



Reproducibility & performance: Why choose?

Ludovic Courtès


Forum ORAP

15 March 2024

Inria

- 
- ▶ Guix started in **2012**
 - ▶ tools for **reproducible software deployment**
 - ▶ runs standalone (Guix System) or atop a **GNU/Linux** distro
 - ▶ **50,000+ packages**
 - ▶ **100+ monthly contributors**

<https://hpc.guix.info>

- 
- ▶ Guix started in 2012
 - ▶ tools for **reproducible software deployment**
 - ▶ runs standalone (Guix System) or atop a **GNU/Linux** distro
 - ▶ **50,000+ packages**
 - ▶ **100+ monthly contributors**
 - ▶ **Guix-HPC effort (Inria, MDC, UBC, UTHCS) started in 2017**

<https://hpc.guix.info>

- ▶ **PlaFRIM** (FR): Inria Bordeaux (3,000+ cores)
- ▶ **GriCAD** (FR): Grenoble (1,000+ cores)
- ▶ **GLICID** (FR): Nantes (4,000+ cores)
- ▶ **Grid'5000** (FR): 8 sites (12,000+ cores)
- ▶ **Max Delbrück Center** (DE): 250-node cluster + workstations
- ▶ **UMC Utrecht** (NL): 68-node cluster (1,000+ cores)
- ▶ ...



<https://www.acm.org/publications/policies/artifact-review-badging>



Software Heritage



guix shell python python-numpy

```
guix shell python python-numpy \  
  -- python3 -c 'import numpy'
```



```
guix shell --manifest=my-packages.scm
```

```
(specifications->manifest  
  '("gcc-toolchain" "openmpi"  
    "scotch" "mumps"))
```

```
guix shell --container \  
  --manifest=my-packages.scm
```

```
(specifications->manifest  
  '("gcc-toolchain" "openmpi"  
    "scotch" "mumps"))
```

Reproducible environments: 2 files, 2 commands

1. `guix describe -f channels > channels.scm`
2. `guix time-machine -C channels.scm -- \`
`shell -m manifest.scm`

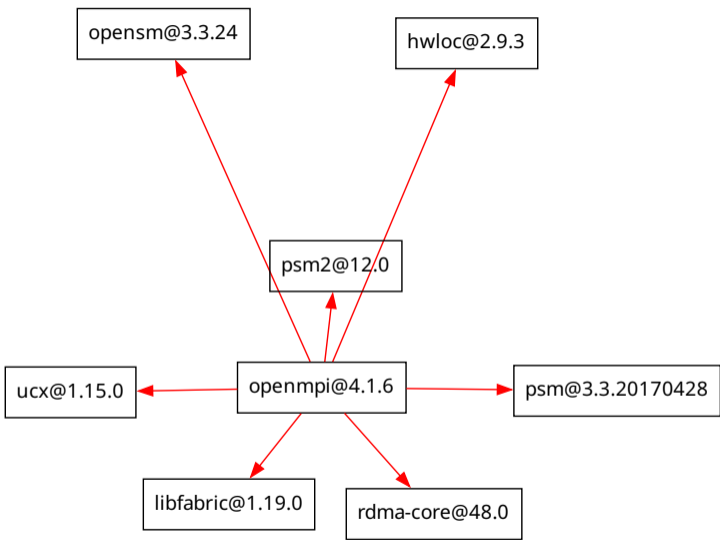
A photograph of a server room aisle. A woman in a light-colored sweater is crouching on the metal grating floor, looking at a server rack. In the background, two other people are standing and talking. The room is filled with rows of server racks, and the lighting is a warm, yellowish-green. The text "Two obsessions: MPI and AVX." is overlaid in white, bold font across the center of the image.

Two obsessions: MPI and AVX.

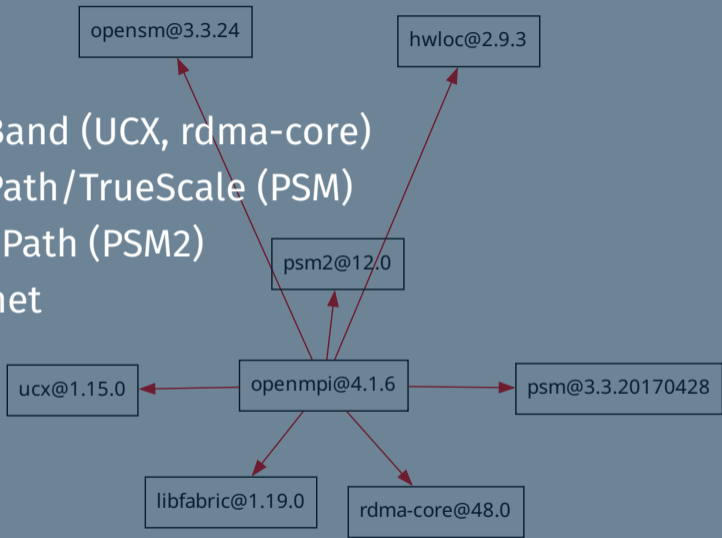
Some fairly common (but questionable) assumptions made by package managers (conda, pip, apt, etc.)

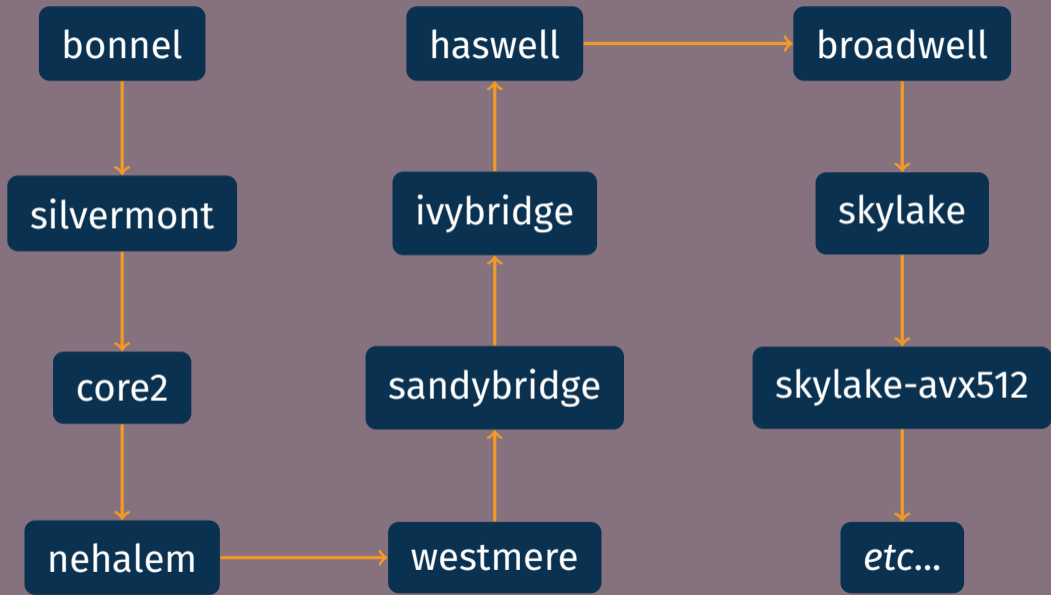
- **1:1 relationship between source code and binary (per platform)**
 - Good for reproducibility (e.g., Debian)
 - Bad for performance optimization
- **Binaries should be as portable as possible**
 - What most distributions do
 - Again, bad for performance
- **Toolchain is the same across the ecosystem**
 - One compiler, one set of runtime libraries
 - Or, no compiler (for interpreted languages)

Todd Gamblin (Spack)



- ▶ InfiniBand (UCX, rdma-core)
- ▶ InfiniPath/TrueScale (PSM)
- ▶ Omni-Path (PSM2)
- ▶ Ethernet
- ▶ ...






```
$ spack spec netcdf-c
```

```
Concretized
```

```
-----  
netcdf-c@4.8.1... arch=linux-ubuntu18.04-skylake_avx512  
  ^ hdf5@1.10.7... arch=linux-ubuntu18.04-skylake_avx512  
    ^ cmake@3.21.4... arch=linux-ubuntu18.04-skylake_avx512  
      ^ ncurses@6.2... arch=linux-ubuntu18.04-skylake_avx512  
        ^ pkgconf@1.8.0... arch=linux-ubuntu18.04-skylake_avx512  
          ^ openssl@1.1.1l... arch=linux-ubuntu18.04-skylake_avx512  
            ^ perl@5.34.0... arch=linux-ubuntu18.04-skylake_avx512  
              ^ berkeley-db@18.1.40... arch=linux-ubuntu18.04-skylake_avx512  
                ^ bzip2@1.0.8... arch=linux-ubuntu18.04-skylake_avx512  
                  ^ diffutils@3.8... arch=linux-ubuntu18.04-skylake_avx512  
                    ^ libiconv@1.16... arch=linux-ubuntu18.04-skylake_avx512  
                      ^ gdbm@1.19... arch=linux-ubuntu18.04-skylake_avx512  
                        ^ readline@8.1... arch=linux-ubuntu18.04-skylake_avx512  
                          ^ zlib@1.2.11... arch=linux-ubuntu18.04-skylake_avx512
```

```
$ spack spec netcdf-c
```

```
Concretized
```

```
-----  
netcdf-c@4.8.1... arch=linux-ubuntu18.04-skylake_avx512  
  ^ hdf5@1.10.7... arch=linux-ubuntu18.04-skylake_avx512  
    ^ cmake@3.21.4... arch=linux-ubuntu18.04-skylake_avx512  
      ^ ncurses@6.2... arch=linux-ubuntu18.04-skylake_avx512  
        ^ pkgconf@1.8.0... arch=linux-ubuntu18.04-skylake_avx512  
          ^ openssl@1.1.1l... arch=linux-ubuntu18.04-skylake_avx512  
            ^ perl@5.34.0... arch=linux-ubuntu18.04-skylake_avx512  
              ^ berkeley-db@18.1.40... arch=linux-ubuntu18.04-skylake_avx512  
                ^ bzip2@1.0.8... arch=linux-ubuntu18.04-skylake_avx512  
                  ^ diffutils@3.8... arch=linux-ubuntu18.04-skylake_avx512  
                    ^ libiconv@1.16... arch=linux-ubuntu18.04-skylake_avx512  
                      ^ gdbm@1.19... arch=linux-ubuntu18.04-skylake_avx512  
                        ^ readline@8.1... arch=linux-ubuntu18.04-skylake_avx512  
                          ^ zlib@1.2.11... arch=linux-ubuntu18.04-skylake_avx512
```

```
$ spack spec netcdf-c
```

```
Concretized
```

```
-----  
netcdf-c@4.8.1... arch=linux-ubuntu18.04-skylake_avx512  
  ^ hdf5@1.10.7... arch=linux-ubuntu18.04-skylake_avx512  
    ^ cmake@3.21.4... arch=linux-ubuntu18.04-skylake_avx512  
      ^ ncurses@6.2... arch=linux-ubuntu18.04-skylake_avx512  
        ^ pkgconf@1.8.0... arch=linux-ubuntu18.04-skylake_avx512  
          ^ openssl@1.1.1l... arch=linux-ubuntu18.04-skylake_avx512  
            ^ perl@5.34.0... arch=linux-ubuntu18.04-skylake_avx512  
              ^ berkeley-db@18.1.40... arch=linux-ubuntu18.04-skylake_avx512  
                ^ bzip2@1.0.8... arch=linux-ubuntu18.04-skylake_avx512  
                  ^ diffutils@3.8... arch=linux-ubuntu18.04-skylake_avx512  
                    ^ libiconv@1.16... arch=linux-ubuntu18.04-skylake_avx512  
                      ^ gdbm@1.19... arch=linux-ubuntu18.04-skylake_avx512  
                        ^ readline@8.1... arch=linux-ubuntu18.04-skylake_avx512  
                          ^ zlib@1.2.11... arch=linux-ubuntu18.04-skylake_avx512
```

```
$ spack spec netcdf-c
```

```
Concretized
```

```
-----  
netcdf-c@4.8.1... arch=linux-ubuntu18.04-skylake_avx512  
  ^ hdf5@1.10.7... arch=linux-ubuntu18.04-skylake_avx512  
    ^ cmake@3.21.4... arch=linux-ubuntu18.04-skylake_avx512  
      ^ ncurses@6.2... arch=linux-ubuntu18.04-skylake_avx512  
        ^ pkgconf@1.8.0... arch=linux-ubuntu18.04-skylake_avx512  
          ^ openssl@1.1.1l... arch=linux-ubuntu18.04-skylake_avx512  
            ^ perl@5.34.0... arch=linux-ubuntu18.04-skylake_avx512  
              ^ berkeley-db@18.1.40... arch=linux-ubuntu18.04-skylake_avx512  
                ^ bzip2@1.0.8... arch=linux-ubuntu18.04-skylake_avx512  
                  ^ diffutils@3.8... arch=linux-ubuntu18.04-skylake_avx512  
                    ^ libiconv@1.16... arch=linux-ubuntu18.04-skylake_avx512  
                      ^ gdbm@1.19... arch=linux-ubuntu18.04-skylake_avx512  
                        ^ readline@8.1... arch=linux-ubuntu18.04-skylake_avx512  
                          ^ zlib@1.2.11... arch=linux-ubuntu18.04-skylake_avx512
```

```
https://spack.readthedocs.io/en/latest/getting\_started.html
```

```
$ spack spec netcdf-c
```

```
Concretized
```

```
-----  
netcdf-c@4.8.
```

```
^ hdf5@1.1
```

```
^ cmake
```

THE #1 PROGRAMMER EXCUSE
FOR LEGITIMATELY SLACKING OFF:

"MY CODE'S COMPILING."

HEY! GET BACK
TO WORK!

COMPILING!

OH. CARRY ON.



```
:512
```

```
lake_avx512
```

```
avx512
```

```
_avx512
```

```
ylake_avx512
```

```
lake_avx512
```

```
ylake_avx512
```

```
ylake_avx512
```

```
lake_avx512
```

```
ylake_avx512
```

```
lake_avx512
```

GNU libc

Libgcrypt

Function
multi-versioning
New

OpenBLAS

BLIS

FFTW

GMP

Julia

Rust

<https://hpc.guix.info/blog/2018/01/pre-built-binaries-vs-performance/>

```
$ guix shell eigen-benchmarks -- \
    benchBlasGemm 240 240 240
240 x 240 x 240
cblas: 0.20367 (16.289 GFlops/s)
eigen : 0.285149 (11.635 GFlops/s)
```

```
$ guix shell eigen-benchmarks -- \
  benchBlasGemm 240 240 240
240 x 240 x 240
cblas: 0.20367 (16.289 GFlops/s)
eigen : 0.285149 (11.635 GFlops/s)
```

**Package
multi-versioning**

```
$ guix shell --tune eigen-benchmarks -- \
  benchBlasGemm 240 240 240
guix shell: tuning for CPU micro-architecture skylake
240 x 240 x 240
cblas: 0.203131 (16.333 GFlops/s)
eigen : 0.0929638 (35.688 GFlops/s)
```

<https://hpc.guix.info/blog/2022/01/tuning-packages-for-a-cpu-micro-architecture/>

AMD 
ROCm

```
laptop$ guix pack -RR hpcg -S /bin=bin
```

```
laptop$ guix pack -RR hpcg -S /bin=bin
```



```
adastra$ tar xf pack.tar.gz
```

```
adastra$ ./bin/mpirun -n 8 ... \  
./bin/rochpcg 280 280 280 180
```

<https://hpc.guix.info/blog/2024/01/hip-and-rocm-come-to-guix/>

- ▶ P. Swartvagher, *On the Interactions between HPC Task-based Runtime Systems and Communication Libraries*, PhD thesis, Dec. 2022
- ▶ M. Felšöci, *Fast Solvers for High-Frequency Aeroacoustics*, PhD thesis, Feb. 2023
- ▶ N. Vallet *et al.*, *Toward practical transparent verifiable and long-term reproducible research using Guix*, Nature Scientific Data, Oct. 2022

Reproducible deployment
can be achieved
without sacrificing performance.



Let's add

reproducible deployment

to our best practices book.



<https://hpc.guix.info>

ludovic.courtes@inria.fr | [@civodul@toot.aquilenet.fr](https://toot.aquilenet.fr/@civodul)

Copyright © 2012–2024 Ludovic Courtès ludo@gnu.org.

GNU Guix logo by Luis Felipe, CC-BY-SA 4.0, <https://guix.gnu.org/graphics>.

Tandem picture by Jules Beau, public domain, <https://images.bnf.fr/#/detail/1535157/9>

LLNL supercomputer picture by US DoE, public domain,
[https://commons.wikimedia.org/wiki/File:U.S._Department_of_Energy_-_Science_-_477_018_010_\(9563440651\).jpg](https://commons.wikimedia.org/wiki/File:U.S._Department_of_Energy_-_Science_-_477_018_010_(9563440651).jpg)

Copyright of other images included in this document is held by their respective owners.

This work is licensed under the [Creative Commons Attribution-Share Alike 4.0](https://creativecommons.org/licenses/by-sa/4.0/) License. To view a copy of this license, visit <https://creativecommons.org/licenses/by-sa/4.0/> or send a letter to Creative Commons, 171 Second Street, Suite 300, San Francisco, California, 94105, USA.

At your option, you may instead copy, distribute and/or modify this document under the terms of the [GNU Free Documentation License, Version 1.3 or any later version](https://www.gnu.org/licenses/gfdl.html) published by the Free Software Foundation; with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts. A copy of the license is available at <https://www.gnu.org/licenses/gfdl.html>.

The source of this document is available from <https://git.sv.gnu.org/cgiit/guix/maintenance.git>.