

# Package Managers, through a lens or two

Jake Trevor

## Disclaimers:

- Not my main line of work
  - There are some holes
- Not a PL talk - just PL themed

# Package managers are Ubiquitous

How many can you think of?

# Package managers are Ubiquitous

Some are familiar:

- utilities from the system PM (apt, brew, yum...)
- language specific dependencies (cargo, npm, pip, etc.)

Some are a little more exotic:

- ready-made actions in your CI (github actions)
- docker images on GHCR
- games on your iPad's app store
- minecraft mods via a mod manager (other games available)
- music from spotify

# Hygiene factors

- Despite how ubiquitous PMs are, no one spends any time thinking about them
- ... because if you are thinking about your package manager, something has gone wrong
- A good package manager is a *hygiene factor* - you notice it only by its absence

No one notices a clean bathroom

Similarly, no one ever wants to think about **pip**

## What should you take out of this talk?

- How can we characterise good behaviour for package managers?
  - both informally
  - and formally
- One potential idea for how to formalise good behaviour
- Maybe some learning about lenses

# What is a package manager?

There are three characteristic problems package managers solve:

1. **Install** packages
2. Handle **dependencies**
3. (or perhaps, 2.a) Manage **versions**

Almost no one gets #1 wrong

But the rest is a lottery!

**What is a good package manager?**

**How should a well behaved package manager behave?**

**Idea: Lenses**

## Lenses

A lens  $\mathcal{L} \Gamma \Delta$  is a structure on two types  $\Gamma, \Delta$

$\Gamma$  is the *source* type

$\Delta$  is the *target* type - it's a *view* on the source

# Lenses

A lens is just a pair of functions:

$$\text{Get} : \Gamma \rightarrow \Delta \quad \text{Put} : \Delta \times \Gamma \rightarrow \Gamma$$

## Lens Laws:

A lens is *well behaved* if it obeys the lens laws:

$$\forall (i : \Gamma) (x : \Delta),$$

$$\text{Get (Put } x \text{ } i) = x$$

$$\text{Put (Get } i) i = i$$

## Lens Laws:

A lens is *very well behaved* if it is also idempotent:

$$\text{put-put}$$
$$\text{Put } x (\text{Put } x i) = \text{Put } x i$$

## Lenses?

Morally, this seems like it might be a good model ...

- We have some notion of a source type ( $\Gamma$ ) (the remote package repository)
- We want a *view* on that source type ( $\Delta$ ) (our local installation of packages)

but the devil is in the details

## Package lenses:

Let's say that we have:

- a type  $P$  of package names
- and for each package  $p : P$ , a body of type  $B$

A package manager is then a collection of lenses

## Package lenses:

The source type is a map of names to bodies:

$$\Gamma = P \rightarrow B$$

For each subset of package names  $P' \subseteq P$ , we have a target type, which is a map from names to bodies:

$$\Delta_{P'} = P' \rightarrow B$$

**Can we mechanise this?**

## Package lenses:

In lean, we define subsets by their *characteristic predicate*:

e.g.  $hP' = \text{is } p \text{ in the subset } P' \subseteq P?$

```
-- given h : P -> Prop  
type P' = { p : P // h p }  
type P' = Subtype h
```

The *package manager* itself can be understood as a function from the choice of packages to the corresponding lens:

$$\text{PLens} : (hP' : P \rightarrow \text{Prop}) \rightarrow \mathcal{L} \Gamma \Delta_{P'}$$

## Package lenses:

In this formalism, since  $P' <: P$ ,  $\Gamma <: \Delta_{P'}$

So `Get` is just a cast from  $\Gamma$  to  $\Delta_{P'}$

```
get : (P -> Prop) -> (Subtype h -> Prop)
get := fun γ p => γ p
```

## Package lenses:

Put *overwrites*  $\Gamma_P$  with values from  $\Delta_{P'}$  whenever they exist:

```
put : ((Subtype h -> Prop) × (P -> Prop)) -> (P -> Prop)
put := fun (δ, γ) p =>
  if h p
  then δ p
  else γ p
```

## Package lenses

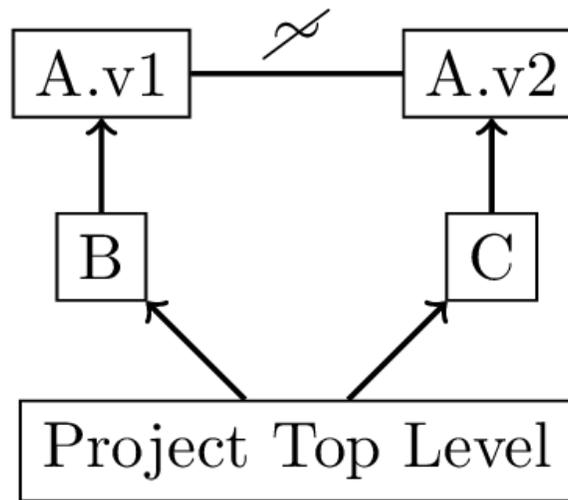
- Models the intuitive behaviour of installing packages
- All the lenses are *very well behaved*
  - I have lean proofs to back this up

# Dependencies

## The Problem: Dependency war

- Let's say that I use package B
- and B depends on A version 1
- Let's say I also want to use C
- which depends on A version 2
- Some combinations of packages are *impossible*
- How should a well behaved package manager deal with this?

# Dependency war



## Dependencies: Internal

- Sometimes, a package *completely encapsulates* its dependencies
  - e.g. if there's no way to tell that A depends on B from the A's interface alone
- In this case, the dependencies are **internal** to the package

→ We can use the package without knowing about its dependencies

# Dependencies: External

Packages can be *leaky*

- What if types from a dependency appear in the public interface of the package?
- Then the user of the package is *exposed* to the dependencies
- These are called **external** dependencies

## To be clear:

- External dependencies make solving dependency war very hard

## My view:

- Dependency war between external dependencies in packages is *not* a technical problem
- It's fundamental to the implementation of the packages

Some sets of packages are *incompatible*  
Nothing a package manager can do will solve this

## Aside: Can we fix it?

If you allow importing multiple versions of the same package, e.g. by naming them differently ...

then you can expose both interfaces without a namespace collision

but doing so is contrary to the *spirit* of versioning

→ Versions provide different ways of doing the same thing  
∴ it's weird to use two versions of the same thing

**However**, dependency war between internal dependencies should not be an issue

**Solution: Dependency Locality**

## Local Dependencies:

The key idea is that we can load dependencies in a *localised* way:

1. When we load a dependency, it must not have side effects on some global environment
2. To support this, packages need to support *dependency injection*

Python does not have this facility; packages are referred to by name, as objects existing on the path

This is not a feature of the package manager, but of the underlying language/system

It's not pip's fault

**Pip sucks because python sucks**

## Module systems:

There's an extensive body of work on module systems, such as:

- **MixML** for ML (Rossberg & Dreyer)
- **backpack** for haskell (Kilpatrick & al.)
- **II** - a language agnostic module calculus; unpublished (Florisson & Mycroft)

These are much more complicated than I have time to go into here, and also a little out of scope

# Versions

## Solution 1: new version, new name

- e.g. `pkg-A-v1` and `pkg-A-v2` are different packages
- This works very nicely with the theory we presented!
- It's also completely useless
  - in that it fails to model the actual problem
- *Dependency war* is just *namespace collision*
- If we really want to think about versions, we need to treat them as more than just 'new names'

## How should versions work?

- Versioning software is a solved problem  
→ Just do what git does!
- What are the semantics of git?

## VCS semantics

We have:

- $S$  - a type of the *items* being versioned
- $H\langle S \rangle$  - the type of histories
- Hash - an associated type of *hashes*

## VCS semantics

We don't care about the content of the types  
so long as we can define the following two functions:

$$\begin{aligned} \text{Commit} & : H\langle S \rangle \times \text{Hash} \times S \rightarrow H\langle S \rangle \times \text{Hash} \\ \text{Checkout} & : H\langle S \rangle \times \text{Hash} \rightarrow S \end{aligned}$$

## VCS semantics

This might look familiar ...

$$\Gamma = H\langle S \rangle \times \text{Hash}$$

**Get = Checkout**

$$\Delta = S$$

**Put = Commit**

## VCS semantics

There's an infinite family of version control lenses (VCLs), constructed by the function:

$$\text{VCL} : (S : \text{Type}) \rightarrow \mathcal{L} (H \langle S \rangle \times \text{Hash}) S$$

## VCS: Is it a lens?

- If I checkout what I just committed, I should get the same thing back:

$$\text{Checkout } (\text{Commit } x \ i) = x$$

- If I check something out, and try to commit it, nothing will happen

→ since Git will not (by default) allow empty commits

$$\text{Commit } (\text{Checkout } i) \ i = i$$

- By the same fact, committing the same thing twice does nothing

$$\text{Commit } x \ (\text{Commit } x \ i) = \text{Commit } x \ i$$

**So, in fact, this kind of abstract VCS is a lens, and it's *very well behaved!***

- Formalising this turned out to be hard
- I have a (very ugly) lean definition for VCLs
- But the proofs are not finished

## What do we want, morally?

- Let's say that we have a *history of bodies* for each package
- For a package  $p : P$ , let  $V_p$  be the set of its versions  
→ i.e. hashes in its history
- Then for a set of packages  $P'$ , a *choice of versions*  $V_{P'}$  is an object with type

$$V_{P'} = (p : P') \rightarrow V_p$$

## What do we want, morally?

Our source type should be a *map of histories* of package bodies:

$$\Gamma = P \rightarrow H\langle B \rangle$$

As before, for a choice of packages  $P'$  versions  $V_{P'}$ , we have a target type - a map of names to bodies:

$$\Delta_{V_{P'}} = P' \rightarrow B$$

We want a family of lenses, constructed by a function:

$$\text{mkPkgMan} = P' \times V_{P'} \rightarrow \mathcal{L} (P \rightarrow H\langle B \rangle) (P' \rightarrow B)$$

## What do we have?

We can construct a package lens from names to histories

$$\text{PLens ref} : \mathcal{L} (P \rightarrow H\langle B \rangle) (\text{Subtype ref} \rightarrow H\langle B \rangle)$$

We can construct a VCL on those histories:

$$\text{VCL } B : \mathcal{L} (H\langle B \rangle \times \text{Hash}) B$$

**Here are two other lenses:**

$$ID : \mathcal{L} A A$$

$$\text{Concentrate} : \mathcal{L} ((B \rightarrow C) \times (B \rightarrow D)) (B \rightarrow (C \times D))$$

These are both vwb. (proofs in lean)

We will need these later

## Constructions on lenses:

Lenses can be tensored (parallel composition):

$$\_ \otimes \_ : \mathcal{L} A B \rightarrow \mathcal{L} C D \rightarrow \mathcal{L} (A \times C) (B \times D)$$

Composed (sequential composition):

$$\_ \circ \_ : \mathcal{L} A B \rightarrow \mathcal{L} B C \rightarrow \mathcal{L} A C$$

We can split a lens over a function in the second argument:

$$\_ \text{split} \_ : \mathcal{L} A (B \rightarrow C) \rightarrow \mathcal{L} C D \rightarrow \mathcal{L} A (B \rightarrow D)$$

These constructions are *behaviour preseving* - their output is vwb. when the input lenses are.

**We can use these lenses and constructions to define a formalism for *versioned package lenses***

- This is a little complex, so we will go through it slowly
- It's OK if you don't follow this completely

## What we get: Versioned Package Lenses

$$\text{VPL ref} : (P \rightarrow \text{Prop}) \rightarrow \mathcal{L} ((P \rightarrow H\langle B \rangle) \times V_{P'}) (P' \rightarrow B)$$

$$\text{VPL ref} = ((\text{PLens ref} \otimes \text{ID}) \circ \text{Concentrate}) \text{ split } (\text{VCL } B)$$

$$\text{VPL ref} = ((\text{PLens ref} \otimes \text{ID}) \circ \text{Concentrate}) \text{ split } (\text{VCL } B)$$

$$(\text{PLens ref} \otimes \text{ID}) : \mathcal{L} ((P \rightarrow H\langle B \rangle) \times V_{P'}) \quad ((P' \rightarrow H\langle B \rangle) \times V_{P'})$$

$$\text{VPL ref} = ((\text{PLens ref} \otimes \text{ID}) \circ \text{Concentrate}) \text{ split } (\text{VCL } B)$$

Recall that  $V_{P'}$  is just  $(p : P') \rightarrow V_p$

$$(\text{PLens ref} \otimes \text{ID}) \circ \text{Concentrate}$$

$$: \mathcal{L} ((P \rightarrow H \langle B \rangle) \times V_{P'}) \quad (P' \rightarrow (H \langle B \rangle \times V_p))$$

$$\text{VPL ref} = ((\text{PLens ref} \otimes \text{ID}) \circ \text{Concentrate}) \text{ split } (\text{VCL } B)$$

$H\langle B \rangle \times V_p$  is exactly the source type of our version control lens

$$\begin{aligned} & ((\text{PLens ref} \otimes \text{ID}) \circ \text{Concentrate}) \text{ split } (\text{VCL } B) \\ & : \mathcal{L} ((P \rightarrow H\langle B \rangle) \times V_{P'}) \quad (P' \rightarrow B) \end{aligned}$$

## What we get: Versioned Package Lenses

$$\begin{aligned} \text{VPL ref} &= ((\text{PLens ref} \otimes \text{ID}) \circ \text{Concentrate})\text{split} (\text{VCL } B) \\ &: (P \rightarrow \text{Prop}) \rightarrow \mathcal{L} ((P \rightarrow H\langle B \rangle) \times V_{P'}) (P' \rightarrow B) \end{aligned}$$

Since this is formed only by vwb. lenses and behaviour preserving constructions, the result is also vwb.

**If you didn't follow, the key take away is:**

**What we get isn't quite what we wanted**

$$\mathcal{L} ((P \rightarrow H\langle B \rangle) \times V_{P'}) (P' \rightarrow B)$$

The target type is correct

$$\Delta = (P' \rightarrow B)$$

but the source type is a little different:

Wanted:

$$\Gamma = P \rightarrow H\langle B \rangle$$

Got:

$$\Gamma = (P \rightarrow H\langle B \rangle) \times V_{P'}$$

# The difference

The source type includes our choice of versions  $V_P$

Intuitively: Put can change our choice of versions

- From a certain point of view, this makes sense
  - The version you just pushed is not the one you checked out
- From another, it does not:
  - If versions can change with a push, why not the set of packages?
  - i.e. what's the story for 'pushing a new package'?

It may well be that a different approach might yield a more pleasing result

wherein versions are names are treated more similarly

Food for thought

## The wrinkles

- Since the formalisation of VCL is not done, I don't have a verified proof that this is all OK just yet
- At present, while 'tensoring' is morally right, it's not sufficient
- The type of Hashes depends on the value of history, so
- We need a dependent tensor - where the second argument can depend on the first

## Take care!

Notice that for us, the package registry is a *map of histories*:

$$P \rightarrow H \langle B \rangle$$

But Nix instead has a *history of maps*:

$$H \langle P \rightarrow B \rangle$$

nixpkgs (the central repository) is a versioned map of names to recipes

Which is the right way around?

## **My view: Nix does it wrong**

With Nix' system, we refer to packages in a particular version of nixpkgs

**Problem 1:** bumping a version might not change your installation at all

→ i.e. if the packages were not changed in that commit

**Problem 2:** Versioned versions

## Versioned Versions

Nix allows you to refer to different nixpkgs versions to allow you to refer to specific builds of a given package

But this can be cumbersome

∴ different versions of a package are often distributed as *different packages*

e.g. php82 and php83 are different packages.

php82 might exist in 2 commits

You can refer to *different versions of php82*

→ Ridiculous

## Conclusions

- We can model package managers with lenses
  - Lots of the nuance is nicely captured
- Particularly, simple unversioned package managers
- For good dependency behaviour, the underlying system needs a good module system
  - Even if we get everything else right, we can still get trapped by non-local dependencies
- This formalism has some weaknesses, space for development
  - A more equal treatment of versions and package names would be nice

## Future work?

- What does proper behaviour for **shared dependencies** look like?
  - What's the next best thing after full locality?
  - What about version constraints?
- **Extra features**
  - e.g. declarative/imperative interfaces, reversibility and uninstallation
    - some of this is already captured by the models I've already presented
- **Security**
  - Push should fail sometimes - is that worth modelling?

Some references:

- [1] Florisson & Mycroft, *Towards a Theory of Packages unpublished draft* (2016)
- [2] Rossberg & Dreyer, *Mixin' Up the ML Module System* (2013)
- [3] Kilpatrick & al, *Backpack: Retrofitting Haskell with Interfaces* (2014)

## **Over cakes:**

What's your most hated package manager?

Can you think of bad behaviour that isn't captured by this model?