation.
- "Swap out" entire process if there are insufficient page frames for working set.