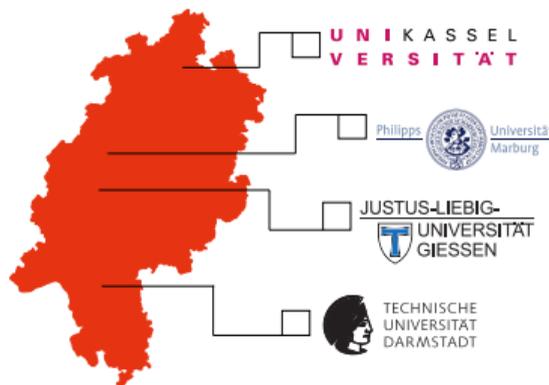


Building Software in Linux

Hessisches Kompetenzzentrum für Hochleistungsrechnen (HKHLR)

Tim Jammer / Dr. Marcel Giar

26.04.2023



HKHLR is funded by the Hessian Ministry of Sciences and Arts



- ▶ Not all software is pre-Installed on an HPC system
- ▶ How can you install software yourself?
- ▶ As Linux software is often open source, compiling it from source code is quite common
- ▶ Errors can happen at different stages of compilation
 - ⇒ Basic understanding of the compilation process will help to resolve potential errors quickly

- 1 Software build steps
 - Preliminaries
 - Pre-processing
 - Compilation and Linking
- 2 Frequently used compiler options
- 3 Makefile introduction
- 4 Variables in Makefiles
- 5 Makefile rules
- 6 Miscellaneous
- 7 Build systems: GNU configure



Definition (Compilation)

The transformation of code written in one programming language (source) into another programming language (target).

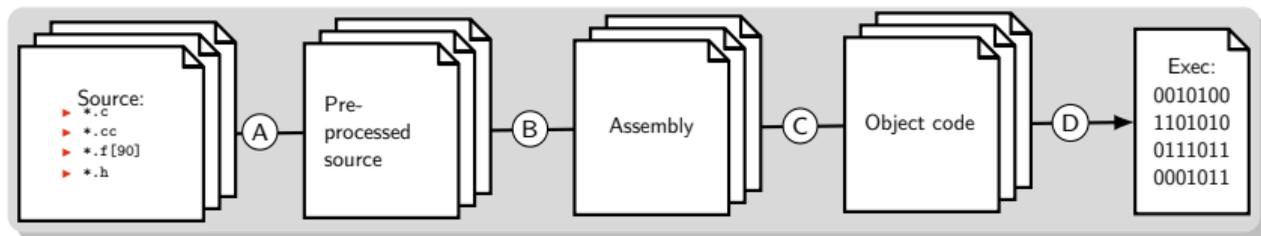
- ▶ To generate a binary, all source code must be available, i.e.:
 - ▶ All functions (methods, member-functions, etc.) declared
 - ▶ All functions defined
- ▶ Compilation: A four-step process

(A) Pre-processing

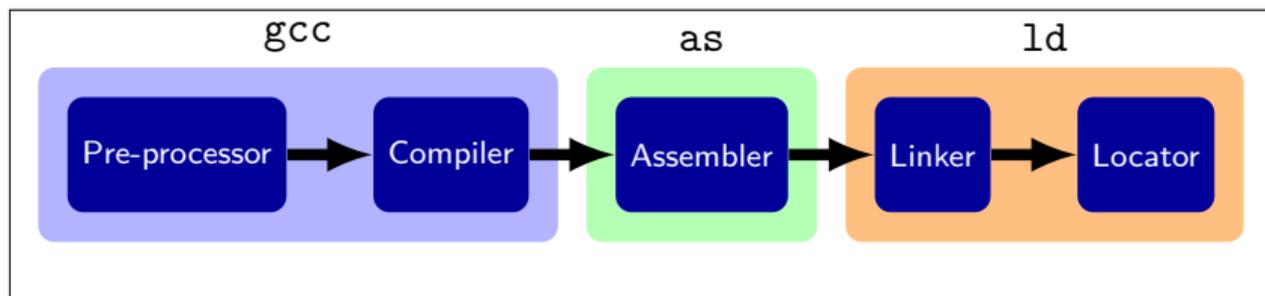
(B) Compilation "proper"

(C) Assembling

(D) Linking



Complete tool-chain can be accessed with gcc frontend



- ▶ GNU project ships complete set of tools to manage the build process
 - ▶ GNU Compiler Collection (GCC)
 - ★ Pre-processor: `cpp`
 - ★ Compiler: `gcc`
 - ▶ GNU Binutils (<https://www.gnu.org/software/binutils/>)
 - ★ Assembler: `as`
 - ★ Linker: `ld`

Invoking a compiler from the command line

Like many other UNIX utilities, compilers must be passed options and files as parameters:

```
1 <CompilerName> <Options> <InputFile> [-o <OutputFile>]
```

An example for generating an object file from a source file:

```
1 gcc -Wall -O0 -g -c SourceFileName.c -o SourceFileName.o
```

Command	"Value"
CompilerName	gcc
Options	-Wall -O0 -g -c
InputFile	SourceFileName.c
-o OutputFileName	-o SourceFileName.o



- ▶ Pre-processing: "Simple" text-manipulation
- ▶ Invoke the pre-processor from the command line:
 - ▶ Call `cpp`: `cpp <filename>.c -o <filename>.i`
 - ▶ Compiler frontend: `gcc -E <filename>.c -o <filename>.i`
- ▶ Pre-processor handles source code lines starting with "#"

```

1 #include <stdio.h>
2 #include <stdlib.h>
3 #include "myfft.h"
4
5 #ifdef __USE_FFTW
6 #define FFT_SIZE 1024
7 #endif
8
9 int main() {
10     // code
11 }

```

▶ Includes:

- ▶ Compiler argument: `-I@PATH@`
- ▶ `#include <...>` and `#include "..."`

★ These behave differently in where to look for include files (see next slide)

▶ Defines:

- ▶ Compiler/Pre-processor argument `-D@DEFINE@=@VALUE@`
- ▶ `#ifdef ... #endif`

`cd Demos/01_includes-and-defines`

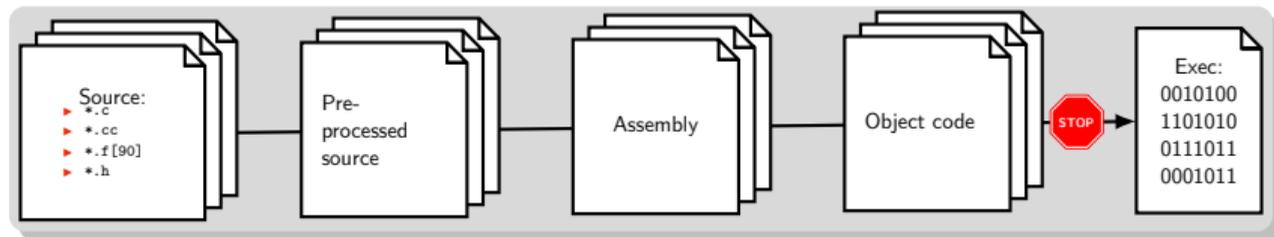
> make DEMO



- ▶ Use pre-processor as `gcc -E` OR `cpp`
 - ▶ Includes header files into source files (system header + user's headers)
 - ▶ Evaluated macro definitions (using argument `-D@DEFINE@=@VALUE@`)
- ▶ Include styles: `#include "..."` VS. `#include <...>`
 - ▶ `#include <...>`
 - ★ Used for system header files
 - ★ Self-defined paths are searched first; then system paths
 - ▶ `#include "..."`
 - ★ Used for *user's own* header files
 - ★ *Current* directory (where the source file is located) is searched first; then system paths
- ▶ Augment search paths for includes with `-I@PATH@` option
- ▶ Monitor include paths with `gcc -E -v` OR `cpp -v`



Compile to object



- ▶ Object files:
 - ▶ "Portable" partial compiled output of a compiler
 - ▶ Code is transformed but function/method names remain unresolved
 - ▶ Enabled by the `-c` option
- ▶ Object files are used to:
 - ▶ Speed up compilation process
 - ▶ Enable parallelization of compilation
 - ▶ Enable cross- language compilation
 - ▶ **Attention:** Compilers are only **mostly** binary compatible

- ▶ Object files must be combined to build a binary
- ▶ Linker processes all symbols and resolves addresses
 - ▶ Object files (*.o), generated by compiling with `-c`
 - ▶ Archive files (*.a), generated by `ar`
 - ▶ Dynamic libraries:
 - ★ Search paths are added via `-L@PATH@`
 - ★ Environment variable `LD_LIBRARY_PATH` contains runtime library paths
 - ★ When given `-l@LIBNAME@`, the linker will look for the file "`lib@LIBNAME@.so`"
 - ★ A library must be present both in the compiler's/linker's search path (`-L`) for compilation, and in the dynamic linker path (`LD_LIBRARY_PATH`) for running
- ▶ Compiler runtime environment is **not** portable

```
cd Demos/04_missing-symbols
```

```
> make DEMO
```

```
undefined reference to 'vectorf_from_array'
```

- ▶ Linker was not able to resolve reference to `vectorf_from_array`
 - ▶ Object file with function definition missing
 - ▶ Dynamic or static library with function definition missing
- ▶ Possible solution: Provide library to linker \Rightarrow `-L/path/to/lib -l@LIBNAME@`

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

- ▶ *Dynamic linker* failed to load shared libraries for binary

```
1 > LD_LIBRARY_PATH=/path/to/lib:${LD_LIBRARY_PATH} ./program
2 > # or
3 > export LD_LIBRARY_PATH=/path/to/lib:${LD_LIBRARY_PATH}
4 > ldd ./program # check dynamically linked dependencies
```



- ▶ Optimizations `-O[0-3]`, `-Ofast`, `-Og`
 - ▶ `-O0`: Mostly disable optimizations, reduce compilation time and make debugging produce the expected results
 - ▶ `-O2`: Good compromise of optimizations and compile time, should be numerically identical to `-O0`
 - ▶ `-O3`: Aggressive optimizations, **slight changes in numerics possible**
 - ▶ `-Ofast`: Most aggressive optimizations, **no guarantee for standard IEEE floating point semantics compliance**
 - ▶ `-Og`: Minimal optimizations, designed for minimal interference with debugger
- ▶ Debugging: `-g` and either `-Og` or `-O0`

Optimization during development

Always use `-O0 -g` during development when checking correctness of your code (during the "edit-compile-debug" cycle).

Rule of thumb

Code should compile equally well with `-O0` and `-O3`.



- ▶ Demand ISO C/C++ standards: `-pedantic` (and `-pedantic-errors`); see <https://gcc.gnu.org/onlinedocs/gcc/Warning-Options.html>
 - ▶ Issue warnings demanded by strict ISO C/C++
 - ▶ ISO C: Follows value of `-std=<value>` flag
 - ▶ Rejects GNU extensions
- ▶ Warnings: `-Wall` and `-Wextra`, e.g.:
 - ▶ `-Wunused-function`
 - ▶ `-Wunused-variable`
- ▶ Turn all warnings into errors: `-Werror` (**aborts compilation!**)
- ▶ Architecture-specific optimizations (e.g. AVX, AVX2 or AVX512): See <https://gcc.gnu.org/onlinedocs/gcc/x86-Options.html>
 - ▶ `-march=native -mtune=native` OR `-mavx -mavx2 -mavx512f`
- ▶ Pointer aliasing: `-fstrict-aliasing` and `-Wstrict-aliasing`

```
cd Demos/05_compiler-flags
```

```
> make DEMO
```



- ▶ Enabling optimization (e.g. `-O3` vs. `-O0`) can yield **much** faster code
 - ▶ Obtainable speedup depends on problem!
- ▶ Compiler applies source code transformations while optimizing
 - ▶ Assembler code differs between `-O0` and `-O3`
- ▶ Using `-g`: Additional symbols in assembler code
 - ▶ Needed by debugger
 - ▶ Should always be used during development, irrespective of optimization level!
- ▶ Architecture optimizations can yield performance gain
 - ▶ Oftentimes library routines (e.g. from Eigen) are written correspondingly
 - ▶ **Attention:** In case the architecture on the compilation host differs from the architecture of the execution host (e.g. if the cluster headnode has a different CPU type than the compute nodes), `-march=native -mtune=native` might be less optimal or even incompatible

- ▶ Common choice for build system in the Linux command line environment: GNU `make`
- ▶ *Not* part of the GCC toolchain \Rightarrow Can be used with any compiler!
- ▶ The `make` utility helps us compile and link a software in order to use it on Linux (e.g. on a HPC cluster)
- ▶ `make` will take care of all steps of compilation:
 - (A) Pre-Processing
 - (B) Compilation proper
 - (C) Assembling
 - (D) Linking
- ▶ `make` offers an extra level of abstraction on top of the frontends provided by compiler toolchains (e.g. GCC)
- ▶ We can configure the build process by using *Makefiles*



```
1 # SPDX-FileCopyrightText: 2021 Competence Center for High Performance Computing in He
2 # SPDX-License-Identifier: MIT
3
4 all: program.exe
5
6 program.exe: array.c helper.c linalg.c main.c
7     gcc -c -O0 -Wall -Wextra -pedantic array.c -o array.o
8     gcc -c -O0 -Wall -Wextra -pedantic helper.c -o helper.o
9     gcc -c -O0 -Wall -Wextra -pedantic linalg.c -o linalg.o
10    gcc -c -O0 -Wall -Wextra -pedantic main.c -o main.o
11    gcc -o program.exe array.o helper.o linalg.o main.o
```

- ▶ Line 3: Executable depends on the source files (a rule)
- ▶ Lines 4-7: Generate object files from C sources (a recipe)
- ▶ Line 8: Link the object files into an executable (a recipe)

Too much repetitive content in this file!

make configuration

To configure the `make` build system, we must add things to our Makefile. These ingredients comprise a set of instructions specifying how to generate an executable from our source files.

- ▶ Variables:
We must specify ...
 - ▶ Source files
 - ▶ Compiler and compiler flags
 - ▶ Libraries to link and headers to include
- ▶ Rules:
We must specify *what* to build
- ▶ Recipes:
We must specify *how* to build it



File: Demos/00_toy-project/makefile

```

1 .SUFFIXES:
2 .SUFFIXES: .c .h
3
4 all: cprogram
5 sync_files:
6     bash ./setup_demo.sh --prepare
7
8 CSources = array.c helper.c linalg.c main.c
9 CObjects = $(CSources:%.c=%.o)
10 ExeFile = program.exe
11
12 CC      = gcc
13 CFLAGS = -Wall -Wextra -pedantic -O0 -g
14
15 vpath %.c ./
16 vpath %.h ./
17
18 cprogram: $(ExeFile)
19
20 $(ExeFile): $(CObjects)
21     $(CC) $(CFLAGS) -o $@ $^
22
23 %.o : %.c
24     $(CC) $(CFLAGS) -c $< -o $@
25
26 .PHONY: clean
27 clean:
28     @rm -fv *.o $(ExeFile)
29     @bash ./setup_demo.sh --clean

```

- ▶ A simple text file ⇒ Can track changes in a version control system (e.g. git)
- ▶ More compact but harder to understand
- ▶ Variables: Set compiler and flags
- ▶ Rules and recipes: Customize the build process
- ▶ Contains all commands to compile applications

Important notice

The `make` utility processes Makefiles in a nonlinear fashion!

Variables in Makefiles

Variables are defined in a Makefile by assigning *values* to them.

- ▶ There are some obvious advantages:
 - ▶ Enhance legibility of Makefiles → More self-explanatory
 - ▶ Allows quick adaption of compile options
 - ▶ Increase portability of whole build system
- ▶ Without using variables, commands become cluttered, e.g.:

```
1 program.exe: array.o helper.o array_ops.o main.o
2 gcc -O3 -mtune=native -march=native -Wall -Wextra -g -o program.exe \
3 array.o helper.o array_ops.o main.o -lm
```

- ▶ Using variables can make the same operation look compact and concise:

```
1 $(EXE): $(OBJS)
2 $(CC) $(CFLAGS_OPT) $(CFLAGS) -o $@ $^ $(LIBS)
```

We can assign values to variables ourselves

```
variable_name = string
```

▶ Setting compilers:

```
1 # Fortran compiler
2 FC = gfortran
3 # C compiler
4 CC = gcc
```

▶ Defining object files:

```
1 # Object files
2 OBJS = utils.o mathfunc.o main.o
```

▶ Setting names of executables (choose suffix as you like):

```
1 # Name of compiled executable
2 EXE = program.x
```



Accessing the value of a variable

Use `$(variable_name)` to access the content of variables

For example (details of this syntax will be explained later):

```
1 $(EXE) : $(OBJS)
2   $(CC) -o $(EXE) $(OBJS)
```

It can be used to define additional variables:

```
1 CFLAGS_BASE = -Wall -std=c11 -O2
2 CFLAGS_DEBUG = -g
3 CFLAGS = $(CFLAGS_BASE) $(CFLAGS_DEBUG)
```

Case sensitivity

`$(variable_name)` is **not** the same as `$(VARIABLE_NAME)`!

Explicit rules[5]

A "rule" within a Makefile connects *targets* and *prerequisites* by specifying a *recipe* of how to (re-)create the target files from the prerequisite files (e.g. compilation, linking, building an archive).

General rule syntax:

```
1 targets: prerequisite1 prerequisite2 prerequisite3 ...
2   command1 # this is a <tab> indentation!
3   command2
4   command3
5   ...
```

- ▶ Targets: List of file names separated by spaces (usually only one per rule)
- ▶ Prerequisites: List of file names separated by spaces (oftentimes more than one per rule)
- ▶ Recipe: List of commands indented with a <tab> character, prescription for how to make the target (interpreted by the shell)



Recipes

A recipe of a corresponding rule consists of `shell` commands that are executed one after the other.

In Makefiles, we really have two types of syntax combined:

- ▶ `make` syntax used for rules and defining variables
- ▶ `shell` syntax for recipes (default: interpreted by `/bin/sh`)

How to distinguish recipes from the rest of a Makefile

Recipes must *always* start with a `<tab>` character. Anything indented with a `<tab>` in a Makefile will be considered part of a recipe and therefore is passed to the `shell`.



Pattern rules

These rules contain a % in the target; the target is a file name matching the pattern.

General syntax: Object files from sources (SUFFIX \rightarrow .o):

```
1 %.o: %$(SUFFIX)
2   command1
3   command2
4   ...
```

- ▶ % in a prerequisite *must* match the same substring as in the target's pattern (e.g. `main.o : main.c`)

Availability of prerequisites

In order to be applicable, the prerequisites must either exist or there must be rules for making them.



► Compile sources and link to executable:

```
1 sources := $(wildcard $(SRCPATH)/*.c)
2 objects := $(subst $(SRCPATH)/,, $(sources:%.c=%.o))
3
4 $(executable): $(objects) # make executable from objects
5     $(CC) $(CFLAGS) -o $(executable) $(objects)
6
7 %.o : %.c # make objects from C sources (pattern rule)
8     $(CC) $(CFLAGS) $(INCS) -c $< -o $@
```

Assuming we have defined the following rules:

```
1 program: array.o linalg.o main.o
2   $(CC) $(CFLAGS) -o $@ $^
3 %.o: %.c
4   $(CC) $(CFLAGS) -c $< -o $@
```

How can we tell `make` to build the target `program`?

```
1 > make program # Pass target to make
```

- ▶ If all `*.o` files exist, the rule `$(CC) $(CFLAGS) -o $@ $^` will be executed
- ▶ If `*.o` files are missing: Build object targets for all `*.c` files first, e.g.:

```
1 array.o: array.c
2   $(CC) $(CFLAGS) -c array.c -o array.o
```



In the previous examples, we have made use of so-called *automatic variables*[1].

Automatic variables

Variables whose values are computed anew whenever the rule involving them is executed; they are based on the target or the prerequisites used in the rule. **They are only valid within the scope of the corresponding recipe.**

Automatic variable	Meaning
\$@	File name of the target in a rule.
\$<	The name of the first prerequisite.
\$*	The substring with which the pattern rule matches.
\$^	The name of <i>all</i> prerequisites, separated by spaces.

The VPATH variable[2]

VPATH specifies *all* directories in which `make` shall search for the files needed in the prerequisites used in the rules defined in the Makefile. The order in which directories are listed is the order used by `make` in its search.

Generic syntax for defining VPATH:

```
1 VPATH = directory_1:directory_2:...
```

Example: Source tree that has different subdirectories

```
1 VPATH = src/tools:src/parser:src/math
2
3 %.o : %.c # with VPATH: Equivalent to %.o : src/tools/%.c:...
4 $(CC) $(CFLAGS) -c $< -o $@
```

⇒ `make` knows that it must search `*.c`-files in the directories specified in the `VPATH` variable.



The `vpath` variable[2]

A more fine-grained configuration of search paths for prerequisites can be obtained with the `vpath` variable (lower case!) by specifying the patterns for file names.

Generic syntax for defining `vpath`:

```
1 vpath <pattern> <directories> # dir_1:dir_2:...
```

Example: Tell `make` to look for `*.c` and `*.h` files in these directories:

```
1 vpath %.c src/c
2 vpath %.h src/headers
```

⇒ Can be useful if different source types are placed in different directories in the source tree.



Goal when using .PHONY targets

Force certain rules to be executed *every time* they are explicitly invoked.

Example[4]: `make clean` (shall *always* execute recipe for `clean`!)

```
1 clean:
2   rm -vf *.o
3   rm -vf *.exe
```

⇒ `make clean` will not work if a file named "`clean`" exists

⇒ `clean` would always be considered up to date since it has no prerequisites that change over time

```
1 .PHONY: clean # make it a prerequisite of built-in .PHONY target
2 clean:
3   rm -vf *.o
4   rm -vf *.exe
```



- ▶ Program name variables should also have a flags variable:

```
1 CC = ...           # C compiler
2 CXX = ...          # C++ compiler
3 FC = ...           # Fortran compiler
4 CFLAGS = ...       # Flags for C compiler
5 FFLAGS = ...       # Flags for Fortran compiler
6 CXXFLAGS = ...     # Flags for C++ compiler
```

- ▶ Use extra variables for compiler options that are mandatory for properly building the application: `CFLAGS_BASE = ...`
- ▶ Place `CFLAGS` variable as last compilation command (usually `CFLAGS` is user-defined; override other variables):

```
1 %.o : %.c
2 $(CC) $(CFLAGS_BASE) $(CFLAGS) -c $< -o $@
```



`CFLAGS` should occur in every invocation of `$(CC)` (compilation *and* linking!)

- ▶ `make` by default searches for files named (in order!):

(1) GNUmakefile

(2) makefile

(3) Makefile

Some useful command line arguments are:

- ▶ Manually specify the Makefile to use:

```
> make -f|--file <filename>
```

⇒ For the default names, no command line argument is needed!

- ▶ Make a "dry run": Do not actually run the recipes, just print the commands that would be executed:

```
> make -n|--dry-run
```

- ▶ Change the directory to `DIRECTORY` before doing anything:

```
> make -C|--directory DIRECTORY
```



Use `make` to build certain targets

By passing the name of specific *targets* to `make` we can execute the *recipes* needed to build them: `make <target>`

Example:

```
1 all: kkrlin kkrpara # first target is compiling everything
2 kkrlin:
3   <recipe_for_kkrlin>
4 kkrpara:
5   <recipe_for_kkrpara>
```

- ▶ `make kkrlin`: Executes the recipe for rule `kkrlin`
- ▶ `make kkrpara`: Executes the recipe for rule `kkrpara`



Commonly used targets in Makefiles:

- ▶ `all`: Compile the entire program. Should be the first target in the Makefile. As a default, `make` should compile with the `-g` option to be able to debug the program if it crashes.
- ▶ `clean`: Delete all files that were created by executing `make` (with some additional options).
- ▶ `distclean`: Delete all files created by configuring or building the program. This option becomes particularly interesting if the program to be built comes with a `configure` file.
- ▶ `install`: Install the program to some destination directory after building.
- ▶ `check`: Run self-tests (if there are any). The program must be built first before executing the recipe for this target.



Overriding defined variables

Variables defined in Makefiles can also be passed or changed from the command line when calling `make`.
These values take precedence over the variables defined inside the Makefile.

The typical syntax is as following:

```
1 > make var_1=value_1 var_2=value_2 ...
```

Example:

```
1 CC = gcc
2 CFLAGS = -O3 -march=native -mtune=native
```

Override values from the command line when calling `make`:

```
1 > make CC=clang CFLAGS=-O0
```

The following ranking shows the priorities of variable definitions in the scope of `make` (first entry overrides later entries):

- 1 Variables passed via the command line
- 2 Self-defined variables inside the Makefile
- 3 Environment variables (using `export <var>=<value>`)
- 4 Pre-defined variables (`make` internal defaults)

How to interchange ranks 2 and 3:

```
1 > make -e|--environment-overrides
```

Important

- ▶ Pre-defined variables can (and should!) *always* be overridden

- ▶ Software packages oftentimes need to be buildable on different platforms
 - ⇒ Makefile needs to be adjusted
 - ⇒ Cumbersome and inefficient if done manually!
- ▶ Solution: Use a script to adjust Makefile to current needs
- ▶ GNU project: `configure` script
 - ▶ Standardized process of customizing the build options
 - ▶ Simplest scenario: `./configure && make && make install`
- ▶ Autotools:
 - ▶ Create the whole build system for current software package
 - ▶ Autoconf: `configure` script
 - ▶ Automake: Makefile
- ▶ For detailed introduction see: https://www.gnu.org/software/automake/manual/html_node/Autotools-Introduction.html



The `configure` script

A GNU `configure` script aims to ensure a working compilation on a specific target system with a specific software environment, without the user having to manually edit the Makefile.

- ▶ Most `configure` scripts offer lots of options for customization of the build process (e.g. which location to install into, etc.)
- ▶ There are some "quasi-standards" (see next slide) w.r.t. parameter names and available options that are adhered to by **most** developers
- ▶ `configure --help` displays an exhaustive list of possible parameters and environment variables with explanations on how to use them



When using a GNU `configure` script that honors a few common standards, the following is expected to work:

- ▶ `configure --prefix=<install_location>` sets the install location (i.e. where compiled files get copied to when typing `make install`)

Make sure that you have write permissions in the prefix path!

- ▶ `--enable-<feature>` and `--disable-<feature>` include or exclude optional components in the compilation
- ▶ `--with-<library>[=<library_path>]` and `--without-<library>` include and set paths to external dependencies, or try to build without them if possible
- ▶ Common GNU Makefile environment (e.g. `CC`, `CFLAGS`, `LD_FLAGS`) are taken into account by `configure`



It is recommended to build **outside of the source tree**, i.e. creating a new empty directory to contain files produced by the build process (e.g. object files).

- 1 Create a build directory and change into it, e.g.:

```
mkdir $HOME/build/<some_software> && cd $HOME/build/<some_software>
```

- 2 Check configure options: `<path_to_source>/configure --help`

- 3 Execute configure with appropriate options: `<path_to_source>/configure <parameters>`

- 4 Build: `make`

- 5 If available, build and execute self-checks: `make check` OR `make test`

- 6 Install into the directory set by `--prefix`: `make install`



The source tarball for the open source interactive plotting software Gnuplot can be found on <https://sourceforge.net/projects/gnuplot/files/gnuplot/5.2.6/>

Demos/configure-Gnuplot/build.sh

```
1 # SPDX-FileCopyrightText: 2021 Competence Center for High Performance Computing in He
2 # SPDX-License-Identifier: MIT
3
4 #!/bin/bash
5
6 tar xvf gnuplot-5.2.6.tar.gz # unpack and untar the sources
7 cd gnuplot-5.2.6
8 # module purge && module load gcc # default gcc module (on a cluster)
9 mkdir build && cd build # create and change into the build directory
10 mkdir -vp ../install/ # create an install directory
11 # ../configure --help # check the options for configure script
12 ../configure CC=gcc CXX=g++ CPP='gcc -E' CXXCPP='g++ -E' \
13     --prefix=$(pwd)/../install --without-latex
14 # less config.log # inspect configure output
15 make -j4 # compile
16 make install # install
```

- ▶ The Modulesystem works with Environment Variables
- ▶ Compilers will automatically use this variables to determine where to search for files
- ▶ Some Important Environment Variables are:
 - ▶ `$PATH`: Location of Executables
 - ▶ `$CPATH` Location of header files
 - ▶ `$LIBRARY_PATH` Location of Libraries (compile time)
 - ▶ `$LD_LIBRARY_PATH` Location of Libraries (run time)
- ▶ Each Variable is a colon separated list of directories
- ▶ Append it by using `export PATH=$PATH:/your/new/path`



- ▶ You can even write your own modules
- ▶ name the directory after the module and the file itself after the version number
- ▶ Use them by appending the `$MODULEPATH` environment variable

libxml2/2.9.7.lua

```
1  whatis("library for xml parsing")
2  whatis("http://www.xmlsoft.org/")
3  whatis("Full module name: libxml2/2.9.7")
4  whatis("Module File: /home/tj75qeje/modules/modulefiles/libxml2/2.9.7.lua")
5  whatis("Application Root: /home/tj75qeje/modules/software/libxml2/2.9.7")
6  -- Environment
7  prepend_path("PATH", "/home/tj75qeje/modules/software/libxml2/2.9.7/bin")
8  prepend_path("CPATH", "/home/tj75qeje/modules/software/libxml2/2.9.7/include")
9  prepend_path("LIBRARY_PATH", "/home/tj75qeje/modules/software/libxml2/2.9.7/lib")
10 prepend_path("LD_LIBRARY_PATH", "/home/tj75qeje/modules/software/libxml2/2.9.7/lib")
11 pushenv("LIBXML2_ROOT", "/home/tj75qeje/modules/software/libxml2/2.9.7")
12
13 -- Messages
14 LmodMessage("Lmod: ", mode().."ing", myModuleName(), myModuleVersion())
```



- ▶ Now try it yourself!
- ▶ build the NetCDF library from source
 - ▶ `https://downloads.unidata.ucar.edu/netcdf/`
 - ▶ The HDF5 library is needed as a prerequisite.
Therefore you have to load this module!
- ▶ build an example program and link it against your NetCDF library
 - ▶ `https://www.unidata.ucar.edu/software/netcdf/examples/programs/simple_xy_wr.c`
- ▶ Can you build the NetCDF library with support for Parallel IO?
- ▶ Can you also write a modulefile for your own build of the NetCDF library?

- ▶ NetCDF can also be built with cmake (not explained in this course)



Automatic Variables. URL: <https://www.gnu.org/software/make/manual/make.html#Automatic-Variables>
(visited on 03/07/2019) (cit. on p. 27).



Directory-Search. URL: <https://www.gnu.org/software/make/manual/make.html#Directory-Search>
(visited on 02/14/2019) (cit. on pp. 28, 29).



Makefile Conventions. URL: <https://www.gnu.org/software/make/manual/make.html#Makefile-Conventions>
(visited on 03/12/2019) (cit. on p. 31).



Phony Targets. URL: <https://www.gnu.org/software/make/manual/make.html#Phony-Targets>
(visited on 02/14/2019) (cit. on p. 30).



Writing Rules. URL:

<https://www.gnu.org/software/make/manual/make.html#Rules> (visited on 02/14/2019) (cit. on p. 22).