

GADTs meet Subtyping

Gabriel Scherer, Didier Rémy

Gallium – INRIA

2014

A reminder on GADTs

GADTs are algebraic data types that may carry *type equalities*. Think of the following simple type:

```
type expr =  
  | Int of int  
  | Bool of bool
```

A reminder on GADTs

GADTs are algebraic data types that may carry *type equalities*. Think of the following simple type:

```
type expr =  
  | Int of int  
  | Bool of bool
```

It can be turned into the more finely typed:

```
type  $\alpha$  expr =  
  | Int of int with  $\alpha = \text{int}$   
  | Bool of bool with  $\alpha = \text{bool}$ 
```

```
type  $\alpha$  expr =  
  | Int : int -> int expr  
  | Bool : bool -> bool expr
```

A reminder on GADTs

GADTs are algebraic data types that may carry *type equalities*. Think of the following simple type:

```
type expr =  
  | Int of int  
  | Bool of bool
```

It can be turned into the more finely typed:

```
type  $\alpha$  expr =  
  | Int of int with  $\alpha = \text{int}$   
  | Bool of bool with  $\alpha = \text{bool}$   
type  $\alpha$  expr =  
  | Int : int  $\rightarrow$  int expr  
  | Bool : bool  $\rightarrow$  bool expr
```

We can now write the following:

```
let eval :  $\forall \alpha. \alpha \text{ expr} \rightarrow \alpha = \text{function}$   
  | Int n  $\rightarrow$  n      (*  $\alpha = \text{int}$  *)  
  | Bool b  $\rightarrow$  b  (*  $\alpha = \text{bool}$  *)
```

Motivating variance

Subtyping: $\sigma \leq \tau$ means “all values of σ are also values of τ ”.
Checked by set of decidable and incomplete inference rules.

$$\frac{\sigma_1 \geq \sigma'_1 \quad \sigma_2 \leq \sigma'_2}{(\sigma_1 \rightarrow \sigma_2) \leq (\sigma'_1 \rightarrow \sigma'_2)}$$

Variance annotations lift subtyping to type parameters.

type $(-\alpha, =\beta, +\gamma) \mathbf{t} = (\alpha * \beta) \rightarrow (\beta * \gamma)$

$$\frac{\alpha \geq \alpha' \quad \beta = \beta' \quad \gamma \leq \gamma'}{(\alpha, \beta, \gamma) \mathbf{t} \leq (\alpha', \beta', \gamma') \mathbf{t}}$$

For simple types, this is easy to check.

Variance for GADT: harder than it seems

Ok?

```
type +α expr =  
  | Val : ∀α. α → α expr  
  | Prod : ∀βγ. β expr * γ expr → (β * γ) expr
```

Variance for GADT: harder than it seems

Ok?

```
type + $\alpha$  expr =  
  | Val :  $\forall \alpha. \alpha \rightarrow \alpha$  expr  
  | Prod :  $\forall \beta \gamma. \beta$  expr *  $\gamma$  expr  $\rightarrow (\beta * \gamma)$  expr
```

And this one?

```
type file_descr = private int (* file_descr  $\leq$  int *)  
val stdin : file_descr
```

```
type + $\alpha$  t =  
  | File : file_descr -> file_descr t  
let o = File stdin in  
let o' = (o : file_descr t :> int t)
```

Variance for GADT: harder than it seems

Ok?

```
type + $\alpha$  expr =  
  | Val :  $\forall \alpha. \alpha \rightarrow \alpha$  expr  
  | Prod :  $\forall \beta \gamma. \beta$  expr *  $\gamma$  expr  $\rightarrow (\beta * \gamma)$  expr
```

And this one?

```
type file_descr = private int (* file_descr  $\leq$  int *)  
val stdin : file_descr
```

```
type + $\alpha$  t =  
  | File : file_descr -> file_descr t  
let o = File stdin in  
let o' = (o : file_descr t :> int t)
```

Breaks abstraction!

```
let project :  $\forall \alpha. \alpha$  t  $\rightarrow (\alpha \rightarrow$  file_descr) = function  
  | File _ -> (fun x -> x)  
project o' : int -> file_descr
```


Proving an example correct

type $+\alpha$ expr =

| Val : $\forall \alpha. \alpha \rightarrow \alpha$ expr

| Prod : $\forall \beta \gamma. \beta$ expr * γ expr $\rightarrow (\beta * \gamma)$ expr

When $\sigma \leq \sigma'$, I know it's safe to assume σ expr $\leq \sigma'$ expr.

Because I could *almost* write that conversion myself.

Proving an example correct

```
type +α expr =  
  | Val : ∀α. α → α expr  
  | Prod : ∀βγ. β expr * γ expr → (β * γ) expr
```

When $\sigma \leq \sigma'$, I know it's safe to assume $\sigma \text{ expr} \leq \sigma' \text{ expr}$.
Because I could *almost* write that conversion myself.

```
let coerce (α ≤ α') : α expr ≤ α' expr = function  
  | Val (v : α) -> Val (v :> α')  
  | Prod β γ ((b, c) : β expr * γ expr) ->  
    (* α = (β * γ), α ≤ α'; Prod? *)
```

Proving an example correct

```
type +α expr =  
  | Val : ∀α. α → α expr  
  | Prod : ∀βγ. β expr * γ expr → (β * γ) expr
```

When $\sigma \leq \sigma'$, I know it's safe to assume $\sigma \text{ expr} \leq \sigma' \text{ expr}$.
Because I could *almost* write that conversion myself.

```
let coerce (α ≤ α') : α expr ≤ α' expr = function  
  | Val (v : α) -> Val (v :> α')  
  | Prod β γ ((b, c) : β expr * γ expr) ->  
    (* α = (β * γ), α ≤ α'; Prod? *)  
    (* if β * γ ≤ α', then α' is of the form  
       β' * γ' with β ≤ β' and γ ≤ γ' *)  
    Prod β' γ' ((b :> β' expr), (c :> γ' expr))
```

Proving an example correct

```
type +α expr =  
  | Val : ∀α. α → α expr  
  | Prod : ∀βγ. β expr * γ expr → (β * γ) expr
```

When $\sigma \leq \sigma'$, I know it's safe to assume $\sigma \text{ expr} \leq \sigma' \text{ expr}$.
Because I could *almost* write that conversion myself.

```
let coerce (α ≤ α') : α expr ≤ α' expr = function  
  | Val (v : α) -> Val (v :> α')  
  | Prod β γ ((b, c) : β expr * γ expr) ->  
    (* α = (β * γ), α ≤ α'; Prod? *)  
    (* if β * γ ≤ α', then α' is of the form  
       β' * γ' with β ≤ β' and γ ≤ γ' *)  
    Prod β' γ' ((b :> β' expr), (c :> γ' expr))
```

Upward closure for $\tau[\bar{\alpha}]$:

If $\tau[\bar{\sigma}] \leq \tau'$, then τ' is also of the form $\tau[\bar{\sigma}']$ for some $\bar{\sigma}'$.

Holds for $\alpha * \beta$, but fails for `file_descr = private int`.

In the general case

Consider a GADT $\alpha \text{ t}$ with a constructor of the form

| K of $\exists \bar{\beta} [\alpha = T[\bar{\beta}]] . \tau[\bar{\beta}]$

Imagine I have a value v of type $\sigma \text{ t}$, and I know $\sigma \leq \sigma'$. Can I convert this $\sigma \text{ t}$ into a $\sigma' \text{ t}$? Let's write the coercion code again:

```
match v :  $\sigma \text{ t}$  with
```

```
...
```

```
| K arg -> (arg :  $\tau[\bar{\rho}]$  :>  $\tau[?]$ )
```

We can type-check this coercion term when $\sigma \leq \sigma'$ if and only if

$$\forall \bar{\rho}, \quad \sigma = T[\bar{\rho}] \implies \exists \bar{\rho}', \quad \sigma' = T[\bar{\rho}'] \wedge \tau[\bar{\rho}] \leq \tau[\bar{\rho}']$$

This extends both upward-closure and the usual variance check on τ .

The semantic criterion

A GADT declaration for $\overline{\nu\alpha} \text{ t}$ is correct if, for each constructor K of type $\exists\overline{\beta}[D[\overline{\alpha}, \overline{\beta}]].\tau[\overline{\beta}]$, we have

$$\forall\overline{\sigma}\overline{\sigma}'\overline{\rho}, \quad \overline{\sigma} \text{ t} \leq \overline{\sigma}' \text{ t} \wedge D[\overline{\sigma}, \overline{\rho}] \implies \exists\overline{\rho}', \quad D[\overline{\sigma}', \overline{\rho}'] \wedge \tau[\overline{\rho}] \leq \tau[\overline{\rho}']$$

How can we check this?

What does it even mean?

Our job: get something *syntactic* out of this semantic criterion, that compilers *and* humans can understand and use.

The plan

We will first explain how to check variance of type variables by a judgment

$\Gamma \vdash \tau : v$.


Resembles previous work with a twist.

We will then extend it to a judgment $\Gamma \vdash \tau : v \Rightarrow v'$ to check closure properties.

From there it's not too hard (but not too easy either) to derive the final syntactic formulation of the correctness criterion.

Variances

- + : only positive occurrences
- - : only negative occurrences




- = : both positive and negative

$$\begin{aligned}\sigma \prec_+ \tau &::= \sigma \leq \tau \\ \sigma \prec_- \tau &::= \sigma \geq \tau \\ \sigma \prec_= \tau &::= \sigma = \tau\end{aligned}$$

Variances

- + : only positive occurrences
- - : only negative occurrences




- = : both positive and negative
- ⊗ : no occurrence at all

$$\begin{aligned}\sigma \prec_+ \tau &::= \sigma \leq \tau \\ \sigma \prec_- \tau &::= \sigma \geq \tau \\ \sigma \prec_= \tau &::= \sigma = \tau \\ \sigma \prec_{\otimes} \tau &::= \mathbf{true}\end{aligned}$$

Variances

- + : only positive occurrences
- - : only negative occurrences



- = : both positive and negative
- ⊗ : no occurrence at all

$$\begin{aligned}\sigma \prec_+ \tau &::= \sigma \leq \tau \\ \sigma \prec_- \tau &::= \sigma \geq \tau \\ \sigma \prec_= \tau &::= \sigma = \tau \\ \sigma \prec_{\otimes} \tau &::= \text{true}\end{aligned}$$

If α has variance v in $(\alpha \text{ t})$, and β variance w in $(\beta \text{ u})$,
what is the variance of α in $((\alpha \text{ t}) \text{ u})$?

| $v.w$ | = | + | - | ⊗ | w |
|-------|---|---|---|---|-----|
| = | = | = | = | ⊗ | |
| + | = | + | - | ⊗ | |
| - | = | - | + | ⊗ | |
| ⊗ | ⊗ | ⊗ | ⊗ | ⊗ | |

v

$\Gamma \vdash \tau : \nu$

We manipulate contexts Γ of variables with variances: $(-\alpha, =\beta, +\gamma)$.
 $\Gamma \vdash \tau : \nu$ means that “if the variables vary according to their variance, τ varies along ν ”.

$$-\alpha, =\beta, +\gamma \vdash (\alpha * \beta) \rightarrow (\beta * \gamma) : (+)$$

$$=\alpha, =\beta, =\gamma \vdash (\alpha * \beta) \rightarrow (\beta * \gamma) : (=)$$

$\Gamma \vdash \tau : v$

We manipulate contexts Γ of variables with variances: $(-\alpha, =\beta, +\gamma)$.
 $\Gamma \vdash \tau : v$ means that “if the variables vary according to their variance, τ varies along v ”.

$$-\alpha, =\beta, +\gamma \vdash (\alpha * \beta) \rightarrow (\beta * \gamma) : (+)$$

$$=\alpha, =\beta, =\gamma \vdash (\alpha * \beta) \rightarrow (\beta * \gamma) : (=)$$

$$\frac{w\alpha \in \Gamma \quad w \geq v}{\Gamma \vdash \alpha : v}$$

$$\frac{\Gamma \vdash \overline{w\alpha} \ t \quad \forall i, \Gamma \vdash \sigma_i : v.w_i}{\Gamma \vdash \overline{\sigma} \ t : v}$$

For instance, in the arrow case:

$$\frac{\Gamma \vdash \sigma_1 : v.- \quad \Gamma \vdash \sigma_2 : v.+}{\Gamma \vdash (\sigma_1 \rightarrow \sigma_2) : v}$$

Closure properties in depth

In our system, $\alpha * \beta$ is upward-closed. This is because the head type constructor, $(*)$, is closed.

For $\alpha \tau \rightarrow (\beta * \gamma)$ to be upward-closed, $\alpha \tau$ must be downward-closed. In the general case, we recursively check closure, according to variance.

What about variables?

What about variables?

[Reminder] *Upward closure*: If $\tau[\bar{\sigma}] \leq \tau'$, then $\tau' = \tau[\bar{\sigma}']$ for some $\bar{\sigma}'$.

$\beta * \beta$ is not closed : $(\text{file_descr} * \text{file_descr}) \leq (\text{file_descr} * \text{int})$.

Repeating a variable twice is dangerous.

What about variables?

[Reminder] *Upward closure*: If $\tau[\bar{\sigma}] \leq \tau'$, then $\tau' = \tau[\bar{\sigma}']$ for some $\bar{\sigma}'$.

$\beta * \beta$ is not closed : $(\text{file_descr} * \text{file_descr}) \leq (\text{file_descr} * \text{int})$.

Repeating a variable twice is dangerous.

Yet, $(\beta \text{ ref}) * (\beta \text{ ref})$ is closed... because *all* occurrences are invariant.

What about variables?

[Reminder] *Upward closure*: If $\tau[\bar{\sigma}] \leq \tau'$, then $\tau' = \tau[\bar{\sigma}']$ for some $\bar{\sigma}'$.

$\beta * \beta$ is not closed : $(\text{file_descr} * \text{file_descr}) \leq (\text{file_descr} * \text{int})$.

Repeating a variable twice is dangerous.

Yet, $(\beta \text{ ref}) * (\beta \text{ ref})$ is closed... because *all* occurrences are invariant.

We capture those subtleties through a partial variance operation $v_1 \hat{\wedge} v_2$.
Defined only when two occurrences at variances v_1 and v_2 can be soundly combined.

| | | | | | |
|--------------------|---|---|---|---|-----|
| $v \hat{\wedge} w$ | = | + | - | ⊗ | w |
| = | = | | | = | |
| + | | | | + | |
| - | | | | - | |
| ⊗ | = | + | - | ⊗ | |
| v | | | | | |

$$\Gamma \vdash \tau : v \Rightarrow v'$$

We can finally extend the judgment $\Gamma \vdash \tau : v$ to capture closure properties. We want to say that $\Gamma \vdash \tau$ is v -closed if:

$$\forall \tau', \bar{\sigma}, \tau[\bar{\sigma}] \prec_v \tau' \implies \exists \bar{\sigma}', \tau[\bar{\sigma}'] = \tau'$$

We need a generalization:

$$\forall \tau', \bar{\sigma}, \tau[\bar{\sigma}] \prec_v \tau' \implies \exists \bar{\sigma}', \tau[\bar{\sigma}'] \prec_{v'} \tau'$$

This is our $\Gamma \vdash \tau : v \Rightarrow v'$ judgment.

Inference rules for the show

They rely on closure information for type constructors, and \wedge to merge contexts of subterms.

Inference rules for the show

They rely on closure information for type constructors, and λ to merge contexts of subterms.

$$\frac{\text{TRIV} \quad v \geq v' \quad \Gamma \vdash \tau : v}{\Gamma \vdash \tau : v \Rightarrow v'}$$

$$\frac{\text{VAR} \quad w\alpha \in \Gamma \quad w = v}{\Gamma \vdash \alpha : v \Rightarrow v'}$$

$$\frac{\text{CONSTR} \quad \Gamma \vdash \overline{w\alpha} \mathbf{t} : v\text{-closed} \quad \Gamma = \lambda_i \Gamma_i \quad \forall i, \Gamma_i \vdash \sigma_i : v.w_i \Rightarrow v'.w_i}{\Gamma \vdash \overline{\sigma} \mathbf{t} : v \Rightarrow v'}$$

How do you check those rules?

Theorem

$\Gamma \vdash \tau : v$ holds if and only if $\forall \bar{\sigma}, \bar{\sigma}', \bar{\sigma} \prec_{\Gamma} \bar{\sigma}' \implies \tau[\bar{\sigma}] \prec_v \tau[\bar{\sigma}']$

Simple.

How do you check those rules?

Theorem

$\Gamma \vdash \tau : v$ holds if and only if $\forall \bar{\sigma}, \bar{\sigma}', \bar{\sigma} \prec_{\Gamma} \bar{\sigma}' \implies \tau[\bar{\sigma}] \prec_v \tau[\bar{\sigma}']$

Simple.

Theorem

$\Gamma \vdash \tau : v \implies (=)$ holds if and only if $\Gamma \vdash \tau : v$ and $\forall \bar{\rho}, \tau', \tau[\bar{\rho}] \prec_v \tau' \implies \exists \bar{\rho}', \tau[\bar{\rho}'] = \tau'$

Soundness (syntactic \implies semantic): routine.

Completeness (semantic \implies syntactic): surprisingly hard.

Back to the check

From these primitives, we can devise a syntactic check for our initial semantic criterion. Assume we have the variances $\overline{v\alpha}$ and a constructor declaration of the form $(\exists \overline{\beta} [\bigwedge_{i \in I} \alpha_i = T_i[\overline{\beta}]] . \tau[\overline{\beta}])$. Remember the *sophisticated* semantic criterion:

$$\forall \overline{\sigma}, \overline{\sigma}', \overline{\rho}, \quad \overline{\sigma} \mathbf{t} \leq \overline{\sigma}' \mathbf{t} \wedge (\sigma_i = T_i[\overline{\rho}])_{i \in I} \implies \\ \exists \overline{\rho}', (\sigma'_i = T_i[\overline{\rho}'])_{i \in I} \wedge \tau[\overline{\rho}] \leq \tau[\overline{\rho}']$$

Back to the check

From these primitives, we can devise a syntactic check for our initial semantic criterion. Assume we have the variances $\overline{v\alpha}$ and a constructor declaration of the form $(\exists \overline{\beta} [\bigwedge_{i \in I} \alpha_i = T_i[\overline{\beta}]] . \tau[\overline{\beta}])$. Remember the *sophisticated* semantic criterion:

$$\forall \overline{\sigma}, \overline{\sigma}', \overline{\rho}, \quad \overline{\sigma} \text{ t} \leq \overline{\sigma}' \text{ t} \wedge (\sigma_i = T_i[\overline{\rho}])_{i \in I} \implies \\ \exists \overline{\rho}', (\sigma'_i = T_i[\overline{\rho}'])_{i \in I} \wedge \tau[\overline{\rho}] \leq \tau[\overline{\rho}']$$

Theorem (Algorithmic criterion)

The soundness criterion above is equivalent to

$$\exists \Gamma, (\Gamma_i)_{i \in I}, \quad \Gamma \vdash \tau : (+) \wedge \Gamma = \bigwedge_{i \in I} \Gamma_i \wedge \forall i \in I, \Gamma_i \vdash T_i : v_i \Rightarrow (=)$$

(Oral explanation)

Phew !

That was the formal side of things.

Phew !

That was the formal side of things.

This criterion raises interesting design issues: `private` definitions make all OCaml types non-downward-closed.

Should we restrict those? The opposite of OOP's `final` keyword.

Another solution

We showed, through hard work, how to check that equality constraints are upward-closed.

With subtyping in constructor types, variance is easy to check

```
type + $\alpha$  expr =  
| Val :  $\forall \beta \geq \alpha. \beta \rightarrow \alpha$  expr  
| Prod :  $\forall \beta \gamma [\alpha \geq (\beta * \gamma)]. \beta$  expr *  $\gamma$  expr  $\rightarrow \alpha$  expr
```

Now pattern-matching only knows a subtyping relation:

```
let eval :  $\forall \alpha. \alpha$  expr  $\rightarrow \alpha$  = function  
| Int n -> (n :>  $\alpha$ )          (*  $\alpha \geq$  int *)  
| Bool b -> (b :>  $\alpha$ )        (*  $\alpha \geq$  bool *)
```

Less convenient: use of subtyping must be annotated explicitly, while equations where implicit.

Future work:

- Completeness of the S&P criterion: type inhabitation.
- Verify behavior through type abstraction.
- Does this also happen with inductive dependent types?

Conclusion

GADT variance checking: suprisingly less obvious than we thought.

Not anecdotal: raises deeper design questions.

We have a sound criterion that can be implemented easily in a type checker.

Bonus Slide: Variance and the value restriction

```
type (= 'a) ref = { mutable contents : 'a }
```

In a language with mutable data, generalizing any expression is unsafe (because you may generalize data locations):

```
# let test = ref [];;  
val test : '_a list ref
```

Solution (Wright, 1992): only generalize values (fun () -> ref [], or []).

Painful when manipulating polymorphic data structures:

```
let test = id [] (* not generalized? *)
```

OCaml relies on variance for the *relaxed value restriction* covariant data is immutable, so covariant type variables may be safely generalized. Very useful in practice.

```
# let test = id [];;  
val test : 'a list = []
```