

On Software Component Co-Installability

Jérôme Vouillon
PPS
University Paris Diderot and CNRS
jerome.vouillon@pps.jussieu.fr

Roberto Di Cosmo
PPS and IRILL
University Paris Diderot and INRIA
roberto@dicosmo.org

ABSTRACT

Modern software systems are built by composing components drawn from large *repositories*, whose size and complexity is increasing at a very fast pace. A fundamental challenge for the maintainability and the scalability of such software systems is the ability to quickly identify components that cannot be installed together: this is the *co-installability* problem, which is related to boolean satisfiability and is known to be algorithmically hard. This paper develops a novel theoretical framework, based on formally certified, semantic preserving graph-theoretic transformations, that allows to associate to each concrete component repository a much smaller one with a simpler structure, but with equivalent co-installability properties. This smaller repository can be represented graphically, giving a concise view of the co-installability issues in the original repository, or used as a basis for various algorithms related to co-installability, like the efficient computation of strong conflicts between components. The proofs contained in this work have been machine checked in Coq.

1. INTRODUCTION

The mainstream adoption of free and open source software (FOSS) has widely popularised component-based software architectures, maintained in a distributed fashion and evolving at a very quick pace. Components are typically made available via a *repository*, and each of these components is equipped with metadata, such as *dependencies* and *conflicts*, used to specify concisely the contexts in which a component can or cannot be installed.

A typical example of the metadata attached to a component, taken from the Debian GNU/Linux distribution, is shown in Example 1: the logical language used for expressing dependencies and conflicts is quite powerful, as it allows conjunctions (symbol ‘,’), disjunctions (symbol ‘|’) and version constraints.

EXAMPLE 1. *The inter-package relationships of postfix, an Internet mail transport agent in the Debian GNU/Linux distribution (<http://www.debian.org>) currently reads:*

```
1 Package: postfix
2 Version: 2.5.5-1.1
3 Depends: libc6 (>= 2.7), libdb4.6, ssl-cert,
4   libsasl2-2, libssl0.9.8 (>= 0.9.8f-5),
5   debconf (>= 0.5) | debconf-2.0,
6   netbase, adduser (>= 3.48), dpkg (>= 1.8),
7   lsb-base (>= 3.0-6)
8 Conflicts: libnss-db (<< 2.2-3), smail,
9   mail-transport-agent, postfix-tls
10 Provides: mail-transport-agent, postfix-tls
```

In most frameworks, determining whether a single component can be installed at all is already an NP-Complete problem, albeit the concrete instances arising in real-world systems, like GNU/Linux distributions, Eclipse plugins or OSGI component repositories, turn out to be tractable [1, 2, 3]. For the maintenance of component repositories, though, more sophisticated analyses are required. This includes identifying for each component the other components in the repository that it absolutely needs [4], and those that it can never be installed with [5].

More generally, a fundamental challenge for component based software maintainability is the study of the problem of *co-installability* of components, that involves identifying and visualising the relevant sets of components that cannot be installed together.

Indeed, from a maintenance point of view, one needs to identify which components cannot be installed together, in order to check whether these incompatibilities are justified or due to erroneous dependencies. The dependency graph is also too rich for end users, which are interested in having a given set of functionalities provided by some components (for instance, they want a word processor and a Web browser) but do not care about the additional components also installed. This is illustrated by the fact that some package managers keep track of these additional components and automatically remove them when they are no more needed.

The sheer size of current mainstream repositories, with tens of thousands of components and hundreds of thousands of relations, makes it completely unfeasible to study such properties of a repository directly: visualising such large graphs is both technically challenging, and of little interest, as one would need to follow recursively a large number of dependency and conflict relations to understand how components relate to one other.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...\$10.00.




Figure 1: Graphical depiction of a repository

In this paper, we develop a novel theoretical framework, based on formally certified, semantic preserving graph-theoretic transformations, that allows to associate to each concrete component repository a much smaller repository with a simpler structure, but with equivalent co-installability properties. This smaller repository can be represented graphically, giving a concise view of the co-installability issues in the original repository, or used as a basis for various algorithms related to co-installability, like the efficient computation of the strong conflicts defined in [5].

We identified several bugs in the Debian distribution using the present work. For instance, the `hardened-servers` package, meant to prevent the installation of unsafe packages, was in conflict with packages `proftpd` et `sendmail` but did not prevent the installation of the corresponding binaries which were in fact in packages `proftpd-base` et `sendmail-bin` (the first two packages were actually transitional packages used to ease the upgrade from a previous release).

The paper is organised as follows: Section 2 recalls the basic notions about packages and dependencies, and overviews the repository transformation developed in the paper, which achieves impressive results on real-world GNU/Linux distributions (Section 3). The technical development follows: repositories are equipped with a partial order in Section 4, put into a flattened form in Sections 5, 6 and 7, simplified by removing irrelevant dependencies and conflicts in Section 8 and quotiented in Section 9. Finally, we show how to draw a simplified graph of the kernel of a repository in Section 11. We discuss related works, and conclude in Section 13.

2. OVERALL APPROACH

While the concrete details may vary from one technology to the other, the core metadata associated to component based systems always allows to express a few fundamental properties: a component, called *package* in the following, may *depend* on a combination of components, expressed as a conjunction of disjunctions of components, and a component may *conflict* with a combination of components, expressed as a conjunction of components.

Extra properties like *provides* (e.g. `postfix-tls` in Example 1), or *versioned constraints* (e.g. `libc6 (>= 2.7)` in Example 1) can be easily preprocessed out [1], so that one can focus on a core dependency system that contains a binary symmetric conflict relation, and a dependency function $D(\pi) = \{\{\pi_1^1, \dots, \pi_{n_1}^1\}, \dots, \{\pi_1^k, \dots, \pi_{n_k}^k\}\}$ that is satisfied when for each i at least one of the π_j^i is installed.

An example repository in this core dependency system is depicted in Figure 1, that also introduces the graphical notation used in the paper: package a has two dependencies, and can be installed only if, first, either package b or package c is installed, and second, package f is installed; package c conflicts with b and f : neither packages b and c , nor packages c and f , can be installed simultaneously; package d depends

on package e , which in turn depends on both f and g .

2.1 Repositories

We follow the notations of [1, 6, 7, 8], that we recall here. A *repository* is a tuple $R = (P, D, C)$ where P is a finite set of *packages*, $D : P \rightarrow \mathcal{P}(\mathcal{P}(P))$ is the *dependency function* (we write $\mathcal{P}(X)$ for the set of subsets of X), and C , a symmetric irreflexive relation over P , is the *conflict relation*. An *installation* I of a repository (P, D, C) is a subset of P . An installation I is *healthy* when the following holds:

- *abundance*: every package has what it needs. Formally, for every package $\pi \in I$, and for every dependency $d \in D(\pi)$, we have $d \cap I \neq \emptyset$.
- *peace*: no two packages conflict, that is, $C \cap (I \times I) = \emptyset$.

We call *dependency* a set of packages d included in $D(\pi)$ for some package $\pi \in P$.

One can give a logical interpretation of the dependency function and the conflict relation. The logical variables are the packages $\pi \in P$. A set of packages $d \in \mathcal{P}(P)$ is interpreted as a disjunction:

$$\llbracket d \rrbracket = \bigvee_{\pi \in d} \pi.$$

A set of set of packages $D(\pi) \in \mathcal{P}(\mathcal{P}(P))$ is interpreted as a conjunction:

$$\llbracket D(\pi) \rrbracket = \bigwedge_{d \in D(\pi)} \llbracket d \rrbracket.$$

A dependency function D is then interpreted as the set of formulas of the shape:

$$\pi \implies \llbracket D(\pi) \rrbracket$$

where π ranges over P , which can also be written:

$$\pi \implies \bigwedge_{1 \leq i \leq n} \bigvee_{1 \leq j \leq m_i} \pi_{i,j}$$

with $D(\pi) = \{d_i \mid 1 \leq i \leq n\}$ and $d_i = \{\pi_{i,j} \mid 1 \leq j \leq m_i\}$. A conflict relation C is interpreted as the set of formulas $\neg \pi \vee \neg \pi^0$ for $(\pi, \pi^0) \in C$. A healthy installation is an assignment which simultaneously satisfies all these formulas.

2.2 Co-Installability

A package π is *installable* in a repository if it is included in a healthy installation I of this repository. A set of packages Π are *co-installable* in a repository if they are all included in some healthy installation I of the repository.

Checking co-installability has been shown equivalent to SAT by taking advantage of the logical interpretation of repositories [1]. However, our experience is that this problem is easy in practice: SAT-solver based tools are currently used routinely to identify non-installable components on repositories that contains dozens of thousands of packages, and hundreds of thousands of dependencies and conflicts.

2.3 Extracting a Co-Installability Kernel

Identifying all the sets of components that cannot be installed together is way more complex: even if we limit ourselves to the simplest case of sets of non co-installable components of size 2 (also known as *strong conflicts*), testing all possible pairs of packages is not a viable option, as the package number is in the tens of thousands. Even using all the

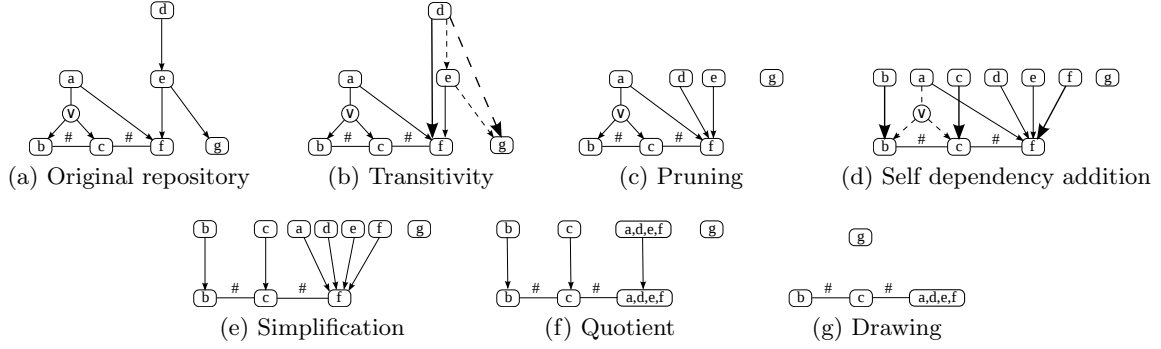


Figure 2: Transformations of a repository (added dependencies are in bold, dotted ones are removed in the next phase)




Figure 3: Repository that blows-up when flattened

optimisations described in [5], the computation takes almost a week on a modern workstation.

In the present work, we lay down the essential theoretical basis and algorithmic insight for tackling this *non co-installability* problem: the fundamental idea is to extract from the component repository a *kernel* repository which is equivalent to the original one, as far as co-installability is concerned, but which turns out in practice to be orders of magnitude smaller, and easily manageable.

The key steps of our approach are now summarised. The effect of each step on the example repository of Figure 1 is shown in Figure 2.

Flattening.

The recursive nature of dependencies is convenient for package developers, as it allows them to describe the dependencies among the different packages very concisely, in a modular way. To study the properties of a repository, though, it is way more convenient to use only a special *flattened* form \widehat{D} of dependency functions that *directly* describes all dependencies of each package: if

$$\widehat{D}(\pi) = \{\{\pi_1^1, \dots, \pi_1^{k_1}\}, \dots, \{\pi_n^1, \dots, \pi_n^{k_n}\}\},$$

then the packages π_i^j are all the packages relevant for installing package π , and only them.

Any dependency function can be converted in this form by a sort of transitive closure that expands the dependencies of each intermediate package, and then converts the result again in a conjunction of disjunctions using distributivity: on our running example, this amounts to adding a dependency from d to f , and one from d to g (Figure 2b).

This transformation has some similarity with the conversion of logical formulae to conjunctive normal form, and is likewise subject to exponential blow-up (see for instance [9]): the dependency function of the repository in Figure 3 has size $3n$, and when expanded, gives rise to a repository whose dependency function has size 2^n .

This is a strong limiting result, but we are only inter-

ested in studying co-installability of packages, so we need not fully maintain the logical equivalence of repositories. In particular, we can prune the expanded dependency function by removing any dependency containing a package with no conflicts without changing the co-installability property. In practice, this suffices to avoid the exponential blow-up. On our running example, this pruning phase removes the dependencies from d to e , from d to g and from e to g , leading to the repository of Figure 2c.

We take a further action to render the repository more homogeneous: we add a self dependency to each package with conflicts. This sort of reflexive closure will be very useful when quotienting the repository, in a later phase. (There is no point in adding self dependencies to other packages as they would be removed by pruning.) We then find it convenient to draw the repository using a two-level structure: on the top, we have all packages; on the bottom, we have packages with conflicts; dependencies connects the top layer to the bottom layer; conflicts are between packages on the bottom layer. On our running example, this leads to the repository of Figure 2d.

The three phases of expansion, pruning and addition of self dependencies can be performed in a single pass, and we thereafter use the term *flattening* to denote all of them.

Elimination of irrelevant dependencies and conflicts.

As a second phase, we identify several classes of dependencies and conflicts that are irrelevant as far as co-installability is concerned, and remove them. In Figure 4, we can see some interesting examples (more are given in Section 8), in all of which the disjunctive dependency connecting package a to packages b and c can be dropped:

- (a) if some branches of a disjunction are forced by a stronger dependency, all other branches can be dropped;
- (b) a package with no conflict can be added to any installation, so dependencies on such a package are always satisfiable and all disjunctive dependencies containing it can be simplified out (this is the pruning defined above);
- (c) if a package has a disjunctive dependency containing a package (here, b) that conflicts only with other packages in this dependency (here, c), this dependency is always satisfiable and can be dropped: either a package conflicting with package b is installed, or package b can be installed; in both cases, the dependency can be satisfied.




Figure 4: Some irrelevant dependencies



Figure 5: A redundant conflict

In Figure 5, the conflict between packages a and b is implied by the conflict between packages c and d and can be dropped.

Proving the soundness of such simplifications is far from trivial: in general, one has to rely on a peculiar structure of the repository to justify that a dependency can be removed, but removing a dependency may modify this very structure. Therefore, a suitable invariant has to be found, that allows to remove most, if not all, irrelevant dependencies.

On our example repository, the disjunctive dependency from package a to packages b and c , corresponding to Figure 4c, can be removed, yielding the repository of Figure 2e.

Quotienting equivalent packages.

It is now quite evident looking at Figure 2e, that packages a , d , e and f are, as far as co-installability is concerned, really equivalent: they share the very same set of dependencies (notice that this fact is easily detectable on the graph thanks to the self dependency of package f introduced during the flattening phase).

As many packages in a repository share the same behaviour with respect to co-installability, it is useful to quotient the final repository, identifying these packages; this step contributes greatly in reducing the size of the repository, as can be seen on our running example in Figure 2f.

After removing self dependencies, one gets the final repository of Figure 2g, where it is now quite easy to see which package can be installed with which other package, and which package cannot.

3. EXPERIMENTAL RESULTS

The transformations described in this paper have been proven correct, and all the proofs have been certified in Coq [10]. An OCaml program implementing these transformations has been run on several mainstream GNU/Linux distributions: Debian testing (full suite, amd64, snapshot taken August 22, 2010), Ubuntu 10.10 (main suite, x386) and Mandriva 2010.1 (main suite, x386). Running time were measured on a machine using a Intel Core 2 Duo Processor E6600 at 2.4GHz. The relevant statistics of the results are given on Table 1.

We can notice that the number of packages is greatly reduced: many packages share the same behaviour as far as co-installability is concerned, and the quotienting phase identifies them. In particular, many packages can *always* be installed, which is good news for the GNU/Linux distribu-

Table 1: Repository sizes

	Debian		Ubuntu		Mandriva	
	before	after	before	after	before	after
Packages	28919	1038	7277	100	7601	84
Dependencies	124246	619	31069	29	38599	8
Conflicts	1146	985	82	60	78	62
Median cone size	38	1	38	1	59	1
Avg. cone size	66	1.7	84	1.3	153	1.1
Max. cone size	1134	15	842	4	1016	5
Running time (s)		10.6		1.19		11.6




Figure 6: Distribution of cone size (Debian testing)

tions; they are thus mapped to a single equivalence class. The number of dependencies is reduced even more: many dependencies are not relevant to co-installability, and are removed by our transformations. The simplification shown in Figure 4c turns out to be essential; for example, in Debian, thousands of packages depend on `debconf` which depends on either `debconf-i18n` or `debconf-english`, these two last packages being mutually exclusive: removing the disjunctive dependency hugely reduces the size of the final repository.

As for conflicts, we notice that distributions contain only few of them, which explains that flattening is practical; most of them cannot be removed.

Finally, to each package p in a repository one can associate its *cone*, the set of packages that are reachable from p by following the dependency relations; the cone size of a package is typically quite large (Figure 6): two third of the packages have a cone of more than a hundred packages. On the other hand, after simplification, two third of the package equivalence classes have a cone of size one, meaning that they do not depend on any other package than themselves.

After simplification, the Ubuntu distribution fits on a letter size paper (see Figure 14 in the Appendix) and can be easily inspected visually for errors. The corresponding graph for the Debian distribution is much larger, but it is our experience that it can still be displayed in a usable way with a suitable graph viewer: the visualisation of the Debian co-installability kernel can be tested online at http://ocsigen.org/js_of_ocaml/graph/.

As a consequence of the small size of the kernel extracted from a repository, many analyses can be performed very quickly. For example, it is easy to compute on the co-installability kernel of a repository the pairs of packages that can never be installed together, known as *strong conflicts* [5], which is the simplest case of co-installability. On the same data as [5], the computation takes a few seconds, instead of the several days reported there. This time is in fact included in the running times of Table 1 as we use this information when drawing the simplified repository to emphasise packages that prevent the installation of many other packages.




Figure 7: An example of dependency functions related by the preorder

4. ORDERING REPOSITORIES

A large part of our work consists in finding constraints that can be removed while leaving co-installability invariant. We make the idea of removing constraints precise by defining a preorder on repositories.

We first define a preorder on dependency functions: $D \sqsubseteq D^0$ iff for every package π , for every dependency $d \in D(\pi)$, there exists a dependency $d^0 \in D^0(\pi)$ such that $d^0 \subseteq d$. As an example, the dependency function D of the repository shown in Figure 7a is subsumed by the dependency function D^0 of the repository shown in Figure 7b.

This preorder corresponds exactly to the converse of logical implication: $D \sqsubseteq D^0$ if and only if the formulas $\llbracket D(\pi) \rrbracket \implies \llbracket D^0(\pi) \rrbracket$ can be derived for all packages $\pi \in P$.

The preorder induces an equivalence relation over dependency functions, and a canonical representative of an equivalence class can be defined by taking the element which is point-wise the smallest. Given a dependency function D , its canonical representative can be explicitly defined as:

$$D^0(\pi) = \{d \in D(\pi) \mid \forall d^0 \in D(\pi), d^0 \subseteq d \implies d^0 = d\}.$$

This provides a first way to simplify the dependency function (illustrated by Figure 4a), and our implementation aggressively put all dependencies in canonical form.

We can now define a preorder on repositories: $(P, D, C) \sqsubseteq (P^0, D^0, C^0)$ if and only if:

$$P = P^0, \quad D \sqsubseteq D^0, \quad C \subseteq C^0.$$

Given two repositories related under this preorder, going from the right hand side one to the left hand side one consists in removing conflicts or relaxing dependencies, thus making it easier to install packages.

THEOREM 2. *When $(P, D, C) \sqsubseteq (P^0, D^0, C^0)$, any healthy installation I of repository (P^0, D^0, C^0) is also a healthy installation of repository (P, D, C) .*

(Proof p.)

5. FLATTENING DEPENDENCIES

The *flattened* form of a dependency function, whose intuition has been given in the introduction, is formally defined as follows; given a repository (P, D, C) , the *flattened* dependency function \widehat{D} is the smallest function (with respect to point-wise inclusion) such that:

$$\frac{\text{REFL} \quad (\pi, \pi^0) \in C}{\{\pi\} \in \widehat{D}(\pi)} \quad \frac{\text{TRANS} \quad \begin{array}{c} \{\pi_1, \dots, \pi_n\} \in D(\pi) \\ d_1 \in \widehat{D}(\pi_1) \quad \dots \quad d_n \in \widehat{D}(\pi_n) \end{array}}{\bigcup_{i=1}^n d_i \in \widehat{D}(\pi)}$$

As the above rules are monotonous, such a function \widehat{D} exists.

The rule TRANS expands the intermediate dependencies of a package π and converts the result back into a conjunction of disjunctions; this rule silently drops circular dependencies: adding a dependency $\{\pi\} \in D(\pi)$ has no effect on $\widehat{D}(\pi)$.

The rule REFL is designed to capture precisely the two properties we have outlined informally in the introduction: on one hand, we want to keep in $\widehat{D}(\pi)$ only dependencies on packages with at least a conflict (we prove in the next section that they are enough for keeping co-installability invariant); on the other hand, we want $\widehat{D}(\pi)$ to contain *explicitly* all the packages that are needed to install package π (so, if π has a conflict, it will also be an explicit dependency of itself).

The application of the transformation on the repository of Figure 1 gives the following result (illustrated by Figure 2e):

	D	\widehat{D}
a	$\{\{b, c\}, \{f\}\}$	$\{\{b, c\}, \{f\}\}$
b	\emptyset	$\{\{b\}\}$
c	\emptyset	$\{\{c\}\}$
d	$\{\{e\}\}$	$\{\{f\}\}$
e	$\{\{f\}, \{g\}\}$	$\{\{f\}\}$
f	\emptyset	$\{\{f\}\}$
g	\emptyset	\emptyset

Packages d , e and f now have the same dependencies, which reflects the intuition that they behave the same way as far as co-installability is concerned.

6. STRONGLY FLAT REPOSITORIES

A flattened repository satisfies two properties (Theorem 3 below): a reflexivity property (if a package has a conflict, then it depends on itself), and a transitivity property (dependencies are stable under composition). Co-installability in such repositories, that we call *strongly flat*, can be shown equivalent to a more convenient property, that we call *weak co-installability* (Theorem 4); this is instrumental to prove the key result of this section: a set of packages are co-installable in a repository if and only if they are co-installable in the corresponding flattened repository (Theorem 6).

We define precisely these two properties and study repositories that satisfies them.

We define the *composition* $D; D^0$ of two dependency functions over a set of packages P as the smallest function (with respect to point-wise inclusion) such that for every package $\pi \in P$, for every set $\{\pi_1, \dots, \pi_n\} \in D(\pi)$, for every sets $d_i \in D^0(\pi_i)$, we have $\bigcup_{i=1}^n d_i \in (D; D^0)(\pi)$.

To any conflict relation C , we associate a dependency function Δ_C defined as follows:

$$\Delta_C(\pi) = \begin{cases} \{\{\pi\}\} & \text{if } (\pi, \pi^0) \in C \text{ for some } \pi^0 \in P \\ \emptyset & \text{otherwise} \end{cases}$$

A repository (P, D, C) is *strongly flat* when the following conditions hold:

- *reflexivity*: $\Delta_C \sqsubseteq D$ (every package with conflict depends on itself);
- *transitivity*: $D; D \sqsubseteq D$ (dependencies are closed under composition).

The flattening transformation produces strongly flat repositories.

THEOREM 3. *Let (P, D, C) be a repository. Let \widehat{D} be the corresponding flattened dependency function. The repository (P, \widehat{D}, C) is strongly flat.*

(Proof p.)

Intuitively, strongly flat repositories have a two level structure. Looking for instance at Figure 2e, we find all packages on the top layer, and typically only packages with conflicts (that we refer to as *features* in the following) at the bottom layer: thanks to transitivity, everything a package π may need to be installed is fully described by $D(\pi)$, without recursive traversal of dependencies; and thanks to reflexivity, conflicts need only be considered on the image of D .

We can take advantage of this to define a more convenient way of capturing co-installability:

A *configuration* is a pair (I, F) of a set I of packages and a set F of features; we say that it is *healthy* when the following conditions hold:

- *abundance*: every package has what it needs. Formally, for every package $\pi \in I$, and for every dependency $d \in D(\pi)$, we have $d \cap F \neq \emptyset$.
- *peace*: no two features conflict, that is, $C \cap (F \times F) = \emptyset$.

This is subtly different from the homonymous definitions regarding installations. Conflicts are only checked in F , and abundance only checked for packages in I w.r.t. F ; the sets F and I might as well be disjoint, here.

A set of packages Π are *weakly co-installable* if there exists a set of features $F \subseteq P$ such that the configuration (Π, F) is healthy. In general, this is a weaker notion. In strongly flat repositories, though, the two notions are equivalent.

THEOREM 4. *Any set of packages Π weakly co-installable in a strongly flat repository are also co-installable.*

(Proof p.)

It is interesting to remark that the result of the flattening operation can be mathematically characterised as follows.

LEMMA 5. *The flattened dependency function \widehat{D} associated to a repository (P, D, C) is the least dependency function D^0 (for preorder \sqsubseteq , and up to equivalence) such that:*

- $\Delta_C \sqsubseteq D^0$
- $D ; D^0 \sqsubseteq D^0$.

(Proof p.)

The essential result of this section is that co-installability is left invariant by flattening.

THEOREM 6. *Let (P, D, C) be a repository. Let \widehat{D} be the corresponding flattened dependency function. Let Π be a set of packages. The following propositions are equivalent:*

1. Π is co-installable in (P, D, C) ;
2. Π is weakly co-installable in (P, \widehat{D}, C) ;
3. Π is co-installable in (P, \widehat{D}, C) .

(Proof p.)

7. FLAT REPOSITORIES

In this section, we focus on a particular class $\nabla_C \subseteq \mathcal{P}(P)$ of dependencies that can be safely removed (Theorem 11). Removing these dependencies may destroy the *strongly flat*




Figure 8: Illustration of monotony requirement

structure of a repository, but we introduce the weaker notion of *flat* repository which is preserved (Theorem 11), and for which co-installability and weak co-installability still coincide (Theorem 10), thus enabling further simplifications introduced in Section 8.

We want to capture in the class ∇_C a set of dependencies that have the following two key properties:

always satisfiable any healthy configuration in the repository can be extended to satisfy these dependencies, so they are irrelevant for co-installability and we can remove them; formally, this means that if $d \in \nabla_C$, then for all $F \in \mathcal{P}(P)$ maximal with respect to set inclusion such that $C \cap (F \times F) = \emptyset$, we have $d \cap F \neq \emptyset$;

monotony a dependency which is in ∇_C must still be always satisfiable even if we remove some conflicts from the repository; formally, if $C^0 \subseteq C$, then $\nabla_C \subseteq \nabla_{C^0}$.

The *monotony* property is necessary, because in the next section we introduce further simplifications that remove redundant conflicts. We want to be sure that removing a conflict later on does not invalidate the decision taken here of removing a dependency, as illustrated in Figure 8: the disjunctive dependency on packages b and c can always be satisfied because the conflict between d and e prevents the simultaneous installation of d and e ; but this dependency is not in ∇_C since it is no longer satisfiable if the conflict between d and e is removed. This leads to the following:

DEFINITION 7. *Given a repository (P, D, C) , the set ∇_C is the the largest set such that $d \in \nabla_C$ if and only if, for all $C^0 \subseteq C$, for all $F \in \mathcal{P}(P)$ maximal with respect to set inclusion such that $C^0 \cap (F \times F) = \emptyset$, we have $d \cap F \neq \emptyset$.*

We can give a more explicit characterisation of the elements of ∇_C : these are *exactly* the dependencies that contain at least a package having only *internal* conflicts, like in Figure 4c.

THEOREM 8. *Let (P, D, C) be a repository. The set ∇_C is the set of dependencies d such that there exists a package $\pi \in d$ such that, for all pairs $(\pi, \pi^0) \in C$, we have $\pi^0 \in d$.*

(Proof p.)

Notice that, if a dependency d contains a package π with no conflict, then it is in ∇_C ; so ∇_C also contains the redundant dependencies shown in Figure 4b.

As we shall see, weak co-installability is left invariant by the transformation that removes the elements of set ∇_C from a dependency function D of a strongly flat repository. But, in general, the strongly flat property is lost, so we need a weaker notion that is preserved by this simplification. We start by defining a coarser preorder on dependency functions that ignores dependencies in ∇_C :

$D \prec_C D^0$ if and only if for every package π , for every dependency $d \in D(\pi)$, either $d \in \nabla_C$ or there exists a dependency $d^0 \in D^0(\pi)$ such that $d^0 \subseteq d$.

A repository (P, D, C) is *flat* when it satisfies the following properties:

- *reflexivity*: $\Delta_C \prec_C D$;
- *transitivity*: $D ; D \prec_C D$.

Flat repositories have a series of good properties: they include strongly flat repositories, co-installability and weak co-installability still coincide, removing ∇_C preserves flatness and keeps co-installability invariant.

LEMMA 9. *Any strongly flat repository is flat.*

(Proof p.)

THEOREM 10. *Any set of packages Π weakly co-installable in a flat repository are co-installable in this repository.*

(Proof p.)

THEOREM 11. *Let (P, D, C) be a flat repository and D^0 be the dependency function such that $D^0(\pi) = D(\pi) \setminus \nabla_C$ for all $\pi \in P$. The repository (P, D^0, C) is flat, and co-installability is left invariant by this transformation.*

(Proof p.)

One can still reason on flat repositories, as far as co-installability is concerned, as if their dependency function was transitive: just choose installations (I, F) where F is maximal, and then any dependency obtained by composition is satisfied, even when it is not explicitly in the dependency function. (For strongly flat repositories, this holds for arbitrary sets F .)

8. IRRELEVANT CONSTRAINTS

We review now several classes of dependencies and conflicts that are redundant and can be simplified out.

8.1 Clearly Irrelevant Dependencies

The results of the previous section let us remove the dependencies in ∇_C from a flat repository while leaving weak co-installability invariant and keeping the repository flat.

8.2 Conflict Covered Dependencies

An example of another very interesting class of irrelevant dependencies is shown in Figure 9, where the dependency for package a can always be satisfied despite the conflict between packages e and f (we assume the other packages in this dependency also have conflicts, not explicitated, so the dependency cannot be obviously removed): indeed, for this conflict to be relevant for the dependency, f needs to be installed; but if f is installed, at least one of packages c and d is installed as well, and thus the dependency is satisfied without needing to install e . This generalizes to the case where package e is in conflict with several packages with the same property as package f .

More formally, we say that a dependency d is *conflict covered* at π if it contains a package π such that for all $(\pi, \pi^0) \in C$, there exists a dependency $d^0 \in D(\pi^0)$ such that $d^0 \subseteq d \setminus \{\pi\}$. Removing one such dependency leave co-installability invariant.

LEMMA 12. *Let (P, D, C) be a flat repository, d be a conflict covered dependency, and D^0 be the dependency function obtained by removing d from D . Any set of packages weakly co-installable in (P, D^0, C) is weakly co-installable in (P, D, C) .*




Figure 9: Dependency covered by the conflict requirements.




Figure 10: Redundant conflict belonging to a clique

(Proof p.)

Unfortunately, removing such dependencies may destroy the flatness of the repository, so we remove them one after another, in a greedy way, and only after checking that flatness is preserved by using the following result.

LEMMA 13. *Let (P, D, C) be a flat repository. Let $\pi \in P$ and $d \in D$. Let D^0 be the dependency function D where the dependency d of package π has been removed. If the following two conditions hold, then (P, D^0, C) is flat.*

- $d \not\subseteq \{\pi\}$;
- for all $d^0 \in (D^0 ; D^0)(\pi)$, we have $d \not\subseteq d^0$.

(Proof p.)

In practice, it can be simpler to get all possible dependencies d^0 above by taking all dependencies $d^0 \in D(\pi) \setminus \{d, \{\pi\}\}$ and composing them with dependencies in D .

The graph reduction rule defined by this simplification has unsolvable critical pairs, so the result of this simplification could depend on the order of removal. In practice, though, we remove all instances present in the initial repository.

8.3 Redundant Conflicts

We consider some of the conflicts that can be removed from a repository while leaving co-installability invariant. A conflicting pair $(\pi_1, \pi_1^0) \in C$ is *redundant* if there exists a dependency $d \in D(\pi_1)$ such that for all $\pi_2 \in d$, there exists a package π_2^0 such that:

- $(\pi_2, \pi_2^0) \in C$;
- $\{\pi_1, \pi_1^0\} \neq \{\pi_2, \pi_2^0\}$;
- there exists $d^0 \in D(\pi_2)$ such that $d^0 \subseteq \{\pi_2^0\}$.

Redundant conflicts can be removed, but only one at a time: for example, if two conflicts are considered redundant thanks to the existence of one another, then removing both of them simultaneously is incorrect.

LEMMA 14. *Let (P, D, C) be a repository. Let (π_1, π_2) be a redundant conflict in this repository. Any healthy installation of repository $(P, D, C \setminus \{(\pi_1, \pi_2), (\pi_2, \pi_1)\})$ is healthy in repository (P, D, C) .*

(Proof p.)

Removing redundant conflicts involves a trade off. On one side, it may allow to remove some additional dependencies; on the other, it can also break some interesting structures. In Figure 10, the conflict between b and c is redundant, but removing it breaks the clique a, b, c , which is useful when drawing a simplified graph.

8.4 Dependence on Conflicting Packages

A special configuration may surface in the repository during simplification when the initial repository contains broken




Figure 11: Dependence on conflicting package

packages, as depicted in Figure 11: clearly, package a cannot be installed, and leaving such a configuration in the repository would pollute the graphical representation. In this case, we mark explicitly package a as broken by replacing its dependencies by the empty dependency \emptyset . All conflicts involving a can then be also removed, as they are redundant (Section 8.3). The transformation preserves healthiness.

LEMMA 15. *Let π be a package not installable in some repository (P, D, C) . Let D^\emptyset be the dependency function that coincide with D for all packages but π and such that $D^\emptyset(\pi) = \{\emptyset\}$. Any healthy installation of repository (P, D, C) is also healthy in repository (P, D^\emptyset, C) .*

(Proof p.)

The flatness of the repository may be destroyed, as some of the removed dependencies may be involved in transitivity. Thus, after applying such simplification, flattening should be performed again.

9. QUOTIENTING THE SET OF PACKAGES

In real-world repositories, many packages share the same behaviour as far as co-installability is concerned: for example, a lot of packages can always be installed, and some groups of packages only conflicts with a single other package.

In this section, we define an equivalence relation between packages, and show that the quotient w.r.t. this relation preserves all the good properties of a repository. We define two packages as equivalent in a repository (P, D, C) if they have the same dependencies:

$$\pi \equiv \pi^\emptyset \text{ if and only if } D(\pi) = D(\pi^\emptyset).$$

We write $[\pi]$ for the equivalence class of package π , and extend this definition to set of packages: $[\Pi] = \{[\pi] \mid \pi \in \Pi\}$. The *quotient repository* $(P^\emptyset, D^\emptyset, C^\emptyset)$ of a repository (P, D, C) is naturally defined as follows:

- P^\emptyset is the set of all equivalence classes: $P^\emptyset = P/\equiv = \{[\pi] \mid \pi \in P\}$;
- the dependency function D^\emptyset is such that $D^\emptyset([\pi]) = \{[d] \mid d \in D(\pi)\}$ for all $\pi \in P$;
- the conflict relation C^\emptyset is defined by

$$C^\emptyset = \{([\pi], [\pi^\emptyset]) \mid (\pi, \pi^\emptyset) \in C\}.$$

If the original repository does not contain dependencies of the form of Figure 11 nor redundant conflicts, then the quotient repository is indeed a repository (the key point to check is irreflexivity of the conflict relation).

LEMMA 16. *Let (P, D, C) be a flat repository such that, for all $\pi \in P$ and for all $d \in D(\pi)$, if $d \subseteq \{\pi^\emptyset\}$ for some $\pi^\emptyset \in P$, then $(\pi, \pi^\emptyset) \notin C$. Its quotient is indeed a repository.*

(Proof p.)

Quotienting preserves flatness and keeps co-installability invariant.

THEOREM 17. *A set of packages Π is weakly co-installable in the flat repository (P, D, C) if and only if the set $[\Pi]$ is weakly co-installable in the associated quotient repository.*

repeat

$(P, D, C) \leftarrow \text{flatten}(P, D, C)$

$(P, D, C) \leftarrow \text{canonise}(P, D, C)$

$(P, D, C) \leftarrow (P, D, C) \setminus \nabla_C$

$(P, D, C) \leftarrow \text{remove-clearly-broken}(P, D, C)$

$(P, D, C) \leftarrow \text{remove-redundant-conflicts}(P, D, C)$

until the last two steps above have no effect

$(P, D, C) \leftarrow \text{remove-conflict-covered-deps}(P, D, C)$

return $\text{quotient}(P, D, C)$

Figure 12: Simplifying the repository.

(Proof p.)

THEOREM 18. *If a repository is flat, then the corresponding quotiented repository is flat as well.*

(Proof p.)

10. REFLEXIVE TRANSITIVE REDUCTION

It would not be suitable to graph directly a repository after flattening as it would be polluted by dependencies which are not informative: due to reflexivity, we have packages π where $\{\pi\} \in D(\pi)$, and some dependencies can be deduced from others by transitivity. Thus, we perform a kind of reflexive transitive reduction of the dependency function: given a repository (P, D, C) , we find a minimal dependency function D^\emptyset with the same flattening (that is, $\widehat{D} = \widehat{D^\emptyset}$).

Because of disjunctive dependencies, the complexity of finding an optimal solution is high [11, 12], in contrast with the case of reflexive transitive reduction for graphs. As this is mostly a cosmetic issue for us, we use a simple non-optimal algorithm. As a first step, we iteratively remove dependencies which are implied from other dependencies by transitivity, in a greedy way. The second step is to remove all self dependencies, that is, dependencies $d \in D(\pi)$ such that $\pi \in d$. Co-installability is left invariant by these operations.

LEMMA 19. *Let (P, D, C) be a repository. Let $\pi \in P$ be a package and $d \in D(\pi)$ be a dependency of this package. Let $D^\emptyset = D \setminus \{\pi \mapsto d\}$ be the dependency function D where the dependency has been removed. If $d \in (D^\emptyset; D^\emptyset)(\pi)$ then any healthy installation I of repository (P, D^\emptyset, C) is a healthy installation of repository (P, D, C) .*

(Proof p.)

LEMMA 20. *Let (P, D, C) be a repository. Let D^\emptyset be the dependency function defined by $D^\emptyset(\pi) = \{d \in D \mid \pi \notin d\}$. Any healthy installation I of repository (P, D^\emptyset, C) is also a healthy installation of repository (P, D, C) .*

(Proof p.)

11. PUTTING ALL TOGETHER

We now have all the ingredients at hand to perform on any repository (P, D, C) the transformations that allow to produce the final repository, which is suitable both for drawing a simplified graph, or performing efficiently various analysis related to co-installability.

11.1 Extracting a Co-Installability Kernel

The complete algorithm is shown in Figure 12. We first flatten the initial repository (Section 5), canonise the dependency function (Section 4), and remove the clearly irrelevant



(a) Explicit representation (b) Compact representation

Figure 13: Conflict clique

dependencies in ∇_C (Section 7). In our implementation, all these operations are performed simultaneously: this is significantly more efficient, as we have less dependencies to consider while flattening. Then, we set the dependencies of broken packages of the form of Figure 11 to the empty dependency \emptyset (Section 8.4), and we remove redundant conflicts (Section 8.3). As changing the dependencies of broken packages may break flatness and removing may grow ∇_C , these five steps are iterated until no change occurs. The process terminates as at each iteration either $D(\pi)$ is set to $\{\emptyset\}$ for a package π or a conflict is removed. In practice, only two iterations are performed: more iterations are only needed in unlikely configurations where dealing with a broken package exposes another package as broken. Finally, we remove the conflict covered dependencies that can be safely dropped (Section 8.2) and the repository is quotiented (Section 9).

By combining the results of the previous sections, we obtain the fundamental result on the simplification performed by the algorithm.

THEOREM 21. *The transformation performed by the simplification algorithm leaves co-installability invariant. This algorithm produces a flat repository.*

(Proof p.)

As noticed above, on a repository with no broken package, it is not necessary to iterate the flattening phase, so the algorithm could run slightly faster; but finding all broken packages is slower than performing the whole simplification, as it requires to call a SAT solver repeatedly on large problems. On the other hand, repositories with good quality control should contain no broken packages, and a simpler version of the simplification algorithm could be used on them.

11.2 Drawing a Simplified Graph

Before drawing the final repository, we perform the transitive reflexive reduction of Section 10. The structure of the graph is then passed as input to the `dot` program of the Graphviz toolkit [13] that performs the layout.

It is important to name nodes using meaningful representatives of each equivalence class: we give preference to packages π that are directly involved in conflicts, as they have more chances to be relevant for the repository maintainers; these are easy to find by checking if $\{\pi\} \in D(\pi)$.

There can be many packages all mutually in conflict. For instance, this is the case of all mailer agents in Debian. We identify maximal such cliques and draw them in a more concise way, as shown in Figure 13.

We compute strong conflicts [5] (non-coinstallable pairs of packages) and use this information to emphasise packages that prevent the installation of many other packages.

12. READING THE UBUNTU SIMPLIFIED GRAPH

In Figure 14 we can see the result of applying the simplifications described in this paper to the main section of version 10.10 of the Ubuntu GNU/Linux distribution. Of the dozens of thousands of packages, and hundreds of thousands of relationships, only a handful is left, and it is possible to *read* on the graph quite a lot of interesting information; solid arrows indicate dependencies, dotted lines indicate conflicts, and conflict cliques are represented with a node containing a `#` connected with dotted lines to all packages in the conflict clique.

First of all, all packages not equivalent to those present on the drawing are *always* installable, so we need not care for them as far as co-installability is concerned.

Distribution maintainers can now review the independent connected components looking for potential quality issues; we give here just a couple of examples of what can be easily discovered at a glance without needing to sift through hundreds of hyperlinked web pages.

We look first at the five conflict cliques: they all correspond to incompatible implementations of libraries which are compiled with different backends; but the clique for *grub* has an apparent anomaly : an old package, *grub-legacy-ec2*, is in conflict with just one of the nodes in the clique, while it seems reasonable that it should be in conflict with all the others; this is an anomaly that the quality assurance team needs to consider, and fix by adding the missing conflicts to the metadata for *grub-legacy-ec2*.

On the other hand, the package *ubuntu-desktop* depends on a particular implementation of the `libSDL` multimedia library, *libSDL1.2debian-pulseaudio*, which is incompatible with the four other implementations available in the distribution. This seems also to be clearly an oversight in the definition of the dependencies of *ubuntu-desktop*, that should not make assumption on the particular sound library backend available on the machine.

13. RELATED AND FUTURE WORK

There are many relevant research issues, coming from the concrete and widespread applications of software component repositories: identifying components that can never be installable has been shown NP-complete, but tractable in practice [1], and there is currently very active research on finding installations that optimise a given objective function, organised around the Mancoosi International Solver Competition (<http://www.mancoosi.org/misc-2010/>); determining what other components a package will always need [4] and what pairs of packages are incompatible [5] have been shown to be relevant for quality assurance in package repositories; finally, since feature diagrams, used in software product lines, can be encoded as component repositories [14], all the problems related to configuration management can be equivalently stated in terms of repositories.

Component repositories have been in widespread use for quite a long time, but surprisingly little research work has been devoted to study the *formal* properties of component repositories: this situation is similar to what happened with the ubiquitous tool `make`, whose optimisations were only formalised and proven in [15]. Connections between component repositories and boolean satisfiability and constraint solving have been made only a few years ago in the framework of GNU/Linux distributions [1, 2] and the Eclipse platform [3], but these connections, and other recent developments such

as [4, 5] do not exploit the *special structure* of the dependencies and conflicts found in a repository.

The underlying structure of software component repositories exposed here can also be seen as a generalisation of some known mathematical structures : prime event structure [16] correspond to repositories without loops and disjunctions; directed hypergraphs [11] correspond to repositories without conflict arcs, and Dual Horn theories correspond to repositories with conflicts [17].

In this paper, we have developed a theory and algorithms to extract from a repository a co-installability kernel, which can be seen as a minimal representation of the dependency and conflict relations: despite the apparent simplicity of the definition of the problem, and the intuitive appealing of the hypergraph transformations we have developed, the proofs of the crucial properties turned out to be surprisingly complex, so we decided to machine check them using Coq [10].

The results presented here pave the way to attacking significantly more complex problems concerning software component repositories. More generally, we believe this work clearly shows the interest of the mathematical objects underlying software repositories, which turn out to be amenable to an elegant formal treatment and of high practical interest.

14. REFERENCES

- [1] F. Mancinelli, J. Boender, R. Di Cosmo, J. Vouillon, B. Durak, X. Leroy, and R. Treinen, "Managing the complexity of large free and open source package-based software distributions." in *ASE*. IEEE Computer Society, 2006, pp. 199–208.
- [2] C. Tucker, D. Shuffelton, R. Jhala, and S. Lerner, "Opium: Optimal package install/uninstall manager," in *ICSE*. IEEE Computer Society, 2007, pp. 178–188.
- [3] D. L. Berre and A. Parrain, "On sat technologies for dependency management and beyond," in *SPLC (2)*, S. Thiel and K. Pohl, Eds. Lero Int. Science Centre, University of Limerick, Ireland, 2008, pp. 197–200.
- [4] P. Abate, J. Boender, R. Di Cosmo, and S. Zacchiroli, "Strong dependencies between software components," in *ESEM*. IEEE Press, Oct. 2009, pp. 89–99. [Online]. Available: <http://ieeexplore.ieee.org:80/search/wrapper.jsp?arnumber=5316017>
- [5] R. Di Cosmo and J. Boender, "Using strong conflicts to detect quality issues in component-based complex systems," in *ISEC '10: Proceedings of the 3rd India software engineering conference*. New York, NY, USA: ACM, 2010, pp. 163–172.
- [6] R. Treinen and S. Zacchiroli, "Solving package dependencies: from EDOS to Mancoosi," in *DebConf8 (9th annual conference of the Debian project)* *DebConf8*, Argentine, 08 2008, pp. 18–43. [Online]. Available: <http://hal.archives-ouvertes.fr/hal-00340581/en/>
- [7] R. Di Cosmo, F. Mancinelli, J. Boender, J. Vouillon, B. Durak, X. Leroy, D. Pinheiro, P. Trezentos, M. Morgado, T. Milo, T. Zur, R. Suarez, M. Lijour, and R. Treinen, "Report on formal mangement of software dependencies," EDOS, Tech. Rep., Apr. 2006, eDOS project Deliverable 2.2, available as <http://www.edos-project.org/xwiki/bin/download/Main/Deliverables/edos-wp2d2.pdf>.
- [8] R. Di Cosmo, B. Durak, X. Leroy, F. Mancinelli, and J. Vouillon, "Maintaining large software distributions: new challenges from the FOSS era," in *Proceedings of the FRCSS 2006 workshop*, 2006, eASST Newsletter.
- [9] H. K. Buning and T. Lettmann, "Propositional logic: Deduction and algorithms," *Studia Logica*, vol. 71, pp. 247–258, 2002.
- [10] The Coq Development Team, *The Coq Proof Assistant Reference Manual – Version V8.2*, 2008. [Online]. Available: <http://coq.inria.fr>
- [11] G. Ausiello, A. D’Atri, and D. Saccà, "Graph algorithms for functional dependency manipulation," *J. ACM*, vol. 30, no. 4, pp. 752–766, 1983.
- [12] G. Ausiello, A. D’Atri, and D. Saccà, "Minimal representation of directed hypergraphs," *SIAM J. Comput.*, vol. 15, no. 2, pp. 418–431, 1986.
- [13] J. Ellson, E. Gansner, L. Koutsofios, S. North, and G. Woodhull, "Graphviz - open source graph drawing tools," in *Graph Drawing*, ser. Lecture Notes in Computer Science, P. Mutzel, M. Jünger, and S. Leipert, Eds. Springer Berlin / Heidelberg, 2002, vol. 2265, pp. 594–597. [Online]. Available: http://dx.doi.org/10.1007/3-540-45848-4_57
- [14] R. Di Cosmo and S. Zacchiroli, "Feature diagrams as package dependencies," in *SPLC*, ser. Lecture Notes in Computer Science, J. Bosch and J. Lee, Eds., vol. 6287. Springer, 2010, pp. 476–480.
- [15] C. A. Gunter, "Abstracting dependencies between software configuration items," *ACM Trans. Softw. Eng. Methodol.*, vol. 9, pp. 94–131, January 2000. [Online]. Available: <http://doi.acm.org/10.1145/332740.332743>
- [16] G. Winskel, "Event structures," in *Petri Nets: Applications and Relationships to Other Models of Concurrency*, ser. Lecture Notes in Computer Science, W. Brauer, W. Reisig, and G. Rozenberg, Eds. Springer Berlin / Heidelberg, 1987, vol. 255, pp. 325–392. [Online]. Available: http://dx.doi.org/10.1007/3-540-17906-2_31
- [17] W. F. Dowling and J. H. Gallier, "Linear-time algorithms for testing the satisfiability of propositional horn formulae," *J. Log. Program.*, vol. 1, no. 3, pp. 267–284, 1984.

APPENDIX

This appendix contains all the proofs of the lemmas and theorems stated in the body of the article, as well as some auxiliary lemmas which are needed in the proofs. Auxiliary material appearing only in the appendix has numbers that are higher than in the paper, and appear hence out of order: this is necessary to avoid disrupting the numbering of lemmas and theorems in the body of the article.

A. PROOF OF THE MAIN RESULTS

REMARK 22. Let (P, D, C) and (P^0, D^0, C^0) be two repositories. If any healthy installation I of repository (P, D, C) is a healthy installation of repository (P^0, D^0, C^0) , then any co-installable set of packages Π in repository (P, D, C) is a co-installable set of packages in repository (P^0, D^0, C^0) .

Proof: Immediate □

The preorder \sqsubseteq is coarser than point-wise inclusion:

REMARK 23. If, for all $\pi \in P$, $D(\pi) \subseteq D^0(\pi)$, then $D \sqsubseteq D^0$.

We write \sim for the equivalence relation associated to the relation \sqsubseteq .

Proof of Theorem 2

We suppose that I is a healthy installation of (P^0, D^0, C^0) and prove that I is a healthy installation of (P, D, C) .

We first prove abundance. Let $\pi \in I$ and $d \in D(\pi)$. As $D \sqsubseteq D^0$, there exists $d^0 \in D^0(\pi)$ such that $d^0 \subseteq d$. As I is a healthy installation of (P^0, D^0, C^0) , by abundance, $d^0 \cap I \neq \emptyset$. Hence, $d \cap I \neq \emptyset$ as wanted.

We now prove peace. As I is a healthy installation of (P^0, D^0, C^0) , we have $C^0 \cap (I \times I) = \emptyset$. On the other hand, $C \subseteq C^0$. Hence, $C \cap (I \times I) = \emptyset$ as wanted. ■

Proof of Theorem 3

Reflexivity, that is $\Delta_C \sqsubseteq \widehat{D}$, is a direct consequence of rule REFL.

We now show transitivity, that is $\widehat{D}; \widehat{D} \sqsubseteq \widehat{D}$. Let $\pi \in P$ and $d \in (\widehat{D}; \widehat{D})(\pi)$. By definition of composition, there exists $d^0 = \{\pi_1, \dots, \pi_n\} \in \widehat{D}(\pi)$ and sets $d_i \in \widehat{D}(\pi_i)$ such that $d = \bigcup_{i=1}^n d_i$. We need to show that there exists $d^{00} \in \widehat{D}(\pi)$ such that $d^{00} \subseteq d = \bigcup_{i=1}^n d_i$. The proof is by induction on a derivation of $d^0 \in \widehat{D}(\pi)$. We have two cases.

- Case REFL. There exists $\pi^0 \in P$ such that $(\pi, \pi^0) \in C$, and $d^0 = \{\pi\}$. We have, $d = d_1 \in \widehat{D}(\pi_1) = \widehat{D}(\pi)$, hence we can take $d^{00} = d$.

- Case TRANS. There exists a set $\{\pi_1, \dots, \pi_m\} \in \widehat{D}(\pi)$ and sets $d_j^0 \in \widehat{D}(\pi_j^0)$ such that $d^0 = \bigcup_{j=1}^m d_j^0$. Besides, by induction hypothesis, for all j , if one can find packages π_j^k and dependencies d_j^k such that $d_j^0 = \{\pi_j^1, \dots, \pi_j^{p_j}\}$, and $d_j^k \in \widehat{D}(\pi_j^k)$, then there exists $d_j^{00} \in \widehat{D}(\pi_j^0)$ such that $d_j^{00} \subseteq \bigcup_{k=1}^{p_j} d_j^k$. We have $d_j^0 \subseteq d^0$. Hence we can choose for each indices j and k an index i such that $\pi_j^k = \pi_i$. Then, we take $d_j^k = d_i$. We therefore have dependencies $d_j^{00} \in \widehat{D}(\pi_j^0)$ such that $d_j^{00} \subseteq \bigcup_{k=1}^{p_j} d_j^k$. We take $d^{00} = \bigcup_{j=1}^m d_j^{00}$. By rule TRANS, $d^{00} \in \widehat{D}(\pi)$. Besides, $d^{00} \subseteq \bigcup_{j=1}^m \bigcup_{k=1}^{p_j} d_j^k$. But, as $d^0 = \bigcup_{j=1}^m d_j^0$, the packages π_j^k ranges over all the distinct packages π_i . Hence, the dependencies d_j^k ranges over all dependencies d_i , and we have $d^{00} \subseteq \bigcup_{i=1}^n d_i$ as wanted. □
-

REMARK 24. The two following properties hold:

- an installation I is healthy iff the configuration (I, I) is healthy;
- if (Π, F) is healthy and $\Pi^0 \subseteq \Pi$, then (Π^0, F) is healthy.

Proof: Immediate. □

LEMMA 25. If a set of packages Π are co-installable, then they are also weakly co-installable.

Proof: Suppose the set of packages Π are co-installable. By definition, there exists a healthy installation F such that $\Pi \subseteq F$. We terminate the proof by using in turn both properties in Remark 24. The configuration (F, F) is healthy. Hence, by anti-monotony of the definition of healthiness, (Π, F) is healthy. □

We present here two technical lemmas needed in the proof of Theorem 6.

LEMMA 26. Let (P, D, C) be a repository such that $\Delta_C \sqsubseteq D$. Let (I, F) be a healthy configuration. For every package $\pi \in I$, if $(\pi, \pi^0) \in C$ for some $\pi^0 \in P$, then $\pi \in F$.

Proof: We consider a package $\pi \in I$ such that $(\pi, \pi^0) \in C$ for some package $\pi^0 \in P$. As $\Delta_C \sqsubseteq D$, there exists $d \in D(\pi)$ such that $d \subseteq \Delta_C(\pi) = \{\pi\}$. By abundance, $d \cap F \neq \emptyset$. Hence, $\pi \in F$. □

LEMMA 27. Let (P, D, C) be a repository. Let D^0 be a dependency function such that $\Delta_C \sqsubseteq D^0$ and $D; D^0 \sqsubseteq D^0$. If a set of packages Π are weakly co-installable in (P, D^0, C) , then they are co-installable in (P, D, C) .

Proof: We consider a set of packages Π which are weakly co-installable in (P, D^0, C) . This means that there exists a set of packages F such that the configuration (Π, F) is

healthy in (P, D^0, C) . We then consider a maximal set of packages I such that $\Pi \subseteq I$ and the configuration (I, F) is healthy in (P, D^0, C) . We can conclude if we prove that I is healthy in (P, D, C) .

We first prove abundance. We take $\pi \in I$ and $d = \{\pi_1, \dots, \pi_n\} \in D(\pi)$ and need to prove that $d \cap I \neq \emptyset$. We want to show that we can find $\pi^{00} \in d$ such that $(I \cup \{\pi^{00}\}, F)$, is healthy. Indeed, then, by maximality of I , we have $\pi^{00} \in I$. Hence, $d \cap I \neq \emptyset$ as wanted. The package π^{00} should be such that, for all $d^{00} \in D(\pi^{00})$, we have $d^{00} \cap F \neq \emptyset$. The proof is by contradiction. Suppose that for all i , there exists $d_i \in D^0(\pi_i)$ such that $d_i \cap F = \emptyset$. As $D; D^0 \sqsubseteq D^0$, there exists $d^0 \in D^0(\pi)$ such that $d^0 \subseteq \bigcup_{1 \leq i \leq n} d_i$. We have $d^0 \cap F \subseteq \bigcup_{1 \leq i \leq n} d_i \cap F = \emptyset$. This contradicts the fact that $\pi \in I$ and (I, F) is healthy. Hence the existence of a suitable package π^{00} which allow us to finish the proof of abundance.

We now prove peace by contradiction. Suppose that $C \cap (I \times I) \neq \emptyset$, that is, there exists $\pi \in I$ and $\pi^0 \in I$ such that $(\pi, \pi^0) \in C$. As (I, F) is healthy, we have $\pi \in F$ and $\pi^0 \in F$ by Lemma 26, and then $(\pi, \pi^0) \notin C$ by peace. We reach a contradiction, hence $C \cap (I \times I) = \emptyset$. \square

Proof of Theorem 4

This is an immediate consequence of Lemma 27, taking $D = D^0$. \blacksquare

Proof of Lemma 5

Clearly, the flattened dependency function satisfies these equations. Let D^0 be another dependency function satisfying the equations. We want to prove that, for all $\pi \in P$, $\widehat{D}(\pi) \subseteq D^0(\pi)$. By unfolding the definition of the preorder, this can be rephrased as: for all $\pi \in P$, for all $d \in \widehat{D}(\pi)$, there exists $d^0 \in D^0(\pi)$ such that $d^0 \subseteq d$. The proof is by induction on a derivation of $d \in \widehat{D}(\pi)$. We have two cases.

- Case REFL. There exists $\pi^0 \in P$ such that $(\pi, \pi^0) \in C$, and $d = \{\pi\}$. From the first assumption, there exists $d^0 \in D^0(\pi)$, such that $d^0 \subseteq \{\pi\}$. Hence, $d^0 \subseteq d$ as wanted.
- Case TRANS. There exists a set $\{\pi_1, \dots, \pi_n\} \in D(\pi)$ and sets $d_i \in \widehat{D}(\pi_i)$ such that $d = \bigcup_{1 \leq i \leq n} d_i$. Besides, by induction hypothesis, for all i , there exists $d_i^0 \in D^0(\pi_i)$ such that $d_i^0 \subseteq d_i$. From the second assumption, there exists $d^0 \in D^0(\pi)$, such that $d^0 \subseteq \bigcup_{1 \leq i \leq n} d_i^0$. We have, $d^0 \subseteq \bigcup_{1 \leq i \leq n} d_i^0 \subseteq \bigcup_{1 \leq i \leq n} d_i = d$ as wanted. \square

Proof of Theorem 6

(3) implies (2) by Lemma 25. (2) implies (3) by Lemmas 4 and 3. (2) implies (1) by Lemmas 27 and 5. We now assume (1) and prove (3).

By definition, there exists a healthy installation I of repository (P, D, C) such that $\Pi \subseteq I$. We can just show that I is also a healthy installation of repository (P, \widehat{D}, C) . Peace is immediate. We prove abundance. We consider $\pi \in P$ and $d \in \widehat{D}(\pi)$. We need to show that if $\pi \in I$, then $d \cap I \neq \emptyset$. The proof is by induction on a derivation of $d \in \widehat{D}(\pi)$. We have two cases.

- Case REFL. We have $d = \{\pi\}$. Hence, if $\pi \in I$ then $d \cap I = \{\pi\} \neq \emptyset$.

- Case TRANS. We have $\{\pi_1, \dots, \pi_n\} \in D(\pi)$ and $d = \bigcup_{1 \leq i \leq n} d_i$ where $d_i \in \widehat{D}(\pi_i)$ and, if $\pi_i \in I$, then $d_i \cap I \neq \emptyset$ (induction hypothesis). Let us assume $\pi \in I$. By abundance, as $\{\pi_1, \dots, \pi_n\} \in D(\pi)$, there must exists some i such that $\pi_i \in I$. Then, $d_i \cap I \neq \emptyset$, and thus $d \cap I \neq \emptyset$. \square

Proof of Theorem 8

We first assume that there exists no package $\pi \in d$ such that, for all pairs $(\pi, \pi^0) \in C$, we have $\pi^0 \in d$. We show that then $d \notin \nabla_C$. Let C^0 be the set of conflicts in C that cross d :

$$C^0 = \{(\pi, \pi^0) \in C \mid \#(\{\pi, \pi^0\} \cap d) = 1\}.$$

Clearly, this is a conflict relation (it is symmetric and ir-reflexive). We consider the set of features $F = P \setminus d$. We have $C^0 \cap (F \times F) = \emptyset$. If we can now prove that the set F is a maximal set satisfying this property, we can conclude. Indeed, then, as $d \cap F = \emptyset$, the set d is not in ∇_C . We consider a strictly larger set of features F^0 and show that it does not satisfies the property. The set F^0 must contain at least a package π in dependency d . But then, by assumption, there exists $\pi^0 \in P \setminus d = F$ such that $(\pi, \pi^0) \in C$. By definition of C^0 , we have $(\pi, \pi^0) \in C^0$. Hence, $C^0 \cap (F^0 \times F^0) \neq \emptyset$, as wanted.

We now assume that there exists a package $\pi \in d$ such that, for all pairs $(\pi, \pi^0) \in C$, we have $\pi^0 \in d$. We show that then $d \in \nabla_C$. Let C^0 be a conflict relation such that $C^0 \subseteq C$. Let F be a maximal set of features such that $C^0 \cap (F \times F) = \emptyset$. We need to show that $d \cap F \neq \emptyset$. If we assume that there exists no $\pi^0 \in F$ such that $(\pi, \pi^0) \in C^0$, then, by maximality of F , we must have $\pi \in F$. Hence, $d \cap F \neq \emptyset$ as wanted. Otherwise, there exists $\pi^0 \in F$ such that $(\pi, \pi^0) \in C^0$. But then, $\pi^0 \in d$. Hence, $d \cap F \neq \emptyset$ as well. \blacksquare

REMARK 28. Let (P, D, C) be a repository. The set ∇_C is upward-closed: if $d \subseteq d^0$ and $d \in \nabla_C$, then $d^0 \in \nabla_C$.

Proof: This is immediate thanks to Theorem 8. \square

This preorder can be extended to repositories by setting $(P, D, C) \prec (P^0, D^0, C^0)$ iff $P = P^0$, $D \prec_C D^0$, and $C \subseteq C^0$. We write \approx for the associated equivalence relation.

REMARK 29. The following implications hold:

- if $C^0 \subseteq C$ and $D \prec_C D^0$, then $D \prec_{C^0} D^0$;
- if $D \subseteq D^0$, then $D \prec_C D^0$;
- if $(P, D, C) \sqsubseteq (P^0, D^0, C^0)$, then $(P, D, C) \prec (P^0, D^0, C^0)$.

Proof: Immediate. \square

Proof of Lemma 9

This is a direct consequence of Remark 29. \blacksquare

REMARK 30. Let (P, D, C) be a flat repository. Let $\pi \in P$ such that $(\pi, \pi^0) \in C$ for some $\pi^0 \in P$. Then, there exists $d \in D(\pi)$ such that $d \subseteq \{\pi\}$.

Proof: Let $\pi \in P$ such that $(\pi, \pi^0) \in C$ for some $\pi^0 \in P$. By reflexivity $(\Delta_{C^0} \prec_C D)$, we have either $\{\pi\} \in \nabla_C$ or there exists $d \in D(\pi)$ such that $d \subseteq \{\pi\}$. By Theorem 8, $\{\pi\} \notin \nabla_C$. Hence the result. \square

LEMMA 31. If $(P, D, C) \prec (P^0, D^0, C^0)$, $(P^0, D^0, C^0) \sqsubseteq (P, D, C)$, and the repository (P, D, C) is flat, then the repository (P^0, D^0, C^0) is flat.

Proof: We assume that $(P, D, C) \prec (P^0, D^0, C^0)$, $(P^0, D^0, C^0) \sqsubseteq (P, D, C)$, and (P, D, C) is flat. We show that (P^0, D^0, C^0) is flat.

We first show reflexivity, that is, $\Delta_{C^0} \prec_{C^0} D^0$. We have $\Delta_C \prec_C D \prec_C D^0$ by hypothesis. Hence, $\Delta_C \prec_{C^0} D \prec_{C^0} D^0$ by Remark 29. On the other hand, we have $C^0 \subseteq C$. Hence, for all $\pi \in P$, we have $\Delta_{C^0}(\pi) \subseteq \Delta_C(\pi)$. Thus, $\Delta_{C^0} \sqsubseteq \Delta_C$ by Remark 23, and finally $\Delta_{C^0} \prec_{C^0} D^0$ by Remark 29. We conclude by transitivity of \prec_{C^0} .

We now show transitivity, that is, $D^0; D^0 \prec_{C^0} D^0$. Let $\pi^0 \in P$ and $d^0 \in (D^0; D^0)(\pi^0)$. We show that either $d^0 \in \nabla_{C^0}$ or there exists a dependency $d \in D^0(\pi^0)$ such that $d \subseteq d^0$. By definition of composition, there exists a dependency $d^0 = \{\pi_1, \dots, \pi_n\} \in D^0(\pi^0)$ and dependencies $d_i^0 \in D^0(\pi_i)$ such that $d^0 = \bigcup_{i=1}^n d_i^0$. As $D^0 \sqsubseteq D$, there exists $d \in D(\pi^0)$ such that $d \subseteq d^0$. Similarly, for all i , there exists $d_i \in D(\pi_i)$ such that $d_i \subseteq d_i^0$. By transitivity in (P, D, C) , either $\bigcup_{i=1}^n d_i \in \nabla_C$ or there exists $d^1 \in D(\pi^0)$ such that $d^1 \subseteq \bigcup_{i=1}^n d_i$. In the first case, we have $\bigcup_{i=1}^n d_i \subseteq \bigcup_{i=1}^n d_i^0 = d^0$. Then, as ∇_C is upward-closed (Remark 28) and $C^0 \subseteq C$, we have $d^0 \in \nabla_{C^0} \subseteq \nabla_{C^0}$ as wanted. In the second case, we have $d^1 \subseteq \bigcup_{i=1}^n d_i \subseteq \bigcup_{i=1}^n d_i^0 = d^0$. As $D \prec_C D^0$, either $d^1 \in \nabla_{C^0}$ or there exists $d^2 \in D^0(\pi^0)$ such that $d^2 \subseteq d^1$. If $d^1 \in \nabla_{C^0}$, then $d^0 \in \nabla_{C^0}$ as well, as ∇_{C^0} is upward-closed. This is what we wanted to prove. Otherwise, we have $d^2 \in D^0(\pi^0)$ such that $d^2 \subseteq d^1 \subseteq d^0$, which is again what is needed to conclude. \square

Proof of Theorem 10

The proof follow the proof of Lemma 10, with minor modifications.

Let (P, D, C) be a flat repository. Let Π be a set of packages weakly co-installable in (P, D, C) . There exists a set of packages F such that the configuration (Π, F) is healthy in (P, D, C) . We assume that F is a maximal set of packages satisfying this property. We then consider a maximal set of packages I such that $\Pi \subseteq I$ and the configuration (I, F) is healthy in (P, D, C) . We can conclude if we prove that I is healthy in (P, D, C) .

We first prove abundance. We take $\pi \in I$ and $d = \{\pi_1, \dots, \pi_n\} \in D(\pi)$ and need to prove that $d \cap I \neq \emptyset$. We want to show that we can find $\pi^{00} \in d$ such that $(I \cup \{\pi^{00}\}, F)$, is healthy. Indeed, then, by maximality of I , we have $\pi^{00} \in I$. Hence, $d \cap I \neq \emptyset$ as wanted. The package π^{00} should be such that, for all $d^{00} \in D(\pi^{00})$, we have $d^{00} \cap F \neq \emptyset$. The proof is by contradiction. Suppose that for all i , there exists $d_i \in D(\pi_i)$ such that $d_i \cap F = \emptyset$. As $D; D \prec_C D$, either $\bigcup_{i=1}^n d_i \in \nabla_C$ or there exists $d^0 \in D(\pi)$ such that $d^0 \subseteq \bigcup_{i=1}^n d_i$. In the first case, by maximality of F and definition of ∇_C , $(\bigcup_{i=1}^n d_i) \cap F \neq \emptyset$. This contradicts the assumption that $d_i \cap F = \emptyset$ for all i . In the second case, we have $d^0 \cap F \subseteq \bigcup_{i=1}^n d_i \cap F = \emptyset$. This contradicts the fact that $\pi \in I$ and (I, F) is healthy. Hence the existence of a suitable package π^{00} which allow us to finish the proof of abundance.

We now prove peace by contradiction. Suppose that $C \cap (I \times I) \neq \emptyset$, that is, there exists $\pi \in I$ and $\pi^0 \in I$ such that $(\pi, \pi^0) \in C$. As (I, F) is healthy, we have $\pi \in F$ and $\pi^0 \in F$ by Lemma 26, and then $(\pi, \pi^0) \notin C$ by peace. We reach a contradiction, hence $C \cap (I \times I) = \emptyset$. \blacksquare

LEMMA 32. If $(P, D, C) \prec (P^0, D^0, C^0)$, then any set of packages Π weakly co-installable in repository (P^0, D^0, C^0) is weakly co-installable in repository (P, D, C) .

Proof: We assume that $(P, D, C) \prec (P^0, D^0, C^0)$ and that the set of packages Π is weakly co-installable in (P^0, D^0, C^0) . There exists a set of packages F such that (Π, F) is a healthy configuration in (P^0, D^0, C^0) . In particular, we have $C^0 \cap (F \times F) = \emptyset$. As $C \subseteq C^0$, we also have $C \cap (F \times F) = \emptyset$. There exists thus a set of packages F^0 maximal such that $C \cap (F^0 \times F^0) = \emptyset$ and $F \subseteq F^0$. We can conclude by showing that (Π, F^0) is a healthy configuration in (P, D, C) . We have peace by definition of F^0 . We show abundance. Let $\pi \in \Pi$ and $d \in D(\pi)$. We need to show that $d \cap F^0 \neq \emptyset$. As $D \prec_C D^0$, either $d \in \nabla_C$ or there exists $d^0 \in D^0(\pi)$ such that $d^0 \subseteq d$. In the first case, by definition of ∇_C and maximality of F^0 , we have $d \cap F^0 \neq \emptyset$ as wanted. In the second case, as (Π, F) is a healthy configuration, we have $d^0 \cap F \neq \emptyset$, and therefore $d \cap F^0 \neq \emptyset$ as wanted. \square

LEMMA 33. Let (P, D, C) be a repository. We define a simplified dependency function D^0 by: $D^0(\pi) = D(\pi) \setminus \nabla_C$ for all $\pi \in P$. Then:

- $(P, D^0, C) \sqsubseteq (P, D, C)$;
- $(P, D^0, C) \approx (P, D, C)$.

Proof: For all $\pi \in P$, we have $D^0(\pi) \subseteq D(\pi)$. Hence, by Remark 23, $D^0 \sqsubseteq D$, and therefore, $(P, D^0, C) \sqsubseteq (P, D, C)$. Then, by Remark 29, $(P, D^0, C) \prec (P, D, C)$.

We now prove $(P, D, C) \prec (P, D^0, C)$. It is sufficient to show that $D \prec_C D^0$. Let $\pi \in P$ and $d \in D(\pi)$. We need to prove that either $d \in \nabla_C$ or there exists $d^0 \in D^0(\pi) = D(\pi) \setminus \nabla_C$ such that $d^0 \subseteq d$. This is clearly satisfied, taking $d^0 = d$ in the second case. \square

Proof of Theorem 11

Flatness is a consequence of Lemmas 31 and 33. Invariance is a consequence of Theorem 10, and Lemmas 32 and 33. \blacksquare

The following technical lemma is a crucial tool for reasoning on flat repositories. It states that a stronger form of transitivity where dependencies are not fully composed holds for flat repositories.

LEMMA 34. *Let (P, D, C) be a flat repository, $\pi \in P$, $d \in D(\pi)$. Let $\Pi = \{\pi_1, \dots, \pi_n\}$ be a subset of d . We assume that, for all i , there exists $d_i \in D(\pi_i)$. Then, there exists $d^0 \in D(\pi) \cup \nabla_C$ such that $d^0 \subseteq (d \setminus \Pi) \cup \bigcup_{i=1}^n d_i$.*

Proof: Suppose that there exists $\pi^0 \in d \setminus \Pi$ such that $\{\pi^0\} \in \nabla_C$ (that is, by Theorem 8, package π^0 conflicts with no package). Then we can take $d^0 = \{\pi^0\}$ and conclude. Thus, we can assume that, for all $\pi^0 \in d \setminus \Pi$, there exists $\pi^{00} \in P$ such that $(\pi^0, \pi^{00}) \in C$.

We show that for all $\pi^0 \in d$ we can find $d^0 \in D(\pi^0)$ such that $d^0 \subseteq (d \setminus \Pi) \cup \bigcup_{i=1}^n d_i$. We can then use transitivity ($D; D \prec_C D$) to conclude. Indeed, by transitivity, there exists $d^{00} \in D; D$ such that $d^{00} \subseteq (d \setminus \Pi) \cup \bigcup_{i=1}^n d_i$ and either $d^{00} \in \nabla_C$ or there exists $d^{000} \in D(\pi)$ such that $d^{000} \subseteq d^{00}$. We can then conclude by taking either d^{00} or d^{000} .

Let $\pi^0 \in d$. If $\pi^0 = \pi_i$ for some i , we can take $d^0 = d_i$. Otherwise, $\pi^0 \in d \setminus \Pi$. We have assumed that there exists π^{00} such that $(\pi^0, \pi^{00}) \in C$. By Remark 30, there exists $d^0 \in D(\pi^0)$ such that $d^0 \subseteq \{\pi^0\}$. This dependency is suitable as $\pi^0 \in d \setminus \Pi$. \square

LEMMA 35. *Let (P, D, C) be a flat repository. Let C^0 be a conflict relation such that $C^0 \subseteq C$. The repository (P, D, C^0) is flat.*

Proof: We first prove reflexivity: $\Delta_{C^0} \prec_C D$. By assumption, $\Delta_C \prec_C D$. Hence, by transitivity, it is sufficient to prove that $\Delta_{C^0} \prec_C \Delta_C$. Clearly, as $C^0 \subseteq C$, we have $\Delta_{C^0}(\pi) \subseteq \Delta_C(\pi)$ for all π . We conclude by Remark 23 and 29.

We now prove transitivity. By assumption, we have $D; D \prec_C D$. Hence, by Remark 29, as $C^0 \subseteq C$, we have $D; D \prec_{C^0} D$ as wanted. \square

In the following, we will write $D \setminus \{\pi \mapsto d\}$ for the dependency function D where the dependency d of package π has been removed. Formally,

$$\begin{aligned} (D \setminus \{\pi_0 \mapsto d_0\})(\pi_0) &= D(\pi_0) \setminus \{d_0\} \\ (D \setminus \{\pi_0 \mapsto d_0\})(\pi) &= D(\pi) \quad \text{when } \pi \neq \pi_0. \end{aligned}$$

Proof of Lemma 12

Let $d_0 \in D(\pi_0)$ be one of these dependencies. There exists $\pi_1 \in d_0$ such that for all $(\pi_1, \pi) \in C$, there exists a dependency $d \in D(\pi)$ such that $d \subseteq d_0 \setminus \{\pi_1\}$. Let $D^0 = D \setminus \{\pi_0 \mapsto d_0\}$ be the dependency function with

this dependency removed. Let Π be a weakly co-installable set of packages in the repository (P, D^0, C) . There exists a set of features F such that (Π, F) is a healthy configuration of (P, D^0, C) . We consider a maximal such a set F (with respect to inclusion). Clearly, the set F is then maximal such that $C \cap (F \times F) = \emptyset$. If $d_0 \cap F \neq \emptyset$, then (Π, F) is a healthy configuration of (P, D, C) as well, as wanted. We now consider the case where $d_0 \cap F = \emptyset$. Let $F^0 = (F \setminus \{\pi \mid (\pi_1, \pi) \in C\}) \cup \{\pi_1\}$. We show that (Π, F^0) is a healthy configuration of (P, D, C) . Peace is clear. We need to show abundance.

Let $\pi \in I$ and $d \in D(\pi)$. We want to show that $d \cap F^0 \neq \emptyset$. If $\pi_1 \in d$, this is the case. We can thus assume that $\pi_1 \notin d$. In particular, this means that $d \neq d_0$, and therefore, by abundance, $d \cap F \neq \emptyset$. The remainder of the proof is by contradiction. We assume that $d \cap F^0 = \emptyset$. We consider the packages π^i such that $d \cap F = \{\pi^1, \dots, \pi^n\}$. We have $\pi^i \in F \setminus F^0$. Thus, for all i , we have $(\pi_1, \pi^i) \in C$, and therefore, by assumption, there exists $d^i \in D(\pi^i)$ such that $d^i \subseteq d_0 \setminus \{\pi_1\}$. By Lemma 34, there exists a dependency d^0 such that $d^0 \in D(\pi) \cup \nabla_C$ and $d^0 \subseteq (d \setminus (d \cap F)) \cup \bigcup_{i=1}^n d^i$. From this last inclusion, we get $d^0 \subseteq (d \setminus F) \cup (d_0 \setminus \{\pi_1\})$. From this, on the one hand, as $\pi_1 \notin d$, we can see that $\pi_1 \notin d^0$, and therefore $d^0 \neq d$. Hence, $d^0 \in D^0(\pi) \cup \nabla_C$. By abundance and maximality of F , we get $d^0 \cap F \neq \emptyset$. On the other hand, $d^0 \cap F \subseteq ((d \setminus F) \cup (d_0 \setminus \{\pi_1\})) \cap F = (d_0 \setminus \{\pi_1\}) \cap F = ((d_0 \cap F) \setminus \{\pi_1\}) = \emptyset$. We reach a contradiction, which completes the proof. \blacksquare

Proof of Lemma 13

We first prove reflexivity, that is, $\Delta_C \prec_C D^0$. Let $\pi^0 \in P$ and $d^0 \in \Delta_C(\pi^0)$. We show that either $d^0 \in \nabla_C$, or there exists $d^{00} \in D^0(\pi^0)$ such that $d^{00} \subseteq d^0$. By definition, we must have $d^0 = \{\pi^0\}$. By reflexivity in repository (P, D, C) , we have either $d^0 \in \nabla_C$, or there exists $d^{000} \in D(\pi^0)$ such that $d^{000} \subseteq d^0$. In the first case, we can conclude immediately. If $d^{000} \in D^0(\pi^0)$, we can conclude as well by taking $d^{00} = d^{000}$. Otherwise, we must have $\pi^0 = \pi$ and $d^{000} = d$. But, then, $d = d^{000} \subseteq d^0 = \{\pi^0\} = \{\pi\}$. This is not possible due to the first assumption.

We now prove transitivity, that is, $D^0; D^0 \prec_C D^0$. Let $\pi^0 \in P$ and $d^0 \in (D^0; D^0)(\pi^0)$. We show that either $d^0 \in \nabla_C$, or there exists $d^{00} \in D^0(\pi^0)$ such that $d^{00} \subseteq d^0$. By definition, there exists a dependency $\{\pi_1, \dots, \pi_n\} \in D^0(\pi^0)$ and dependencies $d_i \in D^0(\pi_i)$ such that $d^0 = \bigcup_{i=1}^n d_i$. By transitivity in repository (P, D, C) , either $d^0 \in \nabla_C$ or there exists $d^{000} \in D(\pi^0)$ such that $d^{000} \subseteq d^0$. In the first case, we can conclude immediately. If $d^{000} \in D^0(\pi^0)$, we can conclude as well by taking $d^{00} = d^{000}$. Otherwise, we must have $\pi^0 = \pi$ and $d^{000} = d$. But, then, $d = d^{000} \subseteq d^0 \in D^0; D^0$. This is not possible due to the second assumption. \blacksquare

Proof of Lemma 14

Let Π be a healthy set of packages of repository $(P, D, C \setminus \{(\pi_1, \pi_2), (\pi_2, \pi_1)\})$. We show that it is healthy in repository (P, D, C) . Abundance is clear. We prove peace, that is, $C \cap (\Pi \times \Pi) = \emptyset$. We have by hypothesis $(C \setminus \{(\pi_1, \pi_2), (\pi_2, \pi_1)\}) \cap (\Pi \times \Pi) = \emptyset$. Thus, we can conclude if packages π_1 and π_2 are not both in Π . The proof is by contradiction: we assume the packages are in Π and reach a contradiction. By abundance in the initial repository, for all a dependency $d \in D(\pi_1)$, there exists a package $\pi_1^0 \in d \cap \Pi$. Then, by definition of redundant conflicts, there exists a dependency $d_1 \in D(\pi_1)$,

a package π_1^0 in $d_1 \cap \Pi$, a dependency $d_2 \in D(\pi_2)$ and a package $\pi_2^0 \in \Pi$ such that $(\pi_1^0, \pi_2^0) \in C \setminus \{(\pi_1, \pi_2), (\pi_2, \pi_1)\}$ and $d_2 \subseteq \{\pi_2^0\}$. By abundance, $d_2 \cap \Pi \neq \emptyset$, hence $\pi_2^0 \in \Pi$. This contradicts peace in the initial repository. Hence the result. ■

Proof of Lemma 15

Let I be an healthy installation of repository (P, D, C) . We show that it is an healthy installation of repository (P, D^0, C) . Clearly, peace is satisfied and abundance is satisfied for all packages distinct from π . We now show abundance for package π . As $D^0(\pi) = \emptyset$, this reduces to showing that $\pi \notin I$. This is indeed the case as π is not installable in the initial repository. ■

Proof of Lemma 16

Let (P^0, D^0, C^0) be the tentative quotiented repository associated to repository (P, D, C) , as defined above. The result is clear, except for irreflexivity of C^0 , which we show below. We assume that $(\pi^0, \pi^0) \in C^0$ for some package $\pi^0 \in P^0$ and reach a contradiction. By definition of C^0 , there exists $\pi_1 \in P$ and $\pi_2 \in P$ such that $\pi^0 = [\pi_1] = [\pi_2]$ and $(\pi_1, \pi_2) \in C$. By Remark 30, there exists a dependency $d \in D(\pi_2)$ such that $d \subseteq \{\pi_2\}$. By definition of quotienting, $D(\pi_1) = D(\pi_2)$, and thus $d \in D(\pi_1)$. But then, by hypothesis, we should have $(\pi_1, \pi_2) \notin C$. Hence a contradiction, as wanted. ■

REMARK 36. *Let (P, D, C) be a repository. Let (I, F) be a healthy configuration of repository (P, D, C) . We assume that the set F is a maximal set (with respect to inclusion) satisfying this property. Then, for all packages π in $P \setminus F$, there exists a package π^0 such that $(\pi, \pi^0) \in C$.*

Proof: We have $\{\pi\} \notin \nabla_C$. Indeed, otherwise, we would have, by definition, $\{\pi\} \cap F \neq \emptyset$, and thus $\pi \in F$. Then, by Theorem 8, there exists π^0 such that $(\pi, \pi^0) \in C$. ■

LEMMA 37. *Let (P, D, C) be a flat repository. Let (I, F) be a healthy configuration of repository (P, D, C) . We assume that the set F is a maximal set (with respect to inclusion) satisfying this property. Let $\pi \in P \setminus F$. Let $\pi^0 \in P$ such that $D(\pi^0) = D(\pi)$ (in other words, $\pi^0 \in [\pi]$). Then, there exists $d \in D(\pi^0)$ such that $d \subseteq \{\pi\}$.*

Proof: By Remark 36, there exists π^{00} such that $(\pi, \pi^{00}) \in C$. By Remark 30, there exists $d \in D(\pi)$ such that $d \subseteq \{\pi\}$. We have $D(\pi^0) = D(\pi)$, hence the result. ■

LEMMA 38. *Let (P, D, C) be a repository. Let (P^0, D^0, C^0) be the corresponding quotient repository. Let (I, F) be a healthy configuration of repository (P, D, C) . We assume that the set F is a maximal set (with respect to inclusion) satisfying this property. Then $([I], \overline{[F]})$ is a healthy configuration of repository (P^0, D^0, C^0) . (We write \overline{F} for the complement $P \setminus F$ of set F .)*

Proof: Let $F^0 = \overline{[F]} = \{[\pi] \in P \mid \forall \pi^0 \in [\pi], \pi^0 \in F\}$.

We first prove abundance. Let $\pi_0^0 \in [I]$ and $d_0^0 \in D^0(\pi_0^0)$. There exists $\pi_0 \in I$ and $d_0 \in D(\pi_0)$ such that $\pi_0^0 = [\pi_0]$ and $d_0^0 = [d_0]$. By abundance in repository (P, D, C) , we

have $d_0 \cap F \neq \emptyset$. Suppose there exists $\pi \in d_0 \cap F$ such that $[\pi] \in F^0$. Then, as $[\pi] \in d_0^0$, we have $d_0^0 \cap F^0 \neq \emptyset$ as wanted. Otherwise, for all package $\pi \in d_0 \cap F$, there exists a package $\pi^0 \in [\pi] \setminus F$. By Lemma 37, there also exists a dependency $d \in D(\pi)$ such that $d \subseteq \{\pi^0\}$, and therefore $d \cap F = \emptyset$. We can apply Lemma 34 to compose d_0 with all these dependencies. Hence, there exists $d^0 \in D(\pi_0) \cup \nabla_C$ such that $d^0 \cap F = \emptyset$. But by abundance and maximality of F , we should have $d^0 \cap F \neq \emptyset$. We reach a contradiction. Hence, this second case is actually not possible.

We now show peace. The proof is by contradiction. Suppose that peace does not hold. Then, there exists $\pi_1^0 \in F^0$ and $\pi_2^0 \in F^0$ such that $(\pi_1^0, \pi_2^0) \in C^0$. By definition of C^0 , there exists a pair $(\pi_1, \pi_2) \in C$ such that $\pi_1^0 = [\pi_1]$ and $\pi_2^0 = [\pi_2]$. By definition of F^0 , $\pi_1 \in F$ and $\pi_2 \in F$. Then, by peace in repository (π, D, C) , $(\pi_1, \pi_2) \notin C$. We reach a contradiction, hence the result. ■

LEMMA 39. *Let (P, D, C) be a repository. Let (P^0, D^0, C^0) be the corresponding quotient repository. Let (I, F) be a healthy configuration of repository (P^0, D^0, C^0) . We pose $I^0 = \{\pi \mid [\pi] \in I\}$ and $F^0 = \{\pi \mid [\pi] \in F\}$. Then (I^0, F^0) is a healthy configuration of repository (P, D, C) .*

Proof: We first prove abundance. Let $\pi \in I^0$ and $d \in D(\pi)$. We need to show that $d \cap F^0 \neq \emptyset$. We have $[\pi] \in I$. By definition of D^0 , we have $[d] \in D^0([\pi])$. By abundance, we thus have $[d] \cap F \neq \emptyset$, that is, there exists a package $\pi \in d$ such that $[\pi] \in F$. We thus have $\pi \in F^0$, hence the result.

We now prove peace, that is $C \cap (I^0 \times I^0) \neq \emptyset$. The proof is by contradiction. Let π_1 and π_2 be two packages in F^0 such that $(\pi_1, \pi_2) \in C$. We have $[\pi_1]$ and $[\pi_2]$ in F and, by definition of C^0 , $([\pi_1], [\pi_2]) \in C^0$. This contradicts peace in the quotient repository. Hence the result. ■

Proof of Theorem 17

This is a direct consequence of Lemmas 38 and 39. ■

Proof of Theorem 18

Let (P, D, C) be a flat repository. We show that its quotient repository (P^0, D^0, C^0) is also flat.

We first prove reflexivity, that is, $\Delta_C \prec_C D^0$. It is sufficient to prove that $\Delta_C \sqsubseteq D^0$. Let $\pi_1, \pi_2 \in P$ such that $([\pi_1], [\pi_2]) \in C^0$. We need to prove that there exists a dependency $d^0 \in D^0([\pi_1])$ such that $d^0 \subseteq \{[\pi_1]\}$. By definition of C^0 , there exists a pair of packages $(\pi_3, \pi_4) \in C$ such that $[\pi_3] = [\pi_1]$ and $[\pi_4] = [\pi_2]$. By reflexivity in the initial repository, either $\{\pi_3\} \in \nabla_C$ or there exists $d \in D(\pi_3)$ such that $d \subseteq \{\pi_3\}$. By Theorem 8, $\{\pi_3\} \notin \nabla_C$. Hence, there exists a dependency $[d]$ in $D^0([\pi_3]) = D^0([\pi_1])$ such that $[d] \subseteq \{[\pi_3]\} = \{[\pi_1]\}$, as wanted.

We now prove transitivity, that is, $D^0; D^0 \prec_C D^0$. Let π be a package. Let $d = \{\pi_1, \dots, \pi_n\}$ a dependency such that $[d] \in D^0([\pi])$ (that is, $d \in D(\pi)$).

Let d_i be dependencies such that $[d_i] \in D^0([\pi_i])$ (that is, $d_i \in D(\pi_i)$) for each i . We need to prove that either $\bigcup_i [d_i] \in \nabla_C^0$ (where ∇_C^0 is our class of always satisfiable dependencies for the quotiented repository) or there exists $d^0 \in D^0([\pi])$ such that $d^0 \subseteq \bigcup_i [d_i]$. We assume that $\bigcup_i [d_i] \notin \nabla_C^0$ and show we are in the second case. By Theorem 8,

for all $[\pi] \in \bigcup_i [d_i]$, there exists a package $\pi^0 \in P^0 \setminus \bigcup_i [d_i]$ such that $([\pi], \pi^0) \in C^0$. This means that we can define a function f which associates to each $[\pi] \in \bigcup_i [d_i]$ a package $f([\pi]) \in [\pi]$ such that there exists a package π^{00} such that $[\pi^{00}] \in P^0 \setminus \bigcup_i [d_i]$ and $(f([\pi]), \pi^{00}) \in C$. Note that $\pi^{00} \in P \setminus f(\bigcup_i [d_i])$. Hence, for any dependency d^{00} , if $d^{00} \subseteq f(\bigcup_i [d_i])$, then $d^{00} \notin \nabla_C$. Now, by reflexivity and Theorem 8 both applied to $f(\pi^{00})$, for each $\pi^{00} \in d_i$, there exists a dependency $d^{00} \in D(f([\pi^{00}]) = D(\pi^{00}))$ such that $d^{00} \subseteq \{f([\pi^{00}])\}$. Thus, for all i , by transitivity applied to d_i and the corresponding d^{00} , there exists a dependency $d_i^{00} \in D(\pi_i)$ such that $d_i^{00} \subseteq f(\bigcup_j [d_j])$. By transitivity again, applied to dependency d and dependencies D_i^{00} , there exists a dependency $d^{00} \in D(\pi)$ such that $d^{00} \subseteq f(\bigcup_i [d_i])$. We have $[d^{00}] \in D^0([\pi])$ and $[d^{00}] \subseteq [f(\bigcup_i [d_i])] = \bigcup_i [d_i]$ as wanted. ■

Proof of Lemma 19

Let I be a healthy installation of (P, D^0, C) . We show that it is a healthy installation of (P, D, C) . Peace is clear. We prove abundance.

Let $\pi^0 \in I$ and $d^0 \in D(\pi^0)$. We need to show that $d^0 \cap I \neq \emptyset$. If $\pi^0 \neq \pi$ or $d^0 \neq d$, then $d^0 \in D^0(\pi^0)$ and the result holds by abundance in (P, D^0, C) . We now show that $d \cap I \neq \emptyset$. As $d \in (D^0; D^0)(\pi)$, there exists a dependency $d^{00} = \{\pi_1, \dots, \pi_n\} \in D^0(\pi)$ and, for all i , a dependency $d_i \in D^0(\pi_i)$ such that $d = \bigcup_{i=1}^n d_i$. By abundance, as $\pi \in I$, we have $d^{00} \cap I \neq \emptyset$, that is, there exists i such that $\pi_i \in I$. Then, by abundance again, we have $d_i \cap I \neq \emptyset$. Hence, $d \cap I \neq \emptyset$ as wanted. ■

Proof of Lemma 20

Let I be a healthy installation of (P, D^0, C) . We show that it is a healthy installation of (P, D, C) . Peace is clear. We prove abundance.

Let $\pi \in I$ and $d \in D(\pi)$. We need to show that $d \cap I \neq \emptyset$. If $\pi \in d$, this is the case. Otherwise, $d \in D^0(\pi)$, and we can conclude by healthiness of I in (P, D^0, C) . ■

Proof of Theorem 21

First, the quotienting operation performed at the end indeed produces a repository, as the hypothesis of Lemma 16 are satisfied: if $\emptyset \in D(\pi)$, then package π has no conflicts thanks to redundant conflict removal; if $\{\pi^0\} \in D(\pi)$, then $(\pi, \pi^0) \notin C$ thanks to the dependency strengthening of clearly broken packages. These two properties are preserved by dependency removal.

We now show that once the loop in Figure 12 is exited, the repository remains flat. Indeed, flattening produces a flat repository (Theorem 3 and Lemma 9), which remains flat by canonisation and removal of dependencies in ∇_C (Remark 29, Lemmas 31 and 33). Then, the removal of conflict cover dependencies is defined so as to preserve flatness, and quotienting preserves flatness (Lemma 18).

To show that co-installability is left invariant, we rely heavily on the equivalence between co-installability and weak co-installability in flat repositories (Lemma 25 and Theorem 10), which let us use weak co-installability preservation results to prove co-installability preservation. We also use the fact that healthiness preservation implies co-installability preservation (Remark 22). Then, co-installability is left invariant by flattening (Theorem 6), canonisation (Theorem 2), removal of dependencies in ∇_C (Theorem 11), de-

pendency strengthening of clearly broken packages (Lemma 15 and Theorem 2), removal of redundant conflicts (Lemma 14 and Theorem 2), removal of conflict covered dependencies (Lemma 12) and Theorem 2), and quotienting (Theorem 17). ■