

Union and Intersection Contracts Are Hard, Actually

Teodoro Freund
Universidad de Buenos Aires
Buenos Aires, Argentina
tfreund@dc.uba.ar

Yann Hamdaoui
Tweag
Paris, France
yann.hamdaoui@tweag.io

Arnaud Spiwack
Tweag
Paris, France
arnaud.spiwack@tweag.io

Abstract

Union and intersection types are a staple of gradually typed languages such as TypeScript. While it's long been recognized that union and intersection types are difficult to verify statically, it may appear at first that the dynamic part of gradual typing is actually pretty simple.

It turns out however, that in presence of higher-order contracts union and intersection are deceptively difficult. The literature on higher-order contracts with union and intersection, while keenly aware of the fact, doesn't really explain why. We point and illustrate the problems and trade-offs inherent to union and intersection contracts, via example and a survey of the literature.

CCS Concepts: • **General and reference** → *Surveys and overviews*; • **Software and its engineering** → **Language features**; *Software verification and validation*.

Keywords: contracts, higher-order contracts, union, intersection

ACM Reference Format:

Teodoro Freund, Yann Hamdaoui, and Arnaud Spiwack. 2021. Union and Intersection Contracts Are Hard, Actually. In *Proceedings of the 17th ACM SIGPLAN International Symposium on Dynamic Languages (DLS '21), October 19, 2021, Chicago, IL, USA*. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3486602.3486767>

1 Introduction

Union types, meaning a type $A \cup B$ containing values which belong either to a type A or B , are a popular tool when adding

static types to a dynamic language. In particular, both TypeScript [9] and MyPy [1], use union types to model the frequent practice to use the value `null` (None in Python) to represent an absent optional value. This is why the gradual typing literature, concerned with formalising the interplay between static and dynamic type systems, has been quite interested in union types [11, 17, 20, 23, 25].

On the other hand, unions are not a common feature of static type systems, mostly because they are quite difficult to verify statically. So unions are really only worth it in gradually typed language where they formalise existing dynamically typed patterns.

Surely, for dynamic tests, unions ought to be really easy: they are simply the Boolean disjunction of two dynamic tests, that fail whenever one of those tests fail. Unfortunately, as we document in this article, as soon as you extend dynamic checks to *contracts* [14], unions become actually pretty difficult, and threaten desirable properties of your language.

1.1 Configuration Languages

To motivate contracts and the problem caused by unions, let's make a detour through configuration languages. A configuration language is a language concerned with describing the configuration of an application. In traditional configuration languages, such as YAML, TOML, or JSON, the configuration is fully, and explicitly, spelt out.

However, with the advent of DevOps, configurations have been extended to describe the entire state of a computer, or even a fleet of computers. For instance, with Kubernetes you need to configure a large fleet of (possibly replicated) docker containers. To describe this sort of configurations, you really want to be able to re-use and abstract parts of the configuration, like traditional programming languages do. To meet this need, languages such as Cue [2], Dhall [6], Jsonnet [7], or Nickel [3], where configurations are generated rather than spelt out, were created.

Another example is continuous integration systems: it's fairly typical to need a matrix of jobs, wherein the same tests are run on different infrastructures, or with different versions of a compiler. Traditional configuration would have you copy the same steps for each infrastructure. This is tedious, hard to maintain, and error prone. It's much better, instead, to write the steps once, and instantiate them for each infrastructure. Continuous integration systems typically do this using a templating system layered on top of

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. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

DLS '21, October 19, 2021, Chicago, IL, USA

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9105-4/21/10...\$15.00

<https://doi.org/10.1145/3486602.3486767>

YAML. Each of the configuration-generating languages above allow such job-matrix definition natively.

1.2 Nickel

In this article, we will use the Nickel language [3] as illustration and motivation. At its core, Nickel is the JSON data model, augmented with abstraction mechanisms, and it comprises:

- dictionaries, written as ¹:
`{field1 = value1, ..., fieldn = valuen}`
- arrays:
`[x1, x2, ..., xn]`
- functions:
`fun arg1 ... argn ⇒ body`
- and let-definitions:
`let id = value in exp`

A Nickel configuration is then evaluated to an explicit configuration, e.g. in JSON, which can then be consumed by an application. Therefore a design constraint of Nickel is any Nickel data must have a straightforward interpretation in JSON.

1.3 Contracts

A useful feature of a configuration language is to provide facilities for schema validation. That is, help answer questions like: does our configuration have all the required fields? does the `url` field indeed contains a URL?

These are inherently dynamic questions, as they are all questions about the evaluated configuration. To this effect, Nickel lets us annotate any expression with a dynamic schema check: `exp | C`. There is also syntactic sugar to annotate definitions: `let id | C = value in exp` stands for `let id = (value | C) in exp`.

Let us pause for a moment and consider the following: it is Nickel's ambition to be able to manipulate configurations like Nixpkgs. With over 50 000 packages, it is one of the largest repository of software packages in existence [8]. Concretely, Nixpkgs is a dictionary mapping packages to build recipes. That is, a massive, over-50 000-key-value-pair wide dictionary. It is absolutely out of the question to evaluate the entirety of this dictionary every time one needs to install 10 new packages: this would result in a painfully slow experience.

To be able to support such large dictionaries, Nickel's dictionaries are *lazy*, that is, the values are only evaluated when explicitly required. For instance, when evaluating the expression `nixpkgs.hello`, only the `hello` package gets evaluated, even if `nixpkgs` contained a `world` package as well.

¹Note that, unlike JSON, Nickel assigns values in dictionaries using `=`, to keep `:` for type annotations.

But let's consider now writing something like `nixpkgs | packages`, to guarantee that all the packages conform to the desired schema. If this were a simple Boolean test, it would have to evaluate all 50 000 package to check their validity, hence breaking the laziness of dictionaries. Do we have to choose between laziness and schema validation? Fortunately, we don't! Enter *contracts* [14]: dynamic checks which can be partially delayed, yet errors can be reported accurately. Contracts can respect laziness of dictionaries, and they can be used to add schema validation to functions as well (in fact functions were the original motivation for contracts).

There is no Boolean function which can check that a value has type `Str → Str`. Instead, a contract for `Str → Str` checks for each call of the function whether

1. the argument has type `Str`, otherwise the caller of the function is faulty
2. if so, that the returned value has type `Str`, otherwise the implementation of the function is faulty

Like in the case of lazy dictionaries, the checks are delayed. Contracts keep track of whether the caller or the implementation is at fault for a violation, hence it can report precise error messages. Contracts are said to *blame* either the caller or the implementation. Compare Figure 1a and Figure 1b: in Figure 1a an error is reported inside the `catHosts` function, but `catHosts` is, in fact, correct, as is made clear by Figure 1b, where `catHosts` is decorated with the `Str → Str` contract, and correctly reports that the caller failed to call `catHosts` with a string argument.

As we shall show, the delayed check of contract, while essential to ensuring that schema validation doesn't affect performance (or, indeed, is possible at all on functions), make union contracts (and their less appreciated sibling, intersection contracts) quite problematic. While usual contracts require only one witness to show the invalidation of a contract, the introduction of unions makes the number of witnesses not bounded.

1.4 Contributions

Our contributions are as follows

- We describe the fundamental difficulties caused by presence of union and intersection contracts in a language, which are kept implicit in the literature (Section 4)
- We survey the various trade-offs which appear in implemented languages and in the academic literature to work around these difficulties (Section 5)

2 A Typology of Language Features

Union contracts are not only difficult to implement, their unrestricted presence is incompatible with potentially desirable properties of the language. In this section we present some of these properties; we will show how these properties interact with union contracts in Sections 4 and 5.

```

let catHosts = fun last =>
  let hosts = ["foo.com", "bar.org"] in
  lists.fold (fun val acc =>
    val ++ ", " ++ acc) hosts last in

let makeHost = fun server ext =>
  server ++ "." ++ ext in

catHosts (makeHost "google")

error: Type error
3 | [...] ", " ++ acc) hosts last in
  |           ^^^
  | This expression has type Fun,
  | but Str was expected
4 |
5 | let makeHost = fun server ext => [...] in
  |           -----
  |                   evaluated to this
  |                   = ++, 2nd argument

```

(a) Error reporting without contract

```

let catHosts | Str → Str = fun last =>
  let hosts = ["foo.com", "bar.org"] in
  lists.fold (fun val acc =>
    val ++ ", " ++ acc) hosts last in

let makeHost = fun server ext =>
  server ++ "." ++ ext in

catHosts (makeHost "google")

error: Blame error: contract broken by the caller.
  | Str → Str
  | --- expected type of the argument [...]
[...]
1 | let catHosts | Str → Str = fun last =>
  |           ^^^^^^^^^^^^^ bound here
[...]
6 | catHosts (makeHost "google")
  | ----- (2) calling <func>

```

(b) Error reporting with contract

Figure 1. Contracts improve error messages

2.1 User-Defined Contracts

A strength of dynamic checking is that we can easily check properties which are impractical to check statically. For instance that a string represents a well-formed URL, or a number is a valid port.

This same property is desirable of contracts as well, otherwise we lose an important benefit of dynamic checking. Preferably, we want to be able to extend the universe of contracts with user-defined predicates.

For instance, Figure 2 shows the definition of a contract for valid ports in Nickel syntax. User-defined contracts can be combined with other contracts normally: $\text{Int} \rightarrow \text{Port}$ is a contract verified by functions which, given an integer returns a valid port.

```

let Port = contracts.fromPred (fun p =>
  num.isInt p && 0 ≤ p && p ≤ 65535) in
80 | Port

```

Figure 2. A contract for valid ports

This type of contracts are present in many different languages, for instance, the Eiffel programming language [19], the precursor of the Design by Contract philosophy, makes it possible to assert these kinds of expression as pre- and post-conditions on functions and as invariants on classes [5].

The Racket programming language also has a system to work with contracts, powerful enough to define user-defined contracts, and to compose them with other kinds of dynamic

checks, like higher order contracts or a lightweight take on union and intersection contracts [4].

2.2 Referential Transparency

The performance of modern programs heavily relies on the optimizations performed by the compiler or the interpreter. Even more so for functional languages, whose execution model is often far removed from the hardware, causing naive execution to exhibit unacceptable slowdowns.

One such important optimization is inlining (Figure 3). Functional programs tend to make heavy use of functions, and a function call is not a free operation: it usually involves a number of low-level operations such as saving and storing registers, pushing a new stack frame and jumping to and back from the function's body. Inlining eliminates a function call by directly substituting the function for its definition at compile time (or before execution, for interpreted language). This is especially efficient for small functions that are called repeatedly.

While inlining expands an expression by substituting a definition for its value, an opposite transformation can be beneficial when a composite expression occurs at several places. In this case, the same expression is wastefully recomputed at each occurrence. Common subexpression elimination (CSE) consists in storing the expression in a variable that is then used in place of the original occurrences (Figure 4), thus evaluating the expression once and for all.

Beyond CSE, other optimizations such as loop-invariant code motion or let-floating [21] apply the same principle of

```

let elem = fun elt  $\Rightarrow$ 
  lists.any (fun x  $\Rightarrow$  x == elt) in

let subList = fun l1 l2  $\Rightarrow$ 
  elem (lists.head l1) l2
  && subList (list.tail l1) l2
  Source program

let subList = fun l1 l2  $\Rightarrow$ 
  lists.any (fun x  $\Rightarrow$  x == (lists.head l1))
  l2
  && subList (list.tail l1) l2
  Optimized program

```

Figure 3. Inlining

```

let elemAtOrLast = fun index list  $\Rightarrow$ 
  if index > lists.length list - 1 then
    lists.elemAt (lists.length list - 1)
    list
  else
    lists.elemAt index list
  Source program

let elemAtOrLast = fun index list  $\Rightarrow$ 
  let l = lists.length list - 1 in
  if index > l then
    lists.elemAt l list
  else
    lists.elemAt index list
  Optimized program

```

Figure 4. Common subexpression elimination

```

let f = fun x  $\Rightarrow$  g y (x + 1)
  Source program

let g' = g y in
let f' = fun x  $\Rightarrow$  g' (x + 1)
  Optimized program

```

Figure 5. Let-floating

extracting out an invariant expression to avoid recomputing it (respectively across loop iterations and function calls).

For example, take the code of Figure 5. The partial application $g\ y$ is recomputed each time f is called. This may be costly, in particular in the presence of contracts: if the

first argument of g must be a list with elements of a specific kind and if that precondition is enforced by a function contract (say $g \mid \text{List Odd} \rightarrow \text{Even} \rightarrow \text{List Odd}$), the additional cost is linear in the size of y . A sensible thing to do is to factor $g\ y$ out of f as in Figure 5, which is something a let-floating transformation could indeed do (given g is pure, as detailed below).

The soundness of these optimizations is tied to the validity of specific program equivalences. Inlining requires that one can replace the application of a function by its body, which is basically β -reduction: as long as the arguments are evaluated following the language's strategy, this is usually a valid transformation. However, Section 4 exposes that the question of inlining a function with a contract attached is more subtle.

A CSE-like transformation on a term M requires on the other hand an equivalence of the form:

$$M[N/x] \simeq \text{let } x = N \text{ in } M \quad (1)$$

$M[N/x]$ stands for the substitution of x for the term N in the term M . This equation clearly fails in presence of side-effects, as demonstrated in Figure 6. In that example, $(f\ 1, f\ 1)$ prints "hi" two times while $\text{let } y = f\ 1 \text{ in } (y, y)$ only prints it once.

```

let f x = print "hi"; (x+1)
  Effectful function

(f 1, f 1)  $\neq$  let y = f 1 in (y, y)
  Invalid expansion

```

Figure 6. Counter-example to (1) in presence of side-effects

However, (1) does hold for *pure* terms, that are terms without side-effects. In a pure language, all these transformations are valid. In impure languages, the situation varies: in some case a large subset of pure terms can be identified (in languages with effects tracking such as PureScript) to be safely transformed. Otherwise, the compiler must stay conservative and only apply CSE to expressions it can prove are without side-effects (arithmetic expressions, for example). Strikingly, we will see in Section 4 that the introduction of union and intersection contracts *breaks referential transparency* and make (1) *unsound*², preventing the kind of optimization of Figure 5 to fire in general.

3 Union & Intersection

Let us now consider union and intersection contracts, before we explain in Section 4 how they can compromise the properties that we described in Section 2.

²(1) holds in Nickel despite non-termination and contract-checking errors because of lazy evaluation. In strict languages (1) needs to be restricted, nevertheless union and intersection contracts make it worse.

```
interface Circle {
  kind: "circle";
  radius: number; }

interface Square {
  kind: "square";
  sideLength: number; }

type Shape = Circle | Square;
```

Figure 7. A sum type as a tagged union

3.1 Unions

A union type $A \cup B$ is a type of values which are either of type A or of type B : literally the union of A and B . Union types are popular in gradual typed systems such as TypeScript [9] and MyPy [1].

In Gradually Typed Systems. The problem that these practical gradual type systems are trying to solve is to capture, in static types, as many programming patterns as possible from the underlying dynamically typed language (JavaScript for TypeScript and Python for MyPy). One such pattern is heterogeneous collections. For instance, in TypeScript, an array which can contain both strings and numbers would have type `Array<string|number>`.

A probably even more common pattern is a variable which can contain either a value of type A (say, a number) or the `null` value. So much so, in fact, that MyPy defines a type alias `Optional[A]` for `Union[A, None]` (`None` is how Python renders the null value).

Yet another application of union types is, rather than capturing a pattern from JavaScript or Python, to capture a pattern from traditional statically typed language: sum types. In statically typed languages, values of sum types are usually thought of as being built out of constructors. But neither JavaScript nor Python have such constructors. So instead, sums are construed as “tagged unions” (or discriminated unions), that is, quite literally, the union of two types which contain a discriminating tag. See Figure 7 for an example from the TypeScript documentation: there the `kind` field is the tag, and its type in both alternatives is a singleton type which contains only the specified string.

Union contracts. In the academic gradual type literature, it is common to use contracts as a glue between static and dynamic types. Therefore, the question of bringing union to contracts is natural, and have indeed been studied (e.g. [17, 25]).

Like for static types, a value which satisfies the contract $A \cup B$ is a value which satisfies either contract A or contract B (though in Section 5 we will see that it may be desirable to weaken this definition).

Nickel is a language built from scratch with contracts, so it may be less clear why unions are useful. However, Nickel’s ambition is to have its data model canonically interpretable in common serialization formats, in particular JSON. It means that it is very convenient to represent optional value by the `null` value like in JavaScript. It also means that Nickel doesn’t have built-in constructors: constructors don’t have a canonical representation in JSON. So it would be quite natural to represent optional contracts and sum contracts as unions.

3.2 Intersections

An intersection type $A \cap B$ is satisfied by values which satisfy both contract A and contract B .

Intersection contracts (and types) are probably less prevalent than union in practical type systems. However, a function from a union is equivalent to an intersection. That is $(A \cup B) \rightarrow C \simeq (A \rightarrow C) \cap (B \rightarrow C)$. So in a system with functions and unions, intersections are already morally present (and, for that matter, in a system with functions and intersections, unions are morally present). Some of our examples in Sections 4 and 5 are better expressed in terms of intersections, so it’s best to include them.

Figure 8 gives a concrete example of this phenomenon. The function `appendDate` appends an element to `list`, whose type is only known to be the union of two lists, each using a different representation. Because the return type is the same as the input type, `appendDate` must preserve this representation (`Date` and `DateWeek` cannot be mixed in a same list). Both alternatives support falling back to a simple string for unparsed dates. Faced with these two possibilities, `appendDate` can only append a value which fits both types: this is precisely the intersection $(\text{Date} \cup \text{Str}) \cap (\text{DateWeek} \cup \text{Str})$, that is, `Str`.

```
let Date = {day | Num, month | Num, year | Num} in
let DateWeek = {dayOfWeek | Num, week | Num, year | Num} in
let appendDate | (List (Date ∪ Str)
  ∪ List (DateWeek ∪ Str))
  → (List (Date ∪ Str)
  ∪ List (DateWeek ∪ Str)) =
  fun list ⇒ lists.cons "01/01/2021" list
```

Figure 8. Adding an element to a union of two arrays

Yet, intersection are useful in their own right: they can be used to combine dictionaries in the style of object-oriented multiple inheritance. For instance, in Figure 9 two types are defined `Animal` and `Pet`, and a variable that is compatible with both types is declared, with type `Animal ∩ Pet`. This particular application is supported, for instance, by TypeScript.


```

let Animal =
  { species | Str, breed | Str, name | Str }
  in
let Pet = { owner | Str, name | Str } in
let myDog | Animal ∩ Pet =
  { species = "Canis Lupus",
    breed   = "Australian Cattle Dog",
    owner   = "Anonymous Author",
    name    = "Juno" }

```

Figure 9. An animal that is also a pet

Another application of intersections shows up when intersecting functions: it can be used to encode overloading. For instance, take a look at Figure 10, where the function `duplicate` works both as a function to duplicate arrays, as well as a function to duplicate strings. This is particularly useful when using unions, since it's a good way to express that a function can deal with different shapes of data. For instance, in the same figure, `duplicate` is (correctly) called on a value of type $(\text{List Str}) \cup \text{Str}$.

```

let duplicate
  | (List Str → List Str)
  ∩ (Str → Str) =
  fun x ⇒ x ++ x in
let text | (List Str) ∪ Str = ... in
duplicate text

```

Figure 10. Duplicating an array of Strings or a String

4 Incompatibilities

However appealing union and intersection contracts may be, they happen to be either hard to combine or even fundamentally incompatible with the desirable language features from Section 2. At least in their full-blown form: in Section 5 we will discuss pragmatic restrictions of union and intersection contracts to recover some or all of the features.

4.1 Union Contracts as a Side-Effect

In Nickel, the failure of a function contract can always be traced back to a single call. For example, take the function `f` with a simple contract attached of Figure 11. The whole program fails with a contract error blaming `f` because the return value of the second call `f 5` violates the `Positive` contract. The first call to `f` does not matter, and `f 5` is a single and independent witness of the contract violation. The user is pointed to this one location in practice.

This single witness property can be justified as follows. Apart from the error reporting part (although this is the crucial bit in practice!), the current contract system of Nickel

can be implemented purely as a library, requiring only a `fail` primitive to abort the execution. In practice, applying a function contract to `f` replaces it with an `f'` that performs the additional checks. Thus, since the core language is pure (albeit partial, if only because `fail`), the failure of `f' 5` must be independent of its environment and of any previous call to `f'`.

```

let f | Positive → Positive
    = fun x ⇒ x - 7 in
(f 10) + (f 5)

```

Figure 11. Simple contract violation

Union contracts are different. Consider the program presented in Figure 12. The same `f` is now given a union contract. `f` is violating this contract once again, as it neither maps all positive numbers to positive numbers nor to non-positive numbers.

```

let f | (Positive → Positive)
      ∪ (Positive → NonPositive)
    = fun x ⇒ x - 7 in
(f 10) + (f 5)

```

Figure 12. Union contract violation

This program must fail, because `f 10` is a witness of `f` failing the contract `Positive → NonPositive`, and `f 5` is a witness of `f` failing `Positive → Positive`. But, as opposed to the example from Figure 11, removing only one of the calls makes the program succeed! Indeed, each call only unveils the violation of one component of the union. In this example, a single call to `f` that would be the witness of the violation of the whole contract doesn't even exist: a minimum of two are always needed.

This behavior shows that union contracts introduce side-effects. The result of `f 5` now depends on the previous execution and more specifically on any prior call to `f`. This behavior of union contracts breaks referential transparency, as well as the property 1 introduced in Section 2.2, that is required to perform CSE-like optimizations.

Figure 13 illustrates this point further. It contains an original program and an optimized version where the common subexpression `f 1` has been eliminated. While equivalent in a pure language with only non-termination or plain higher-order contracts, these two programs behave differently because of unions:

- The original version returns `(1, "False")` without failing.
- The optimized version fails with a contract violation.

In the original version, each partial application $f\ 1$ gives rise to a fresh instance of the contract $\text{Bool} \rightarrow \text{Num} \cup \text{Bool} \rightarrow \text{Str}$. These instances are independent, and can pick a different component of the union to satisfy. Although f doesn't actually respect the contract, these calls are not enough to prove so. In the optimized version, g is endowed with a single contract, that must pick one of the two components of the union. There, the two calls refer to the same union contract, and shows that f does violate its initial contract.

```
let f | Num → (Bool → Num ∪ Bool → Str)
    = fun x y ⇒ if y then x else "False"
in (f 1 true, f 1 false)
```

Original

```
let f | Num → (Bool → Num ∪ Bool → Str)
    = fun x y ⇒ if y then x else "False"
let g = f 1 in
(g true, g false)
```

Optimized

Figure 13. Equivalent programs with CSE applied

To sum up, the addition of union contracts introduce side-effects in a pure language. Side-effects have well-known pitfalls:

- For the programmer, they are hard to reason about. They prevent local reasoning. In our previous examples, removing or adding a function call somewhere can toggle a failure in a call at a totally different location.
- For the interpreter (or compiler), side-effects inhibit many optimizations and program transformations.

4.2 Intersection with User-Defined Contracts

A natural — but naive — implementation of intersection contracts could be the following: to apply a contract $A \cap B$, apply both contracts A and B sequentially, resulting in the naive decomposition rule of Figure 14.

This intuition works for simple contracts: checking that $x \mid \text{Natural} \cap \text{Odd}$ amounts to check that $x \mid \text{Natural}$ and $x \mid \text{Odd}$. Unfortunately, this doesn't scale to higher-order contracts. The overloaded identity example of Figure 14 illustrates the use of an intersection to model a simple overloading of the identity function. If we were to apply the naive decomposition, the argument 1 would fail the $\text{Str} \rightarrow \text{Str}$ contract and abort the execution. Perhaps the exchange rule given in Figure 14, which is a direct consequence of the naive decomposition, illustrates the issue better. It is clear that this exchange law isn't the right semantics for overloading. With this law, the contract for overloaded identity of Figure 14 would always fail because no argument can satisfy $\text{Num} \cap \text{Str}$.

$$M \mid A \cap B \simeq (M \mid A) \mid B$$

Naive decomposition

$$(A \rightarrow B) \cap (C \rightarrow D) \simeq (A \cap C) \rightarrow (B \cap D)$$

Exchange law

```
let g | Num → Num ∩ Str → Str
    = fun x ⇒ x in
g 1
```

Overloaded identity

Figure 14. Naive implementation of intersection

In a higher-order intersection contract, blame is raised when:

Faulty caller The argument fails *both* components.

Faulty implementation The function fails at least *one* component that the argument previously satisfied.

To fix the naive implementation, the interpreter can share state between the sub-contracts, in order to decide if blame must be raised or not when a sub-contract fails:

$$x \mid A \cap B \simeq (x \mid A[1]) \mid B[1]$$

Shared state is represented by the label 1 . Such a shared state is in essence the approach proposed by Williams, Morris, and Wadler in [25].

```
let C = contracts.fromPred (fun f ⇒
  f 0 == 0) in
let g | (Str → Str) ∩ C
    = fun x ⇒ x
in g 0
```

Figure 15. Intersection and user defined contracts

However, this shared-state approach has a major drawback: it isn't straightforwardly compatible with user-defined contracts (introduced in Section 2.1). The issue is similar to our initial issue with higher-order contracts and the naive decomposition: user-defined contracts may apply functions and thus make a sub-contract of the intersection fail, but this failure shouldn't always result in raising blame. An example is given on Figure 15. Decomposing using the shared state approach, we end up with:

$$((\text{fun } x \Rightarrow x) \mid (\text{Str} \rightarrow \text{Str})[1]) \mid C[1]$$

Stateful decomposition

where 1 represents the shared state. At this point, applying the C contract results in evaluating:

$$((\text{fun } x \Rightarrow x) \mid (\text{Str} \rightarrow \text{Str})[1])\ 0 == 0.$$

Applying a function wrapped in a $\text{Str} \rightarrow \text{Str}$ contract to \emptyset fails negatively. This is not the expected behavior, since the identity function does respect semantically both contracts. As opposed to built-in higher-order contracts, user-defined contracts are black-box from the interpreter's point of view, and it is thus not obvious how to extend the shared state approach to handle user-defined contracts.

Once again, intersection contracts introduce side-effects in the picture. What's more, these side-effects interact with user-defined contracts in a non-trivial way, while they are an important feature for validation.

5 Pragmatic Trade-Offs

Despite the difficulties of Section 4, union and intersection contracts are still sought after. In this section we turn to existing systems with union and intersection contracts in the literature and in implementations.

These systems all make trade-offs, sacrificing some features of union and intersection contracts to preserve language features. We survey and discuss those trade-offs and their implications.

5.1 A Coinductive Semantics

In order to give a precise definition to what values ought to satisfy union and intersection contracts, [Keil and Thiemann](#) [17] give a coinductively defined semantics inspired by union and intersection type systems. The key innovation of their work is recognizing that giving a semantics to higher-order contracts requires defining not only what values satisfy a contract, but also what *contexts* satisfy the contract. This models the situation where context may violate a contract by calling a function with an inappropriate argument.

Concretely, given a contract C , [Keil and Thiemann](#) introduce the two sets $\llbracket C \rrbracket^+$ and $\llbracket C \rrbracket^-$ of values and contexts, respectively, satisfying the contract. They are defined by mutual induction and coinduction.

This semantics has limited support for overloading. Consider the example in Figure 16: it could evaluate to the pair $(1, 1)$. But [Keil and Thiemann](#)'s coinductive semantics rejects it as a contract violation. This can be phrased pithily as the fact that the coinductive semantics doesn't validate the property $A \rightarrow B \cap A \rightarrow C \simeq A \rightarrow (B \cap C)$.

```
let f = fun x y => x in
let g = f | (Num → Num → Num)
        ∩ (Num → Bool → Num) in
let h = g 1 in
(h 1, h true)
```

Figure 16. Intersection contracts don't distribute

A solution, for the programmer, is to use an uncurried function $\text{fun } (x, y) \Rightarrow x$. So one way to think about this limitation is that currying function is not fully supported.

Note that if the last two lines of Figure 16 had read $(g\ 1\ 1, g\ 1\ \text{true})$ instead, then the coinductive semantics would accept the example. It implies that under the coinductive semantics, common-subexpression elimination (see Section 2.2) is quite perilous.

5.2 A First Realization

In Section 4, we've seen that different calls to a function with a union contract must share information: the behavior of one call is influenced by the previous ones, as the function must pick one component of the union to satisfy across all usages. Conversely, following the semantics of overloading, each application of a function with an intersection contract can select a different branch and is thus independent from the others. A general contract composed of nested unions, intersections and higher-order contracts appears to require complex book-keeping in order to correctly raise blame.

All of this still holds true of the coinductive semantics described in Section 5.1. Nevertheless, [Keil and Thiemann](#) give an algorithmic system which is complete for their coinductive semantics. In a remarkable technical tour de force, their algorithmic system allows for user-defined contracts (see Section 4.2).

A key aspect of the approach of [Keil and Thiemann](#) is to rewrite nested union and intersection contracts into a disjunctive normal form using the De Morgan's law $A \cap (B \cup C) \simeq (A \cap B) \cup (A \cap C)$. The goal is to be able to delay the choice of branch in intersections as much as possible.

To implement contract verification, [Keil and Thiemann](#) resort to specific reduction rules for unions and intersections which perform this rewriting on the fly. This aspect is critiqued in [Williams et al.](#) [25]: "the monitoring semantics for contracts of intersection and union types given by Keil and Thiemann are not uniform. (...) If uniformity helps composition, then special cases can hinder composition."

A cost of this approach is that the De Morgan's law $A \cap (B \cup C) \simeq (A \cap B) \cup (A \cap C)$ duplicates contract A , which will cause some contracts to be checked several times. This can be an issue with user-defined contracts which may include costly tests.

Efficiency is also affected another way: each time a function with a contract attached is applied, the whole *context* must be traversed to check for a compatibility property.

The algorithmic system of [Keil and Thiemann](#) is, on balance, a technically impressive realization of the coinductive semantics that supports user-defined contracts, though it is fairly complex and probably difficult to implement efficiently.

5.3 Monitoring Properties

Another realization of the coinductive semantics described in Section 5.1 is given in Williams et al. [25], which aims at simplifying the algorithmic system proposed by Keil and Thiemann and described in Section 5.2.

A key ingredient of Williams et al. is to disallow user-defined contracts. This choice gives the authors more freedom in the quest of a more uniform operational semantics. This is sensible trade-off in the context of gradual typing à la TypeScript: the problem is to match contracts with static types, and user-defined contracts don't have a static type equivalent. On the other hand, the cost would probably not be worth it for a configuration language like Nickel.

As a means of proving the correctness of their simplified system, the authors introduce what they call *sound monitoring properties*. Here is the sound monitoring property for contexts of intersection contracts:

$$K \in \llbracket A \cap B \rrbracket^- \text{ if } K \in \llbracket A \rrbracket^- \vee K \in \llbracket B \rrbracket^-$$

This reads as: a context K satisfies the intersection of A and B if it satisfies at least one of the two. Morally, the K s in $\llbracket A \cap B \rrbracket^-$ should be the ones that can have their hole filled with a term satisfying $A \cap B$ without violating the $A \cap B$ contract.

Although sound, this interpretation is weaker than what the coinductive semantics permits. Consider the two contexts presented on Figure 17. The first one is a context satisfying $\text{Num} \rightarrow \text{Num}$, applying the hole to a number. Similarly, the second context from the same figure satisfies $\text{Bool} \rightarrow \text{Bool}$.

```
□ 3
    Num → Num context

□ true
    Bool → Bool context
```

Figure 17. Two different contexts in Nickel

Now, combining these two contexts as in Figure 18 gives a context that doesn't satisfy $\text{Num} \rightarrow \text{Num}$ nor $\text{Bool} \rightarrow \text{Bool}$. According to the sound monitoring property of intersection, Figure 18 thus doesn't satisfy $\text{Num} \rightarrow \text{Num} \cap \text{Bool} \rightarrow \text{Bool}$.

```
let f = □ in
(f 3, f true)
```

Figure 18. Combined context

The consequence is that Williams et al. don't prove their system complete for the coinductive semantics. It's probably just an oversight in the proof: we believe their system to be

indeed complete for the coinductive semantics. But it does speak to the intrinsic complexity of union and intersection contracts: it is remarkably easy to get details wrong. This difficulty contrasts with the standard framework of higher-order contracts where satisfaction is much more straightforward.

5.4 Racket

Racket is a language based on the Scheme dialect of Lisp. Among established languages, Racket is probably the one with the most comprehensive contract system [4]. Regarding union and intersection, Racket provides the **and/c** and **or/c** combinators for contract.

The **and/c** combinator corresponds to the naive interpretation of intersection described in Section 4.2: applying contract **(and/c A B)** is like applying contract A then applying contract B . In particular **and/c** doesn't model overloading. As in Figure 14, the example given in Figure 19 always fail because no argument satisfies both number? and string? .

```
(define/contract overload
  (and/c (→ number? number?)
        (→ string? string?)))
(lambda (x) x))
```

Figure 19. **and/c** and overloading

The union combinator **or/c** is similarly simple. It must be able to decide immediately which branch holds: **or/c** is a simple Boolean disjunction. For instance, when higher-order contracts are combined using **or/c**, Racket imposes that contracts must be distinguishable by their arity. Doing so, there is at most one candidate that can be selected directly. This is illustrated in Figure 20 whose program is accepted. On the other hand, the program of Figure 21 is rejected.

```
(define/contract united
  (or/c (→ number? number?)
        (→ string? string? string?)))
(lambda (x) x))
```

Figure 20. Accepted use of **or/c** with higher-order contracts

```
(define/contract united
  (or/c (→ number? number?)
        (→ even? even?)))
(lambda (x) x))
```

Figure 21. Rejected use of **or/c** with higher-order contracts

case→. To compensate for the fact that **and/c** doesn't support overloading, Racket provides a second intersection-like combinator: **case→**. As for the **or/c**, the candidate contracts must have distinct arities to avoid ambiguity, unlike **or/c** the alternative chosen when the function is called rather than when the contract is applied to the function. An example is provided in Figure 22. The resulting possibilities are similar to static overloading, where one function can take additional parameters for example (e.g. as supported for Java methods). On the other hand, it excludes the overloading of generic operations with fixed arity such as equality, comparison, arithmetic operators, and so on.

```
(define/contract overcase
  (case→ (→ string? string?)
         (→ number? number? number?))
  )
  (lambda (x [y 0]) (if (number? x)
                        (+ x y)
                        x)))

(overcase 1 2)

(overcase "hello")
```

Figure 22. Overloading with the **case→** combinator

In conclusion, Racket, a programming language with a large user base, avoids the difficulties of general unions and intersections (Section 4). They make the pragmatic choice of a simple semantics with limited support for higher-order contracts in intersections and unions.

6 (More) Related Work

6.1 Higher-Order Contracts

Enforcing pre- and post-conditions at runtime is a widely established practice. In their foundational paper [14], Findler and Felleisen introduce higher-order contracts, a principled approach to run-time assertion checking that nicely supports functions. They introduce the notion of blame, which is crucial to good error reporting. It became apparent later that their contracts are closely related to the type casts introduced by gradual typing, excluding blame: both [18] and [22] see the value of contracts as a safe interface between typed and untyped code. In [24], the authors precisely introduce a system integrating gradual typing with contracts *à la Findler & Felleisen*. Nickel adopts a similar type system, with both statically typed terms, dynamically typed terms, and first-class contracts. Higher-order contracts are the basis of the work of Keil and Thiemann [17] and Williams et al. [25] that this paper explored extensively.

6.2 Unions and Intersections in Gradual Typing

Castagna and Lanvin [10] introduce a gradual type system based on set-theoretic types [15]. Set-theoretic types feature unions, intersections, negation types together with a notion of subtyping. This work adheres to the *static first school* (see [16]) of gradual typing: the reason for gradual types is to allow for a less precise type information. It follows that their goals and constraints are somehow different, resulting in the absence of first-class contracts. As for any gradual type system, they do have to implement casts — that are very close to contracts — for unions and intersections, but these casts are neither visible nor available to the programmer. Obviously, no contracts mean no user-defined contracts as well.

Castagna and Lanvin use abstract interpretation to derive the semantics of unions and intersections. In their words: “the resulting definitions are quite technical and barely intuitive but they have the properties we seek for [...]”. This makes any comparison with the coinductive semantics of Keil and Thiemann rather difficult. Castagna et al. [11] build on Castagna and Lanvin to add polymorphism in the picture. However, it is at the price of restricting the union and intersection part: it is not possible to assign intersection types to a function anymore.

6.3 Context Sensitive Contracts

The work by Dimoulas et al. [12, 13] presents a complete study on systems that may execute (and potentially invalidate) contracts inside other contracts, in particular through the study of dependent function contracts. Their ideas have the potential to help solve the problems outlined on Sections 4.2 and 5.2, however, it isn't clear how to leverage their techniques for intersections.

7 Conclusion

Despite the fact that union and intersection of dynamic properties may at first appear like an easy task, as soon as they are combined with higher-order contracts for increased accuracy of error messages, they really aren't.

The problem of union and intersection contracts is that they are not orthogonal to apparently independent features of programming languages. The mere presence of intersection and union contract induces computational effects, and can make it quite difficult to perform simple program optimizations such as inlining.

Designing a language with union and intersection contracts necessarily means making difficult choices: some features of union and intersection contracts and of the rest of the language must be abandoned. Various trade-offs can be made, but it is worth mentioning that implementing a system with union and intersection contracts appears to be pretty complex a task when unions and intersections are fairly complete.

It's hard not to have sympathy for the minimalist end of this spectrum, where, like Racket, a language only has a very simple notion of unions and intersections. Many applications of unions and intersections are not possible in such a context, but the presence of union and intersection contracts doesn't interact with the rest of the language. It's probably more manageable, and it's the approach that we are currently considering for Nickel.

To conclude, let us make clear that we do not think union and intersection contracts are fundamentally broken, that they can not be implemented correctly, or that they do not bear any value (quite the contrary). They may still make sense to have in a language, and some apparent difficulties in the implementation could be lifted some day. But as often, there are gaps between the theoretical foundation, a proof-of-concept, a prototype, and the integration in an actual language. We hope that our attempt may serve as a cautionary tale: for union and intersection contracts, these gaps may be larger than they appear.

References

- [1] 2016. *Optional types and the None type*. https://web.archive.org/web/20210910024940/https://mypy.readthedocs.io/en/latest/kinds_of_types.html#optional-types-and-the-none-type
- [2] 2019. *The CUE Configuration Language*. <https://web.archive.org/web/20210910025022/https://cuelang.org/>
- [3] 2020. *The Nickel repository on GitHub*. <https://archive.softwareheritage.org/swh:1:dir:66ab3b7cda21c78149b1b51eb1c7827d3e35e77b;origin=https://github.com/tweag/nickel;visit=swh:1:snp:9147140b16412b3c288b3d235bf28c481ae46b6d;anchor=swh:1:rev:5cb1ba649eb145fc0fee82eece8c62a6b67b597e>
- [4] 2021?. *Contracts - Racket Documentation*. <https://web.archive.org/web/20210814023650/https://docs.racket-lang.org/reference/contracts.html>
- [5] 2021?. *Design by Contract and Assertions - Eiffel*. https://web.archive.org/web/20210116003625/https://www.eiffel.org/doc/solutions/Design_by_Contract_and_Assertions
- [6] 2021?. *The Dhall configuration language*. <https://web.archive.org/web/20210910025152/https://dhall-lang.org/>
- [7] 2021?. *Jsonnet - The Data Templating Language*. <https://web.archive.org/web/20210910025206/https://jsonnet.org/>
- [8] 2021. *Repology package tracker*. <https://web.archive.org/web/20210910024725/https://repology.org/repositories/statistics/total>
- [9] 2021. *Union Types in TypeScript, TypeScript handbook*. <https://web.archive.org/web/20210910024843/https://www.typescriptlang.org/docs/handbook/2/everyday-types.html#union-types>
- [10] Giuseppe Castagna and Victor Lanvin. 2017. Gradual Typing with Union and Intersection Types. *Proc. ACM Program. Lang.* 1, ICFP, Article 41 (Aug. 2017), 28 pages. <https://doi.org/10.1145/3110285>
- [11] Giuseppe Castagna, Victor Lanvin, Tommaso Petrucciani, and Jeremy G. Siek. 2019. Gradual Typing: A New Perspective. *Proc. ACM Program. Lang.* 3, POPL, Article 16 (Jan. 2019), 32 pages. <https://doi.org/10.1145/3290329>
- [12] Christos Dimoulas, Robert Bruce Findler, Cormac Flanagan, and Matthias Felleisen. 2011. Correct Blame for Contracts: No More Scapegoating. In *Proceedings of the 38th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages* (Austin, Texas, USA) (POPL '11). Association for Computing Machinery, New York, NY, USA, 215–226. <https://doi.org/10.1145/1926385.1926410>
- [13] Christos Dimoulas, Sam Tobin-Hochstadt, and Matthias Felleisen. 2012. Complete Monitors for Behavioral Contracts. In *Programming Languages and Systems*, Helmut Seidl (Ed.). Springer Berlin Heidelberg, Berlin, Heidelberg, 214–233.
- [14] Robert Bruce Findler and Matthias Felleisen. 2002. Contracts for Higher-Order Functions. *SIGPLAN Not.* 37, 9 (Sept. 2002), 48–59. <https://doi.org/10.1145/583852.581484>
- [15] Alain Frisch, Giuseppe Castagna, and Véronique Benzaken. 2008. Semantic subtyping: Dealing set-theoretically with function, union, intersection, and negation types. *J. ACM* 55 (09 2008). <https://doi.org/10.1145/1391289.1391293>
- [16] Michael Greenberg. 2019. The Dynamic Practice and Static Theory of Gradual Typing. In *3rd Summit on Advances in Programming Languages (SNAPL 2019) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 136)*, Benjamin S. Lerner, Rastislav Bodík, and Shriram Krishnamurthi (Eds.). Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 6:1–6:20. <https://doi.org/10.4230/LIPIcs.SNAPL.2019.6>
- [17] Matthias Keil and Peter Thiemann. 2015. Blame Assignment for Higher-Order Contracts with Intersection and Union. In *Proceedings of the 20th ACM SIGPLAN International Conference on Functional Programming (Vancouver, BC, Canada) (ICFP 2015)*. Association for Computing Machinery, New York, NY, USA, 375–386. <https://doi.org/10.1145/2784731.2784737>
- [18] Jacob Matthews and Robert Bruce Findler. 2007. Operational Semantics for Multi-Language Programs. In *Proceedings of the 34th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (Nice, France) (POPL '07)*. Association for Computing Machinery, New York, NY, USA, 3–10. <https://doi.org/10.1145/1190216.1190220>
- [19] B Meyer. 1987. Eiffel: Programming for Reusability and Extendibility. *SIGPLAN Not.* 22, 2 (Feb. 1987), 85–94. <https://doi.org/10.1145/24686.24694>
- [20] Francisco Ortin and Miguel García. 2011. Union and intersection types to support both dynamic and static typing. *Inform. Process. Lett.* 111, 6 (2011), 278–286. <https://doi.org/10.1016/j.ipl.2010.12.006>
- [21] Simon Peyton Jones, Will Partain, and André Santos. 1996. Let-Floating: Moving Bindings to Give Faster Programs. In *Proceedings of the First ACM SIGPLAN International Conference on Functional Programming (Philadelphia, Pennsylvania, USA) (ICFP '96)*. Association for Computing Machinery, New York, NY, USA, 1–12. <https://doi.org/10.1145/232627.232630>
- [22] Sam Tobin-Hochstadt and Matthias Felleisen. 2006. Interlanguage Migration: From Scripts to Programs. In *Companion to the 21st ACM SIGPLAN Symposium on Object-Oriented Programming Systems, Languages, and Applications (Portland, Oregon, USA) (OOPSLA '06)*. Association for Computing Machinery, New York, NY, USA, 964–974. <https://doi.org/10.1145/1176617.1176755>
- [23] Matías Toro and Éric Tanter. 2017. A Gradual Interpretation of Union Types. 382–404. https://doi.org/10.1007/978-3-319-66706-5_19
- [24] Philip Wadler and Robert Bruce Findler. 2009. Well-Typed Programs Can't Be Blamed. In *Proceedings of the 18th European Symposium on Programming Languages and Systems: Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2009 (York, UK) (ESOP '09)*. Springer-Verlag, Berlin, Heidelberg, 1–16. https://doi.org/10.1007/978-3-642-00590-9_1
- [25] Jack Williams, J. Garrett Morris, and Philip Wadler. 2018. The Root Cause of Blame: Contracts for Intersection and Union Types. *Proc. ACM Program. Lang.* 2, OOPSLA, Article 134 (Oct. 2018), 29 pages. <https://doi.org/10.1145/3276504>