

A Tale of Nix and Nickel

YOW! Lambda Jam

Yann Hamdaoui

May 5, 2021

YWEAG

Introduction

A cautionary tale

Once upon a time...


A cautionary tale

There was a brilliant princess




A cautionary tale

One night, she goes to the Kingdom University Ball




A cautionary tale

Where she falls in love with a mysterious girl




A cautionary tale

But when midnight strikes, the girl suddenly runs!





A cautionary tale

She could only get a photo of her shoe




Mercifully, there is the DeepShoe paper!



A cautionary tale

She downloads the repository, which requires Python



A cautionary tale

Well, Python2, actually

```
$python main.py  
Failed to import the shoe  
module...
```

```
$install python2
```

A cautionary tale

And the python bindings to opencv as well

```
$python2 main.py  
Failed to import the  
opencv module
```

```
$pip2 install opencv-python
```

With opencv v1.14

```
$python2 main.py
Undefined symbol:
    "_ZNreadImgRaw12"
```

```
$pack version opencv
package: OpenCV
version: 2.0
```

A cautionary tale

She moves heaven and earth

```
$uninstall opencv --force  
$install opencv-0.92  
No candidate found.
```

To find a solution

```
$curl https...opencv-0.92  
$./install.sh  
Syntax error:...
```

Without success

```
$git clone ...opencv.git .  
$git checkout rel-0.92  
$make install
```

```
$python2 main.py  
Segmentation fault
```

A cautionary tale

The end



Reproducibility

works on my machine ^_^(ツ)_ノ

Reproducibility

works on my machine ↗_(ツ)_↗

Reproducibility

Reproducibility

works on my machine ↗_(ツ)_↗

Reproducibility

1. Concrete and widespread


Reproducibility

works on my machine ↗_(ツ)_↗


Reproducibility

1. Concrete and widespread
2. Not addressed by mainstream tools


The problem



Looks familiar?



Looks familiar?



Functional approach to reproducibility




Nix: the (pure) functional package manager

works on my machine _(ツ)_/



works ~~on my machine~~ everywhere

What if the princess had access to a version packaged by Nix?



Steps

Steps

1. [Describe](#) a package and its dependencies in full

Describing

```
1 {python2WithOpenCV, opencv, stdenv}:
2 stdenv.mkDerivation rec {
3     pname = "gh-from-shoe";
4     version = "2021-04-30";
5
6     buildInputs = [ python2WithOpenCV opencv ];
7
8     installPhase = ''
9         mkdir -p $out/bin
10        cp ${./main.py} $out/bin/gh-from-shoe
11    '';
12};
```

gh-from-shoe/default.nix

Describing

```
1 Derive(  
2   [("out", "/nix/store/qya..-gh-from-shoe", "", "")] ,  
3   [  
4     ("/nix/store/ae4..-python-2-7-10.drv",  
5      ["out"]),  
6     ("/nix/store/78f..-opencv-1-14.drv",  
7      ["out"]),  
8     ...  
9     ["/nix/store/9kr..-default-builder.sh"],  
10    "x86_64-linux",  
11    ...
```

gh-from-shoe-1-0.drv (generated)

Steps

1. Describe a package and its dependencies in full
2. **Build it in isolation**

Building


```
gh-from-shoe-1.0$ nix build
```

1. Pull and build dependencies
(opencv-1-14, python-2-7-10, ...)
2. Create an isolated environment.
3. Run the builder.


Steps

1. Describe a package and its dependencies in full
2. Build it in isolation
3. Put the result in the **store**

/nix/store (read-only)



/nix/store (read-only)



/bin/gh-from-shoe

Steps

1. Describe a package and its dependencies in full
2. Build it in isolation
3. Put the result in the store
4. Profit: find love!

Perks

- Reproducible
- Declarative
- Complete dependencies
- Fearless upgrades: atomic upgrades and rollbacks

Usage examples

- [Nix](#): package management
- [NixOS](#): declarative system configuration
- [Nix shell](#): project-specific environments
- [NixOps](#): Nix-based cloud deployment

Constraints

- Steep learning curve
- Everything has to be "Nixified"


Purely functional package management

Nix	Pure functional programming
Read-only store	Immutability
Hash addressing + sharing	Hash consing
Cleaning	Garbage collection
Reproducibility	Referential transparency

Nix expressions

Epilogue

The princess found love and now wants to use Nix for her own projects



Package as a function

Building a package should be a [pure](#) function: use a functional programming language!

Nix expressions

The Nix language

JSON + λ (higher-order functions)

Nix expressions

```
1 {python2WithOpenCV, opencv, stdenv}:
2 stdenv.mkDerivation rec {
3     pname = "gh-from-shoe-rust";
4     version = "2021-04-30";
5
6     buildInputs = [ python2WithOpenCV opencv ];
7
8     installPhase = ''
9         mkdir -p $out/bin
10        cp ${./myscript.py} $out/bin/myscript
11    '';
12 };
```

gh-from-shoe/default.nix

Derivation: Nix object code

```
1 Derive(
2   [("out", "/nix/store/qya...-gh-from-shoe", "", "")] ,
3   [
4     ("/nix/store/ae4...-python-2-7-10.drv",
5      ["out"]),
6     ("/nix/store/78f...-opencv-1-14.drv",
7      ["out"]),
8     ...
9     ["/nix/store/9kr...-default-builder.sh"],
10    "x86_64-linux",
11    ...
```

/nix/store/27az7...gh-from-shoe-1-0.drv

State of affairs

Nix expressions outgrew their initial scope.

In the wild

- Various user-defined abstractions
- Object systems (kind of): overriding
- A module system: NixOS
- Non-trivial algorithms (e.g. topological sort)

All of this without types!



Nickel

A new take

- Gradual typing
- Run-time contracts
- Recursive records merge system
- Stand-alone language: free us from YAML templating! (Terraform, Kubernetes, build systems, etc.)

A teaser: contract

```
1  let Port = ...
2
3  let Service = {
4      name | doc "Service name"
5          | Str,
6
7      openPorts | doc "Open ports (firewall)"
8          | List #Port
9          | default = [] ,
10     ...
11 }
```

contracts.ncl

A teaser: configuration

```
1  let portToUrl : Str -> Num -> Str =
2      fun host port => ...  in
3
4  {
5      name = "nginx",
6      openPorts = [80, 443],
7      server = "localhost",
8      urls = lists.map
9          (portToUrl server)
10         openPorts,
11     }
12 | #Service
```

nginx.ncl

A teaser: result

```
1  {
2      "name": "nginx",
3      "openPorts": [
4          80,
5          443
6      ],
7      "server": "localhost",
8      "urls": [
9          "http://localhost",
10         "https://localhost"
11     ]
12 }
```

nginx.json

Untyped code

By default, code is **untyped**:

- Terminating & fixed inputs
- Contracts for validation
- JSON interop

Example

```
1 services = [  
2     "init",  
3     {name = "firewall", bin = "/bin/firewall"},  
4     {name = "service", repo = "github.com/johndoe/dns-service"}  
5 ]
```

Heterogeneous values

Typed code

Library code is **statically typed**:

- Triggered by *annotations*
- Scoped
- Type-inference

Example

```
1 map : forall a b. (a -> b) -> List a -> List b
2   = fun f list =>
3     if list == [] then []
4     else
5       let head = lists.head list in
6       let tail = lists.tail list in
7       [f head] @ map f tail
```

Statically typed map

Interaction typed/untyped

Problem

Untyped code can sneak in ill-typed parameters

Example

```
1 let add : Num -> Num -> Num
2     = fun x y => x + y in
3 add "a" 0
```

```
let add : Num -> Num -> Num = fun x y => x + y
                                         ^
```

This expression has type Str, expected Num

Contracts, the invisible glue

Typed code is protected by run-time casts, or *contracts*.

```
1 let safeNum = fun value =>
2     if builtins.isNum value then value
3     else panic! in
4
5 let addSafe = fun x y =>
6     let safeX = safeNum x in
7     let safeY = safeNum y in
8     safeNum (add safeX safeY)
```

Generated code for add (simplified)

Contracts, the invisible glue

```
error: Blame error: contract broken by the caller.
```

```
:1:1
1  Num -> Num -> Num
   --- expected type of the argument provided by the caller
```

```
repl-input-0:2:5
2  add "a" 0
   --- evaluated to this expression
```

```
<unkown> (generated by evaluation):1:1
1  "a"
   --- evaluated to this value
```

note:

```
repl-input-0:1:11
1  let add : Num -> Num -> Num = fun x y => x + y in
   ^^^^^^^^^^^^^^^^^ bound here
```

First-class contracts

```
1  let Url = let pattern = "[-a-zA-Z0-9@:...]" in
2      fun label value =>
3          if builtins.isStr value then
4              if strings.isMatch value pattern then
5                  value
6              else
7                  contracts.blame
8                      (contracts.tag "invalid URL" label)
9          else
10             contracts.blame
11                 (contracts.tag "not a string" label) in
12
13 let mkUrls | {url: #Url, pattern: Str} -> List #Url = ...
```

First-class contracts

```
1  Derivation | doc "A Nix package, in Nickel" = {
2      name | Str,
3      buildInputs | List #NixPackage,
4  },
5
6  NixPackage | doc "Interchange format" = {
7      package | Str,
8      input | Str
9          | default = "nixpkgs",
10     _type = "package",
11 }
```

First-class contracts

Perks

- Can check arbitrary properties
- Composable
- Allow safe typed/untyped interactions
- Built-in error reporting

Limits

- Run-time cost
- Untriggered code paths

Conclusion

Summary

- **Reproducibility** is a concrete and hard problem. Nix helps.
- Nix expressions have shortcomings. We started the **Nickel** language to overcome them.
- There is a design space for alternative type systems. **Gradual typing** and **first-class contracts** is an exciting combo explored in Nickel.

Configuration languages are a worthy area of research.

The 1st Workshop on Configuration Languages

Website <https://2021.splashcon.org/home/conflang-2021>

Deadline Friday 6 August 2021

Duration 1 day

Event October 2021, at SPLASH 2021

The end

Nickel <https://github.com/tweag/nickel/>

Nix <https://nixos.org/>

Tweag's blog <https://www.tweag.io/blog>

Contact

- yann.hamdaoui@tweag.io
- hello@tweag.io

