

Reproducible and User-Controlled Package Management in HPC with GNU Guix

Ludovic Courtès (`ludovic.courtes@inria.fr`)
Ricardo Wurmus (`ricardo.wurmus@mdc-berlin.de`)

Workshop on Reproducibility in Parallel Computing (RepPar)
25 August 2015

“Reproducibility”?

“Reproducibility”?

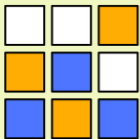
1. **bit-reproducible** builds

“Reproducibility”?

1. **bit-reproducible** builds
2. **isolating** a software environment from changes

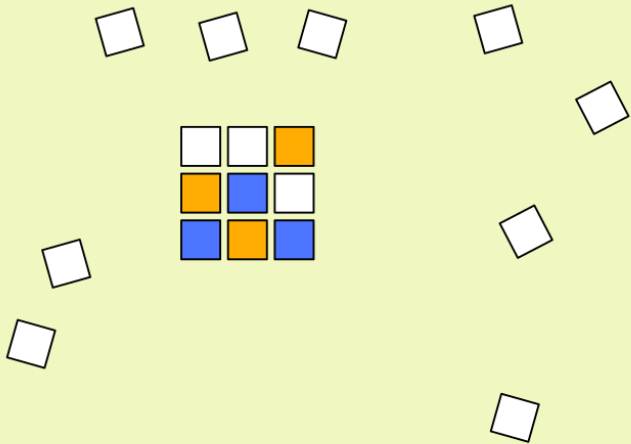
“Reproducibility”?

localhost



“Reproducibility”?

localhost

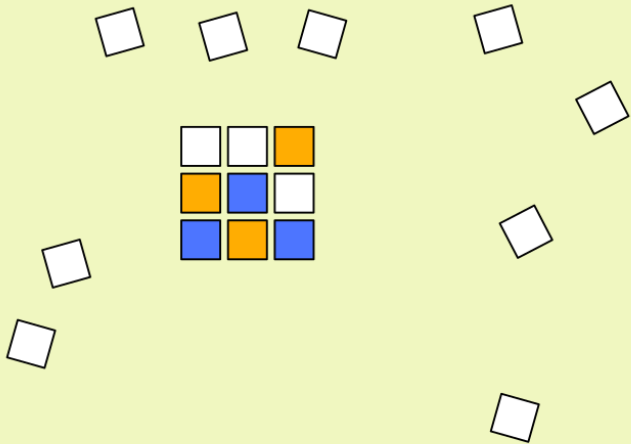


“Reproducibility”?

1. **bit-reproducible** builds
2. **isolating** a software environment from changes
3. **sharing** environments with others

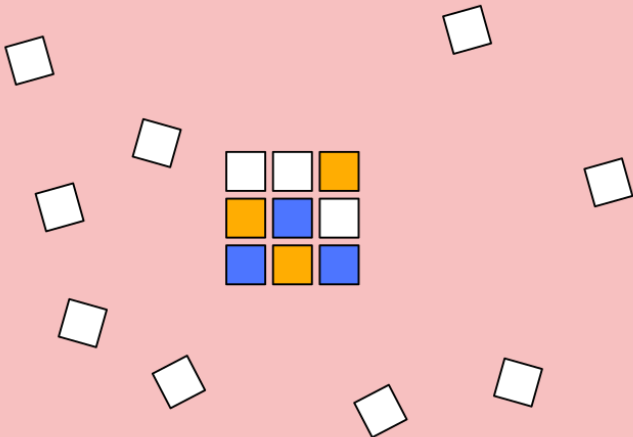
“Reproducibility”?

localhost



“Reproducibility”?

remote

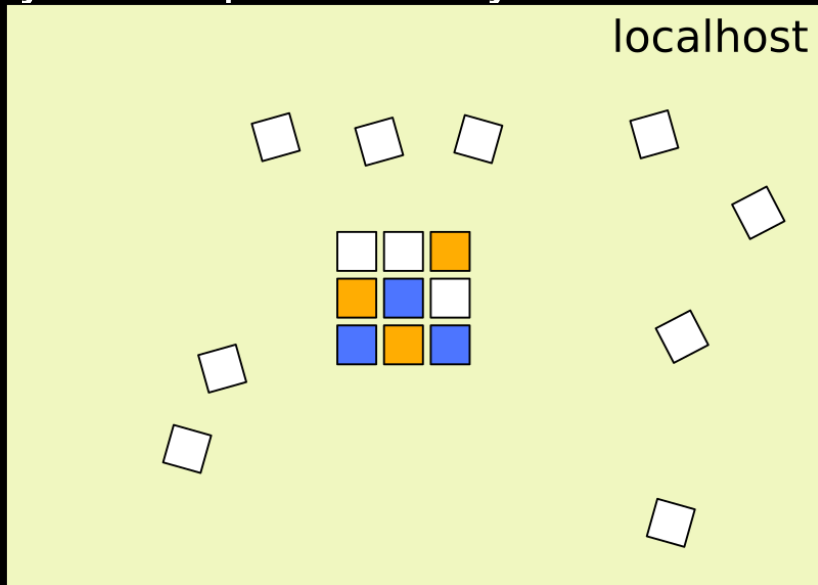


Beyond Reproducibility

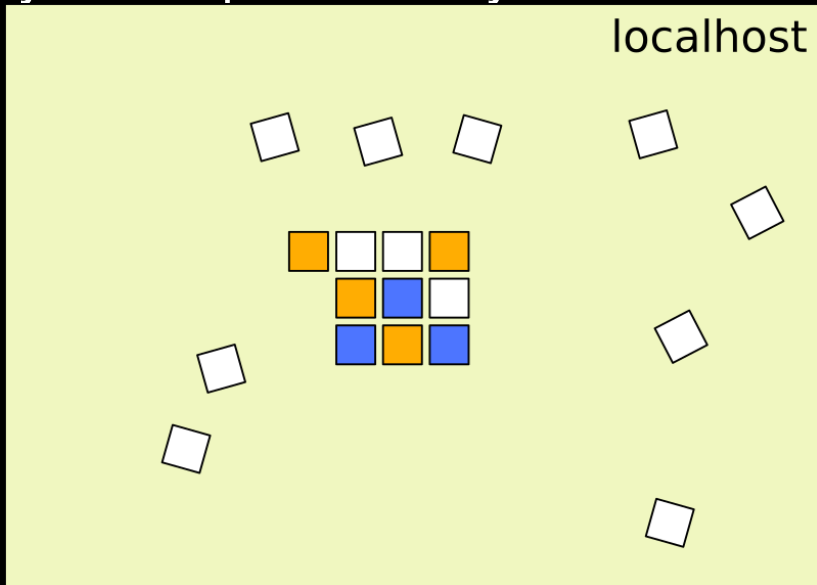
Beyond Reproducibility

- ▶ **user-controlled upgrades** and roll-backs

Beyond Reproducibility



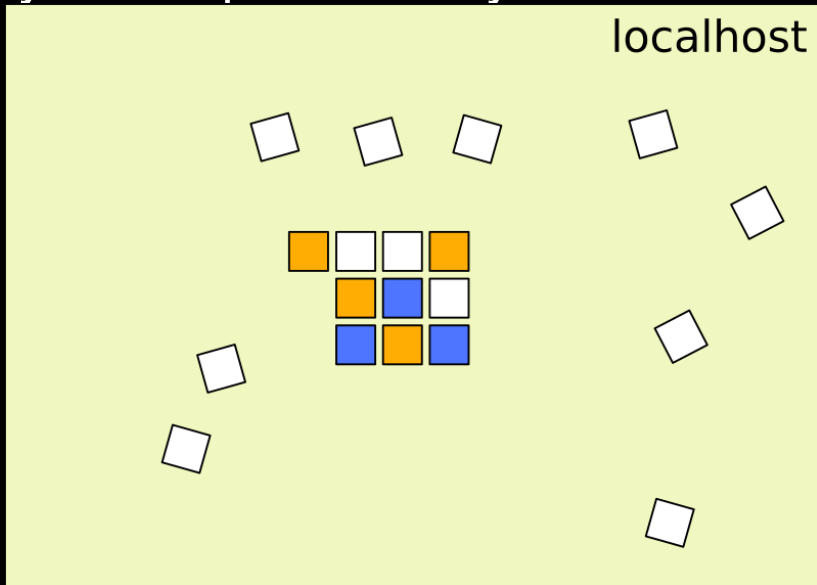
Beyond Reproducibility



Beyond Reproducibility

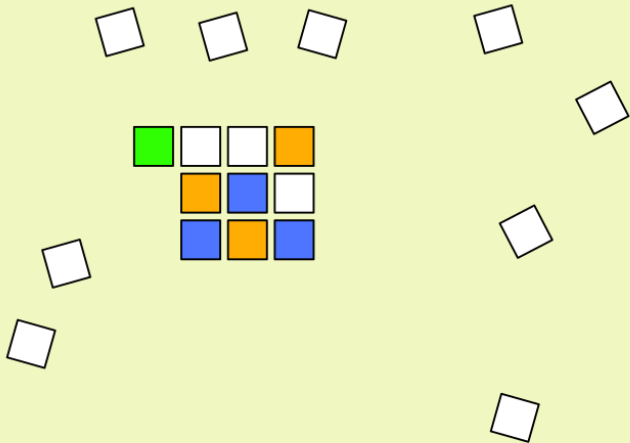
- ▶ **user-controlled upgrades** and roll-backs
- ▶ change **specific parts** of the software stack

Beyond Reproducibility



Beyond Reproducibility

localhost



Beyond Reproducibility

- ▶ **user-controlled upgrades** and roll-backs
- ▶ change **specific parts** of the software stack
- ▶ **hackability**: no black boxes

VMs and Docker

VMs and Docker

Pros:

- ▶ “bit-reproducible”
- ▶ reproducible
anywhere by anyone

VMs and Docker

Pros:

- ▶ “bit-reproducible”
- ▶ reproducible
anywhere by anyone

Problems:

- ▶ VMs are
heavyweight
- ▶ binary images are
opaque
- ▶ not composable

Functional Package Management

Functional Package Management

Regarding the build & installation process of a package as a **pure function**.

`openmpi = f(hwloc, gcc, make, coreutils)`

where `f = ./configure && make && make install`

```
openmpi = f(hwloc, gcc, make, coreutils)
hwloc = g(pciaccess, gcc, make, coreutils)
```



```
openmpi = f(hwloc, gcc, make, coreutils)
hwloc = g(pciaccess, gcc, make, coreutils)
gcc = h(make, coreutils, gcc0)
...
```

```
openmpi = f(hwloc, gcc, make, coreutils)
hwloc = g(pciaccess, gcc, make, coreutils)
gcc = h(make, coreutils, gcc0)
```

...

the complete DAG is captured

References

- ▶ *A Safe and Policy-Free System for Software Deployment* (“Nix”), Dolstra et al., 2003
- ▶ *Functional Package Management with Guix*, Courtès, 2013

Thesis

1. functional package management (FPM)
empowers cluster users
2. FPM is a solid foundation for **reproducible software deployment**
3. **beyond reproducibility**: Guix is programmable, supports experimentation

From the Architecture of Nix...

<http://nixos.org/nix/>

build processes
chroot, separate UIDs

Nix tools

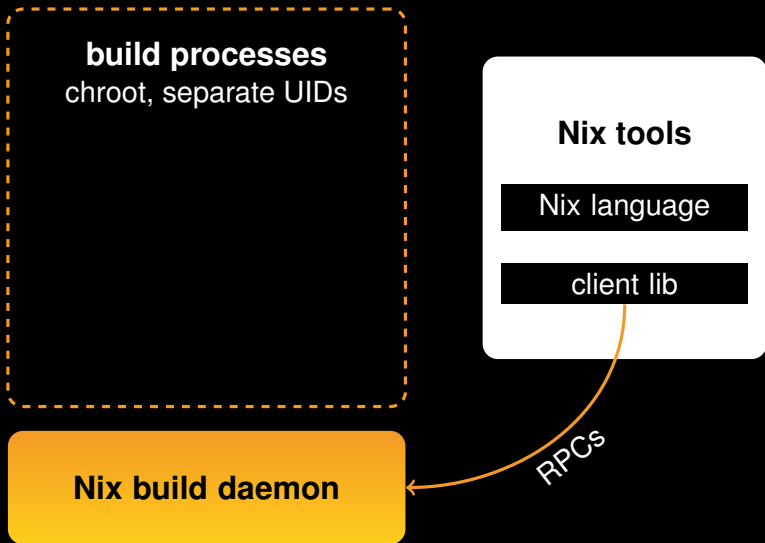
Nix language

client lib

Nix build daemon

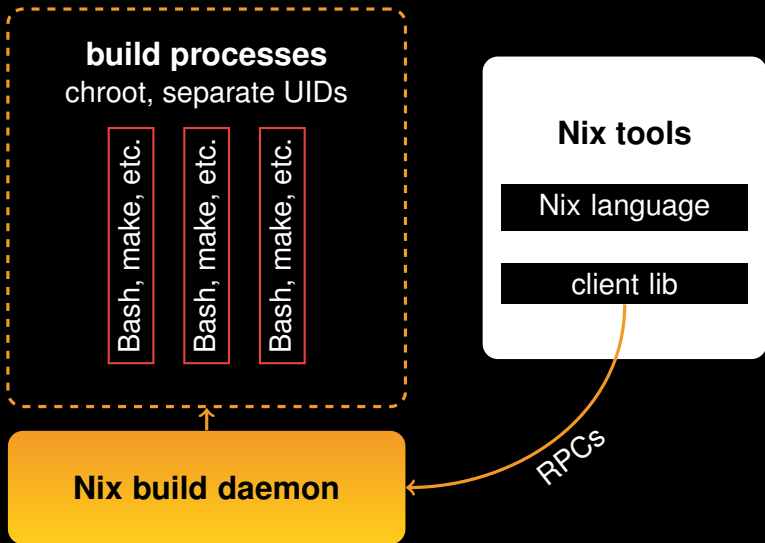
From the Architecture of Nix...

<http://nixos.org/nix/>



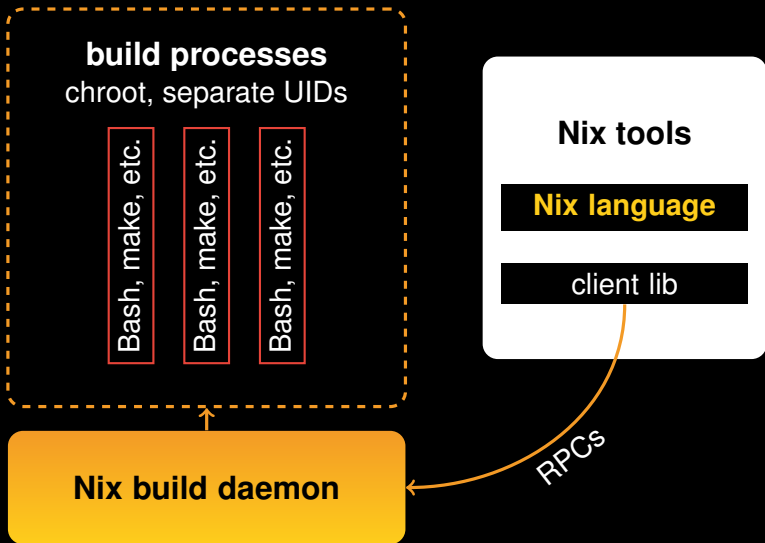
From the Architecture of Nix...

<http://nixos.org/nix/>



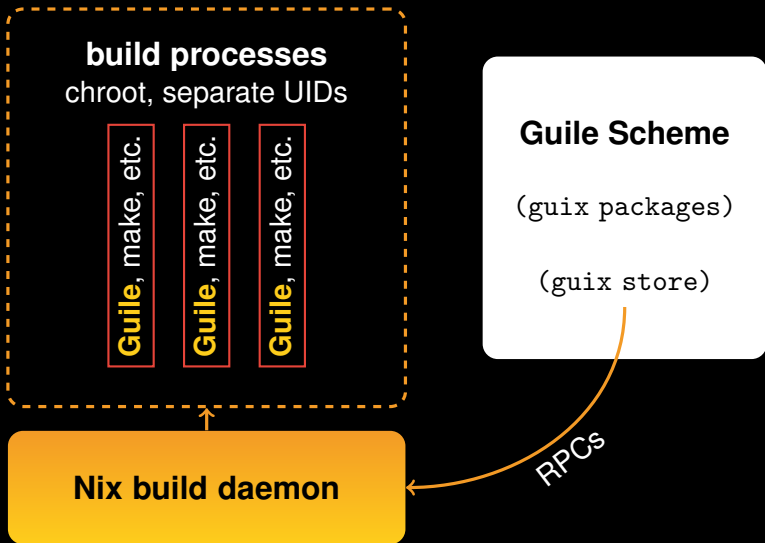
From the Architecture of Nix...

<http://nixos.org/nix/>



... to the Architecture of Guix

<http://gnu.org/s/guix/>



Why Guix?

[Courtès 2013]

1. Scheme is a “programmable programming language” → **tailored EDSLs**
2. **general-purpose language** with compiler, debugger, libraries, etc.
3. **a single language** → more code reuse, unified environment
4. **complete package programming interface**

Bit-Reproducible Builds*

* almost!


```
$ guix build petsc
```

isolated build: chroot, separate name spaces, etc.

Bit-Reproducible Builds*

* almost!

```
$ guix build petsc  
/gnu/store/ h2g4sf72... -petsc-3.6.0
```



hash of **all** the dependencies

Bit-Reproducible Builds*

* almost!

```
$ guix build petsc  
/gnu/store/h2g4sf72... -petsc-3.6.0
```

```
$ guix gc --references /gnu/store/...-petsc-3.6.0  
/gnu/store/...-openmpi-1.8.5  
/gnu/store/...-gfortran-4.9.3-lib  
/gnu/store/...-superlu-4.3  
/gnu/store/...-lapack-3.5.0  
/gnu/store/...-glibc-2.21  
...
```

Bit-Reproducible Builds*

* almost!

```
$ guix build petsc  
/gnu/store/h2g4sf72... -petsc-3.6.0
```

```
$ guix gc --references /gnu/store/...-petsc-3.6.0  
/gnu/store/...-openmpi-1.8.5  
/gnu/store/...-gfortran-4.9.3-lib  
/gnu/store/...-superlu-4.3  
/gnu/store/...-lapack-3.5.0  
/gnu/store/...-g...  
...
```

(nearly) bit-identical for everyone

Reproducible Environments

- ▶ **per-user “profiles”**
- ▶ **non-interference** among users/profiles
- ▶ **transactional upgrades & rollbacks**

Reproducible Environments

```
$ guix package -i gcc-toolchain coreutils sed grep
```

```
...
```

```
$ eval 'guix package --search-paths'
```

```
...
```

```
$ guix package --manifest=my-software.scm
```

```
...
```

```
$ guix environment demo petsc
```

```
...
```

```
$ guix environment --ad-hoc python-ipython python-numpy \  
-E ipython
```

```
...
```


Experience at the Max Delbrück Center, Berlin

- ▶ **Guix deployed on two clusters** (≈ 250 nodes) + workstations
- ▶ used by **bioinformatics** researchers
- ▶ **50+ bioinfo packages** in use (C/C++, Python, etc.)
- ▶ replaces CentOS packages + sysadmin-managed software

Experience at the Max Delbrück Center, Berlin

- ▶ **Guix deployed on two clusters** (≈ 250 nodes) + workstations
- ▶ used by **bioinformatics** researchers
- ▶ **50+ bioinfo packages** in use (C/C++, Python, etc.)
- ▶ replaces CentOS packages + sysadmin-managed software
- ▶ **advantages**: more user control, better resource usage, ...

Beyond Reproducibility: Supporting Experimentation

Fiddling with the HPC Stack

Example from Inria

linear algebra

PaSTiX, Chameleon (solvers)

StarPU (task scheduling)

run-time

hwloc (hardware topology)

MPI (message passing)

compiler's run-time support, C library, etc.

Requirements for an Experimentation-Capable System

1. customize + non-ambiguously **specify package DAG**
2. reliably reproduce **variants of the DAG**

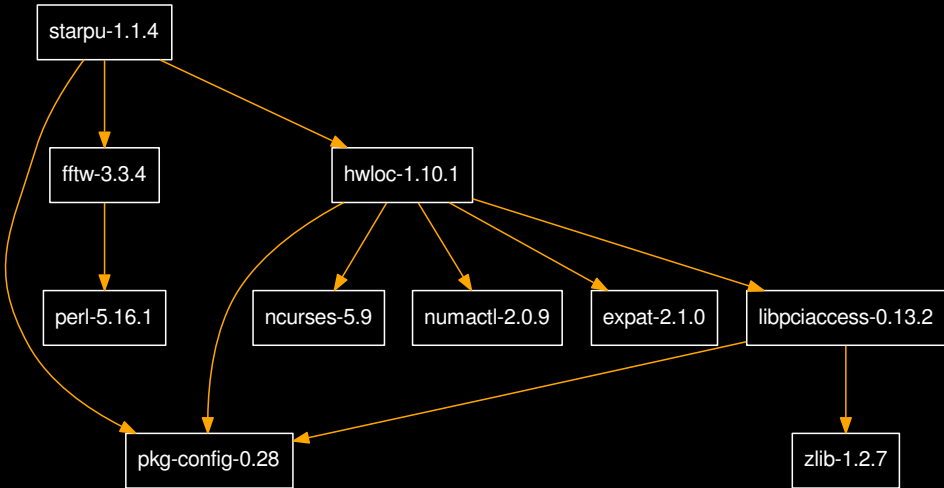
```
(define starpu
  (package
    (name "starpu")
    (version "1.1.4")
    (source (origin
              (method url-fetch)
              (uri "http://...")
              (sha256 (base32 "0zmkw..."))))
    (build-system gnu-build-system )
    (native-inputs '(("pkg-config" ,pkg-config)))
    (inputs '(("fftw" ,fftw)
              ("hwloc" , hwloc )))
    (home-page "http://starpu.gforge.inria.fr/")
    (synopsis "Run-time system for heterogeneous computing")
    (description "Blah...")
    (license lgpl2.1+)))
```

```
(define starpu
  (package
    (name "starpu")
    (version "1.1.4")
    (source (origin
              (method url-fetch)
              (uri "http://...")
              (sha256 (base32 "0zmkw..."))))
    (build-system gnu-build-system )
    (native-inputs '(("pkg-config" ,pkg-config)))
    (inputs '(("fftw" ,fftw)
              ("hwloc" ,hwloc )))
    (home-page "http://starpu.gforge.inria.fr/")
    (synopsis "Run-time system for heterogeneous computing")
    (description "Blah...")
    (license lgpl2.1+)))
```

dependencies

```
(define starpu
  (package
    (name "starpu")
    (version "1.1.4")
    (source (origin
      (uri "http://...")
      (sha256 (base32 "0zmkw..."))))
    (build-system gnu-build-system)
    (native-inputs '("pkg-config" ,pkg-config))
    (inputs '("fftw" ,fftw) depends on gcc, make, bash, etc.
      ("hwloc" , hwloc)))
    (home-page "http://starpu.gforge.inria.fr/")
    (synopsis "Run-time system for heterogeneous computing")
    (description "Blah...")
    (license lgpl2.1+)))
```

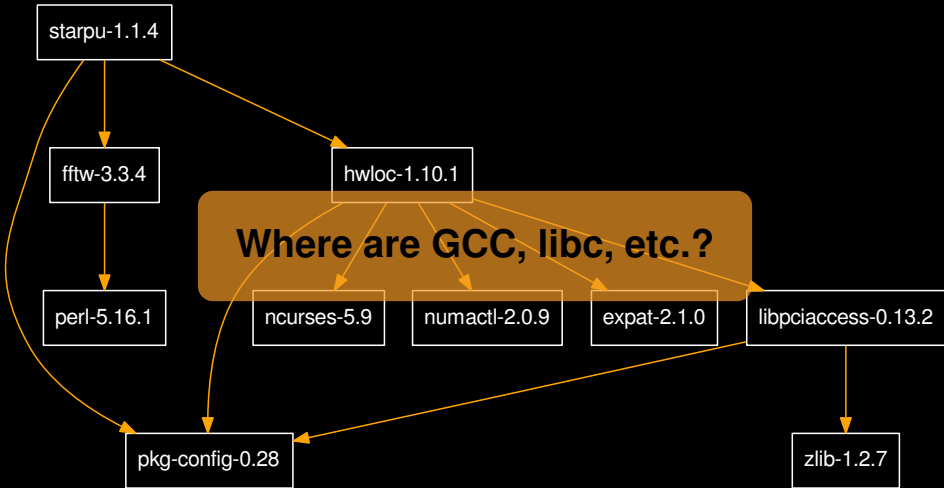

DAG of Package Objects



10 nodes

`guix graph --type=package starpu`

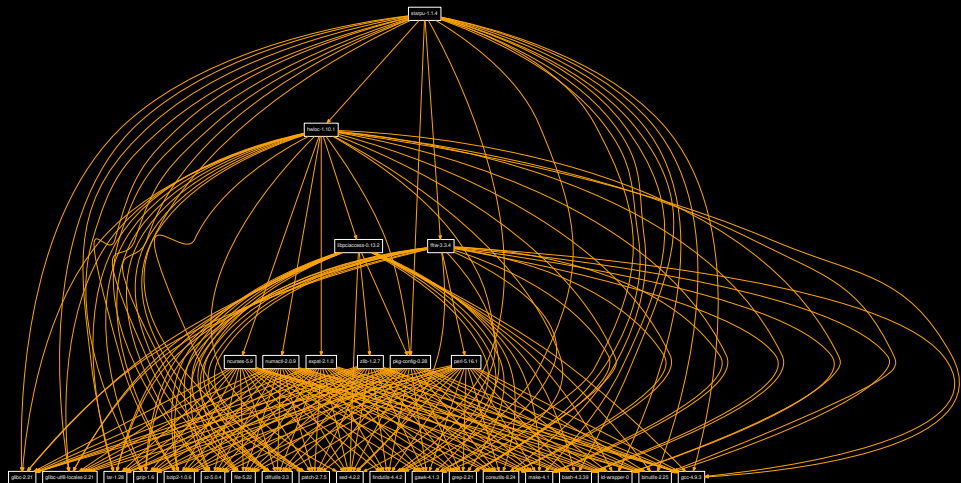
DAG of Package Objects



10 nodes

`guix graph --type=package starpu`

Same DAG, including implicit inputs



29 nodes

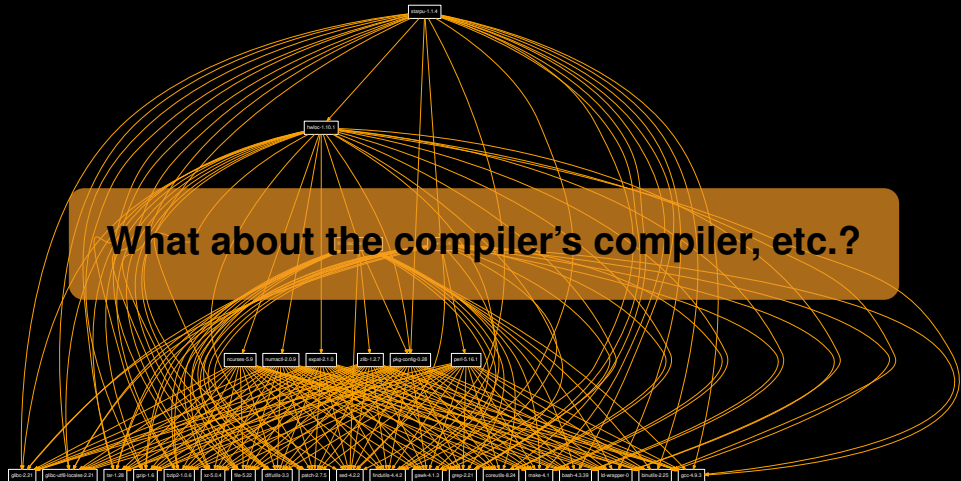
`guix graph --type=bag-emerged starpu`

Same DAG, including implicit inputs

What about the compiler's compiler, etc.?

29 nodes

`guix graph --type=bag-emerged starpu`



Full DAG, including bootstrap

(too big)

321 nodes

```
guix graph --type=bag starpu
```

Defining Package Variants

```
(define starpu-1.2rc                                ;release candidate
  (package (inherit starpu)
    (version "1.2.0rc2")
    (source (origin
      (method url-fetch)
      (uri (string-append "http://.../"
                          "starpu-" version ".tar.gz"))
      (sha256 (base32 "0qgb6y..."))))))))
```


Defining Package Variants

```
(define starpu-with-simgrid
  (package (inherit starpu)
    (name "starpu-with-simgrid")

    ;; Add SimGrid, an optional dependency.
    (inputs '(("simgrid" ,simgrid)
              ,@(package-inputs starpu))))))
```

Package Functions

```
(define (make-chameleon starpu )  
  ;; Return the Chameleon solver linked against  
  ;; this particular variant of StarPU.  
  (package  
    ;; ...  
    (inputs ‘(("starpu" , starpu )  
              ("blas" ,atlas)  
              ("lapack" ,lapack)  
              ("gfortran" ,gfortran-4.8)  
              ("python" ,python-2))))))
```



Package Functions

```
(define (make-chameleon starpu )
  ;; Return the Chameleon solver linked against
  ;; this particular variant of StarPU.
  (package
    ;; ...
    (inputs ‘(("starpu" , starpu )
              ("blas" ,atlas)
              ("lapack" ,lapack)
              ("gfortran" ,gfortran-4.8)
              ("python" ,python-2))))))

(define chameleon
  (make-chameleon starpu))

(define chameleon/starpu-simgrid
  (make-chameleon starpu-with-simgrid))
```

Conclusion

Limitations

- ▶ daemon must run as root to isolate builds
- ▶ non-deterministic build systems
- ▶ non-free software unavailable in Guix

Limitations

- ▶ daemon must run as root to isolate builds
 - ▶ WIP: have daemon rely on **user name spaces** (Linux 3.8+)
- ▶ non-deterministic build systems
- ▶ non-free software unavailable in Guix

Limitations

- ▶ daemon must run as root to isolate builds
 - ▶ WIP: have daemon rely on **user name spaces** (Linux 3.8+)
- ▶ non-deterministic build systems
 - ▶ must be identified & **fixed upstream**
 - ▶ WIP thanks to <http://reproducible.debian.net>
- ▶ non-free software unavailable in Guix

Limitations

- ▶ daemon must run as root to isolate builds
 - ▶ WIP: have daemon rely on **user name spaces** (Linux 3.8+)
- ▶ non-deterministic build systems
 - ▶ must be identified & **fixed upstream**
 - ▶ WIP thanks to <http://reproducible.debian.net>
- ▶ non-free software unavailable in Guix
 - ▶ **would you do chemistry research out of magic potions?**

Limitations

- ▶ daemon must run as root to isolate builds
 - ▶ WIP: have daemon rely on **user name spaces** (Linux 3.8+)
- ▶ non-deterministic build systems
 - ▶ must be identified & **fixed upstream**
 - ▶ WIP thanks to <http://reproducible.debian.net>
- ▶ non-free software unavailable in Guix
 - ▶ would you do chemistry research out of magic potions?
 - ▶ **reproducible research demands free software**

Summary

- ▶ Guix allows **cluster users** to reproduce environments

Summary

- ▶ Guix allows **cluster users** to reproduce environments
- ▶ it provides **the source of software environments**, not just the bits

Summary

- ▶ Guix allows **cluster users** to reproduce environments
- ▶ it provides **the source of software environments**, not just the bits
- ▶ **composability, transparency, and hackability** of software stacks are key to reproducible research



<http://gnu.org/software/guix/>

Copyright © 2010, 2012, 2013, 2014, 2015 Ludovic Courtès ludo@gnu.org.
Copyright © 2015 Ricardo Wurmus ricardo.wurmus@mdc-berlin.de.

GNU Guix logo, GFDL, <http://gnu.org/s/guix/graphics>
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 3.0** License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-sa/3.0/> 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** 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 <http://www.gnu.org/licenses/gfdl.html>.

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