

---

# Enforcing Type-Safe Linking using Inter-Package Relationships

M. Dogguy\* — S. Glondu\* — S. Le Gall\*\* — S. Zacchioli\*<sup>1</sup>

\* *Laboratoire PPS, UMR 7126  
Université Paris Diderot-Paris 7  
Case 7014*

*F-75205 Paris Cedex 13*

`{dogguy,glondu,zack}@pps.jussieu.fr`

\*\* *OCamlCore S.A.R.L.*

`sylvain.le-gall@ocamlcore.com`

---

**Abstract.** *Strongly-typed languages rely on link-time checks to ensure that type safety is not violated at the borders of compilation units. Such checks entail very fine-grained dependencies among compilation units, which are at odds with the implicit assumption of backward compatibility that is relied upon by common library packaging techniques adopted by FOSS (Free and Open Source Software) package-based distributions. As a consequence, package managers are often unable to prevent users to install a set of libraries which cannot be linked together.*

*We discuss how to guarantee link-time compatibility using inter-package relationships. In doing so, we take into account real-life maintainability problems such as support for automatic package rebuild and manageability of ABI (Application Binary Interface) strings by humans.*

*We present the `dh_ocaml` implementation of the proposed solution, which is currently in use in the Debian distribution to safely deploy more than 300 OCaml-related packages.*

**Keywords:** *static typing, separate compilation, linking, free software, FOSS, distribution, package, dependency, OCaml, Debian*

---

1. Partially supported by the European Community FP7, MANCOOSI project, grant agreement n. 214898

## Résumé long

Les langages fortement typés reposent sur des vérifications lors de l'édition de liens afin de garantir que la sûreté du typage reste respectée entre les différentes unités de compilation. Ces vérifications sont très strictes, et incompatibles avec l'usage, dans les distributions de logiciels libres utilisant des paquets, qui est de supposer que les bibliothèques sont rétrocompatibles. Ainsi, les gestionnaires de paquets sont souvent incapables d'empêcher un utilisateur d'installer un ensemble incohérent de bibliothèques.

Une des causes d'incohérence peut apparaître au niveau du typage. Dans le cas de langages fortement typés, ce problème peut se manifester quand deux unités de compilation ont été générées avec deux versions différentes du compilateur. Cette source d'incohérence est détectée par l'éditeur de liens qui effectue les tests nécessaires pour garantir un ensemble cohérent ; mais cela peut persister quand il s'agit d'unités séparées.

Pour des langages de bas niveau, tels que le langage C, l'éditeur de liens effectue quelques vérifications lors de la phase finale:

- chaque symbole est résolvable,
- aucun symbole n'est défini plusieurs fois.

Malheureusement, ces vérifications ne sont pas suffisantes pour garantir la sûreté du typage quand il s'agit de langages fortement typés tels que OCaml et Haskell. Aux vérifications habituelles, il faudra rajouter des vérifications de types entre unités de compilation pour garantir la cohérence de l'ensemble. Ces informations de types sont habituellement représentées dans un format court et exportées dans l'unité de compilation produite. Pour l'éditeur de liens, un symbole devient donc un nom associé à l'information de type. Ceci permet de faire des vérifications plus fines. Ayant cette information pour tous les symboles disponibles, il est possible de définir une version pour l'ensemble des unités de compilation, obtenant ainsi une version de compatibilité binaire pour la bibliothèque, aussi appelée ABI.

En C, une grande attention est portée sur la compatibilité binaire, gérée manuellement, lors d'un changement d'interface. Il est ainsi pos-

sible de rajouter des fonctionnalités dans de nouvelles versions tout en conservant une certaine compatibilité binaire. Par conséquent, une borne inférieure sur la version d'une bibliothèque suffit souvent à assurer le bon fonctionnement d'un paquet en dépendant. Au contraire, en OCaml, la compatibilité binaire est gérée automatiquement par le compilateur et le moindre changement sera considéré comme incompatible par les outils. Ceci amène à repenser le schéma de dépendances utilisé dans les paquets afin de refléter le changement d'ABI pour que ce soit visible à l'utilisateur dès la tentative d'installation.

Le but de la solution proposée est de détecter les incompatibilités d'ABI en analysant les dépendances des différents paquets concernés. Nous établissons donc des critères pour évaluer notre solution et la comparer aux solutions existantes. Les critères établis sont:

- dépendances correctes;
- inférence des dépendances;
- stabilité des dépendances;
- légèreté des dépendances.

Notre solution, appelée "ABI approximation", consiste à calculer un hash cryptographique par paquet à partir des unités de compilation qu'il fournit. Le résultat calculé et la liste des unités fournies sont ensuite enregistrés dans un registre global, et pourront ainsi être utilisés par le mécanisme d'inférence des dépendances pour d'autres paquets. Cette solution a été implémentée en Perl, pour éviter des dépendances circulaires, et vérifie tous les critères établis. Elle a été testée sur plus de 150 paquets sources, produisant près de 353 paquets binaires et plus de 2500 modules. Elle a ensuite été déployée dans la distribution Debian avec succès. Le gain est présent aussi bien pour les utilisateurs, qui ne peuvent donc plus installer un ensemble incohérent de paquets, mais aussi pour les mainteneurs de paquets, qui peuvent désormais identifier facilement l'impact de l'introduction d'une nouvelle version de bibliothèque. Une solution similaire a également été adoptée par les mainteneurs des paquets Haskell.

Dans le cadre de ce travail, nous avons été amenés à écrire de nouveaux outils pour pouvoir lire les informations dont on a besoin à partir des binaires code-octet et des greffons OCaml. Ces outils ont été

soumis aux auteurs de OCaml qui les ont acceptés (et feront partie de la prochaine version).

## 1. Introduction

Type safety is a tricky business, even more so when separate compilation is a desired feature. In the world of system-level languages and linkers [14]—such as the C language and the widespread GNU linker—very few checks are performed at the final linking stage; a bit of formalization will help in understanding them. Given a set of *compilation units*  $\{u_1, \dots, u_n\}$  to be linked together, the linker essentially checks a form of referential integrity, i.e. that all symbols needed by involved compilation units are actually available within the set. We call *requirements* of a compilation unit  $\mathcal{R}(u)$  the set of required symbols and *application binary interface* (ABI) of a compilation unit  $\mathcal{A}(u)$  the set of provided symbols.

The linker notion of “linkability” can now be grasped as follows:<sup>1</sup>

**Definition 1** (Linkability).  $link\{u_1, \dots, u_n\} = \checkmark$  iff:

- 1)  $\bigcup_i \mathcal{R}(u_i) \subseteq \bigcup_i \mathcal{A}(u_i)$
- 2)  $\forall i, j, \quad i \neq j \rightarrow \mathcal{A}(u_i) \cap \mathcal{A}(u_j) = \emptyset$

The first condition is the linker’s view of referential integrity: all referenced symbols must be provided by linked objects; the second condition avoids multiple definitions.<sup>2</sup> No matter the expressivity of the type system, such a linking discipline cannot help in enforcing type safety *across compilation units*, if not relying on name mangling hacks [15] or delegating it to external whole program verification techniques [8].

Moving to the world of functional, statically typed programming languages, such as OCaml and Haskell, link-time checks get more thorough mainly because types come into play. Not only cross-module type compatibility is challenging to verify *per se* [1, 11], but also technical guarantees that ABIs do not change between compile-time of individual units and link-time should be type-aware. The solution adopted

---

1. We do not strive for completeness here, we cover only the linker checks that will help in comparing with the strongly typed language world. For the same reason, we do not distinguish between static and dynamic linking.

2. Actually, in some corner cases, the C linker can allow multiple definitions, but that is uninteresting for our purposes.

Listing 1: foo.ml	Listing 2: bar.ml	Listing 3: main.ml
<pre>let hello () =   Printf.printf     "Hello!\n"</pre>	<pre>let hello () =   Foo.hello ()</pre>	<pre>let _ =   Bar.hello ()</pre>

Figure 1: sample OCaml compilation units

by OCaml is, for each compilation unit, to expose two sets of *module names*, associating each name to a cryptographic hash or *checksum* that captures the type information of that module.

**Example 1.** *Let us consider the OCaml sources of Figure 1.*

*After (bytecode) compilation of bar.ml—which in turn needs a compiled version of foo.ml—the resulting compilation unit contains the following “assumptions”:*

```
$ ocamlc -c foo.ml bar.ml
$ ocamlobjinfo bar.cmo
Unit name: Bar
Interfaces imported:
  807ecd3a1538992580464c03462c9964  Printf
  da00042bb934260afe41d004bc91fe2e  Foo
  9e3404342379641955461e6944482508  Bar
```

*where we can see that bar.cmo exports an ABI consisting of the interface Bar with a specific MD5 checksum and that it requires some other checksum-tagged interfaces. Among them we can spot Foo, provided by foo.ml, and Printf, provided by the OCaml standard library which is linked in by default.*

*If the ABI of foo.cmo changes between the compile-time of bar.cmo and the final link-time, the user will incur in the sadly well-known “inconsistent assumptions” error:*

```
$ ocamlc -c foo.ml bar.ml
$ echo "let gotcha () = ()" >> foo.ml
$ ocamlc -c foo.ml
$ ocamlc foo.cmo bar.cmo main.ml
Files bar.cmo and foo.cmo make inconsistent
```

The additional type-aware checks performed by OCaml already fit in the given simple formalization. It is enough to consider both  $\mathcal{R}(u)$  and  $\mathcal{A}(u)$  to be sets of *pairs*  $\langle m, c \rangle$ , where  $m$  is the name of an OCaml module and  $c$  is its associated checksum. The only additional property checked by the OCaml linker is that, given a set of compilation units, the mapping between module names and checksums is a function, i.e. that a module is not associated with more than one checksum.

Simplifying the characteristics of the standard library *stdlib*, the failure of Example 1 can now be explained as follows:

$$\begin{aligned} \mathcal{R}(\text{stdlib}) &= \emptyset \\ \mathcal{R}(\text{foo.cmo}) &= \{\langle \text{Printf}, 807\text{ec} \dots \rangle\} \\ \mathcal{R}(\text{bar.cmo}) &= \{\langle \text{Printf}, 807\text{ec} \dots \rangle, \\ &\quad \langle \text{Foo}, \text{da}000 \dots \rangle\} \\ \mathcal{A}(\text{stdlib}) &= \{\langle \text{Printf}, 807\text{ec} \dots \rangle\} \\ \mathcal{A}(\text{foo.cmo}) &= \{\langle \text{Foo}, \text{a}4293 \dots \rangle\} \\ \mathcal{A}(\text{bar.cmo}) &= \{\langle \text{Bar}, 9\text{e}340 \dots \rangle\} \end{aligned}$$

$$\begin{aligned} \bigcup_{u \in M} \mathcal{R}(u) \setminus \bigcup_{u \in M} \mathcal{A}(u) &= \{\langle \text{Foo}, \text{da}000 \dots \rangle\} \\ \text{where } M &= \{\text{stdlib}, \text{foo.cmo}, \text{bar.cmo}\} \end{aligned}$$

In order to satisfy Definition 1, the next to last equation was supposed to be the empty set, whereas it contains the pair  $\langle \text{Foo}, \text{da}000 \dots \rangle$ .

The ability of the linker to detect this kind of unsound assumptions comes at a cost: ABIs change more frequently than in the C language case. Indeed, with system-level linking, the ABI of a given unit can inhibit linkability only by removing symbols from it. Backward ABI compatibility—which of course does not imply type safety—can be retained only when adding new symbols, which is an unsurprisingly common practice in the evolution of C libraries.

With type-aware linking, ABIs break at each change in a module,<sup>3</sup> no matter if it is an addition, a removal, or a simple (type) change, be-

3. Here, we do not distinguish among assumptions on interfaces and assumptions on implementations, these details are postponed to Section 6.

cause each such modification will change the checksum of the module. While C libraries offer (type-unsafe) backward binary compatibility by default, OCaml libraries have the converse default: they break binary compatibility at every single change.

Unfortunately for most users of languages such as OCaml and Haskell, packaging systems and techniques used in mainstream FOSS (Free and Open Source Software) distributions have been designed with *implicit backward compatibility* for libraries in mind. Dependencies on library packages are indeed usually expressed according to the following schema:

```
Package : my-app  
Depends : libfoo1 (>= 1.2.3)
```

Such a schema implicitly assumes that future versions of the `foo` library will either be binary backward compatible or change the package name altogether, for instance switching to `libfoo2`.

The advantage of the backward compatibility assumption is that inconsistent software installations, where installed libraries lack the needed objects to be linkable, are detected at the dependency level and can hence be spotted by package managers. Given that there is no reliable mapping between library versions and ABIs (new versions *can* change ABIs, but are not forced to), we observe that with the OCaml linking discipline the advantage of spotting linkability errors at the dependency level is gone.

As most OCaml users on distributions such as Debian<sup>4</sup> have experienced in the recent past, it indeed frequently happens that upgrades involving OCaml libraries might temporarily leave the build toolchain in an inconsistent state,<sup>5</sup> requiring users to: recompile dependent libraries by hand, rollback the upgrade (if possible), or simply wait for fixed packages (of all involved libraries) from the distribution.

In this paper we show how to make the packaging system aware of complex and frequently changing ABIs used by languages such as

---

4. <http://www.debian.org>

5. For example, <http://bugs.debian.org/cgi-bin/bugreport.cgi?bug=238727> (retrieved October 2009); similar reports are very frequent during transitions from one OCaml version to the next, since all libraries need to be rebuilt.



OCaml and Haskell, hence shielding users from the installation of non-linkable libraries. Attempting to do so will amount to trying to install a package whose dependencies cannot be satisfied. We achieve that result by coalescing ABIs in a single human-readable checksum—called *ABI approximation*—which is then reified as a “virtual” package known to the packaging system, *de facto* closing the gap between coarse-grained inter-package relationships and fine-grained inter-unit linking assumptions.

### 1.1. Paper structure

The next section introduces some concepts of package-based FOSS distributions that will be needed throughout the paper. Section 3 states the goals and requirements for the solution we look for; Section 4 reviews the state of the art against them. The proposed solution is described and analyzed in Section 5; its actual implementation—in the context of the Debian distribution and for the OCaml language—is discussed in Section 6.

From now on, we will focus on the OCaml language and its packaging in Debian, but reasoning and choices are general enough to be ported to similar languages (with strong typing and inspectable linking assumptions) and distributions (based on binary packages). In fact, for what concerns Debian, the discussed concepts have recently been adopted by the Haskell library stack as well.

## 2. Packaging basics

### 2.1. Packages

In most FOSS distributions—and, more generally, in component-based systems [16]—software components are managed as *binary packages* [2] which define the granularity at which users can add or remove software.

A binary package is essentially a bundle of data (executables, documentation, media, etc.) to be deployed on the file system. Additionally, binary packages are described by meta-information which include com-

plex *inter-package relationships* that describe the static requirements to run properly on a target system. Requirements are expressed in terms of other binary packages, possibly with restrictions on the desired versions. Both positive requirements (*dependencies*) and negative requirements (*conflicts*) are supported by most packaging systems.

**Example 2.** An excerpt of the meta-data of the *ocamlnet* library in Debian currently reads:

```
1 Package: libocamlnet-ocaml-dev
2 Version: 2.2.9-3
3 Depends: ocaml-nox-3.10.2, ocaml-findlib,
4   libocamlnet-ocaml (= 2.2.9-3),
5   libpcr-ocaml-dev (>= 5.11.1),
6   libcryptgps-ocaml-dev (>= 0.2.1)
7 Provides: libequeue-ocaml-dev,
8   libnetclient-ocaml-dev,
9   librpc-ocaml-dev
10 Conflicts: libequeue-ocaml-dev (<< 2.2.3-1),
11   libnetclient-ocaml-dev (<< 2.2.3-1),
12   librpc-ocaml-dev (<< 2.2.3-1)
13 Description: OCaml application-level Internet
14   libraries - core libraries
```

In this short but representative example we can recognize the distinguishing features of inter-package relationships [7]:

- dependencies on other libraries, tools, and the OCaml interpreter (line 3) which can either be versioned (e.g. `libocamlnet-ocaml`—the runtime part of the package shown) or non-versioned (e.g. `ocaml-findlib`);

- conflicts with some libraries (line 10), in the example with superseded packages now included into `libocamlnet-ocaml-dev` itself;

- *virtual packages* (line 7), i.e. the ability to declare *features* provided by the owning “real” package so that others can depend on (or conflict with) feature names. In the example `libocamlnet-ocaml-dev` provides old library names such as `libequeue-ocaml-dev`; dependencies on it by other packages can be satisfied by installing `libocamlnet-ocaml-dev`. The other typical use case of vir-

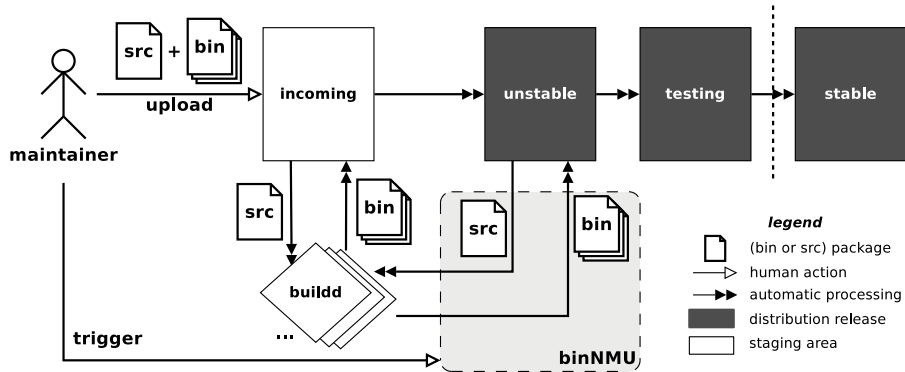


Figure 2: package (auto-)building

tual packages is to add an indirection layer over system services (e.g. mail-transport-agent), which stays in between packages *providing* the service (e.g. postfix or exim4) and packages needing it (e.g. cron).

## 2.2. Auto-building and binNMUs

*Source packages* are a different type of packages which contain the source code. Compiling them produces binary packages.<sup>6</sup>

Source packages are usually manipulated by human package *maintainers* contributing to specific distributions. The mapping between source and binary packages is one to many; for instance, the `ocamlnet` source package produces the package shown in Example 2 together with eight other binary packages. Also, there exists a many-to-many *build-dependency* relation among source packages and the binary packages needed to build them.

The natural work-flow of packages is rather complex [9], the part that mostly concerns us is sketched in Figure 2. Maintainers upload source packages to a package queue, together with the corresponding binary packages; considered as a whole, such an upload is said to be a *sourceful upload*. Given that maintainers usually only own a machine

6. According to established conventions, we use “packages” to refer to binary packages, and explicitly “source packages” for the others.

for a single architecture, *auto-builders* (or *buildds*) pick up the uploaded source package and rebuild the corresponding binary packages for each architecture supported by the distribution (about a dozen, in the case of Debian).

The only exception to the autobuilding scheme are *architecture independent* packages (or `arch:all`) that once built on any given architecture, can then be installed on all architectures, for instance because they contain portable interpreted code or bytecode. Therefore `arch:all` packages do not get rebuilt at all, after the first build. All obtained binary packages flow to the “unstable” *suite* which is ready to be used by developers and testers in preparation of a future stable release.

In certain circumstances (e.g. library transitions), binary packages can be rebuilt without intervention on their source packages. This is needed to fix builds performed in malformed environments (e.g. buggy compiler), to ensure that a package still builds against a new version of a library it depends upon, or precisely to trigger a rebuild against a new (shared) library when the package built against the old library is not binary compatible with the new one. To that end, maintainers can trigger rebuilds by requesting a *binNMU*, a historic and unfortunate name standing for “binary Non-Maintainer Upload”. For our needs, the important aspect of binNMUs is that during these no changes can be performed to the *source* package.

### 2.3. *Dependency inference*

But how can a binary package change if its source package remains unchanged?

The contents of the package (e.g. binary files) can of course change as the toolchain used to build it changes (compiler, libraries, etc.), but it turns out that package meta-information such as dependencies can also change across rebuilds (and hence differ across architectures). In fact, maintainers can exploit the mechanisms of *dependency interpolation* and *dependency inference* to ease the tedious task of maintaining dependency information.

**Example 3.** *The dependency information of the `mldonkey-server` package read as follows in the `mldonkey` source package:*

```
Package: mldonkey-server  
Depends: adduser, mime-support, ucf,  
${shlibs:Depends}, ${misc:Depends}
```

*whereas in the resulting binary package they read:*

```
Package: mldonkey-server  
Depends: adduser, mime-support, ucf, libbz2-1.0,  
libc6 (>= 2.3.2), libjpeg62,  
libfreetype6 (>= 2.2.1), libgcc1 (>= 1:4.1.1),  
zlib1g (>= 1:1.1.4),  
libgd2-noxpm (>= 2.0.36~rc1~dfsg) |  
libgd2-xpm (>= 2.0.36~rc1~dfsg),  
libpng12-0 (>= 1.2.13-4), libstdc++6 (>= 4.2.1),  
debconf | debconf-2.0
```

What is going on in Example 3 is that interpolated variables—expressed in `${this:form}`—get expanded during build according to dependency inference. While maintainers can use custom expansion mechanisms, the common way to expand variables is to rely upon specific *registries*. For instance, the `${shlibs:Depends}` variable expands to all dependencies that can be inferred by C shared library relationships [3].

To that end, each C shared library installed on the system provides a `shlibs` file which is in essence a record of tuples  $\langle lib\ n\ ,\ k\ n\ ,\ si\ n\ di\ t \rangle$  which contributes to the C shared library registry. For instance, the `libpng12-0` package provides a record  $\langle libpng12, libpng12-0, (>= 1.2.13-4) \rangle$ . When the `mldonkey-server` package gets built, all executables it ships are inspected for relationships with C shared libraries (using tools like `objdump`). Since the `/usr/bin/mldonkey-server` executable needs a C library called `libpng12` to be loaded, a lookup for that library name in the registry provides the matching dependency `libpng12-0 (>= 1.2.13-4)` which is then added to the expansion of `${shlibs:Depends}`.

### 3. Requirements

The goal of the solution we are looking for is to detect type-aware linking incompatibilities at the dependency level. As there are several possible solutions to that problem, we define in this section a set of contingent requirements, mostly inherited from the context.

The first obvious requirement is goal fulfillment.

**Requirement 1** (Dependency soundness). *All binary packages shipping compilation units should satisfy the dependency soundness property.*

**Property 1** (Dependency soundness). <sup>7</sup> *A package shipping OCaml compilation units  $\{u_1, \dots, u_n\}$  has sound dependencies with respect to a repository if and only if for all healthy installation  $I$ , it holds that  $\text{link}(\{u \in I \mid I \text{ ships } u\}) = \checkmark$*

Intuitively, a package has sound dependencies if all compilation units shipped by its (transitive) dependencies can be linked together. Dependency soundness is clearly not a local property—i.e. it cannot be decided by considering a package in isolation—but this is not surprising given that our notion of linking is geared towards obtaining a self-contained executable.

**Requirement 2** (Dependency inference). *Given a source package  $p$  and its build-dependencies  $bd = \{d_1, \dots, d_n\}$ , the dependencies that ensure soundness on all binary packages obtained by building  $p$ , should be inferrable on the basis of  $bd$  and its (transitive) dependencies.*

Acknowledging that OCaml ABIs change very frequently, Requirement 2 asks for a mechanism of dependency inference (see Section 2) that relieves maintainers from the error-prone task of maintaining by hand the needed dependencies.

---

7. The notions of *healthy installation* and *generated subrepository*, noted  $R|_{\Pi}$ , are easily explained. An installation is a set of installed packages, it is healthy if all dependencies are satisfied and no conflicts occur among installed packages; a generated subrepository is a repository subset obtained as the closure of the “depends on” relation starting from a given set of packages. Formal details can be found in [12].

**Requirement 3** (binNMU-safety). *If a package has sound dependencies, performing a binNMU on it should not make its dependencies unsound.*

This requirement not only means that inferred dependencies should not be tied to source package version, but also that we cannot rely on sourceful uploads. Hence, solutions requiring changes to be incorporated in source packages at each rebuild (including most notably binNMU-driven rebuilds) are not suitable for the task.

All dependencies ensuring soundness should be recomputed during build, on the basis of the build environment. While binNMU-safety mimics the homonymous accepted best-practice of packaging, it is also a real need to reduce maintenance burden. Since transitions from one compiler version to the next require rebuilding all packages, and given that distributions like Debian contain nowadays more than a hundred OCaml-related source packages, the ability to do transitions via binNMs is crucial for the maintainability of the whole stack in package-based binary distributions.

**Requirement 4** (Light dependencies). *All inter-package relationships needed to ensure dependency soundness should be terse, readable, human-manageable.*

Albeit admittedly very informal, this requirement is meant to ensure that package dependencies remain manageable by humans. Even if there is no clear definition for that, discussions within the Debian project between OCaml maintainers, the release team<sup>8</sup>, and users have given evidence that solutions like a simple dump of all MD5 requirements as dependencies would not be acceptable. Indeed, during release management, quality assurance, or simply broken dependency analysis within a package manager, such dependencies will be too many (one for each ABI requirement) and too unfriendly. Requirement 4 attempts to convey this and similar desiderata.

---

```

Package: libfoo-ocaml-dev
Version: 1

Package: libbar-ocaml-dev
Version: 1
Depends: libfoo-ocaml-dev (>= 1)
Build-Depends: libfoo-ocaml-dev (>= 1)

Package: main
Version: 1
Build-Depends: libbar-ocaml-dev (>= 1)

```

---

Figure 3: Past Debian dependency scheme for OCaml-related packages

#### 4. Related work

Until the adoption of the solution presented in this paper, the Debian distribution was using the natural dependency scheme [10] depicted in Figure 3, where the snippets of Figure 1 are considered shipped by packages with corresponding names.

The drawbacks of that solution are twofold: no dependency inference (the maintainer writes dependencies by hand), no dependency soundness (there is no guarantee that future versions will preserve ABI compatibility). Note that stricter version predicates, such as e.g. `libfoo-ocaml-dev (>= 1)`, `libfoo-ocaml-dev (<< 2)`, (as it was done for Haskell packages in Debian until very recently) will not provide soundness either, because version numbers are densely ordered.

In 2005, a helper tool called `dh_ocaml` was proposed [17] to ease the burden of maintaining OCaml-related dependencies,<sup>9</sup> mimicking the architecture of the `shlibs` registry for C shared libraries. Dependencies generated by (old-style) `dh_ocaml` follow the scheme of Figure 3. The obtained solution hence also fulfills Requirement 2, but not yet Requirement 1 (arguably the most important one).

---

8. The release team is a group of Debian developers who coordinate transitions, goals and deadline for the next stable release.

9. <http://bugs.debian.org/cgi-bin/bugreport.cgi?bug=328422>, retrieved October 2009



```

$ rpm -qRp ocaml-ocamlnet*.rpm
ocaml(Arg) = b6513be035dc9c8a458c189cd8841700
ocaml(Array) = 9c9fa5f11e2d6992c427dde4d1168489
ocaml(Bigarray) = fc2b6c88ffd318b9f111abe46ba99902
ocaml(Buffer) = 23af67395823b652b807c4ae0b581211
... snipped 67 more ocaml(*) deps

$ rpm --provides -qp ocaml-ocamlnet*.rpm
ocaml(Equeue) = 329e036bb2778b249d6763d22407af19
ocaml(Ftp_client) =
  d36822b105eacef219a2b6e0331ba34b
ocaml(Ftp_data_endpoint) =
  f279805dc3b7ced5d8554f92e287c889
ocaml(Generate) = 418dedddda65b04bdc4d0c6e9fb918d4
... snipped 110 more ocaml(*) virtual packages

```

Figure 4: Sample dependencies (on the left) and provided features (on the right) according to the Fedora solution, for the `ocamlnet` library.

The Fedora and Red Hat distributions, and after them other RPM-based distributions such as OpenSUSE and Mandriva, adopt a different solution [4]. At the end of the build, they automatically inspect all OCaml bytecode objects  $u_i$  shipped by each binary package and, for each pair  $\langle m, c \rangle \in \mathcal{R}(u_i)$ , they add a dependency on a virtual package in the OCaml namespace which has the full MD5 checksum as its version; pairs in  $\mathcal{A}(u_i)$  are similarly handled to generate the list of provided virtual packages. No intermediate registry is used.

Fedora’s solution provides soundness (up to a technical detail discussed in Section 6) and implicit dependency inference, however it fails lightness. As it can be seen in Figure 4, it bloats package lists with one virtual package/dependency per exported/required OCaml module; that can also have a performance hit on dependency resolution, in distributions the size of Debian. Such solution ultimately exposes users and members of relevant teams to checksum strings which are both meaningless for humans as well as being overly long.

Table 1: Review of linkability enforcement solutions

<b>Solution</b>	<b>Req. 1 soundness</b>	<b>Req. 2 inference</b>	<b>Req. 3 binNMU</b>	<b>Req. 4 lightness</b>
past Debian <i>status quo</i>	✗	✗	✓	✓
old-style dh_ocaml	✗	✓	✓	✓
current Fedora guidelines	✓	✓	✓	✗
ABI evolution tracking	✓	✓	✗	✓
<b>ABI approximation</b>	✓	✓	✓	✓

The last alternative solution we are aware of is *ABI evolution tracking* [17]: it establishes an injective mapping between package ABIs and integers. The integer ideally represents ABI “version” and is used to establish human-friendly virtual package names (e.g. `libpcre-ocaml-dev-1`) that provide soundness. Unfortunately, this solution defeats binNMU-safety. Indeed, to avoid ABI version clashes the mapping should be either maintained by hand (similarly to what happens with C symbols tracking [5], where unexpected changes at build time trigger build failures), or obtained via automatic monotonic increase of ABI versions. For the latter, we would need to preserve somewhere a history of past ABI numbers, which has no place to stay during binNMUs: network is not accessible to avoid non-deterministic builds and source packages cannot be changed.

Table 1 reviews the state of the art against the requirements of Section 3. The last line corresponds to the solution we propose, described in next section.

## 5. ABI approximation

To overcome the limitations of the discussed approaches, we have devised a solution based on the idea of computing a single approxima-

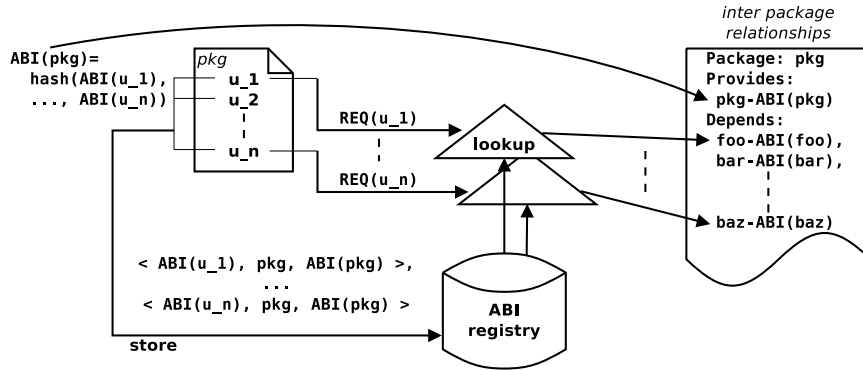


Figure 5: ABI approximation: architecture

tion of the ABIs of all compilation units shipped by a given package. We call this solution *ABI approximation*; its architecture is sketched in Figure 5.

Consider a (binary) package  $pkg$  shipping compilation units  $u_1, \dots, u_n$  (at the top-left of Figure 5). Its ABI approximation—noted  $\mathcal{A}(k)$ —is obtained as a cryptographic hash of all ABI pairs of all compilation units; intuitively:  $\mathcal{A}(k) = s(\mathcal{A}(u_1) \ \mathcal{A}(u_n))$ , where  $s$  is a concatenation operator which is used to incrementally add material on which the hash is taken.<sup>10</sup>

The obtained ABI approximation is then used in two different ways, both contributing to form the inter-package relationships of  $pkg$ . For what concerns the interface exposed by  $pkg$ , we directly add a virtual package obtained by juxtaposing  $pkg$ 's name with the ABI approximation itself. All packages relying on the set of ABIs  $\mathcal{A}(u_1), \dots, \mathcal{A}(u_n)$  will depend on such virtual package.

Let us now assume that all dependencies of  $pkg$  itself already have computed ABI approximations and that they are available at build-time, which is a sound assumption according to our linking discipline. To ensure that we can compute the dependencies on the correct virtual packages—i.e. the virtual packages corresponding to packages able to

10. A more formal description would require more details on the hashing function, on module naming conventions, etc. They are omitted for the sake of conciseness. Section 6 nevertheless provides some details about the current implementation.

satisfy the linking assumptions  $\mathcal{R}(u_1), \dots, \mathcal{R}(u_n)$ —we use an *ABI registry* that stores information about which packages provide which ABI.

The registry is populated by tuples (simplified in Figure 5)  $\langle m, c, k, \mathcal{A}(k) \rangle$  where  $\langle m, c \rangle \in \mathcal{R}(u)$  for some  $u$  shipped by  $pkg$ . Intuitively, each module  $m$  exposed by some compilation unit with checksum  $c$  has a tuple in the registry which relates it with a specific package name and overall ABI approximation. Technically, the tuples are computed at build-time and the registry is constructed incrementally by files shipped by  $pkg$  itself, similarly to what happens with the `shlibs` mechanism.

Using the ABI registry, we can now compute the dependency entries of  $pkg$  which will ensure its linkability. For each  $\langle m, c \rangle \in \mathcal{R}(u_i)$ , we look up  $\langle m, c \rangle$  in the registry to get a pair of package name and ABI approximation; note that  $\langle m, c \rangle$  is expected to be a primary key in any possible configuration. By juxtaposing the obtained package name and approximation as before, we can now emit the appropriate dependency snippets.

**Example 4.** *The resulting dependencies for the `ocamlnet` package, which we have already seen before ABI approximation in Example 2, is as follows.*

```
Package: libocamlnet-ocaml-dev
Version: 2.2.9-7
Depends: ocaml-findlib,
  libcryptgps-ocaml-dev-139d7,
  libocamlnet-ocaml-3rxe6,
  libpcre-ocaml-dev-kh2c0,
  ocaml-nox-3.11.1
Provides: libequeue-ocaml-dev,
  libnetclient-ocaml-dev,
  librpc-ocaml-dev,
  libocamlnet-ocaml-dev-3rxe6
Description: OCaml application-level Internet
  libraries - core libraries
```

*We can notice that the package depends on the external `cryptgps` and `pcre` libraries with specific ABI approximations, and that it exports an ABI approximation itself.*

*By comparison with the Fedora dependencies for the same package (see Figure 4), we observe that the huge amount of provided and expected module checksums is here hidden in the registry and resolved at build time: the package interface just exposes a single extra virtual package and one dependency for each external library needed.*

To review ABI approximation against the requirements of Section 3, we start by observing that soundness (Requirement 1) is granted up to the risk of clashes of different sets of ABIs to the same approximated ABI. Literature on the birthday paradox [13] gives an approximation on the number of different, incompatible, versions of a library to have collision with probability :

$$P(n) \approx \frac{n^2}{2 \times 10^6}$$

where  $n$  denotes the ABI approximation space size. In our practical implementation  $n = 1103$  (see next section), so we reach the negligible probability of 1 % of collision with 1103 versions.

Dependency inference (Requirement 2) is trivially provided by the ABI registry, according to the lookup mechanism described above.

binNMU-safety (Requirement 3) is granted too as all the state, needed to compute both the inferred dependencies and the exposed approximated ABI, is not stored in the source package (and hence does not need sourceful uploads to be changed); rather, the state is kept in the registry which is contributed to by build-dependencies of the package being built.

Finally, we argue that dependency lightness (Requirement 4) is provided too. While ABI approximated strings are not really meaningful to humans, they have the benefit of being short (only 5 alphanumeric characters in the current implementation). Additionally the number of such strings that are needed for a given binary package is bounded: one for the exported interface and one for each package dependency, with the latter set replacing dependencies which would have been needed anyhow, substituting ABI approximations for version numbers.

The most notable limitation of ABI approximation is the “instability” of approximations with respect to backward *compatible* changes. This

is not surprising, given that the solution has been designed precisely to cope with frequent backward *incompatible* changes. Still, sometimes, backward compatible changes can happen in a given package, for example by adding a new module to a given library without touching any other existing module. In our formalism that translates to adding a new couple  $\langle m, c \rangle$  to some  $\mathcal{A}(u_i)$ ; trivially, property 1 is unaffected by such addition. Nevertheless, the resulting ABI approximation will change, since the content on which the hash is computed grows. Practically, the problem is non-existent, given the low frequency of such changes and given that—thanks to binNMU-safety—the affected packages can be fixed by just scheduling their recompilation. We observe that this issue can be fixed only by allowing to export ABI approximations whose sizes are linear with respect to the number of modules shipped by a given package; in essence, that would mean mimicking Fedora’s solution, getting back most of its disadvantages.

## 6. Implementation

ABI approximation has been implemented in the context of the Debian distribution to manage inter-package dependencies of all OCaml-related packages in the distribution. At the time of writing, that amounts to 158 source packages that build-depend on OCaml, producing 353 binary packages, and providing interfaces for 2502 OCaml modules.

Technically, the implementation is based on `dh_ocaml`, a helper tool meant to be compatible with the Debhelper packaging helper suite [6]. Some of its inner workings depends on how OCaml libraries are split into binary packages, which we briefly highlight below.

Compiling a source package results in one or several binary packages. For example, library source packages usually produce (at least) two binary packages:

- 1) a *runtime* package, containing all the objects that might be needed at runtime by software linked to the library;
- 2) a *development* package, containing all the objects that might be needed to compile software that uses the library.

In the OCaml world, the runtime package typically contains stubs to C libraries (the so called “bindings”) that are dynamically loaded when running pure bytecode executables (.so), and the development package contains interfaces in source (.mli or .ml) and compiled (.cmi) form, along with compiled module objects (.cmo, .cma, .cmx, .o, .cmxa, .a), META files and (mostly ocaml-doc-generated) documentation. A library might have no runtime package at all if it is meant to be always statically linked, which happens quite often with pure-OCaml libraries. However, when a library can be dynamically loaded by a program that supports plugins, the runtime package can also provide .cma, .cmxs, and META files.

A library can be further split into several sub-components. This is usually left at the package maintainer’s discretion, and choices can vary between distributions. A typical reason for splitting a library is optional dependencies that might be needed only by some compilation units. Each resulting component might then have its runtime and development packages. Finally, there are also non-library packages shipping OCaml binaries.

Overall, dh\_ocaml distinguishes three classes of (binary) packages: runtime, development, and others. We henceforth consider as a *library* a development package together with its runtime package, if any. For example, the ocamlnet source package produces the following binary packages:

- libocamlnet-ocaml: core library (runtime);
- libocamlnet-ocaml-dev: core library (development);
- libocamlnet-ocaml-doc: documentation;
- libocamlnet-ocaml-bin: miscellaneous tools;
- libocamlnet-gtk2-ocaml-dev: GTK2 layer (development);
- libocamlnet-ssl-ocaml: SSL layer (runtime);
- libocamlnet-ssl-ocaml-dev: SSL layer (development);
- libnethttpd-ocaml-dev: HTTP daemon libraries (development);
- libapache2-mod-ocamlnet: Apache2 module.

Among them, four OCaml libraries can be recognized:

- libocamlnet-ocaml/-dev,
- libocamlnet-gtk2-ocaml-dev (without runtime),
- libocamlnet-ssl-ocaml/-dev, and
- libnethttpd-ocaml-dev (without runtime).

To inspect OCaml compilation units (a required ability to implement ABI approximation), the OCaml standard distribution provides some useful tools:

`ocamlobjinfo` reads assumptions of bytecode objects (`.cmi`, `.cmo` and `.cma` files), as shown in Example 1;

`ocamldumpapprox` reads assumptions of native code objects (`.cmx`, `.cmxa` files).

While OCaml bytecode objects make assumptions only on interfaces, native code objects may also make assumptions on implementations, due to code inlining across compilation units. These assumptions are stored in `.cmx` files, which might hence be useful even if the module they represent are included in a `.cmxa` file.

Most notably, the above tools lack the ability to inspect bytecode executables and native plugins (`.cmxs`). Such abilities are much needed though: bytecode executables might depend on external C stubs and native plugins are special cases of libraries and should hence be considered as such. To overcome these limitations, we have implemented `ocamlbyteinfo` and `ocamlplugininfo`; they are currently shipped as part of the `ocaml` source package and relied upon by `dh_ocaml`.

The solution described in section 5 is then implemented in Debian using two tools:

`ocaml-md5sums` manages the ABI registry. This tool is meant to be generic (not Debian-specific). It inspects a set of object files by calling the most appropriate among the above tools, uses their output to compute their dependencies and its approximated ABI. It considers a library as a whole (runtime and development together);

`dh_ocaml` is a frontend to `ocaml-md5sums`; it looks for installed objects and binaries, passes them to `ocaml-md5sums` and fills-



in substitution variables corresponding to dependency inference variables; it distinguishes the three categories of packages:

- a development package provides a hashed virtual package name, and depends on its own runtime package (that explains the dependency on `libocamlnet-ocaml-3rx6` in Example 4), and possibly on other development packages;
- a runtime package provides a hashed virtual package name, and depends only on other runtime packages;
- any other package depends only on runtime packages.

The actual ABI approximation is computed by considering the union of two different sets of  $\mathcal{A}()$  for each compilation unit: one for assumptions over interfaces, one for assumptions over implementations (no matter where they come from: plugin, object, or binary). On a simple and canonical textual representation of the resulting set, a MD5 checksum is then computed and then taken modulo a fixed hash space. Currently, such hash space relies on 5 alphanumeric, lowercase, plain-ASCII characters, accounting for ( ) different ABI approximations. As discussed in the previous section, that choice gives a negligible clash probability, without sacrificing lightness.

We notice that alternative solutions, such as Fedora’s, are currently not considering assumption on implementations, most likely because the original tools in the OCaml distributions were not able to inspect them. As a consequence they only provide soundness up to implementation incompatibility; unfortunately such incompatibilities are more likely to occur than interface incompatibility, given that not all programmers seem to be aware of (or care about) cross-unit code inlining. This being a minor drawback easily fixable (now that we have provided the missing tools), we also observe that fixing it naïvely would mean doubling the already large number of depends/provides in Figure 4.

**Example 5.** *Here is an excerpt of dependencies from the `ocamlnet` source package (i.e. before expansion by `dh_ocaml`):*

```
Package: libocamlnet-ocaml-dev
Provides: libequeue-ocaml-dev,
libnetclient-ocaml-dev,
librpc-ocaml-dev,
```

```
dh_ocaml:Provides
Depends: ocaml-findlib, dh_ocaml:Depends,
dh_shlibs:Depends, dh_misc:Depends
```

*The result of the substitution performed by `dh_ocaml` in the corresponding binary package can be seen in Example 4.*

Both `ocaml-md5sums` and `dh_ocaml` are implemented in Perl to escape (build-)dependency cycles among the package shipping them and the `ocaml` package itself. Both tools are part of the `dh-ocaml` source package,<sup>11</sup> which is distributed under the terms of the GNU General Public License (version 3) and available from its Git repository.<sup>12</sup>

So far, we have effectively deployed this solution in about 50 source packages (out of 100 library packages). This deployment revealed not only dormant packaging errors such as missing dependencies, but also unrelated errors like libraries re-exporting modules they do not own. The most common example of that has been the embedding of the `Unix` module into third-party libraries. The remaining packages, and most notably libraries, are being ported to `dh_ocaml` at the time of writing, their ported versions are all expected to be shipped with the next stable release of Debian (codename “Squeeze”).

The currently known `dh_ocaml`-specific limitations are as follows:

- it is theoretically possible to change C stubs in a binary incompatible way while keeping the same interface; it is therefore possible to install an outdated version of C stubs while satisfying all dependencies of a bytecode executable. This limitation is inherent to OCaml;

- architecture-independent packages that depend on libraries with runtime cannot currently benefit from this solution, because the ABI approximation string depends on the architecture (given assumptions on implementations are not present on architectures lacking the `ocamlopt` native code compiler). This limitation can be solved by computing two separate ABIs: a native code and a bytecode one; the latter should be the same on all architectures.

---

11. <http://packages.debian.org/sid/dh-ocaml>

12. <http://git.debian.org/?p=pkg-ocaml-maint/packages/dh-ocaml.git>

## 7. Conclusions and future work

In this paper we have given some insights on the gap existing between fine-grained dependencies that type-aware linkers consider to enforce type-safety at compilation units borders and the coarse-grained dependencies that can be expressed among packages in FOSS binary distributions. That gap persisting, package managers will not be able to defend users from installing libraries of strongly typed languages that can not be linked together due to annoying “inconsistent assumptions” link-time errors.

We have introduced a solution called *ABI approximation* that guarantees link-time compatibility as long as inter-package relationships are satisfied. ABI approximation satisfies a set of natural requirements inherited from the context: dependency soundness, inference, and lightness, as well as binNMU-safety.

ABI approximation has been implemented and deployed in the Debian distribution through the `dh_ocaml` helper, aiming at managing the dependencies of more than 300 OCaml-related binary packages. On the premises of this work, also the whole stack of Haskell packages in Debian is currently migrating to an ABI approximation solution.

Future work consists in some technical fixes whose need was revealed by using `dh_ocaml`, such as the observation that computing a single ABI for bytecode and native-code is incompatible with architecture-independent packages. Also, we are considering designing a mechanism for computing automatically OCaml-related ABIs for C stubs whose purpose is interacting with the OCaml runtime. That can be done at C stub compilation-time (where and when the interface of the matching OCaml objects are known) and then stored in some section of the resulting object.

## References

- [1] A.W. Appel and D.B. MacQueen. Separate compilation for Standard ML. In *PLDI 1994: Proceedings of the ACM SIGPLAN 1994 conference on Programming language design and implementation*, pages 13–23. ACM, 1994.

- [2] Roberto Di Cosmo, Paulo Trezentos, and Stefano Zacchiroli. Package upgrades in FOSS distributions: Details and challenges. In *HotSWup'08: Proceedings of First ACM Workshop on Hot Topics in Software Upgrades*. ACM, 2008.
- [3] Ulrich Drepper. How to write shared libraries. <http://people.redhat.com/drepper/dsohowto.pdf>, 2002. Revision August 2006.
- [4] Fedora Project. OCaml packaging guidelines. <http://fedoraproject.org/wiki/Packaging/OCaml>. Retrieved October 2009.
- [5] Raphael Hertzog. dpkg-gensymbols: generate symbols files (shared library dependency information). Manual page dpkg-gensymbols(1), on line at [http://man.cx/dpkg-gensymbols\(1\)](http://man.cx/dpkg-gensymbols(1)), July 2007. Retrieved July 2010.
- [6] Joey Hess. debhelper debian package: helper programs for debian/rules. <http://packages.debian.org/sid/debhelper/>, 2009. Version 7.0.15.
- [7] Ian Jackson and Christian Schwarz. Debian policy manual. <http://www.debian.org/doc/debian-policy/>, 2009.
- [8] S. C. Johnson. Lint, a C program checker. (65), 1978.
- [9] Martin F. Krafft. *The Debian system: concepts and techniques*. Open Source Press, 2005.
- [10] Sylvain Le Gall, Sven Luther, Samuel Mimram, Ralf Treinen, and Stefano Zacchiroli. Debian OCaml packaging policy. [http://pkg-ocaml-maint.alioth.debian.org/ocaml\\_packaging\\_policy.html/](http://pkg-ocaml-maint.alioth.debian.org/ocaml_packaging_policy.html/), 2009.
- [11] Xavier Leroy. Manifest types, modules, and separate compilation. In *POPL'94: Proceedings of the 21st ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pages 109–122. ACM, 1994.
- [12] Fabio Mancinelli, Jaap Boender, Roberto Di Cosmo, Jérôme Vouillon, Berke Durak, Xavier Leroy, and Ralf Treinen. Managing the complexity of large free and open source package-based software distributions. In *ASE 2006: 21st IEEE/ACM International*

- Conference on Automated Software Engineering*, pages 199–208. IEEE CS Press, 2006.
- [13] E.H. McKinney. Generalized birthday problem. *American Mathematical Monthly*, 73(4):385–387, 1966.
  - [14] L. Presser and J.R. White. Linkers and loaders. *Computing Surveys (CSUR)*, 4(3):149–167, 1972.
  - [15] Diomidis Spinellis. Type-safe linkage for variables and functions. *SIGPLAN Notices*, 26(8):74–79, 1991.
  - [16] Clemens Szyperski. *Component Software: Beyond Object-Oriented Programming*. Addison Wesley Professional, 1997.
  - [17] Stefano Zacchiroli. Enforcing OCaml link-time compatibility using Debian dependencies. Debian wiki: <http://wiki.debian.org/Teams/OCamlTaskForce?action=AttachFile&do=get&target=dh-ocaml-design.pdf>, January 2009. Design document, retrieved October 2009.