

# Computer Systems and Organization 2

# themes

automating building software

libraries, taking advantage of incremental compilation

sharing machines

multiple users/programs on one system

parallelism and concurrency

doing two+ things at once

under the hood of sockets

layered design of networks

implementing secure communication

under the hood of fast processors

caching, (hidden) parallelism, avoiding idle time

# themes

## automating building software

libraries, taking advantage of incremental compilation

## sharing machines

multiple users/programs on one system

## parallelism and concurrency

doing two+ things at once

## under the hood of sockets

layered design of networks

implementing secure communication

## under the hood of fast processors

caching, (hidden) parallelism, avoiding idle time

# make

```
$ ./foo.exe
```

```
...
```

```
...
```

```
$ edit readline.c
```

```
$ make
```

```
clang -g -O -Wall -c readline.c -o readline.o
```

```
ar rcs terminal.o readline.o libreadline.a
```

```
clang -o foo.exe foo.o foo-utility.o -L. -lreadline
```

```
$
```

# themes

automating building software

libraries, taking advantage of incremental compilation

sharing machines

multiple users/programs on one system

parallelism and concurrency

doing two+ things at once

under the hood of sockets


layered design of networks

implementing secure communication


under the hood of fast processors

caching, (hidden) parallelism, avoiding idle time

# address translation




# address translation



every address accessed  
instructions *and* data

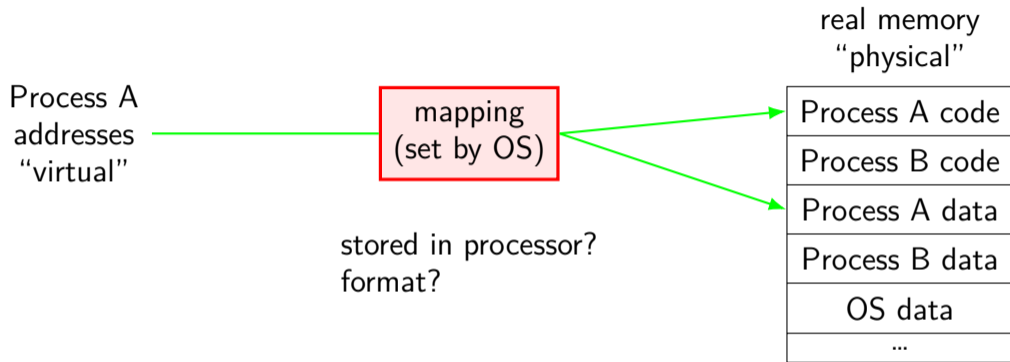
# address translation



program addresses are 'virtual'  
real addresses are 'physical'  
can be **different sizes!**




# address translation



# address spaces

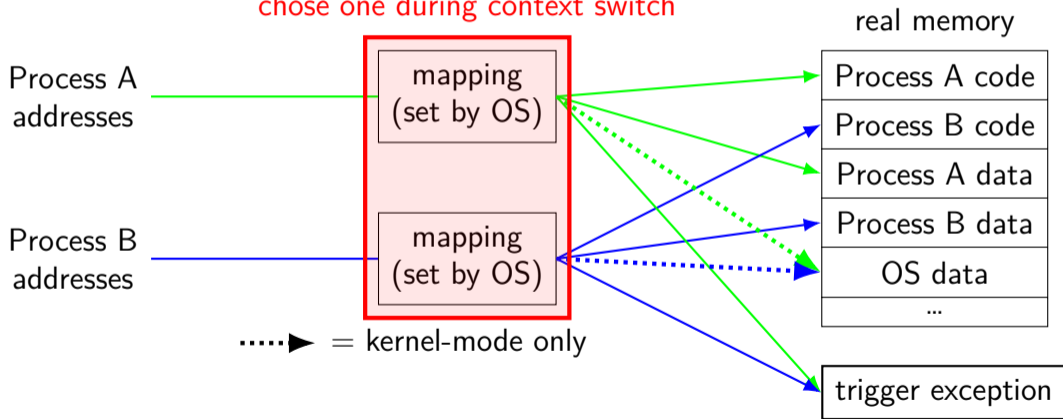
illusion of **dedicated memory**



# address spaces

illusion of **dedicated memory**

chose one during context switch



# themes

automating building software

libraries, taking advantage of incremental compilation

sharing machines

multiple users/programs on one system

parallelism and concurrency

doing two+ things at once

under the hood of sockets


layered design of networks

implementing secure communication


under the hood of fast processors

caching, (hidden) parallelism, avoiding idle time

# keyboard input timeline



# time multiplexing



# time multiplexing



...


```
call get_time  
    // whatever get_time does  
movq %rax, %rbp
```

———— million cycle delay ————

```
call get_time  
    // whatever get_time does  
subq %rbp, %rax
```

...

# time multiplexing



...

```
call get_time
```

```
    // whatever get_time does
```

```
movq %rax, %rbp
```

———— million cycle delay ————

```
call get_time
```

```
    // whatever get_time does
```

```
subq %rbp, %rax
```

...




# multiple cores+threads

core 1:



core 2:




multiple cores? each core still divided up

# multiple cores+threads

core 1:



core 2:



one program with multiple *threads*

# themes

automating building software

libraries, taking advantage of incremental compilation

sharing machines

multiple users/programs on one system

parallelism and concurrency

doing two+ things at once

under the hood of sockets

layered design of networks

implementing secure communication

under the hood of fast processors

caching, (hidden) parallelism, avoiding idle time

# permissions

```
$ ls /u/other/secret  
ls: cannot open directory '/u/other/secret': Permission denied  
$ shutdown  
shutdown: Permission denied
```

# themes

automating building software

libraries, taking advantage of incremental compilation

sharing machines

multiple users/programs on one system

parallelism and concurrency

doing two+ things at once

under the hood of sockets

layered design of networks

implementing secure communication

under the hood of fast processors

caching, (hidden) parallelism, avoiding idle time

# layers

application	HTTP, SSH, SMTP, ...	application-defined meanings
transport	TCP, UDP, ...	reach    correct    program, reliability/streams
network	IPv4, IPv6, ...	reach    correct    machine (across networks)
link	Ethernet, Wi-Fi, ...	coordinate shared wire/radio
physical	...	encode bits for wire/radio

# names and addresses

**name**

logical identifier

variable counter

DNS name `www.virginia.edu`

DNS name `mail.google.com`

DNS name `mail.google.com`

DNS name `reiss-t3620.cs.virginia.edu`

DNS name `reiss-t3620.cs.virginia.edu`

service name `https`

service name `ssh`

**address**

location/how to locate

memory address `0x7FFF9430`

IPv4 address `128.143.22.36`

IPv4 address `216.58.217.69`

IPv6 address `2607:f8b0:4004:80b::2005`

IPv4 address `128.143.67.91`

MAC address `18:66:da:2e:7f:da`

port number `443`

port number `22`

# secure communication?

how do you know who your socket is to?

who can read what's on the socket?

what can you do to restrict this?



# themes

automating building software

libraries, taking advantage of incremental compilation

sharing machines

multiple users/programs on one system

parallelism and concurrency

doing two+ things at once

under the hood of sockets


layered design of networks

implementing secure communication

under the hood of fast processors

caching, (hidden) parallelism, avoiding idle time

# 2004 CPU




Clock Generator



Image: approx 2004 AMD press image of Opteron die;  
approx register location via chip-architect.org (Hans de Vries)

# 2004 CPU

▲ Registers




Clock Generator



Image: approx 2004 AMD press image of Opteron die;  
approx register location via chip-architect.org (Hans de Vries)

# 2004 CPU


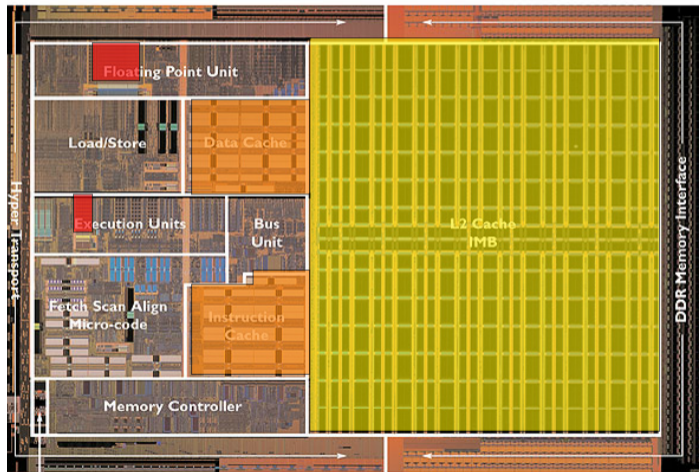


Clock Generator

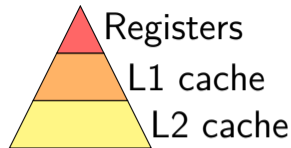
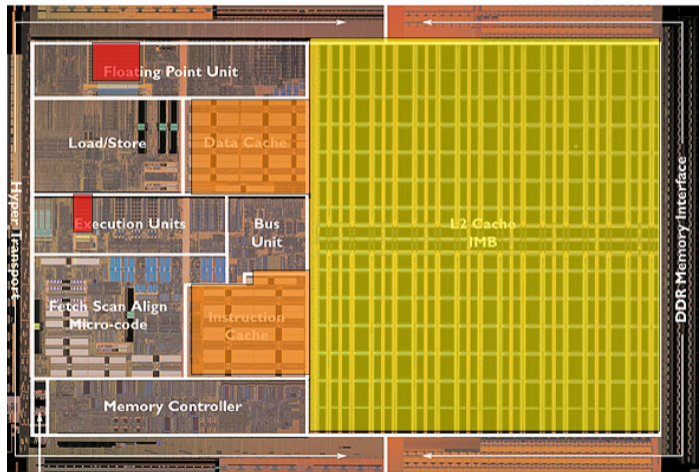


Image: approx 2004 AMD press image of Opteron die;  
approx register location via chip-architect.org (Hans de Vries)


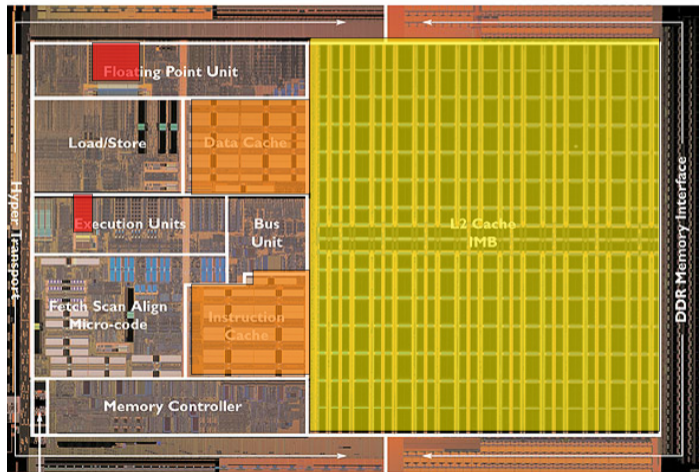
# 2004 CPU





# 2004 CPU



# 2004 CPU



# 2004 CPU



Clock Generator



Image: approx 2004 AMD press image of Opteron die;  
approx register location via chip-architect.org (Hans de Vries)



## some performance examples

```
example1:  
    movq $1000000000000, %rax  
loop1:  
    addq %rbx, %rcx  
    decq %rax  
    jge loop1  
    ret
```

about 30B instructions  
my desktop: approx 2.65 sec

```
example2:  
    movq $1000000000000, %rax  
loop2:  
    addq %rbx, %rcx  
    addq %r8, %r9  
    decq %rax  
    jge loop2  
    ret
```

about 40B instructions  
my desktop: approx 2.65 sec

# some performance examples

```
example1:  
    movq $1000000000000, %rax  
loop1:  
    addq %rbx, %rcx  
    decq %rax  
    jge loop1  
    ret
```

about 30B instructions  
my desktop: approx 2.65 sec

```
example2:  
    movq $1000000000000, %rax  
loop2:  
    addq %rbx, %rcx  
    addq %r8, %r9  
    decq %rax  
    jge loop2  
    ret
```

about 40B instructions  
my desktop: approx 2.65 sec

# logistics

# labs

attend lab in person and get checked off by TA, *or*

(most labs) submit something to submission site and we'll grade it

submit to submission site? don't care if you attend the lab

more strict about submissions without checkoffs

in-person lab checkoff of incomplete lab at least 50% credit

if both checkoff + submission, will use higher score

some labs will basically require attendance

or contact me for other arrangements if you can't (sick, etc.)

logistically won't work otherwise — e.g. code review

# lab collaboration and submissions

please collaborate on labs!

when working with others on lab and submitting code files

please indicate who you worked with in those files  
via comment or similar

# lab space

if labs are full, might kick out students from 'wrong' lab section

for 3:30pm, please come to registered room

for 5pm, 6:30pm, based on registration  
should only need one room

plan: those labs Rice 130 only

will send announcement/have something posted on unused room

# homeworks

several homework assignments

done individually

generally due on Fridays

(tentative dates on schedule)

# homework/lab automatic testing

some homeworks/labs have automatic testing

with some delay after you submit

- usually 10s of minutes

- depending on assignment, number of submissions in queue

- if you submit very early, testing program might not be setup yet

when testing program doesn't understand/can't test something,  
left for manual grading ("not yet graded")

intention is that testing results are not surprises

if you did some manual testing (no hidden requirements, etc.)

if you think testing program made a mistake,  
please submit regrade request



# warmup assignment

first homework

due week from Friday 8 Sep @ 11:59pm

write C function to split a string into array of strings  
with dynamic memory allocation

write C program to call function using input/command-line  
arguments

write Makefile for it (next topic, next week's lab)

# quizzes

released evening after Thursday lecture  
starting *next* week

due 15 minutes before lecture on Tuesdays

about lecture and/or lab from the prior week

4–6 questions

*individual*, open book, open notes, open Internet

okay: looking up resources/tutorials/etc.

not okay: asking Stack Overflow the quiz question

not okay: IMing your friend the quiz question

## on help on quiz questions

I and the TAs won't answer quiz questions...

but we will answer questions about the lecture material, etc.

(and TAs (not you) are responsible for knowing  
what they can't answer

but we'd prefer you don't try to test those limits)

# going over past quizzes

have in past gone over quiz Qs in lecture  
either when a lot missed it  
or on request in lecture

also fine office hour/Piazza question

# readings

in lieu of textbook, have readings

mostly written by Prof Tychnovich (now at UIUC) with edits by me

on website; should be indicated with corresponding lecture

# lecture + assignment sync

generally:

quiz after lecture and/or lab coverage

labs after lecture coverage

homework after lab coverage

means homework (and sometimes quiz)

may be relatively delayed from lecture coverage

# exams

1 final exam

no midterms — instead:

- quizzes count a lot

- slightly more homework/lab than pilot

# development enviroment

official: department machines via SSH or NX (remote desktop)

you can also use your own machines, but...

we will test your code on x86-64 Linux

I haven't checked assignments on a Windows or OS X machine



# getting help

office hours — calendar will be posted on website

mix of in-person and remote, indicated on calendar

remote OH will use Discord + online queue

in-person OH may or may not — indicated on whiteboard, probably

## Piazza

use private questions if homework code, etc.

emailing me (preferably with '3130' in subject)

# collaboration (1)

labs — you can/should work with other students  
everyone should understand the work submitted

homeworks — individual

write your own code / do not share your code  
can ask/look up conceptual questions of others  
others includes other students, Q&A sites, code generation tools, etc.  
**cite** any sources you use (comments in code)

## collaboration (2)

quizzes — individual

but open book+notes+etc.

can/should have help reviewing lecture/readings/etc.

legitimate questions for office hours

don't ask other students, stack overflow, gen AI tools, etc. the quiz questions

don't try to find exactly the quiz question on stack overflow

# feedback

anonymous feedback on Canvas

would appreciate feedback (esp. when I can do something)

(but not a good way to ask for regrades, etc.)

# late policy

no late quizzes

one quiz dropped (unconditionally)

90% credit for 0–72 hours late homeworks

for labs that allow submission only

lab submission due time is 11:59am the next day

90% credit for 0–24 hours late

no late lab checkoffs except by special arrangement

# excused lateness

special circumstances?

illness, emergency, etc.

contact me, we'll figure something out

please don't attend lab/etc. sick!

# attendance

I won't take attendance in lecture


I will attempt to have lecture recordings

sometimes there may be issues with the recording


**building**



# files in building C programs [dynamic linking]

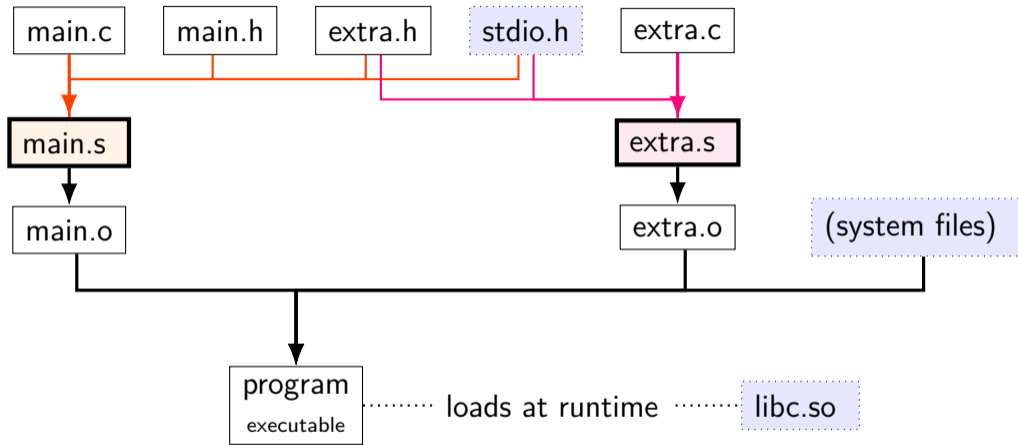


# files in building C programs [dynamic linking]



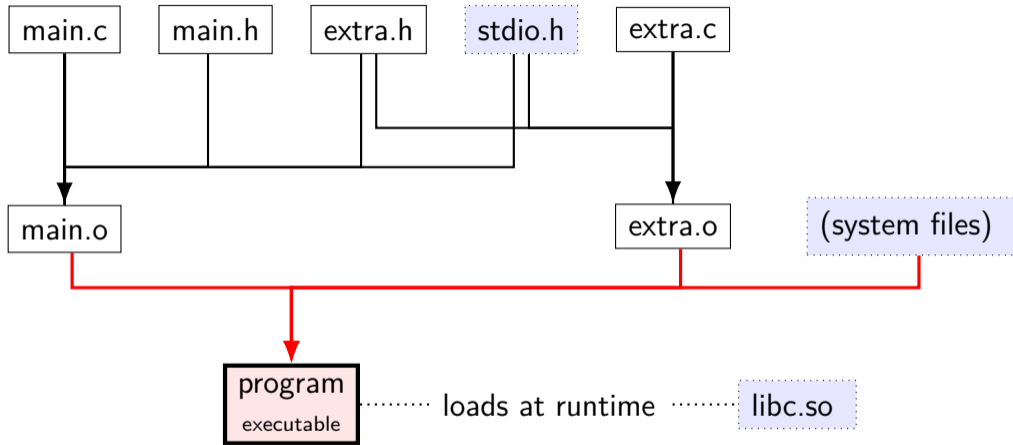
```
clang -c main.c  
clang -c extra.c
```

# files in building C programs [dynamic linking]




```
clang -S -c main.c  
clang -S -c extra.c
```

# files in building C programs [dynamic linking]




```
clang -o program main.o extra.o
```

# files in building C programs [dynamic linking]



`./program ...`

# files in building C programs [static linking]



# file extensions

name

.c	C source code
.h	C header file
.s (or .asm)	assembly file
.o (or .obj)	object file (binary of assembly)
(none) (or .exe)	executable file
.a (or .lib)	statically linked library [collection of .o files]
.so (or .dll or .dylib)	dynamically linked library ['shared object']

# static libraries

Unix-like *static* libraries: libfoo.a

internally: archive of .o files with index

create: `ar rcs libfoo.a file1.o file2.o ...`

use: `cc ... -o program -L/path/to/lib ... -lfoo`

no space between `-l` and library name

`cc` could be `clang`, `gcc`, `clang++`, `g++`, etc.

`-L/path/to/lib` not needed if in standard location



# shared libraries

Linux *shared* libraries: libfoo.so

create:

compile .o files with `-fPIC` (position independent code)

then: `cc -shared ... -o libfoo.so`

use: `cc ...-o program -L/path/to/lib ...-lfoo`

## finding shared libraries (1)

```
$ ls
```

```
libexample.so  main.c
```

```
$ clang -o main main.c -lexample
```

```
/usr/bin/ld: cannot find -lexample
```

```
clang: error: linker command failed with exit code 1 (use -v to show invocation)
```

```
$ clang -o main main.c -L. -lexample
```

```
$ ./main
```

```
./main: error while loading shared libraries:
```

```
libexample.so: cannot open shared object file: No such file or directory
```

# finding shared libraries (1)

```
$ ls
```

```
libexample.so  main.c
```

```
$ clang -o main main.c -lexample
```

```
/usr/bin/ld: cannot find -lexample
```

```
clang: error: linker command failed with exit code 1 (use -v to show invocation)
```

```
$ clang -o main main.c -L. -lexample
```

```
$ ./main
```

```
./main: error while loading shared libraries:
```

```
libexample.so: cannot open shared object file: No such file or directory
```

---

```
$ LD_LIBRARY_PATH=. ./main
```

*or*

```
$ export LD_LIBRARY_PATH=.
```

```
$ ./main
```

*or*

```
$ clang -o main main.c -L. -lexample -Wl,-rpath .
```

```
$ ./main
```

# finding shared libraries (1)

```
cc ...-o program -L/path/to/lib ...-lfoo
```

on Linux: `/path/to/lib` only used to create program  
program contains `libfoo.so` *without full path*

Linux default: `libfoo.so` expected to be in `/usr/lib`, `/lib`, and other 'standard' locations

possible overrides:

`LD_LIBRARY_PATH` environment variable

paths specified with `-Wl,-rpath=/path/to/lib` when creating executable

# exercise (incremental compilation)

program built from main.c + extra.c

main.c, extra.c both include extra.h, stdio.h

```
clang -c main.c           # command 1
clang -c extra.c          # command 2
clang -o program main.o extra.o # command 3
```

What commands need to be rerun if...

Question A: ...main.c changes?

Question B: ...extra.h changes?

# backup slides