CS 3130 intro

changelog

2023-01-19: add missing libfoo.a to ar command line example

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 on networks

under the hood of fast processors

- caching and (hidden) parallelism

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 on networks

under the hood of fast processors

caching and (hidden) parallelism

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 on networks

under the hood of fast processors

- caching and (hidden) parallelism

address translation

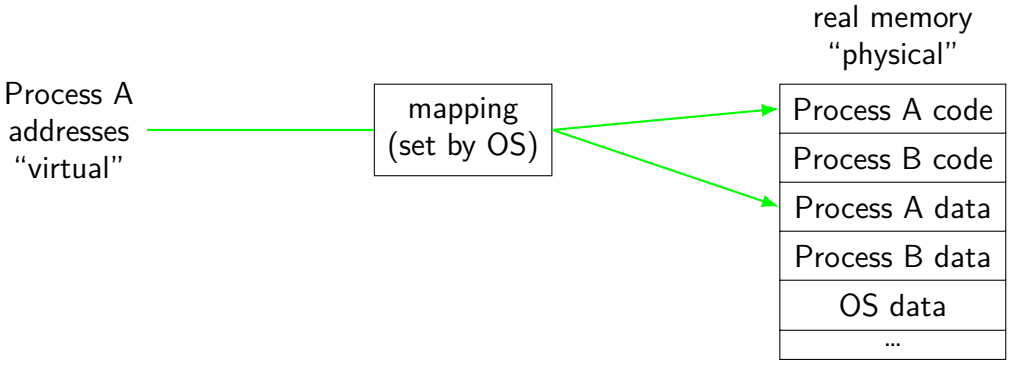

address translation

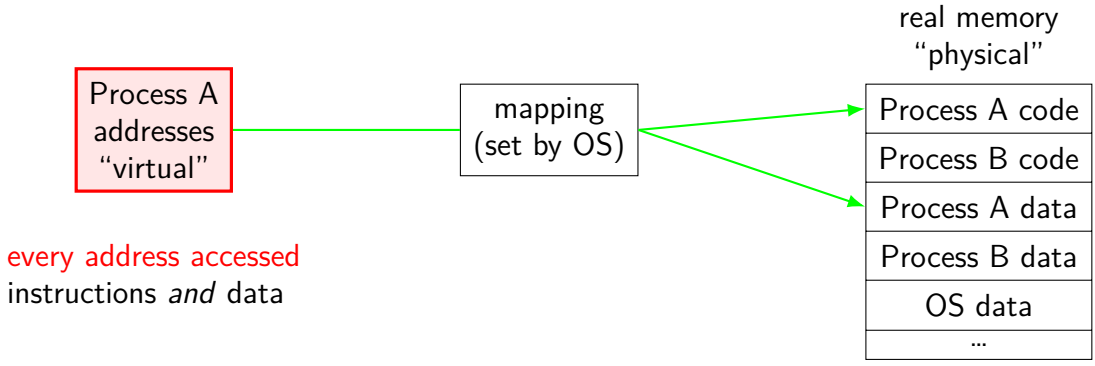

every address accessed
instructions *and* data

address translation

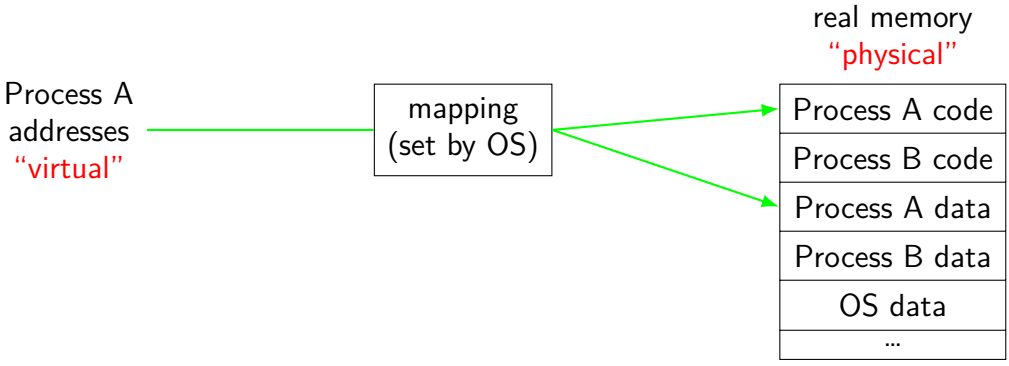

program addresses are 'virtual'
real addresses are 'physical'
can be different sizes!

address translation

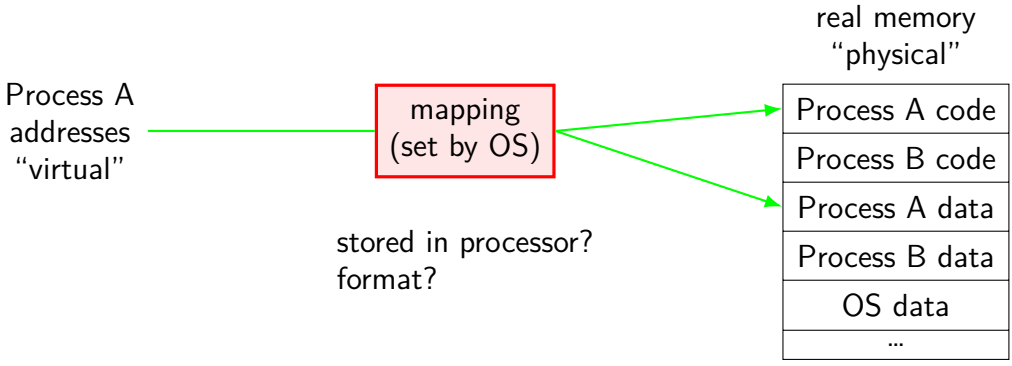

address spaces

illuision of dedicated memory

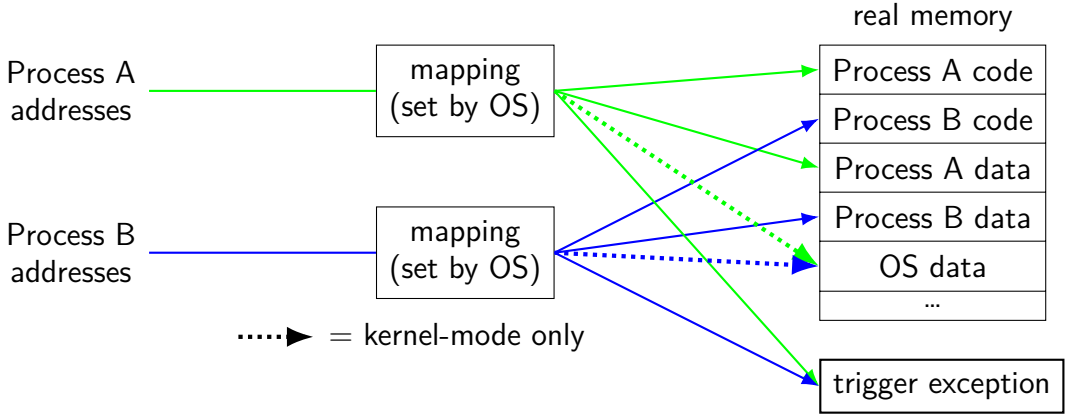

address spaces

illuision of dedicated memory

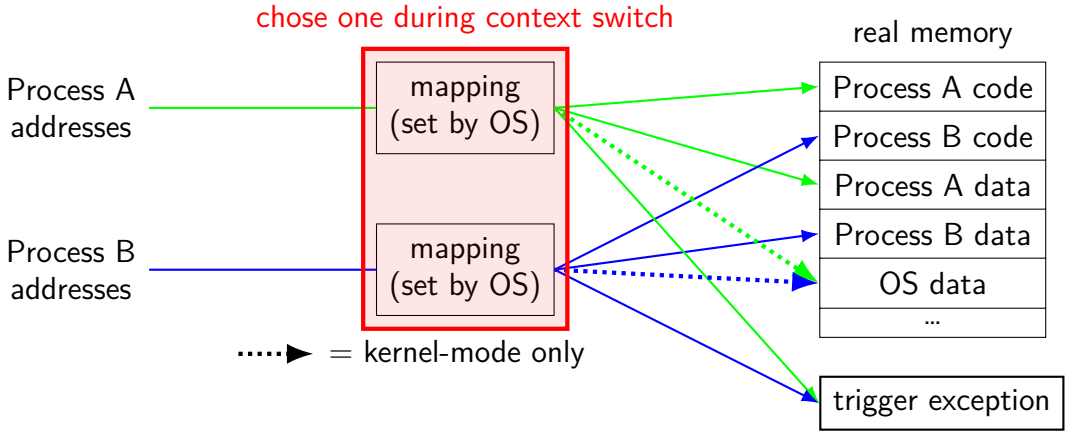

themes

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 on networks

under the hood of fast processors

- caching and (hidden) parallelism

keyboard input timeline

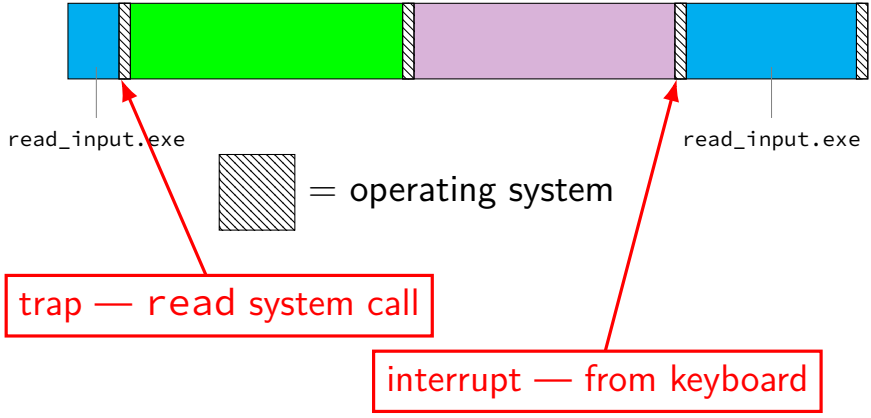

time multiplexing

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

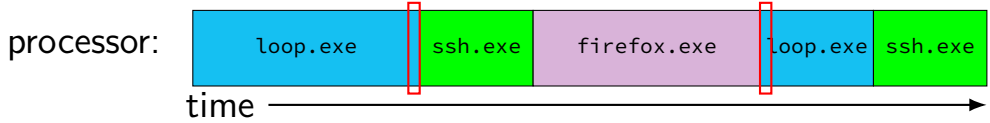


```
...
call get_time
    // whatever get_time does
movq %rax, %rbp
```

million cycle delay

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

multiple cores+threads

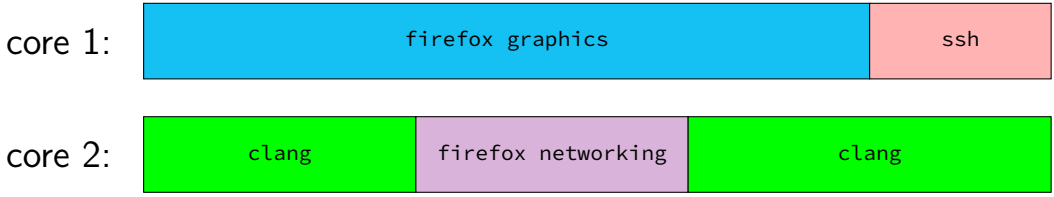

multiple cores? each core still divided up

multiple cores+threads

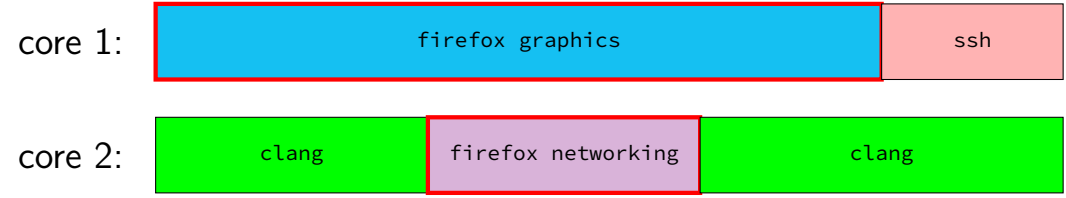

one program with multiple *threads*

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 on networks

under the hood of fast processors

- caching and (hidden) parallelism

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 on networks

under the hood of fast processors

- caching and (hidden) parallelism

layers

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

more than four layers?

sometimes more layers above 'application'

e.g. HTTPS:
HTTP (app layer) on TLS (another app layer) on TCP (network) on ...

e.g. DNS over HTTPS:
DNS (app layer) on HTTP on on TLS on TCP on ...

e.g. SFTP:
SFTP (app layer??) on SSH (another app layer) on TCP on ...

e.g. HTTP over OpenVPN:
HTTP on TCP on IP on OpenVPN on UDP on different IP on ...

names and addresses

name	address
logical identifier	location/how to locate
variable `counter`	memory address `0x7FFF9430`
DNS name `www.virginia.edu`	IPv4 address `128.143.22.36`
DNS name `mail.google.com`	IPv4 address `216.58.217.69`
DNS name `mail.google.com`	IPv6 address `2607:f8b0:4004:80b`
DNS name `reiss-t3620.cs.virginia.edu`	IPv4 address `128.143.67.91`
DNS name `reiss-t3620.cs.virginia.edu`	MAC address `18:66:da:2e:7f`
service name `https`	port number 443
service name `ssh`	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 on networks

under the hood of fast processors

- caching and (hidden) parallelism

2004 CPU

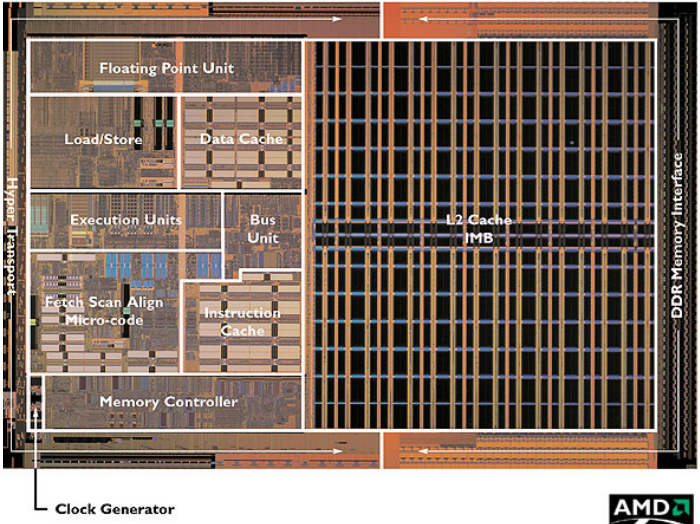

Image: approx 2004 AMD press image of Opteron die;
approx register location via chip-architect.org (Hans de Vries)

2004 CPU

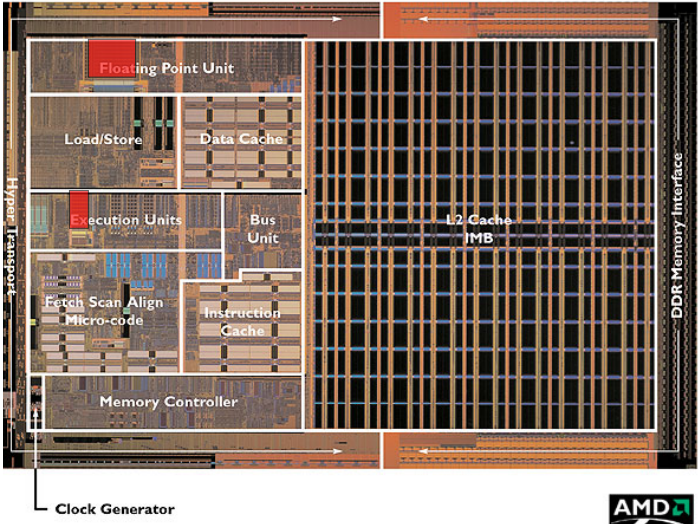

Registers

AMD 64 Opteron

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

2004 CPU

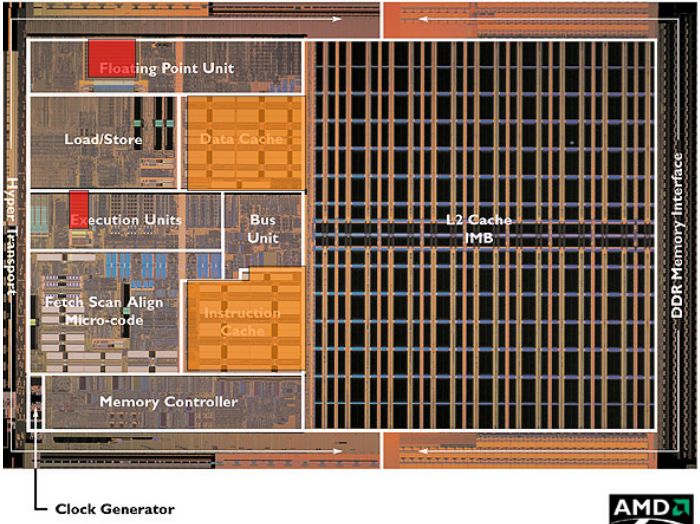

Image: approx 2004 AMD press image of Opteron die;
approx register location via chip-architect.org (Hans de Vries)

2004 CPU

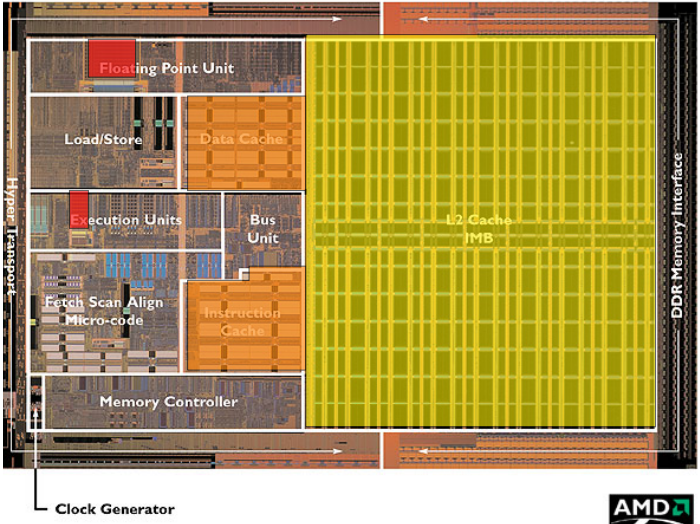

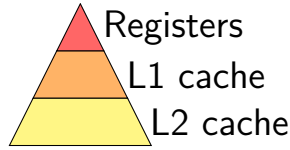

Image: approx 2004 AMD press image of Opteron die;
approx register location via chip-architect.org (Hans de Vries)

2004 CPU

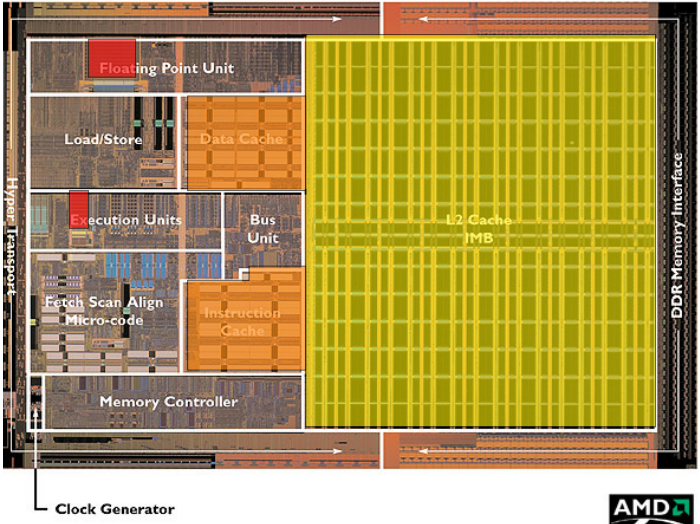

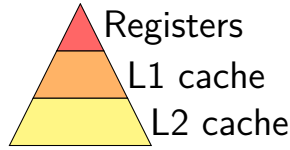

Image: approx 2004 AMD press image of Opteron die;
approx register location via chip-architect.org (Hans de Vries)

2004 CPU

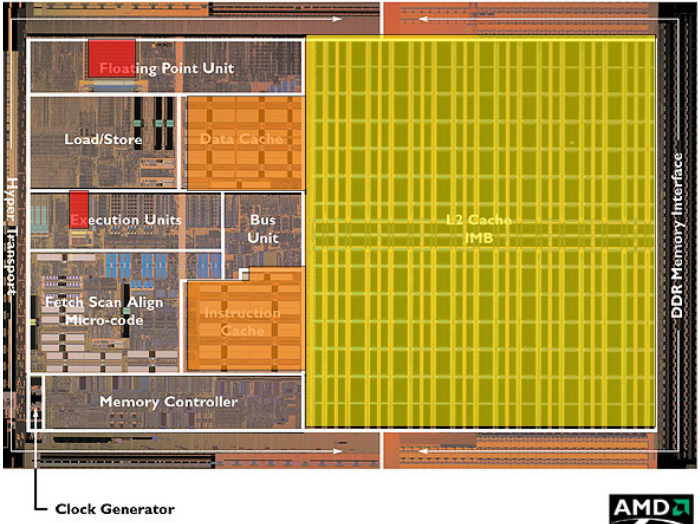

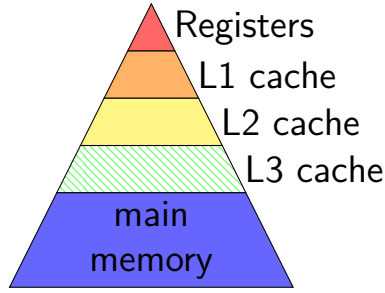

Image: approx 2004 AMD press image of Opteron die;
approx register location via chip-architect.org (Hans de Vries)

2004 CPU

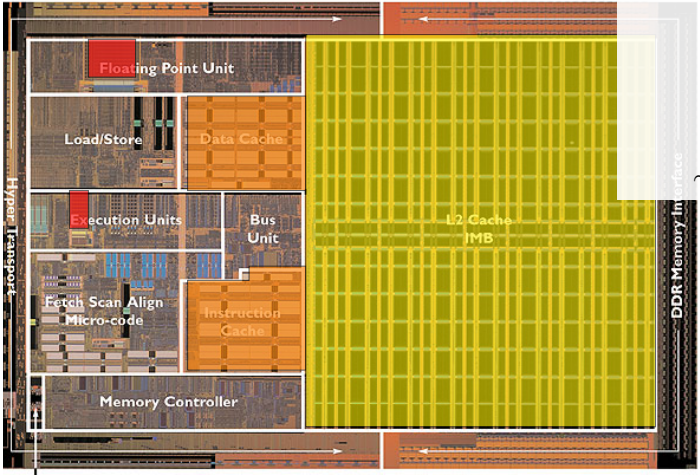

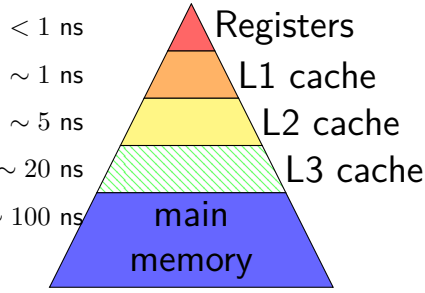

Image: approx 2004 AMD press image of Opteron die;
approx register location via chip-architect.org (Hans de Vries)

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 on networks

under the hood of fast processors

- caching and (hidden) parallelism

some performance examples

```
example1:
    movq $10000000000, %rax
loop1:
    addq %rbx, %rcx
    decq %rax
    jge loop1
    ret
```

about 30B instructions

my desktop: approx 2.65 sec

```
example2:
    movq $10000000000, %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 $10000000000, %rax
loop1:
    addq %rbx, %rcx
    decq %rax
    jge loop1
    ret
```

about 30B instructions

my desktop: approx 2.65 sec

```
example2:
    movq $10000000000, %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 being complete/correct (can't tell how much time you actually spent)
- in-person lab checkoff of incomplete lab at least 50% credit

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

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

asking about 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)

homeworks

several homework assignments

done individually

due before a week's first lab

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)

late policy

no late quizzes

one quiz dropped (unconditionally)

90% credit for 0–48 hours late homeworks

80% credit for 48–72 hours late homeworks

for labs that allow submission only, same policy as homeworks

- lab submission due time is 11:59pm

for other labs, policy on a lab-by-lab basis

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

files in building C programs [dynamic linking]

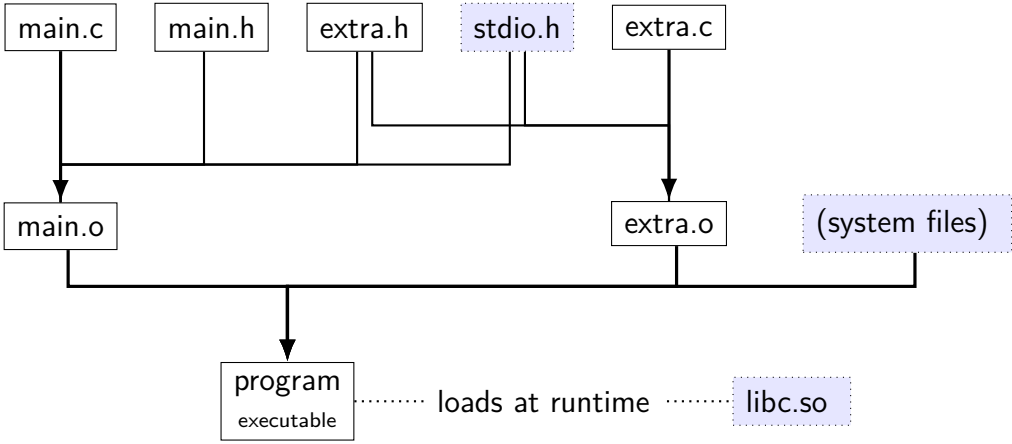

files in building C programs [dynamic linking]

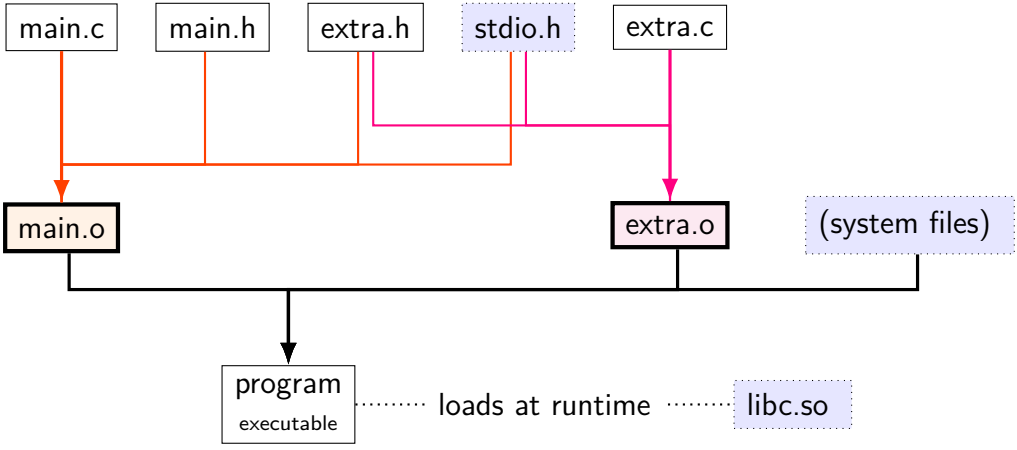


```
clang -c main.c
clang -c extra.c
```

files in building C programs [dynamic linking]

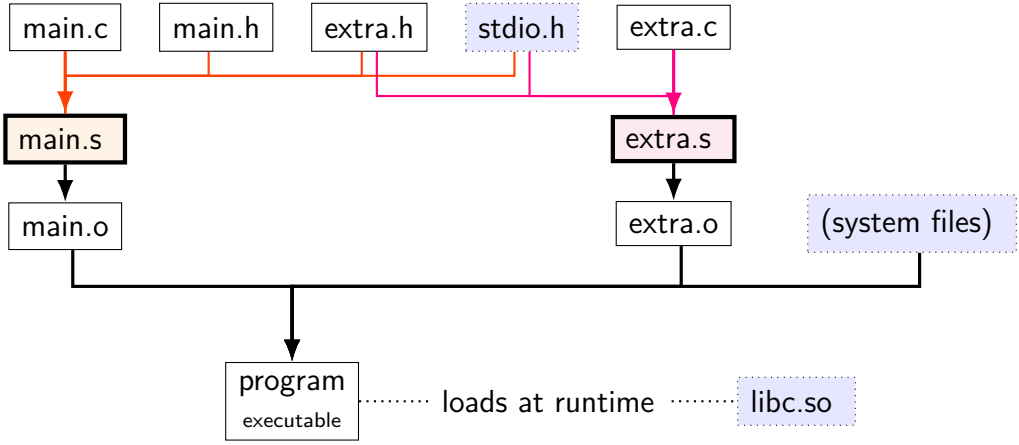


```
clang -S -c main.c
clang -S -c extra.c
```

files in building C programs [dynamic linking]

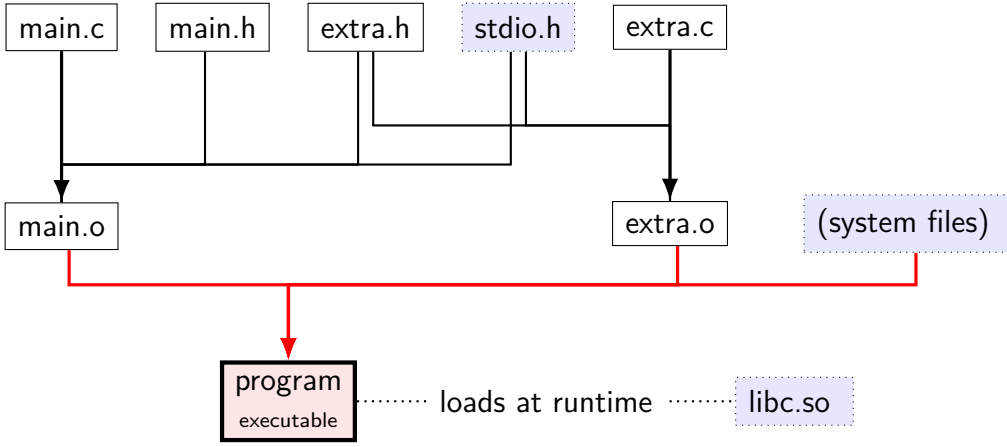


```
clang -o program main.o extra.o
```

files in building C programs [dynamic linking]

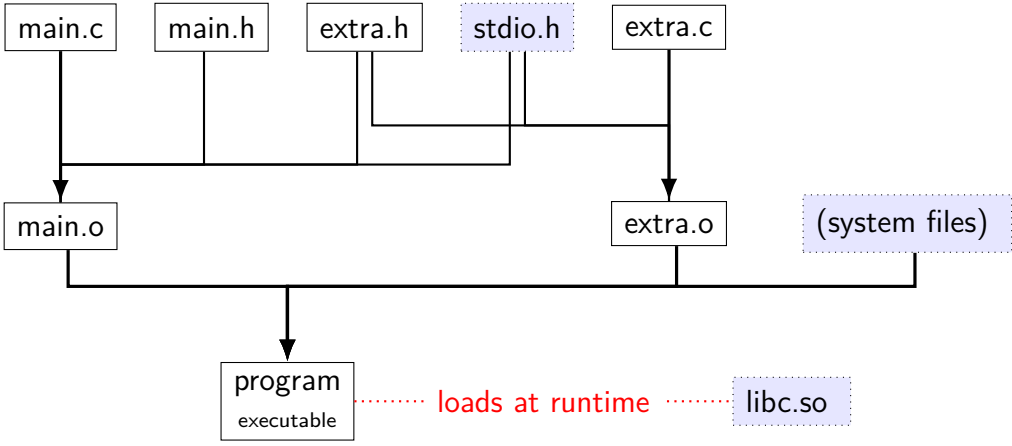


```
./program ...
```

files in building C programs [static linking]

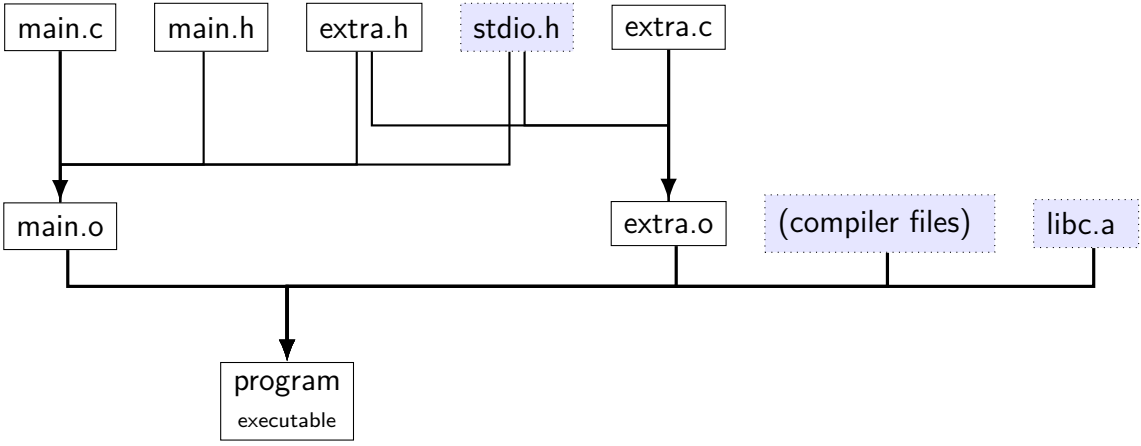

file extensions

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`)	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`

`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

`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