

Norges miljø- og biovitenskapelige universitet



INF205

Resource-efficient programming

- 4 Production
- 4.5 "#define" macros
- 4.6 Shared libraries
- 4.7 Benchmarks
- 4.8 CMake

INF205



Noregs miljø- og biovitskaplege universitet

4 **Production**

4.5 <u>"#define" macros</u>







Norwegian University of Life Sciences

"#define" macros

We have been using the preprocessor directive "**#define**" to protect headers:

#ifndef HEADER_NAME_H
#define HEADER_NAME_H

 \ldots // content of the header

#endif

This particular use is still recommended.

Other uses of **#define** are seen very frequently in legacy code and projects that have existed for a long time ... such as many libraries!

The C/C++ preprocessor accepts defines both as directives or, equivalently, by define flags passed through the compiler, e.g.:

g++ -D**NDEBUG** -D**PARAM=7** -c -o object_file.o code_file.cpp

The "NDEBUG" define also deactivates "assert" and a few other constructions.



Norwegian University of Life Sciences

Case distinctions using macros

C/C++ preprocessor directives are a programming language in its own right.

#ifdef SOME_CONSTANT
 // here comes some code
#else
 // alternative solution
#endif

Exercise

Use macros for a sequential version of the MPI-parallel primes counting code from the same **single code base**.

Overusing preprocessor directives makes code *hard to debug and maintain*!

Features that are turned on through -D... *directives can interact* in unforeseen ways. If there are *n options*, we would need to test 2^{*n*} versions of the code.

If it does not cost much in performance, make case distinctions at runtime.

Where possible, **better use a template** rather than a #define. In the future, **generic programming** will make template-like constructions more powerful.

INF205

Case distinctions using macros



Norwegian University of Life Sciences

```
#ifdef USE_MPI
#include <mpi.h>
#endif
```

. . .

. . .

```
int main(int argc, char** argv)
```

```
#ifdef USE_MPI
MPI_Init(&argc, &argv);
```

```
int size = 0;
MPI_Comm_size(MPI_COMM_WORLD, &size);
```

```
int rank = 0;
MPI_Comm_rank(MPI_COMM_WORLD, &rank);
```

#else

int size = 1;

```
int rank = 0;
```

#endif

```
int64_t limit = std::atoi(argv[1]);
...
int64_t counted_primes = 0;
```

```
for(int64_t n = 6*(rank+1) - 1; n < limit; n += 6*size)
if(is_prime(n)) counted_primes++;</pre>
```

```
for(int64_t n = 6*(rank+1) + 1; n < limit; n += 6*size)
if(is_prime(n)) counted_primes++;</pre>
```

```
int64_t overall_primes = 0;
```

```
#ifdef USE_MPI
    MPI_Reduce(&counted_primes, &overall_primes, ...);
    if(!rank) overall_primes += 2;
    MPI_Finalize();
#else
    overall_primes = counted_primes + 2;
#endif
```

Style advice on #define macros

#define NAMEis equivalent tog++ -DNAME#define NAME valueis equivalent tog++ -DNAME=value

Core guidelines about defines:

- SF.8: Use #include guards for all header files
- ES.30: Don't use macros for program text manipulation
- ES.31: Don't use macros for constants or "functions"

#define PI 3.14
#define SQUARE(a, b) (a * b) (a * b) ("Example, bad", from Core Guidelines)

- ES.32: Use ALL_CAPS for all macro names
- ES.33: If you must use macros, give them unique names

#define MYCHAR /* BAD, will eventually clash with someone else's MYCHAR*/ #define ZCORP_CHAR /* Still evil, but less likely to clash */



Noregs miljø- og biovitskaplege universitet

4 Production

4.5 "#define" macros4.6 Shared libraries





Shared libraries



Norwegian University of Life Sciences

In programming with C/C++ libraries, we have already seen:

- How to work with the C standard library, e.g., <cassert>, <cmath>, ...
- ... with the C++ standard library, with the STL, OpenMPI, and ROS

Technically, libraries are pre-compiled object code that can be reused.

The library needs to be accessed at three stages:

- At compile time, we need to include the library headers.
 - The complete source code for the libraries is unnecessary. It is even possible for the library to be coded in another language.
- During **linking**, the object code is **dynamically linked** against the library.
 - At this stage, the library "static object" (*.so file) is needed.
 - The executable does not contain the library's object code!
- At execution time, the executable and the library are loaded jointly.
 - If the library's **static object** code is gone now, the code will not run!



ImageMagick can deal with

many file formats; we need

graphics format such as BMP

copy colour value of pixels

an uncompressed pixel

Norwegian University

Library use example

The **image-benchmark** code uses the Magick++ API of ImageMagick.^{1, 2}

#include <Magick++.h>

int main(int argc, char** argv)

// charmap input read a "Charmap" object from a pixel std::ifstream pixistrm(argv[1]); diskgraphics::Charmap cm; graphics file, using last week's file format

// image object setup

pixistrm >> cm; pixistrm.close();

Magick::InitializeMagick(*argv); Magick::Image img(Magick::Geometry(cm.get_sizex(), cm.get_sizey()), "white"); img.magick("BMP"); img.monochrome(); img.type(Magick::BilevelType);

// pixel-by-pixel transfer of content

for(int x = 0; x < cm.get sizex(); x++) for(int y = 0; y < cm.get sizey(); y++) img.pixelColor(x, y, Magick::Color(cm.get pixel(x, y) == 0? "black": "white"));

// output in BMP format using one bit per pixel imq.quantize(2); img.write(argv[2]);

with quantize(2) we get one bit per pixel

¹Magick++ API documentation: https://imagemagick.org/Magick++/Documentation.html ²Magick++ Tutorial: https://www.imagemagick.org/Magick++/tutorial/Magick++_tutorial.pdf

INF205

Library use: Compiling and linking

The **image-benchmark** code uses the Magick++ API of ImageMagick.^{1, 2}

#include <Magick++.h> How does the compiler know where to look for this file? For this particular library, there is a tool that helps call g++ with the right flags: g++ -c -std=c++17 -o <name>.o <name>.cpp `Magick++-config --cppflags` -fopenmp -DMAGICKCORE HDRI ENABLE=1 -DMAGICKCORE_QUANTUM_DEPTH=16 -I/usr/local/include/ImageMagick-7 g++ -std=c++17 -o <name> *.o `Magick++-config --libs` -L/usr/local/lib -IMagick++-7.Q16HDRI -IMagickWand-7.Q16HDRI -IMagickCore-7.Q16HDRI More typically, you need to provide this information to the compiler by hand. ¹Magick++ API documentation: https://imagemagick.org/Magick++/Documentation.html ²Magick++ Tutorial: https://www.imagemagick.org/Magick++/tutorial/Magick++_tutorial.pdf

INF205



Norwegian University of Life Sciences

this is a capital i,

not a lower-case L

Library programming

A shared object can be created from an object file using g++ -shared:

g++ -c -o first.o first.cpp g++ -c -o second.o second.cpp g++ -shared -o libname.so first.o second.o

The library header location can be passed to g++ at compile time with -l..., and the shared object is found by the linker with the -L and -l options.

But the library also needs to be found at execution time. For that to work, it must be in the appropriate path, or one of the environment variables for library paths must be set to include the location of the shared object.

this time it is a lower-case L

this can be\$LD_LIBRARY_PATH

INF205



Noregs miljø- og biovitskaplege universitet

4 Production

- 4.5 "#define" macros4.6 Shared libraries
- 4.7 Benchmarks





Norwegian University of Life Sciences

Role of the benchmarks

What are our benchmarks like?

- The benchmark is not a single input file. It is an **input file generator**.
- The **problem size** is controlled through one or multiple parameters.

What to do with the benchmark:

- Record **average runtimes**, maybe worst-case runtimes, on a single core
- Use it to conduct weak scaling tests in parallel (weak scaling: the problem size increases with the number of cores)

What *not* to do with the benchmark:

- The program should not assume that it only receives benchmark input: It should work for the problem in general.
- The benchmark is there so you can document how your code performs and scales. If you tailor the code so specifically to the benchmark that it becomes worse (or incorrect) overall, it has failed its purpose.

INF205

Benchmark scenario: Disk graphics



See **image-benchmark** generator code. Main scenario parameter:

- Graphics edge size a, in pixels, so that the figure has a^2 pixels

Additional benchmark scenario parameter (change only if you have a reason):

- number of disks m from which the figures are generated; default: m = 10

```
float disk_radius = a;
unsigned char disk_colour = 0;
for(int i = 0; i < m; i++) {
    diskgraphics::Disk dsk;
    disk_radius /= 1.125 + 0.125*rand()/(float)RAND_MAX;
    dsk.x = a * (0.75*(rand()/(float)RAND_MAX) + 0.125);
    dsk.y = a * (0.75*(rand()/(float)RAND_MAX) + 0.125);
    dsk.r = disk_radius;
    dsk.colour = disk_colour;
    dv.add_disk(dsk);
```

./generator (from generator.cpp) creates a disk vector with m disks and the figure in charmap format.

./bitmap (from **bitmap.cpp**) uses Magick++ from ImageMagick to convert charmaps to BMP format.

```
if(std::rand() > RAND_MAX/3.0) disk_colour = 255 - disk_colour;
}
```

INF205

Terminal output of ./generator 32 && ./bitmap

random disks: 1 output to: 2 255 255 255 255 255 255 255 255 255 255 255 255	10 benchmark.px1 benchmark.uct 255 255 255 255 255 255 255 255	255 255 255 255	255 255 255 255	255 255 255 0	255 255 255 0	255 255 255 0	255 255 255 0	255 255 255 255	255 255 255 255	255 255 255 255	255 255 255 255	255 255 255 255	255 255 255 255	255 255 255 255	255 255 255	255 255 255 255	255 255 255 255	255 255 255 255	255 255 255 255	255 255 255				F			٦
255 255 255 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0	200 200 0	000000000000000000000000000000000000000	; ;h ;h ;h	ہ 1 جمع	25 103 28	。 8 r 0 r 2 r	i nov nov nov	ŏ + +	15 15 15	1: 1: 1:	3:2 3:2 3:2		ben ben	chr chr	255 1ar 1ar	266 k _ k_+ k_+	px1	265		255 255 255 255 255 255 0 0 0 0 0 0 0 0						-	
205 0 205 255 265 255 275 2	0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0	0 0 255 255 255 255 255 255 255 255 255	0 2 2 2 5 5 2 5 5 2 5 5 2 5 5 2 5 5 2 5 5 2 5 5 2 5 5 2 5 5	255 255 255 255	255 255 255 255 255 255	255 255 255 255 255 255	255 255 255 255 255 255	255 255 255 255 255	fiele 255 255 255	255 255 255 255	0 255 255 255	255 255 255	255 255 255 255 0 0 255	255 255 255 255 255 255 0 0 0 255	a m o 255 255 a m o 255	e 255 255 255 0 0 0 255 255	255 255 255 255 255	255 255 255 255 255 255 0 0 255	0 2555 2555 2555 2555 2555 0 0 2555 2555	0 255 255 255 255 255 255 0 0 255	200 255	200 255	200 255	200 255	200 255	200 255	25t
295 295 295 295 295 295 295 295 295 295 295 295 295 295 295 295 295 295 295 295 295 295 295 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 295 295 295 295 295 295 295 295 295 295 295 295 295 295 295 295 295 <td>2555555 255555 25555 2555 2555 2555 25</td> <td>25555 2555 2555 2555 222 20000000000000</td> <td>255 255 255 255 225 225 225 255 255 255</td> <td>255 255 2000000000000000000000000000000</td> <td>25555 2555 2000000000000000000000000000</td> <td>2555 222200000000000000000000000 2555</td> <td>2555 2555 2000000 200000000000000000000</td> <td>255 255 2255 2255 000 2555 2555 2555 25</td> <td>25555 2555 2555 2555 2555 2555 2555 25</td> <td>25555 2555 2555 2555 2555555</td> <td>25555555555555555555 25555555555555555</td> <td>25555555555555555555555555555555555555</td> <td>255 255 255 255 255 255 255 255 255 255</td> <td>255 255 255 255 255 255 255 255 255 255</td> <td>25555555555555555555555555555555555555</td> <td>255 2555 2555 2555 2555 2555 2555 2555</td> <td>25555555 2555555 2555555 2555555 2555555</td> <td>2555 22555 22555 22555 22555 22555 22555 00000000</td> <td>2555 2555 2555 2555 2555 2555 2555 255</td> <td>2555 2555 2555 2555 2555 2555 2555 255</td> <td>255 255 255 255 255 255 255 255 255 255</td> <td>255 255 255 255 255 255 255 255 255 255</td> <td>25555555555555555555555555555555555555</td> <td>25555555555555555555555555555555555555</td> <td>255 255 255 255 255 255 255 255 255 255</td> <td>25555555555555555555555555555555555555</td> <td>255522552255225522552255225522552255225522552255225522552255225522552252255222552225522255222552225522255222552225522255222552225522255222552225522255222552225222552222</td>	2555555 255555 25555 2555 2555 2555 25	25555 2555 2555 2555 222 20000000000000	255 255 255 255 225 225 225 255 255 255	255 255 2000000000000000000000000000000	25555 2555 2000000000000000000000000000	2555 222200000000000000000000000 2555	2555 2555 2000000 200000000000000000000	255 255 2255 2255 000 2555 2555 2555 25	25555 2555 2555 2555 2555 2555 2555 25	25555 2555 2555 2555 2555555	25555555555555555555 25555555555555555	25555555555555555555555555555555555555	255 255 255 255 255 255 255 255 255 255	255 255 255 255 255 255 255 255 255 255	25555555555555555555555555555555555555	255 2555 2555 2555 2555 2555 2555 2555	25555555 2555555 2555555 2555555 2555555	2555 22555 22555 22555 22555 22555 22555 00000000	2555 2555 2555 2555 2555 2555 2555 255	2555 2555 2555 2555 2555 2555 2555 255	255 255 255 255 255 255 255 255 255 255	255 255 255 255 255 255 255 255 255 255	25555555555555555555555555555555555555	25555555555555555555555555555555555555	255 255 255 255 255 255 255 255 255 255	25555555555555555555555555555555555555	255522552255225522552255225522552255225522552255225522552255225522552252255222552225522255222552225522255222552225522255222552225522255222552225522255222552225222552222



Norwegian University of Life Sciences

255

benchmark.bmp file

Benchmark scenario: Disk graphics

A series of figures is uploaded as **disk-benchmark-series.zip**, including these:



Norwegian University of Life Sciences

16th November 2022

17

Benchmark scenario: ANN

Use the ANN as a compressed representation of multiple benchmark figures.





Norwegian University of Life Sciences

Benchmark scenario: Graph queries

See graph-benchmark generator code. Two main scenario parameters:

- Graph size *n*, the number of nodes in the graph (not number of edges)
- Query size m, length of each of the two paths from the query



Random graphs are generated such that both cases with and without an instance of the graph pattern occur (the aim was to approximate a 50:50 ratio).

INF205



Norwegian University of Life Sciences

Benchmark scenario: Graph queries

Random benchmark graph from "./generate-graph 15 4 kb.dat query.dat":



The benchmark graphs consist of m+1 partitions. Assuming that n and m grow in proportion, the number of edges is in O(n), generating sparse graphs.

INF205

Benchmark scenario: Sphere overlap

Norwegian University of Life Sciences

See **sphere-benchmark** generator code. Main scenario parameter:

- N, the number of spherical particles in the system

Additional benchmark scenario parameters (change only if you have a reason):

- packing fraction ξ , *i.e.*, total sphere volume / box volume; default: $\xi = 7/9$
- ratio $\zeta_{\rm max}$ between largest and smallest sphere diameter; default: ζ = 10/3

Remark: If the spheres are all of the same size, the densest packing (without any overlaps) has the packing fraction 0.7405.

This had been known as one of the "Hilbert problems."

THEOREM 1.1 (The Kepler conjecture). No packing of congruent balls in Euclidean three space has density greater than that of the face-centered cubic packing.

This density is $\pi/\sqrt{18} \approx 0.74$.



T. C. Hales, "A proof of the Kepler conjecture," Ann. Math. 162(3): 1065-1185, doi:10.4007/annals. 2005.162.1065, 2005.

Figure 1.1: The face-centered cubic packing

The proof of this result is presented in this paper. Here, we describe the



Norwegian University of Life Sciences

Benchmark scenario: Sphere overlap

particle number <i>N</i> : packing fraction:	256 0.77778	example output from ./generator 256									
output file name:	benchmark-configuration.dat										
max. size ratio:	3.33333										
	generating 5 spher	res with diameter 3.33333 (occupied volume: 96.9627)									
	generating 2 spheres with diameter 2.33333 (occupied volume: 13.3033)										
	generating 122 spheres with diameter 1.33333 (occupied volume: 151.417										
	generating 127 spher	res with diameter 1 (occupied volume: 66.497)									
	occupied volume:	328.18									
	volume of the box:	421.946									
	size of the box:	7.50042									

Proposal that avoids boring solutions from "stacking" of spheres (placing ≥2 spheres at the same coordinates):



See implementation in **sphere-benchmark**, sphere.cpp, line 51.



Norwegian University of Life Sciences

Benchmark scenario: Sphere overlap





Noregs miljø- og biovitskaplege universitet

4 Production

- 4.5 "#define" macros
- 4.6 Shared libraries
- 4.7 Benchmarks
- <u>4.8</u> <u>CMake</u>



Generating makefiles using CMake



CMake is used by many complex C/C++ projects that require developers or users to compile code on their systems, which may be very diverse. Typically:

cmake . && make && sudo make install

There, **CMake** generates the Makefile that is then used by **GNU make**. We have done this before when we looked into the C++ interface to ROS.

Instructions for CMake are communicated through a file called **CMakeLists.txt**.

- CMake documentation: https://cmake.org/cmake/help/latest/
- CMake tutorial: https://cmake.org/cmake/help/latest/guide/tutorial/

CMake can be helpful if your project has a complex system of dependencies, or compile-time case distinctions are needed beyond what you can implement in a simple way using GNU make; *e.g.*, embedded system cross-compiling.

INF205

25

./CMakeLists.txt

src/CMakeLists.txt

CMakeLists.txt commands

Create a simple CMake project (cmake-dirgraph):

- Copy the "directed-graph" code into the subdirectory ./src of some working folder.
- Remove the original Makefile.
- Write CMakeLists.txt (see commands¹) in the main folder and in the ./src folder.
- Calling "cmake ." in the main working folder generates Makefiles in both folders.

```
)
project(
dirgraph
VERSION 1.0.0
LANGUAGES CXX
)
set(AUTHOR "Martin Horsch")
set(CMAKE_CXX_STANDARD 17)
add subdirectory(src)
```

cmake_minimum_required(

VERSION 3.14

set(EXECUTABLE_OUTPUT_PATH ../bin)
add_executable(dirgraph graph.cpp query.cpp run-graph.cpp)

Now **make** will automatically call g++ with the right options and flags.

¹See: https://cmake.org/cmake/help/latest/manual/cmake-commands.7.html

INF205

16th November 2022



Norwegian University of Life Sciences

CMake support for unit tests

Unit tests are generally a helpful debugging tool in complex development projects. Here they can also help the user verify that everything worked well.

enable_testing() Matches at beginning of input Λ add_test(Matches any single character NAME example_graph [Xy2] Any of the characters X, y, or 2 [^vV] Any character other than v or V COMMAND dirgraph kb.dat query.dat [C-F] Any of the characters C, D, E, or F WORKING DIRECTORY data Preceding pattern occurs ≥ 0 times Preceding pattern occurs >= 1 time +set_tests_properties(Optional (occurs 0 or 1 times) ? Disjunction ("or") example_graph **PROPERTIES** PASS_REGULAR_EXPRESSION "<INF200 2022H>[\t\r\n]*<Rune Grønnevik>" PASS_REGULAR_EXPRESSION "<INF205 2023H>[\t\r\n]*<Trine Næss Henriksen>" PASS_REGULAR_EXPRESSION "<KJM230 2023V>[\t\r\n]*<Heidi Rudi>"

https://cmake.org/cmake/help/latest/manual/cmake-properties.7.html#test-properties

CMakeLists.txt (cmake-dirgraph)



Norwegian University of Life Sciences

CMake GUI

Graphical interface to CMake: cmake-gui

	CMake 3.16.3 - /arc/tr/lehre/2022/inf205/4-production/examples/cmake-dirgraph/bin									
	File Tools Options	<u>l</u> elp								
	Where is the source code:	/arc/tr/lehre/2022/inf205/4-production/examples/cmake-dirgraph	Browse <u>S</u> ource							
	Where to build the binaries:	/arc/tr/lehre/2022/inf205/4-production/examples/cmake-dirgraph/bin	Browse <u>B</u> uild							
File Tools Options Help	Search:	Create Directory	X Remove Entry							
Where is the source code: /arc/tr/lehre/20	Name	Build directory does not exist, should I create it?								
Where to build the binaries: //tmp/dirgraph Search:		Directory: /arc/tr/lehre/2022/inf205/4-production/examples/cmake- dirgraph/bin								
Name CMAKE_CXX_FLAGS_DEBUG CMAKE_CXX_FLAGS_MINSIZEREL CMAKE_CXX_FLAGS_RELEASE	<u>C</u> onfigure <u>C</u>	No Yes rate to generate selected build files.								
CMAKE_CXX_FLAGS_RELWITHDEBIN CMAKE_DLLTOOL CMAKE_EXE_LINKER_FLAGS										
CMAKE_EXE_LINKER_FLAGS_DEBUG CMAKE_EXE_LINKER_FLAGS_MINSIZE CMAKE_EXE_LINKER_FLAGS_RELEASS CMAKE_EXE_LINKER_FLAGS_PEINIT										
CMAKE_EXE_ENRET_FAGS_RELWIT CMAKE_EXPORT_COMPILE_COMMAN CMAKE_INSTALL_PREFIX CMAKE_LINKER	IDS	/usr/local /usr/bin/ld								
CMAKE_MAKE_PROGRAM		/usr/bin/make								
Press Configu	re to update and display new value	es in red, then press Generate to generate selected build files.								
<u>Configure</u> <u>Generate</u> Open <u>Projec</u>	t Current Generator: Unix Makef	les								
Detecting CXX compiler ABI info - Detecting CXX compile features Detecting CXX compile features - of Configuring done Generating done	done									





Noregs miljø- og biovitskaplege universitet



Institutt for datavitskap

Digitalisering på Ås

Conclusion





Update: Presentation schedule

Thursday, 1st December 2022

Black-white images and binary classification problems

10.00 - 10.20: Group #9 10.30 - 10.50: Group #23 11.00 - 11.20: Group #26

Graph querying: Paths and cycles in graphs

12.00 - 12.20: Group #8 12.30 - 12.50: Group #10 13.00 - 13.20: Group #17 13.30 - 13.50: Group #20

Wednesday, 7th December 2022

Spherical particle configurations

14.15 - 14.35: Group #1 14.40 - 15.00: Group #6 15.15 - 15.35: Group #12 15.40 - 16.00: Group #19 Friday, 9th December 2022

Various topics

12.30 - 12.50: Group #2 13.00 - 13.20: Group #3

Thursday, 15th December 2022

Various topics

10.30 - 10.50: Group #4 11.00 - 11.20: Group #5 11.30 - 11.50: Group #7

Black-white images

12.30 - 12.50: Group #13 13.00 - 13.20: Group #16 13.30 - 13.50: Group #25



Norwegian University of Life Sciences



Norges miljø- og biovitenskapelige universitet



INF205

Resource-efficient programming

- 4 Production
- 4.5 "#define" macros
- 4.6 Shared libraries
- 4.7 Benchmarks
- 4.8 CMake

INF205