

System F with Type Equality Coercions

Martin Sulzmann

School of Computing
National University of Singapore
sulzmann@comp.nus.edu.sg

Manuel M. T. Chakravarty

Computer Science & Engineering
University of New South Wales
chak@cse.unsw.edu.au

Simon Peyton Jones

Microsoft Research Ltd
Cambridge, England

Kevin Donnelly

{simonpj,t-kevind}@microsoft.com

Abstract

We introduce System F_C , which extends System F with support for non-syntactic type equality. There are two main extensions: (i) explicit witnesses for type equalities, and (ii) open, non-parametric type functions, given meaning by top-level equality axioms. Unlike System F, F_C is expressive enough to serve as a target for several different source-language features, including Haskell’s `newtype`, generalised algebraic data types, associated types, functional dependencies, and perhaps more besides.

Categories and Subject Descriptors D.3.1 [Programming Languages]: Formal Definitions and Theory—Semantics; F.3.3 [Logics and Meanings of Programs]: Studies of Program Constructs—Type structure

General Terms Languages, Theory

Keywords Typed intermediate language, advanced type features

1. Introduction

The polymorphic lambda calculus, System F, is popular as a highly-expressive typed intermediate language in compilers for functional languages. However, language designers have begun to experiment with a variety of type systems that are difficult or impossible to translate into System F, such as functional dependencies [21], generalised algebraic data types (GADTs) [43, 31], and associated types [6, 5]. For example, when we added GADTs to GHC, we extended GHC’s intermediate language with GADTs as well, even though GADTs are arguably an over-sophisticated addition to a typed intermediate language. But when it came to associated types, even with this richer intermediate language, the translation became extremely clumsy or in places impossible.

In this paper we resolve this problem by presenting System $F_C(X)$, a super-set of F that is both *more foundational* and *more powerful* than adding *ad hoc* extensions to System F such as GADTs or associated types. $F_C(X)$ uses explicit type-equality coercions as witnesses to justify explicit type-cast operations. Like types, coercions are erased before running the program, so they are guaranteed to have no run-time cost.

This single mechanism allows a very direct encoding of associated types and GADTs, and allows us to deal with some exotic functional-dependency programs that GHC currently rejects on the grounds that they have no System-F translation (§2). Our specific contributions are these:

- We give a formal description of System F_C , our new intermediate language, including its type system, operational semantics, soundness result, and erasure properties (§3). There are two distinct extensions. The first, explicit equality witnesses, gives a system equivalent in power to System F + GADTs (§3.2); the second introduces non-parametric type functions, and adds substantial new power, well beyond System F + GADTs (§3.3).
- A distinctive property of F_C ’s type functions is that they are *open* (§3.4). Here we use “open” in the same sense that Haskell type classes are open: just as a newly defined type can be made an instance of an existing class, so in F_C we can extend an existing type function with a case for the new type. This property is crucial to the translation of associated types.
- The system is very general, and its soundness requires that the axioms stated as part of the program text are *consistent* (§3.5). That is why we call the system $F_C(X)$: the “X” indicates that it is parametrised over a decision procedure for checking consistency, rather than baking in a particular decision procedure. (We often omit the “(X)” for brevity.) Conditions identified in earlier work on GADTs, associated types, and functional dependencies, already define such decision procedures.
- A major goal is that F_C should be a *practical* compiler intermediate language. We have paid particular attention to ensuring that F_C programs are robust to program transformation (§3.8).
- It must obviously be *possible* to translate the source language into the intermediate language; but it is also highly desirable that it be *straightforward*. We demonstrate that F_C has this property, by sketching a type-preserving translation of two source language idioms, namely GADTs (Section 4) and associated types (Section 5). The latter, and the corresponding translation for functional dependencies, are more general than all previous type-preserving translations for these features.

System F_C has no new foundational content: rather, it is an intriguing and practically-useful application of techniques that have been well studied in the type-theory community. Several other calculi exist that might in principle be used for our purpose, but they generally do not handle open type functions, are less robust to transformation, and are significantly more complicated. We defer a comparison with related work until §6.

To substantiate our claim that F_C is practical, we have implemented it in GHC, a state-of-the-art compiler for Haskell, including both GADTs and associated (data) types. This is not just a prototype; F_C now *is* GHC’s intermediate language.

F_C does not strive to do everything; rather we hope that it strikes an elegant balance between expressiveness and complexity. While our motivating examples were GADTs and associated types, we believe that F_C may have much wider application as a typed target for sophisticated HOT (higher-order typed) source languages.

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.

TLDI’07 January 16, 2007, Nice, France.

Copyright © 2007 ACM 1-59593-393-X/07/0001...\$5.00.

2. The key ideas

No compiler uses *pure* System F as an intermediate language, because some source-language constructs can only be desugared into pure System F by very heavy encodings. A good example is the algebraic data types of Haskell or ML, which are made more complicated in Haskell because algebraic data types can capture existential type variables. To avoid heavy encoding, most compilers invariably extend System F by adding algebraic data types, data constructors, and `case` expressions. We will use F_A to describe System F extended in this way, where the data constructors are allowed to have existential components [24], type variables can be of higher kind, and type constructor applications can be partial.

Over the last few years, source languages (notably Haskell) have started to explore language features that embody *non-syntactic* or *definitional* type equality. These features include functional dependencies [16], generalised algebraic data types (GADTs) [43, 36], and associated types [6, 5]. All three are difficult or impossible to translate into System F — and yet the alternative of simply extending System F by adding functional dependencies, GADTs, and associated types, seems wildly unattractive. Where would one stop?

In the rest of this section we informally present System F_C , an extension of System F that resolves the dilemma. We show how it can serve as a target for each of the three examples. The formal details are presented in §3. Throughout we use `typewriter` font for source-code, and *italics* for F_C .

2.1 GADTs

Consider the following simple type-safe evaluator, often used as the poster child of GADTs, written in the GADT extension of Haskell supported by GHC:

```
data Exp a where
  Zero :: Exp Int
  Succ :: Exp Int -> Exp Int
  Pair :: Exp b -> Exp c -> Exp (b, c)
```

```
eval :: Exp a -> a
eval Zero      = 0
eval (Succ e)  = eval e + 1
eval (Pair x y) = (eval x, eval y)
```

```
main = eval (Pair (Succ Zero) Zero)
```

The key point about this program, and the aspect that is hard to express in System F, is that in the `Zero` branch of `eval`, the type variable `a` is the same as `Int`, even though the two are syntactically quite different. That is why the `0` in the `Zero` branch is well-typed in a context expecting a result of type `a`.

Rather than extend the intermediate language with GADTs themselves — GHC’s pre- F_C “solution” — we instead propose a general mechanism of parameterising functions with *type equalities*, written $\sigma_1 \sim \sigma_2$, witnessed by *coercions*. Coercion types are passed around using System F’s existing type passing facilities and enable representing GADTs by ordinary algebraic data types encapsulating such type equality coercions.

Specifically, we translate the GADT `Exp` to an ordinary algebraic data type, where each variant is parameterised by a coercion:

```
data Exp :  $\star \rightarrow \star$  where
  Zero :  $\forall a. (a \sim \text{Int}) \Rightarrow \text{Exp } a$ 
  Succ :  $\forall a. (a \sim \text{Int}) \Rightarrow \text{Exp } \text{Int} \rightarrow \text{Exp } a$ 
  Pair :  $\forall abc. (a \sim (b, c)) \Rightarrow \text{Exp } b \rightarrow \text{Exp } c \rightarrow \text{Exp } a$ 
```

So far, this is quite standard; indeed, several authors present GADTs in the source language using a syntax involving explicit equality constraints, similar to that above [43, 10]. However, for us the equality constraints are extra type arguments to the constructor,

which must be given when the constructor is applied, and which are brought into scope by pattern matching. The “ \Rightarrow ” is syntactic sugar, and we sloppily omitted the kind of the quantified type variables, so the type of `Zero` is really this:

$$\text{Zero} : \forall a : \star. \forall (co : a \sim \text{Int}). \text{Exp } a$$

Here a ranges over *types*, of kind \star , while co ranges over *coercions*, of kind $a \sim \text{Int}$. An important property of our approach is that *coercions are types*, and hence, *equalities $\tau_1 \sim \tau_2$ are kinds*. An equality kind $\tau_1 \sim \tau_2$ categorises all coercion types that witness the interchangeability of the two types τ_1 and τ_2 . So, our slogan is *propositions as kinds*, and *proofs as (coercion) types*.

Coercion types may be formed from a set of elementary coercions that correspond to the rules of equational logic; for example, $\text{Int} : (\text{Int} \sim \text{Int})$ is an instance of the reflexivity of equality and $\text{sym } co : (\text{Int} \sim a)$, with $co : (a \sim \text{Int})$, is an instance of symmetry. A call of the constructor `Zero` must be given a type (to instantiate a) and a coercion (to instantiate co), thus for example:

$$\text{Zero } \text{Int } \text{Int} : \text{Exp } \text{Int}$$

As indicated above, regular types like `Int`, when interpreted as coercions, witness reflexivity.

Just like value arguments, the coercions passed to a constructor when it is built are made available again by pattern matching. Here, then, is the code of `eval` in F_C :

```
eval =  $\Lambda a : \star. \lambda e : \text{Exp } a.$ 
  case e of
    Zero (co :  $a \sim \text{Int}$ )  $\rightarrow$ 
      0  $\blacktriangleright$  sym co
    Succ (co :  $a \sim \text{Int}$ ) (e' :  $\text{Exp } \text{Int}$ )  $\rightarrow$ 
      (eval Int e' + 1)  $\blacktriangleright$  sym co
    Pair (b :  $\star$ ) (c :  $\star$ ) (co :  $a \sim (b, c)$ )
      (e1 :  $\text{Exp } b$ ) (e2 :  $\text{Exp } c$ )  $\rightarrow$ 
      (eval b e1, eval c e2)  $\blacktriangleright$  sym co
```

The form $\Lambda a : \star. e$ abstracts over types, as usual. In the first alternative of the `case` expression, the pattern binds the coercion type argument of `Zero` to co . We use the symmetry of equality in $(\text{sym } co)$ to get a coercion from `Int` to a and use that to cast the type of `0` to a , using the *cast expression* $0 \blacktriangleright \text{sym } co$. Cast expressions have no *operational* effect, but they serve to explain to the type system when a value of one type (here `Int`) should be treated as another (here a), and provide evidence for this equivalence. In general, the form $e \blacktriangleright g$ has type t_2 if $e : t_1$ and $g : (t_1 \sim t_2)$. So, $\text{eval } \text{Int} (\text{Zero } \text{Int } \text{Int})$ is of type `Int` as required by `eval`’s signature. We shall discuss coercion types and their kinds in more detail in §3.2.

In a similar manner, the recently-proposed extended algebraic data types [40], which add equality and predicate constraints to GADTs, can be translated to F_C .

2.2 Associated types

Associated types are a recently-proposed extension to Haskell’s type-class mechanism [6, 5]. They offer open, type-indexed types that are associated with a type class. Here is a standard example:

```
class Collects c where
  type Elem c      -- associated type synonym
  empty  :: c
  insert :: Elem c -> c -> c
```

The class `Collects` abstracts over a family of containers, where the representation type of the container, `c`, defines (or constrains) the type of its elements `Elem c`. That is, `Elem` is a type-level function that transforms the collection type to the element type. Just as `insert` is non-parametric – its implementation varies depending on `c` – so is `Elem`. For example, a list container can contain

elements of any type supporting equality, and a bit-set container might represent a collection of characters:

```
instance Eq e => Collects [e] where
  {type Elem [e] = e; ...}
instance Collects BitSet where
  {type Elem BitSet = Char; ...}
```

Generally, type classes are translated into System F [17] by (1) turning each class into a record type, called a dictionary, containing the class methods, (2) converting each instance into a dictionary value, and (3) passing such dictionaries to whichever function mentions a class in its signature. For example, a function of type `negate :: Num a => a -> a` will translate to `negate :: NumDict a -> a -> a`, where `NumDict` is the record generated from the class `Num`.

A record only encapsulates values, so what to do about associated types, such as `Elem` in the example? The system given in [6] translates each associated type into an additional type parameter of the class's dictionary type, provided the class and instance declarations abide by some moderate constraints [6]. For example, the class `Collects` translates to dictionary type `CollectsDict c e`, where e represents `Elem c` and where all occurrences of `Elem c` of the method signatures have been replaced by the new type parameter e . So, the (System F) type for `insert` would now be `CollectDict c e -> e -> c -> c`. The required type transformations become more complex when associated types occur in data types; the data types have to be rewritten substantially during translation, which can be a considerable burden in a compiler.

Type equality coercions enable a far more direct translation. Here is the translation of `Collects` into F_C :

```
type Elem : * -> *
data CollectsDict c =
  Collects {empty : c; insert : Elem c -> c -> c}
```

The dictionary type is as in a translation without associated types. The `type` declaration in F_C introduces a new *type function*. An instance declaration for `Collects` is translated to (a) a dictionary transformer for the values and (b) an equality axiom that describes (part) of the interpretation for the type function `Elem`. For example, here is the translation into F_C of the `Collects BitSet` instance:

```
axiom elemBS : Elem BitSet ~ Char
dCollectsBS : CollectsDict BitSet
dCollectsBS = ...
```

The `axiom` definition introduces a new, named *coercion constant*, `elemBS`, which serves as a witness of the equality asserted by the axiom; here, that we can convert between types of the form `Elem BitSet` and `Char`. Using this coercion, we can `insert` the character 'b' into a `BitSet` by applying the coercion `elemBS` backwards to 'b', thus:

```
('b' ▶ (sym elemBS)) : Elem BitSet
```

This argument fits the signature of `insert`.

In short, System F_C supports a very direct translation of associated types, in contrast to the clumsy one described in [6]. What is more, there are several obvious extensions to the latter paper that cannot be translated into System F at all, even clumsily, and F_C supports them too, as we sketch in Section 5.

2.3 Functional dependencies

Functional dependencies are another popular extension of Haskell's type-class mechanism [21]. With functional dependencies, we can encode a function over types F as a relation, thus

```
class F a b | a -> b
instance F Int Bool
```

However, some programs involving functional dependencies are impossible to translate into System F. For example, a useful idiom in type-level programming is to abstract over the co-domain of a type function by way of an existential type, the `b` in this example:

```
data T a = forall b. F a b => MkT (b -> b)
```

In this Haskell declaration, `MkT` is the constructor of type `T`, capturing an existential type variable `b`. One might hope that the following function would type-check:

```
combine :: T a -> T a -> T a
combine (MkT f) (MkT f') = MkT (f . f')
```

After all, since the type `a` functionally determines `b`, `f` and `f'` must have the same type. Yet GHC rejects this program, because it cannot be translated into System F_A , because `f` and `f'` each have distinct, existentially-quantified types, and there is no way to express their (non-syntactic) identity in F_A .

It is easy to translate this example into F_C , however:

```
type F1 : * -> *
data FDict : * -> * -> * where
  F : ∀ a b. (b ~ F1 a) => FDict a b
axiom fIntBool : F1 Int ~ Bool
data T : * -> * where
  MkT : ∀ a b. FDict a b -> (b -> b) -> T a

combine : T a -> T a -> T a
combine (MkT b (F (co : b ~ F1 a)) f)
  (MkT b' (F (co' : b' ~ F1 a)) f')
  = MkT a b (F a b co) (f . (f' ▶ d2))
where
  d1 : (b' ~ b) = co' ◦ sym co
  d2 : (b' -> b' ~ b -> b) = d1 -> d1
```

The functional dependency is expressed as a type function `F1`, with one equality axiom per instance. (In general there might be many functional dependencies for a single class.) The dictionary for class `F` includes a witness that indeed `b` is equal to `F1 a`, as you can see from the declaration of constructor `F`. When pattern matching in `combine`, we gain access to these witnesses, and can use them to cast `f'` so that it has the same type as `f`. (To construct the witness d_1 we use the coercion combinators `sym ·` and `· ◦`, which represent symmetry and transitivity; and from d_1 we build the witness d_2 .)

Even in the absence of existential types, there are reasonable source programs involving functional dependencies that have no System F translation, and hence are rejected by GHC. We have encountered this problem in real programs, but here is a boiled-down example, using the same class `F` as before:

```
class D a where { op :: F a b => a -> b }
instance D Int where { op _ = True }
```

The crucial point is that the context `F a b` of the signature of `op` constrains the parameter of the enclosing type class `D`. This becomes a problem when typing the definition of `op` in the instance `D Int`. In `D`'s dictionary `DDict`, we have `op : ∀ b. C a b -> a -> b` with `b` universally quantified, but in the instance declaration, we would need to instantiate `b` with `Bool`. The instance declaration for `D` cannot be translated into System F. Using F_C , this problem is easily solved: the coercion in the dictionary for `F` enables the result of `op` to be cast to type `b` as required.

To summarise, a compiler that uses translation into System F (or F_A) must reject some reasonable (albeit slightly exotic) programs involving functional dependencies, and also similar programs involving associated types. The extra expressiveness of System F_C solves the problem neatly.

2.4 Translating newtype

F_C is extremely expressive, and can support language features beyond those we have discussed so far. Another example are Haskell 98's `newtype` declarations:

```
newtype T = MkT (T -> T)
```

In Haskell, this declares T to be isomorphic to $T \rightarrow T$, but there is no good way to express that in System F. In the past, GHC has handled this with an *ad hoc* hack, but F_C allows it to be handled directly, by introducing a new axiom

```
axiom CoT : (T -> T) ~ T
```

2.5 Summary

In this section we have shown that System F is inadequate as a typed intermediate language for source languages that embody non-syntactic type equality — and Haskell has developed several such extensions. We have sketchily introduced System F_C as a solution to these difficulties. We will formalise it in the next section.

3. System $F_C(X)$

The main idea in $F_C(X)$ is that we pass around explicit evidence for type equalities, in just the same way that System F passes types explicitly. Indeed, in F_C the evidence γ for a type equality *is* a type; we use type abstraction for evidence abstraction, and type application for evidence application. Ultimately, we erase all types before running the program, and thereby erase all type-equality evidence as well, so the evidence passing has no run-time cost. However, that is not the only reason that it is better to represent evidence as a *type* rather than as a *term*, as we discuss in §3.10.

Figure 1 defines the syntax of System F_C , while Figures 2 and 3 give its static semantics. The notation \bar{a}^n (where $n \geq 0$) means the sequence $a_1 \cdots a_n$; the “ n ” may be omitted when it is unimportant. Moreover, we use comma to mean sequence extension as follows: $\bar{a}^n, a_{n+1} \triangleq \bar{a}^{n+1}$. We use $fv(x)$ to denote the free variables of a structure x , which may be an expression, type term, or environment.

3.1 Conventional features

System F_C is a superset of System F. The syntax of types and kinds is given in Figure 1. Like F, F_C is impredicative, and has no stratification of types into polytypes and monotypes. The meta-variables $\varphi, \rho, \sigma, \tau, v$, and γ all range over types, and hence also over coercions. However, we adopt the convention that we use ρ, σ, τ , and v in places where we can only have regular types (i.e., no coercions), and we use γ in places where we can only have coercion types. We use φ for types that can take either form. This choice of meta-variables is only a convention to aid the reader; formally, the coercion typing and kinding rules enforce the appropriate restrictions.

Our system allows types of higher kind; hence the type application form $\tau_1 \tau_2$. However, like Haskell but unlike $F\omega$, our system has no type-level lambda, and type equality is syntactic identity (modulo alpha-conversion). This choice has pervasive consequences; it gives a remarkable combination of economy and expressiveness, but leaves some useful higher-kinded types out of reach. For example, there is no way to write the type constructor $(\lambda a. \text{Either } a \text{ Bool})$.

Value type constructors T range over (a) the built-in function type constructor, (b) any other built-in types, such as *Int*, and (c) algebraic data types. We regard a function type $\sigma_1 \rightarrow \sigma_2$ as the curried application of the built-in function type constructor to two arguments, thus $(\rightarrow) \sigma_1 \sigma_2$. Furthermore, although we give the syntax of arrow types and quantified types in an uncurried way, we also

Symbol Classes

a, b, c, co	\rightarrow	\langle type variable
x, f	\rightarrow	\langle term variable
C	\rightarrow	\langle coercion constant
T	\rightarrow	\langle value type constructor
S_n	\rightarrow	$\langle n$ -ary type function
K	\rightarrow	\langle data constructor

Declarations

pgm	\rightarrow	$\overline{decl}; e$
$decl$	\rightarrow	$\mathbf{data} \ T : \overline{\kappa} \rightarrow \star \ \mathbf{where}$ $\quad \overline{K : \forall \bar{a} : \overline{\kappa}. \forall \bar{b} : \bar{\iota}. \bar{\sigma} \rightarrow T \ \bar{a}}$ $\quad \ \mathbf{type} \ S_n : \overline{\kappa}^n \rightarrow \iota$ $\quad \ \mathbf{axiom} \ C : \sigma_1 \sim \sigma_2$

Sorts and kinds

δ	\rightarrow	$\text{TY} \mid \text{CO}$	Sorts
κ, ι	\rightarrow	$\star \mid \kappa_1 \rightarrow \kappa_2 \mid \sigma_1 \sim \sigma_2$	Kinds

Types and Coercions

d	\rightarrow	$a \mid T$	Atom of sort TY
g	\rightarrow	$c \mid C$	Atom of sort CO
$\varphi, \rho, \sigma, \tau, v, \gamma$	\rightarrow	$a \mid C \mid T \mid \varphi_1 \varphi_2 \mid S_n \overline{\varphi}^n \mid \forall a : \kappa. \varphi$ $\quad \ \text{sym } \gamma \mid \gamma_1 \circ \gamma_2 \mid \gamma @ \varphi \mid \text{left } \gamma \mid \text{right } \gamma$ $\quad \ \gamma \sim \gamma \mid \text{rightc } \gamma \mid \text{leftc } \gamma \mid \gamma \blacktriangleright \gamma$	

We use ρ, σ, τ , and v for regular types, γ for coercions, and φ for both.

Syntactic sugar

Types $\kappa \Rightarrow \sigma \equiv \forall _ : \kappa. \sigma$

Terms

u	\rightarrow	$x \mid K$	Variables and data constructors
e	\rightarrow	u	Term atoms
		$\Lambda a : \kappa. e \mid e \varphi$	Type abstraction/application
		$\lambda x : \sigma. e \mid e_1 e_2$	Term abstraction/application
		$\mathbf{let} \ x : \sigma = e_1 \ \mathbf{in} \ e_2$	
		$\mathbf{case} \ e_1 \ \mathbf{of} \ \overline{p} \rightarrow \overline{e_2}$	Cast
		$e \blacktriangleright \gamma$	
p	\rightarrow	$K \ \overline{b} : \overline{\kappa} \ \overline{x} : \overline{\sigma}$	Pattern

Environments

$\Gamma \rightarrow \epsilon \mid \Gamma, u : \sigma \mid \Gamma, d : \kappa \mid \Gamma, g : \kappa \mid \Gamma, S_n : \kappa$
A *top-level environment* binds only type constructors, T, S_n , data constructors K , and coercion constants C .

Figure 1: Syntax of System $F_C(X)$

sometimes use the following syntactic sugar:

$$\begin{aligned} \overline{\varphi}^n \rightarrow \varphi_r &\equiv \varphi_1 \rightarrow \cdots \rightarrow \varphi_n \rightarrow \varphi_r \\ \forall \overline{\alpha}^n. \varphi &\equiv \forall \alpha_1 \cdots \forall \alpha_n. \varphi \end{aligned}$$

An *algebraic data type* T is introduced by a top-level `data` declaration, which also introduces its *data constructors*. The type of a data constructor K takes the form

$$K : \forall \bar{a} : \overline{\kappa}. \forall \bar{b} : \bar{\iota}. \bar{\sigma} \rightarrow T \ \bar{a}^n$$

The first n quantified type variables \bar{a} appear in the same order in the return type $T \ \bar{a}$. The remaining quantified type variables bind either existentially quantified type variables, or (as we shall see) coercions.

Types are classified by *kinds* κ , using the \vdash_{TY} judgement in Figure 2. Temporarily ignoring the kind $\sigma_1 \sim \sigma_2$, the structure of kinds

$\Gamma \vdash_{TY} \sigma : \kappa$		
(TyVar) $\frac{d : \kappa \in \Gamma \quad \Gamma \vdash_k \kappa : TY}{\Gamma \vdash_{TY} d : \kappa}$	(TyApp) $\frac{\Gamma \vdash_{TY} \sigma_1 : \kappa_1 \rightarrow \kappa_2 \quad \Gamma \vdash_{TY} \sigma_2 : \kappa_1}{\Gamma \vdash_{TY} \sigma_1 \sigma_2 : \kappa_2}$	
(TySCon) $\frac{(S_n : \overline{\kappa}^n \rightarrow \iota) \in \Gamma \quad \Gamma \vdash_{TY} \overline{\sigma} : \overline{\kappa}^n}{\Gamma \vdash_{TY} S_n \overline{\sigma}^n : \iota}$	(TyAll) $\frac{\Gamma, a : \kappa \vdash_{TY} \sigma : \star \quad \Gamma \vdash_k \kappa : \delta \quad a \notin fv(\Gamma)}{\Gamma \vdash_{TY} \forall a : \kappa. \sigma : \star}$	
$\Gamma \vdash_{CO} \gamma : \sigma \sim \tau$		
(CoRefl) $\frac{a : \kappa \in \Gamma \quad \Gamma \vdash_k \kappa : TY}{\Gamma \vdash_{CO} a : a \sim a}$	(CoVar) $\frac{g : \sigma \sim \tau \in \Gamma}{\Gamma \vdash_{CO} g : \sigma \sim \tau}$	
(CoAllT) $\frac{\Gamma, a : \kappa \vdash_{CO} \gamma : \sigma \sim \tau \quad \Gamma \vdash_k \kappa : TY \quad a \notin fv(\Gamma)}{\Gamma \vdash_{CO} \forall a : \kappa. \gamma : \forall a : \kappa. \sigma \sim \forall a : \kappa. \tau}$	(CoInstT) $\frac{\Gamma \vdash_{CO} \gamma : \forall a : \kappa. \sigma \sim \forall b : \kappa. \tau \quad \Gamma \vdash_{TY} v : \kappa}{\Gamma \vdash_{CO} \gamma @ v : [v/a]\sigma \sim [v/b]\tau}$	
(SComp) $\frac{\Gamma \vdash_{CO} \overline{\gamma} : \overline{\sigma} \sim \overline{\tau}^n \quad \Gamma \vdash_{TY} S_n \overline{\sigma}^n : \kappa}{\Gamma \vdash_{CO} S_n \overline{\gamma}^n : S_n \overline{\sigma}^n \sim S_n \overline{\tau}^n}$	(Sym) $\frac{\Gamma \vdash_{CO} \gamma : \sigma \sim \tau}{\Gamma \vdash_{CO} \text{sym } \gamma : \tau \sim \sigma}$	(Trans) $\frac{\Gamma \vdash_{CO} \gamma_1 : \sigma_1 \sim \sigma_2 \quad \Gamma \vdash_{CO} \gamma_2 : \sigma_2 \sim \sigma_3}{\Gamma \vdash_{CO} \gamma_1 \circ \gamma_2 : \sigma_1 \sim \sigma_3}$
(Comp) $\frac{\Gamma \vdash_{CO} \gamma_1 : \sigma_1 \sim \tau_1 \quad \Gamma \vdash_{CO} \gamma_2 : \sigma_2 \sim \tau_2 \quad \Gamma \vdash_{TY} \sigma_1 \sigma_2 : \kappa}{\Gamma \vdash_{CO} \gamma_1 \gamma_2 : \sigma_1 \sigma_2 \sim \tau_1 \tau_2}$	(Left) $\frac{\Gamma \vdash_{CO} \gamma : \sigma_1 \sigma_2 \sim \tau_1 \tau_2}{\Gamma \vdash_{CO} \text{left } \gamma : \sigma_1 \sim \tau_1}$	(Right) $\frac{\Gamma \vdash_{CO} \gamma : \sigma_1 \sigma_2 \sim \tau_1 \tau_2}{\Gamma \vdash_{CO} \text{right } \gamma : \sigma_2 \sim \tau_2}$
(CompC) $\frac{\Gamma \vdash_{CO} \gamma : \kappa_1 \sim \kappa_2 \quad \Gamma \vdash_{CO} \gamma' : \sigma_1 \sim \sigma_2 \quad \Gamma \vdash_k \kappa_1 : CO}{\Gamma \vdash_{CO} \gamma \Rightarrow \gamma' : (\kappa_1 \Rightarrow \sigma_1) \sim (\kappa_2 \Rightarrow \sigma_2)}$	(LeftC) $\frac{\Gamma \vdash_{CO} \gamma : \kappa_1 \xrightarrow{\sim} \sigma_1 \quad \kappa_2 \xrightarrow{\sim} \sigma_2}{\Gamma \vdash_{CO} \text{leftc } \gamma : \kappa_1 \sim \kappa_2}$	(RightC) $\frac{\Gamma \vdash_{CO} \gamma : \kappa_1 \xrightarrow{\sim} \sigma_1 \quad \kappa_2 \xrightarrow{\sim} \sigma_2}{\Gamma \vdash_{CO} \text{rightc } \gamma : \sigma_1 \sim \sigma_2}$
(\sim) $\frac{\Gamma \vdash_{CO} \gamma_1 : \sigma_1 \sim \tau_1 \quad \Gamma \vdash_{CO} \gamma_2 : \sigma_2 \sim \tau_2}{\Gamma \vdash_{CO} \gamma_1 \sim \gamma_2 : (\sigma_1 \sim \sigma_2) \sim (\tau_1 \sim \tau_2)}$	(CastC) $\frac{\Gamma \vdash_{CO} \gamma_1 : \kappa \quad \Gamma \vdash_{CO} \gamma_2 : \kappa \sim \kappa'}{\Gamma \vdash_{CO} \gamma_1 \blacktriangleright \gamma_2 : \kappa'}$	
$\Gamma \vdash_e e : \sigma$		
(Var) $\frac{u : \sigma \in \Gamma}{\Gamma \vdash_e u : \sigma}$	(Case) $\frac{\Gamma \vdash_e e : \sigma \quad \overline{\Gamma \vdash_p p \rightarrow e : \sigma \rightarrow \tau}}{\Gamma \vdash_e \text{case } e \text{ of } \overline{p \rightarrow e} : \tau}$	(Let) $\frac{\Gamma \vdash_e e_1 : \sigma_1 \quad \Gamma, x : \sigma_1 \vdash_e e_2 : \sigma_2}{\Gamma \vdash_e \text{let } x : \sigma_1 = e_1 \text{ in } e_2 : \sigma_2}$
(Cast) $\frac{\Gamma \vdash_e e : \sigma \quad \Gamma \vdash_{CO} \gamma : \sigma \sim \tau}{\Gamma \vdash_e e \blacktriangleright \gamma : \tau}$	(Abs) $\frac{\Gamma \vdash_{TY} \sigma_x : \star \quad \Gamma, x : \sigma_x \vdash_e e : \sigma}{\Gamma \vdash_e \lambda x : \sigma_x. e : \sigma_x \rightarrow \sigma}$	(App) $\frac{\Gamma \vdash_e e_1 : \sigma_2 \rightarrow \sigma_1 \quad \Gamma \vdash_e e_2 : \sigma_2}{\Gamma \vdash_e e_1 e_2 : \sigma_1}$
(AbsT) $\frac{\Gamma, a : \kappa \vdash_e e : \sigma \quad \Gamma \vdash_k \kappa : \delta \quad a \notin fv(\Gamma)}{\Gamma \vdash_e \Lambda a : \kappa. e : \forall a : \kappa. \sigma}$	(AppT) $\frac{\Gamma \vdash_e e : \forall a : \kappa. \sigma \quad \Gamma \vdash_k \kappa : \delta \quad \Gamma \vdash_\delta \varphi : \kappa}{\Gamma \vdash_e e \varphi : \sigma[\varphi/a]}$	
$\Gamma \vdash_p p \rightarrow e : \sigma \rightarrow \tau$		
(Alt) $\frac{K : \forall \overline{a} : \overline{\kappa}. \forall \overline{b} : \overline{\iota}. \overline{\sigma} \rightarrow T \overline{a} \in \Gamma \quad \theta = [\overline{v}/\overline{a}] \quad \Gamma, \overline{b} : \overline{\theta}(\overline{\iota}), \overline{x} : \overline{\theta}(\overline{\sigma}) \vdash_e e : \tau}{\Gamma \vdash_p K \overline{b} : \overline{\theta}(\overline{\iota}) \overline{x} : \overline{\theta}(\overline{\sigma}) \rightarrow e : T \overline{v} \rightarrow \tau}$		
$\Gamma \vdash \text{decl} : \Gamma'$		$\Gamma \vdash \text{pgm} : \sigma$
(Data) $\frac{\overline{\Gamma \vdash_{TY} \sigma : \star} \quad \Gamma \vdash_k \kappa : TY}{\Gamma \vdash (\text{data } T : \kappa \text{ where } \overline{K : \sigma}) : (T : \kappa, \overline{K : \sigma})}$		(Pgm) $\frac{\overline{\Gamma \vdash \text{decl} : \Gamma_d} \quad \Gamma = \Gamma_0, \overline{\Gamma_d} \quad \Gamma \vdash_e e : \sigma}{\Gamma_0 \vdash \text{decl}; e : \sigma}$
(Type) $\frac{\Gamma \vdash_k \kappa : TY}{\Gamma \vdash (\text{type } S : \kappa) : (S : \kappa)}$	(Coerce) $\frac{\Gamma \vdash_k \kappa : CO}{\Gamma \vdash (\text{axiom } C : \kappa) : (C : \kappa)}$	

Figure 2: Typing rules for System $F_C(X)$

$\Gamma \vdash_k \kappa : \delta$	
(Star)	$\frac{}{\Gamma \vdash_k \star : \text{TY}}$
(FunK)	$\frac{\Gamma \vdash_k \kappa_1 : \text{TY} \quad \Gamma \vdash_k \kappa_2 : \text{TY}}{\Gamma \vdash_k \kappa_1 \rightarrow \kappa_2 : \text{TY}}$
(EqTy)	$\frac{\Gamma \vdash_{\text{TY}} \sigma_1 : \kappa \quad \Gamma \vdash_{\text{TY}} \sigma_2 : \kappa}{\Gamma \vdash_k \sigma_1 \sim \sigma_2 : \text{CO}}$
(EqCo)	$\frac{\Gamma \vdash_k \gamma_1 : \text{CO} \quad \Gamma \vdash_k \gamma_2 : \text{CO}}{\Gamma \vdash_k \gamma_1 \sim \gamma_2 : \text{CO}}$

Figure 3: Kinding rules for System $F_C(X)$

is conventional: \star is the kind of proper types (that is, the types that a term can have), while higher kinds take the form $\kappa_1 \rightarrow \kappa_2$. Kinds guide type application by way of Rule (TyApp). Finally, the rules for judgements of the form $\Gamma \vdash_k \kappa : \delta$, given in Figure 3, ensure the well-formedness of kinds. Here δ is either TY for kinds formed from arrows and \star , or CO for coercion kinds of form $\sigma_1 \sim \sigma_2$. The conclusions of Rule (EqTy) and (EqCo) appear to overlap, but an actual implementation can deterministically decide which rule to apply, choosing (EqCo) iff γ_1 has the form $\varphi_1 \sim \varphi_2$.

The syntax of terms is largely conventional, as are their type rules which take the form $\Gamma \vdash_e e : \sigma$. As in F , every binder has an explicit type annotation, and type abstraction and application are also explicit. There is a `case` expression to take apart values built with data constructors. The patterns of a case expression are flat — there are no nested patterns — and bind existential type variables, coercion variables, and value variables. For example, suppose

$$K : \forall a : \star. \forall b : \star. a \rightarrow b \rightarrow (b \rightarrow \text{Int}) \rightarrow T \ a$$

Then a `case` expression that deconstructs K would have the form

$$\text{case } e \text{ of } K (b : \star) (v : a) (x : b) (f : b \rightarrow \text{Int}) \rightarrow e'$$

Note that only the existential type variable b is bound in the pattern. To see why, one need only realise that K 's type is isomorphic to:

$$K : \forall a : \star. (\exists b : \star. (a, b, (b \rightarrow \text{Int}))) \rightarrow T \ a$$

3.2 Type equality coercions

We now describe the unconventional features of our system. To begin with, consider the fragment of System F_C that omits type functions (i.e., **type** and **axiom** declarations). This fragment is sufficient to serve as a target for translating GADTs, and so is of interest in its own right. We return to type functions in §3.3.

The essential addition to plain F (beyond algebraic data types and higher kinds) is an infrastructure to construct, pass, and apply *type-equality coercions*. In F_C , a coercion, γ , is a special sort of type whose kind takes the unusual form $\sigma_1 \sim \sigma_2$. We can use such a coercion to cast an expression $e : \sigma_1$ to type σ_2 using the *cast expression* ($e \blacktriangleright \gamma$); see Rule (Cast) in Figure 2. Our intuition for equality coercions is an *extensional* one:

$\gamma : \sigma_1 \sim \sigma_2$ is evidence that a value of type σ_1 can be used in any context that expects a value of type σ_2 , and vice versa.

By “can be used”, we mean that running the program after type erasure will not go wrong. We stress that this is only an intuition; the soundness of our system is proved without appealing to any semantic notion of what $\sigma_1 \sim \sigma_2$ “means”. We use the symbol

“ \sim ” rather than “ $=$ ”, to avoid suggesting that the two types are intensionally equal.

Coercions are types — some would call them “constructors” [25, 12] since they certainly do not have kind \star — and hence the term-level syntax for type abstraction and application ($\Lambda a.e$ and $e \ \varphi$) also serves for coercion abstraction and application. However, coercions have their own kinding judgement \vdash_{CO} , given in Figure 2. The type of a term often has the form $\forall co : (\sigma_1 \sim \sigma_2). \varphi$, where φ does not mention co . We allow the standard syntactic sugar for this case, writing it thus: $(\sigma_1 \sim \sigma_2) \Rightarrow \varphi$ (see Figure 1). Incidentally, note that although coercions are types, they do not classify values. This is standard in F_ω ; for example, there are no values whose type has kind $\star \rightarrow \star$.

More complex coercions can be built by combining or transforming other coercions, such that every syntactic form corresponds to an inference rule of equational logic. We have the reflexivity of equality for a given type σ (witnessed by the type itself), symmetry ‘ $\text{sym } \gamma$ ’, transitivity ‘ $\gamma_1 \circ \gamma_2$ ’, type composition ‘ $\gamma_1 \ \gamma_2$ ’, and decomposition ‘ $\text{left } \gamma$ ’ and ‘ $\text{right } \gamma$ ’. The typing rules for these coercion expressions are given in Figure 2.

Here is an example, taken from §2. Suppose a GADT *Expr* has a constructor *Succ* with type

$$\text{Succ} : \forall a : \star. (a \sim \text{Int}) \Rightarrow \text{Exp Int} \rightarrow \text{Exp } a$$

(notice the use of the syntactic sugar $\kappa \Rightarrow \sigma$). Then we can construct a value of type *Exp Int* thus: *Succ Int Int e*. The second argument *Int* is a regular type used as a coercion witnessing reflexivity — i.e., we have $\text{Int} : (\text{Int} \sim \text{Int})$ by Rule (CoRef). Rule (CoRef) itself only covers type variables and constructors, but in combination with Rule (Comp), the reflexivity of complex types is also covered. More interestingly, here is a function that decomposes a value of type *Exp a*:

$$\begin{aligned} \text{foo} &: \forall a : \star. \text{Exp } a \rightarrow a \rightarrow a \\ &= \Lambda a : \star. \lambda e : \text{Exp } a. \lambda x : a. \\ &\quad \text{case } e \text{ of} \\ &\quad \text{Succ } (co : a \sim \text{Int}) (e' : \text{Exp Int}) \rightarrow \\ &\quad (\text{foo Int } e' \ 0 + (x \blacktriangleright co)) \blacktriangleright \text{sym } co \end{aligned}$$

The `case` pattern binds the coercion co , which provides evidence that a and Int are the same type. This evidence is needed twice, once to cast $x : a$ to Int , and once to coerce the *Int* result back to a , via the coercion ($\text{sym } co$).

Coercion composition allows us to “lift” coercions through arbitrary types, in the style of logical relations [1]. For example, if we have a coercion $\gamma : (\sigma_1 \sim \sigma_2)$ then the coercion *Tree* γ is evidence that *Tree* $\sigma_1 \sim \text{Tree } \sigma_2$, using rules (Comp) and (CoRef) and (CoVar). More generally, our system has the following theorem.

THEOREM 1 (Lifting). *If $\Gamma' \vdash_{\text{CO}} \gamma : \sigma_1 \sim \sigma_2$ and $\Gamma \vdash_{\text{TY}} \varphi : \kappa$, then $\Gamma' \vdash_{\text{CO}} [\gamma/a] \varphi : [\sigma_1/a] \varphi \sim [\sigma_2/a] \varphi$, for any type φ , including polytypes, where $\Gamma = \Gamma', a : \kappa'$ such that a does not appear in Γ' .*

For example, if $\gamma : \sigma_1 \sim \sigma_2$ then

$$\forall b. \gamma \rightarrow \text{Int} : (\forall b. \sigma_1 \rightarrow \text{Int}) \sim (\forall b. \sigma_2 \rightarrow \text{Int})$$

Dually decomposition enables us to take evidence apart. For example, assume $\gamma : \text{Tree } \sigma_1 \sim \text{Tree } \sigma_2$; then, ($\text{right } \gamma$) is evidence that $\sigma_1 \sim \sigma_2$, by rule (Right). Decomposition is necessary for the translation of GADT programs to F_C , but is problematic in earlier approaches [3, 9]. The soundness of decomposition relies, of course, on algebraic types being injective; i.e., $\text{Tree } \sigma_1 = \text{Tree } \sigma_2$ iff $\sigma_1 = \sigma_2$. Notice, too, that *Tree* by itself is a coercion relating two types of higher kind.

Similarly, one can compose and decompose equalities over polytypes, using rules (CoAllT) and (CoInstT). For example,

$$\begin{aligned} \gamma : (\forall a. a \rightarrow \text{Int}) &\sim (\forall a. a \rightarrow b) \\ \vdash_{\text{co}} \gamma @ \text{Bool} : (\text{Bool} \rightarrow \text{Int}) &\sim (\text{Bool} \rightarrow b) \end{aligned}$$

This simply says that if the two polymorphic functions are interchangeable, then so are their instantiations at *Bool*.

Rules (CompC), (LeftC), and (RightC) are analogous to (Comp), (Left), and (Right): they allow composition and decomposition of a type of form $\kappa \Rightarrow \varphi$, where κ is a coercion kind. These rules are essential to allow us to validate this consequence of Theorem 1:

$$\gamma : \sigma_1 \sim \sigma_2 \vdash_{\text{co}} (\gamma \sim \text{Int} \Rightarrow \text{Tree } \gamma) : \begin{array}{l} \sigma_1 \sim \text{Int} \Rightarrow \text{Tree } \sigma_1 \\ \sigma_2 \sim \text{Int} \Rightarrow \text{Tree } \sigma_2 \end{array}$$

Even though $\kappa \Rightarrow \varphi$ is sugar for $\forall _ : \kappa. \varphi$, we cannot generalise (CoAllT) to cover (CompC) because the former insists that the two kinds are identical.

We will motivate the need for rules (\sim) and (CastC) when discussing the dynamic semantics (§3.7).

3.3 Type functions

Our last step extends the power of F_C by adding *type functions* and *equality axioms*, which are crucial for translating associated types, functional dependencies, and the like. A type function S_n is introduced by a top-level **type** declaration, which specifies its *kind* $\bar{\kappa}^n \rightarrow \iota$, but says nothing about its *interpretation*. The index n indicates the *arity* of S . The syntax of types requires that S_n always appears applied to its full complement of n arguments (§3.6 explains why). The arity subscript should be considered part of the name of the type constructor, although we will often elide it, writing *Elem* σ rather than *Elem*₁ σ , for example.

A type function is given its interpretation by one or more equality **axioms**. Each axiom introduces a coercion constant, whose kind specifies the types between which the coercion converts. Thus:

$$\text{axiom } \text{elemBitSet} : \text{Elem BitSet} \sim \text{Char}$$

introduces the named coercion constant *elemBitSet*. Given an expression $e : \text{Elem BitSet}$, we can use the axiom via the coercion constant as in the cast $e \blacktriangleright \text{elemBitSet}$, which is of type *Char*.

We often want to state axioms involving parametric types, thus:

$$\text{axiom } \text{elemList} : (\forall e : \star. \text{Elem } [e]) \sim (\forall e : \star. e)$$

This is the axiom generated from the instance declaration for *Collects* $[e]$ in §2.2. To use this axiom as a coercion, say, for lists of integers, we need to apply the coercion constant to a type argument:

$$\text{elemList Int} : (\text{Elem } [\text{Int}] \sim \text{Int})$$

which appeals to Rule (CoInstT) of Figure 2. We have already seen the usefulness of (CoInstT) towards the end of §3.2, and here we simply re-use it. It may be surprising that we use one quantifier on each side of the equality, instead of quantifying over the entire equality as in

$$\forall a : \star. (\text{Elem } [a] \sim a) \quad \text{-- Not well-formed } F_C!$$

One could add such a construct, but it is simply unnecessary. We already have enough machinery, and when thought of as a logical relation, the form with a quantifier on each side makes perfect sense.

3.4 Type functions are open

A crucial property of type functions is that they are *open*, or *extensible*. A type function S may take an argument of kind \star (or $\star \rightarrow \star$, etc), and, since the kind \star is extensible, we cannot write out all the cases for S at the moment we introduce S . For example, imagine that a library module contains the definition of the *Collects* class (§2.2). Then a client imports this module, defines a new type T (thereby adding a new constant to the extensible kind \star), and wants to make T an instance of *Collects*. In F_C this is easy by simply writing in the client module

```
import CollectsLib
instance Collects T where {type Elem T = E; ...}
```

where we assume that E is the element type of the collection type T . In short, open type functions are absolutely required to support modular extensibility.

We do not argue that *all* type functions should be open; it would certainly be possible to extend F_C with non-extensible kind declarations and closed type functions. Such an extension would be useful; consider the well-worn example of lists parametrised by their length, which we give in source-code form to reduce clutter:

```
kind Nat = Z | S Nat

data Seq a (n :: Nat) where
  Nil  :: Seq a Z
  Cons :: a -> Seq a n -> Seq a (S n)

app :: Seq a n -> Seq a m -> Seq a (Plus n m)
app Nil      ys = ys
app (Cons x xs) ys = Cons x (app xs ys)

type Plus :: Nat -> Nat -> Nat
Plus Z     b = b
Plus (S a) b = S (Plus a b)
```

Whilst we can translate this into F_C , we would be forced to give *Plus* the kind $\star \rightarrow \star \rightarrow \star$, which allows nonsensical forms like *Plus Int Bool*. Furthermore, the non-extensibility of *Nat* would allow induction, which is not available in F_C precisely because kind \star is extensible.

Other closely-related languages support closed type functions; for example LH [25], LX [12], and Ω mega [36]. In this paper, however, we focus on open-ness, since it is one of F_C 's most distinctive features and is crucial to translating associated types.

3.5 Consistency

In System $F_C(X)$, we refine the equational theory of types by giving non-standard equality axioms. *So what is to prevent us declaring unsound axioms?* For example, one could easily write a program that would crash, using the coercion constant introduced by the following axiom:

$$\text{axiom } \text{utterlybogus} : \text{Int} \sim \text{Bool}$$

(where *Int* and *Bool* are both algebraic data types). There are many *ad hoc* ways to restrict the system to exclude such cases. The most general way is this: we require that the axioms, taken together, are *consistent*. We essentially adapt the standard notion of consistency of sets of equations [13, Section 3] to our setting.

DEFINITION 1 (Value type). *A type σ is a value type if it is of form $\forall a. v$ or $T \bar{v}$.*

DEFINITION 2 (Consistency). Γ is consistent iff

1. If $\Gamma \vdash_{\text{co}} \gamma : T \bar{\sigma} \sim v$, and v is a value type, then $v = T \bar{\tau}$.
2. If $\Gamma \vdash_{\text{co}} \gamma : \forall a : \kappa. \sigma \sim v$, and v is a value type, then $v = \forall a : \kappa. \tau$.

That is, if there is a coercion connecting two *value* types — algebraic data types, built-in types, functions, or forall — then the outermost type constructors must be the same. For example, there can be no coercion of type *Bool* \sim *Int*. It is clear that the consistency of Γ is necessary for soundness, and it turns out that it is also sufficient (§3.7).

Consistency is only required of the *top-level* environment, however (Figure 1). For example, consider this function:

$$f = \lambda(g : \text{Int} \sim \text{Bool}). 1 + (\text{True} \blacktriangleright g)$$

It uses the bogus coercion g to cast an Int to a $Bool$, so f would crash if called. But there is no problem, because *the function can never be called*; to do so, one would have to produce evidence that Int and $Bool$ are interchangeable. The proof in §3.7 substantiates this intuition.

Nevertheless, consistency is absolutely required for the top-level environment, but alas it is an undecidable property. That is why we call the system “ $F_C(X)$ ”: it is parametrised by a decision procedure X for determining consistency. There is no “best” choice for X , so instead of baking a particular choice into the language, we have left the choice open. Each particular source-program construct that exploits type equalities comes with its own decision procedure — or, alternatively, guarantees *by construction* to generate only consistent axioms, so that consistency need never be checked. All the applications we have implemented so far take the latter approach. For example, GADTs generate no axioms at all (Section 4); newtypes generate exactly one axiom per newtype; and associated types are constrained to generate a non-overlapping rewrite system (Section 5).

3.6 Saturation of type functions

We remarked earlier that applications of type functions S_n are required to be saturated. The reason for this insistence is, again, consistency. We definitely want to allow abstract types to be non-injective; for example:

axiom $c1 : S_1 Int \sim Bool$
axiom $c2 : S_1 Bool \sim Bool$

Here, both $S_1 Int$ and $S_1 Bool$ are represented by the $Bool$ type. But now we can form the coercion $(c1 \circ (\text{sym } c2))$ which has type $S_1 Int \sim S_1 Bool$, and from that we must not be able to deduce (via right) that $Int \sim Bool$, because that would violate consistency! Applications of type functions are therefore syntactically distinguished so that right and left apply only to ordinary type application (Rules (Left) and (Right) in Figure 2), and not to applications of type functions. The same syntactic mechanism prevents a partial type-function application from appearing as a type argument, thereby instantiating a type variable with a partial application — in effect, type variables of higher-kind range only over injective type constructors.

However, it is perfectly acceptable for a type function to have an arity of 1, say, but a higher kind of $\star \rightarrow \star \rightarrow \star$. For example:

type $HS_1 : \star \rightarrow \star \rightarrow \star$
axiom $c1 : HS_1 Int \sim []$
axiom $c2 : HS_1 Bool \sim Maybe$

An application of HS to one type is saturated, although it has kind $\star \rightarrow \star$ and can be applied (via ordinary type application) to another type.

3.7 Dynamic semantics and soundness

The operational semantics of F_C is shown in Figure 4. In the expression reductions we omit the type annotations on binders to save clutter, but that is mere abbreviation.

An unusual feature of our system, which we share with Cray’s coercion calculus for inclusive subtyping [11], is that values are stratified into *cvalues* and *plain values*; their syntax is in Figure 4. Evaluation reduces a closed term to a *cvalue*, or diverges. A *cvalue* is either a *plain value* v (an abstraction or saturated constructor application), or it is a value wrapped in a single cast, thus $v \blacktriangleright \gamma$ (Figure 4). The latter form is needed because we cannot reduce a term to a plain value without losing type preservation; for example, we cannot reduce $(True \blacktriangleright \gamma)$, where $\gamma : Bool \sim S$ any further without changing its type from S to $Bool$.

However, there are four situations when a *cvalue* will not do, namely as the function part of a type, coercion, or function ap-

plication, or as the scrutinee of a **case** expression. Rules (TPush), (CPush), (Push) and (KPush) deal with those situations, by pushing the coercion inside the term, turning the cast into a plain value. Notice that all four rules leave the *context* (the application or case expression) unchanged; they rewrite the function or case scrutinee respectively. Nevertheless, the context is necessary to guarantee that the type of the rewritten term is a function or data type respectively.

Rules (TPush) and (Push) are quite straightforward. Rule (CPush) is rather like (Push), but at the level of coercions. It is this rule that forces us to add the forms $(\gamma_1 \sim \gamma_2)$, $(\gamma_1 \blacktriangleright \gamma_2)$, $(\text{left } \gamma)$, and $(\text{right } \gamma)$ to the language of coercions. We will shortly provide an example to illustrate this point.

The final rule, (KPush), is more complicated. Here is an example, stripped of the **case** context, where $Cons : \forall a. a \rightarrow [a] \rightarrow [a]$, and $\gamma : [Int] \sim [S Bool]$:

$$(Cons Int e_1 e_2) \blacktriangleright \gamma \longrightarrow Cons (S Bool) \begin{array}{l} (e_1 \blacktriangleright \text{right } \gamma) \\ (e_2 \blacktriangleright ([]) (\text{right } \gamma)) \end{array}$$

The coercion wrapped around the application of $Cons$ is pushed inside to wrap each of its components. (Of course, an implementation does none of this, because types and coercions are erased.) The type preservation of this rule relies on Theorem 1 in Section 3.2, which guarantees that $e_i \blacktriangleright \theta(\rho_i)$ has the correct type.

The rule is careful to cast the *coercion* arguments as well as the *value* arguments. Here is an example, taken from Section 2.3:

$$\begin{array}{l} F : \forall a b. (b \sim F1 a) \Rightarrow FDict a b \\ \gamma : FDict Int Bool \sim FDict c d \\ \varphi : Bool \sim F1 Int \end{array}$$

Now, stripped of the **case** context, rule (KPush) describes the following transition:

$$(F Int Bool \varphi) \blacktriangleright \gamma \longrightarrow F c d (\varphi \blacktriangleright (\gamma_2 \sim F1 \gamma_1))$$

where $\gamma_1 = \text{right } (\text{left } \gamma)$ and $\gamma_2 = \text{right } \gamma$. The coercion argument φ is cast by the strange-looking coercion $\gamma_2 \sim F1 \gamma_1$, whose kind is $(Bool \sim F1 Int) \sim (d \sim F1 c)$. That is why we need rule (\sim) in Figure 2, so that we can type such coercions.

We derived all three “push” rules in a systematic way. For example, for (Push) we asked what e' (involving e and γ) would ensure that $((\lambda x. e) \blacktriangleright \gamma) = \lambda y. e'$. The reader may like to check that if the left-hand side of each rule is well-typed (in the top-level context) then so is the right-hand side.

When a data constructor has a higher-rank type, in which the argument types are themselves quantified, a good deal of book-keeping is needed. For example, suppose that

$$\begin{array}{l} K : \forall a : *. (a \sim Int \Rightarrow a \rightarrow Int) \rightarrow T a \\ \gamma : T \sigma_1 \sim T \sigma_2 \\ e : (\sigma_1 \sim Int) \Rightarrow \sigma_1 \rightarrow Int \end{array}$$

Then, according to rule (KPush) we find (as before we strip off the **case** context)

$$(K \sigma_1 e) \blacktriangleright \gamma \longrightarrow K \sigma_2 (e \blacktriangleright \gamma')$$

where $\gamma' = (\text{right } \gamma \sim Int) \Rightarrow \text{right } \gamma \rightarrow Int$, which is obtained by substituting $[\text{right } \gamma / a]$ in $(a \sim Int) \Rightarrow a \rightarrow Int$. Now suppose that we later reduce the (sub)-expression

$$(e \blacktriangleright \gamma') \gamma''$$

where $e = \Lambda b : (\sigma_1 \sim Int). \lambda x : \sigma_1. x \blacktriangleright b$. Before we can apply rule (CPush) we have to determine the kind of γ' . It is straightforward to deduce that

$$\gamma' : (\sigma_1 \sim Int \Rightarrow \sigma_1 \rightarrow Int) \sim (\sigma_2 \sim Int \Rightarrow \sigma_2 \rightarrow Int)$$

Values:

Plain values $v ::= \Lambda a.e \mid \lambda x.e \mid K \bar{\sigma} \bar{\varphi} \bar{e}$
 Cvalues $cv ::= v \blacktriangleright \gamma \mid v$

Evaluation contexts:

$$\frac{e \longrightarrow e'}{E[e] \longrightarrow E[e']} \quad E ::= [] \mid E e \mid E \tau \mid E \blacktriangleright \gamma \mid \text{case } E \text{ of } \overline{p \rightarrow rhs}$$
Expression reductions:

(TBeta)	$(\Lambda a.e) \varphi$	\longrightarrow	$[\varphi/a]e$
(Beta)	$(\lambda x.e) e'$	\longrightarrow	$[e'/x]e$
(Case)	$\text{case } (K \bar{\sigma} \bar{\varphi} \bar{e}) \text{ of } \dots K \bar{b} \bar{x} \rightarrow e' \dots$	\longrightarrow	$[\varphi/\bar{b}, e'/\bar{x}]e'$
(Comb)	$(v \blacktriangleright \gamma_1) \blacktriangleright \gamma_2$	\longrightarrow	$v \blacktriangleright (\gamma_1 \circ \gamma_2)$
(TPush)	$((\Lambda a:\kappa.e) \blacktriangleright \gamma) \varphi$	\longrightarrow	$(\Lambda a:\kappa.(e \blacktriangleright \gamma @ a)) \varphi$ where $\gamma : (\forall a:\kappa.\sigma_1) \sim (\forall b:\kappa.\sigma_2)$
(CPush)	$((\Lambda a:\kappa.e) \blacktriangleright \gamma) \varphi$	\longrightarrow	$(\Lambda a':\kappa'.((a' \blacktriangleright \gamma_1)/a)e) \blacktriangleright \gamma_2) \varphi$ where $\gamma : (\kappa \Rightarrow \sigma) \sim (\kappa' \Rightarrow \sigma')$ $\gamma_1 = \text{sym}(\text{leftc } \gamma)$ – coercion for argument $\gamma_2 = \text{rightc } \gamma$ – coercion for result
(Push)	$((\lambda x.e) \blacktriangleright \gamma) e'$	\longrightarrow	$(\lambda y.((y \blacktriangleright \gamma_1)/x)e) \blacktriangleright \gamma_2) e'$ where $\gamma_1 = \text{sym}(\text{right}(\text{left } \gamma))$ – coercion for argument $\gamma_2 = \text{right } \gamma$ – coercion for result
(KPush)	$\text{case } (K \bar{\sigma} \bar{\varphi} \bar{e} \blacktriangleright \gamma) \text{ of } \overline{p \rightarrow rhs}$	\longrightarrow	$\text{case } (K \bar{\tau} \bar{\varphi}' \bar{e}') \text{ of } \overline{p \rightarrow rhs}$ where $\gamma : T \bar{\sigma} \sim T \bar{\tau}$ $K : \forall \bar{a}:\bar{\kappa}.\forall \bar{b}:\bar{\iota}.\bar{\rho} \rightarrow T \bar{a}^n$ $\varphi'_i = \begin{cases} \varphi_i \blacktriangleright \theta(v_1 \sim v_2) & \text{if } b_i : v_1 \sim v_2 \\ \varphi_i & \text{otherwise} \end{cases}$ $e'_i = e_i \blacktriangleright \theta(\rho_i)$ $\theta = [\gamma_i/a_i, \varphi_i/b_i]$ $\gamma_i = \text{right}(\underbrace{\text{left} \dots \text{left}}_{n-i} \gamma)$

Figure 4: Operational semantics

Hence, via (CPush) we find that

$$\begin{aligned} & ((\Lambda b:(\sigma_1 \sim \text{Int}). \lambda x:\sigma_1. x \blacktriangleright b) \blacktriangleright \gamma') \gamma'' \\ \longrightarrow & (\Lambda c:(\sigma_2 \sim \text{Int}). (\lambda x:\sigma_1. x \blacktriangleright (c \blacktriangleright \gamma_1)) \blacktriangleright \gamma_2) \gamma'' \end{aligned}$$

where $\gamma_1 = \text{sym}(\text{leftc } \gamma')$, $\gamma_1 : (\sigma_1 \sim \text{Int}) \sim (\sigma_2 \sim \text{Int})$, $\gamma_2 = \text{rightc } \gamma'$ and $\gamma_2 : (\sigma_1 \rightarrow \text{Int}) \sim (\sigma_2 \rightarrow \text{Int})$.

Notice that forms $(\gamma_1 \sim \gamma_2)$, $(\gamma_1 \blacktriangleright \gamma_2)$, $(\text{leftc } \gamma)$, and $(\text{rightc } \gamma)$ only appear during the reduction of F_C programs. In case we restrict F_C types to be rank 1 none of these forms are necessary.

THEOREM 2 (Progress and subject reduction). *Suppose that a top-level environment Γ is consistent, and $\Gamma \vdash_e e : \sigma$. Then either e is a cvalue, or $e \longrightarrow e'$ and $\Gamma \vdash_e e' : \sigma$ for some term e' .*

A proof is given in [37].

COROLLARY 1 (Syntactic Soundness). *Let Γ be consistent top-level environment and $\Gamma \vdash_e e : \sigma$. Then either $e \longrightarrow^* cv$ and $\Gamma \vdash_e cv : \sigma$ for some cvalue cv , or the evaluation diverges.*

We give a call-by-name semantics here, but a call-by-value semantics would be equally easy: simply extend the syntax of evaluation contexts with the form $v E$, and restrict the argument of rule (Beta) to be a cvalue.

In general, evaluation affects *expressions* only, not types. Since coercions are types, it follows that coercions are not evaluated either. This means that we can completely avoid the issue of normalisation of coercions, what a coercion “value” might be, and so on.

3.8 Robustness to transformation

One of our major concerns was to make it easy for a compiler to transform and optimise the program. For example, consider this fragment:

$$\lambda x. \text{case } x \text{ of } \{ T1 \rightarrow \text{let } z = y + 1 \text{ in } \dots; \dots \}$$

A good optimisation might be to float the let-binding out of the lambda, thus:

$$\text{let } z = y + 1 \text{ in } \lambda x. \text{case } x \text{ of } \{ T1 \rightarrow \dots; \dots \}$$

But suppose that $x : Ta$ and $y : a$, and that the pattern match on $T1$ refines a to Int . Then the floated form is type-incorrect, because the let-binding is now out of the scope of the pattern match. This is a real problem for any intermediate language in which equality is implicit. In F_C , however, y will be cast to Int using a coercion that is bound by the pattern match on $T1$. So the type-incorrect transformation is prevented, because the binding mentions a variable bound by the match; and better still, we can perform the optimisation in a type-correct way by abstracting over the variable to get this:

$$\begin{aligned} & \text{let } z' = \lambda g. (y \blacktriangleright g) + 1 \\ & \text{in } \lambda x. \text{case } x \text{ of } \{ T1 g \rightarrow \text{let } z = z' g \text{ in } \dots; \dots \} \end{aligned}$$

The inner let-binding obviously cannot be floated outside, because it mentions a variable bound by the match.

Another useful transformation is this:

$$(\text{case } x \text{ of } p_i \rightarrow e_i) \text{ arg} = \text{case } x \text{ of } p_i \rightarrow e_i \text{ arg}$$

This is valid in F_C , but not in the more implicit language LH, for example [25].

In summary, we believe that F_C 's obsessively-explicit form makes it easy to write type-preserving transformations, whereas doing so is significantly harder in a language where type equality is more implicit.

3.9 Type and coercion erasure

System F_C permits syntactic type erasure much as plain System F does, thereby providing a solid guarantee that coercions impose absolutely no run-time penalty. Like types, coercions simply provide a statically-checkable guarantee that there will be no run-time crash.

Formally, we can define an erasure function e° , which erases all types and coercions from the term, and prove the standard erasure theorem. Following Pierce [32, Section 23.7] we erase a type abstraction to a trivial term abstraction, and type application to term application to the unit value; this standard device preserves termination behaviour in the presence of `seq`, or with call-by-value semantics. The only difference from plain F is that we also erase casts.

$$\begin{array}{ll} x^\circ = x & (\lambda x:\varphi. e)^\circ = \lambda x.e^\circ \\ K^\circ = K & (e_1 e_2)^\circ = e_1^\circ e_2^\circ \\ (\Lambda a:\kappa. e)^\circ = \lambda a.e^\circ & (e \blacktriangleright \gamma)^\circ = e^\circ \\ (e \sigma)^\circ = e^\circ() & (K \bar{a}:\bar{\kappa} \bar{x}:\bar{\varphi})^\circ = K \bar{a} \bar{x} \end{array}$$

$$\begin{array}{ll} (\text{let } x:\sigma = e_1 \text{ in } e_2)^\circ = \text{let } x = e_1^\circ \text{ in } e_2^\circ \\ (\text{case } e_1 \text{ of } \bar{p} \rightarrow e_2)^\circ = \text{case } e_1^\circ \text{ of } \bar{p}^\circ \rightarrow e_2^\circ \end{array}$$

THEOREM 3. *Suppose that a top-level environment Γ is consistent, and $\Gamma \vdash_e e_1 : \sigma$. Then, (a) either e_1 is a cvalue and e_1° is a value or (b) we have $e_1 \rightarrow e_2$ and either $e_1^\circ \rightarrow e_2^\circ$ or $e_1^\circ = e_2^\circ$.*

A proof is given in [37].

COROLLARY 2 (Erasure soundness). *For an well-typed System F_C term e_1 , we have $e_1 \rightarrow^* e_2$ iff $e_1^\circ \rightarrow^* e_2^\circ$.*

The dynamic semantics of Figure 4 makes all the coercions in the program bigger and bigger. This is not a run-time concern, because of erasure, but it might be a concern for compiler transformations. Fortunately there are many type-preserving simplifications that can be performed on coercions, such as:

$$\begin{array}{ll} \text{sym } \sigma & = \sigma \\ \text{left } (Tree \ Int) & = Tree \\ e \blacktriangleright \sigma & = e \end{array}$$

and so on. The compiler writer is free (but not obliged) to use such identities to keep the size of coercions under control.

In this context, it is interesting to note the connection of type-equality coercions to the notion of proof objects in machine-supported theorem proving. Coercion terms are essentially proof objects of equational logic and the above simplification rules, as well the manipulations performed by rules, such as (PushK), correspond to proof transformations.

3.10 Summary and observations

F_C is an intensional type theory, like F: that is, *every term encodes its own typing derivation*. This is bad for humans (because the terms are bigger) but good for a compiler (because type checking is simple, syntax-directed, and decidable). An obvious question is this: could we maintain simple, syntax-directed, decidable type-checking for F_C with less clutter? In particular, a coercion is an explicit proof of a type equality; could we omit the coercions, retaining only their kinds, and reconstructing the proofs on the fly?

No, we could not. Previous work showed that such proofs can indeed be inferred for the special case of GADTs [43, 34, 38]. But our

setting is much more general because of our type functions, which in turn are necessary to support the source-language extensions we seek. Reconstructing an equality proof amounts to unification modulo an equational theory (E-unification), which is undecidable even in various restricted forms, let alone in the general case [2]. In short, dropping the explicit proofs encoded by coercions would render type checking undecidable (see [37] for a formal proof).

Why do we express coercions as *types*, rather than as *terms*? The latter is more conventional; for example, GADTs can be used to encode equality evidence [36], via a GADT of the form

```
data Eq a b where { EQ :: Eq a a }
```

F_C turns this idea on its head, instead using equality evidence to encode GADTs. This is good for several reasons. First, F_C is more foundational than System F plus GADTs. Second, F_C expresses equality evidence in *types*, which permit erasure; GADTs encode equality evidence as *values*, and these values cannot be erased. Why not? Because in the presence of recursion, the mere existence of an expression of type `Eq a b` is not enough to guarantee that `a` is the same as `b`, because \perp has any type. Instead, one must *evaluate* evidence before using it, to ensure that it converges, or else provide a meta-level proof that asserts that the evidence always converges. In contrast, our language of types deliberately lacks recursion, and hence coercions can be trusted simply by virtue of being well-kinded.

4. Translating GADTs

With F_C in hand, we now sketch the translation of a source language supporting GADTs into F_C . As highlighted in §2.1, the key idea is to turn type equalities into coercion types. This approach strongly resembles the dictionary-passing translation known from translating type classes [17]. The difference is that we do not turn type equalities into values, rather, we turn them into types.

We do not have space to present a full source language supporting GADTs, but instead sketch its main features; other papers give full details [43, 10]. We assume that the GADT source language has the following syntax of types:

$$\begin{array}{ll} \text{Polytypes} & \pi \rightarrow \eta \mid \forall a.\pi \\ \text{Constrained types} & \eta \rightarrow \tau \mid \tau \sim \tau \Rightarrow \eta \\ \text{Monotypes} & \tau \rightarrow a \mid \tau \rightarrow \tau \mid T \bar{\tau} \end{array}$$

We deliberately re-use F_C 's syntax $\tau_1 \sim \tau_2$ to describe GADT type equalities. These equality constraints are used in the source-language type of data constructors. For example, the `Succ` constructor from §2.1 would have type

$$\text{Succ} : \forall a.(a \sim \text{Int}) \Rightarrow \text{Int} \rightarrow \text{Exp } a$$

Notice that this already *is* an F_C type.

To keep the presentation simple, we use a non-syntax-directed translation scheme based on the judgement

$$C; \Gamma \vdash_{GADT} e : \pi \rightsquigarrow e'$$

We read it as “assuming constraint C and type environment Γ , the source-language expression e has type π , and translates to the F_C expression e' ”. The translation scheme can be made syntax-directed along the lines of [31, 34, 38]. The constraint C consists of a set of named type equalities:

$$C \rightarrow \epsilon \mid C, c:\tau_1 \sim \tau_2$$

The most interesting translation rules are shown in Figure 5, where we assume for simplicity that all quantified GADT variables are of kind $*$. The Rules (Var), (\forall -Intro), and (\forall -Elim), dealing with variables and the introduction and elimination of polymorphic types, are standard for translating Hindley/Milner to System F [19]. The introduction and elimination rules for constrained

$$\boxed{C; \Gamma \vdash_{GADT} e : \pi \rightsquigarrow e'}$$

$$\begin{array}{c}
(\text{Var}) \quad \frac{(x : \pi) \in \Gamma}{C; \Gamma \vdash_{GADT} x : \pi \rightsquigarrow x} \quad (\text{Eq}) \quad \frac{C; \Gamma \vdash_{GADT} e : \tau \rightsquigarrow e' \quad C \vdash_{\text{co}} \gamma : \tau \sim \tau'}{C; \Gamma \vdash_{GADT} e : \tau' \rightsquigarrow e' \blacktriangleright \gamma} \\
(\forall\text{-Intro}) \quad \frac{C; \Gamma \vdash_{GADT} e : \pi \rightsquigarrow e' \quad a \notin \text{fv}(C, \Gamma)}{C; \Gamma \vdash_{GADT} e : \forall a. \pi \rightsquigarrow \Lambda a : *. e'} \quad (\text{C-Intro}) \quad \frac{C, c : \tau_1 \sim \tau_2; \Gamma \vdash_{GADT} e : \eta \rightsquigarrow e'}{C; \Gamma \vdash_{GADT} e : \tau_1 \sim \tau_2 \Rightarrow \eta \rightsquigarrow \Lambda(c : \tau_1 \sim \tau_2). e'} \\
(\forall\text{-Elim}) \quad \frac{C; \Gamma \vdash_{GADT} e : \forall a. \pi \rightsquigarrow e'}{C; \Gamma \vdash_{GADT} e : [\tau/a] \pi \rightsquigarrow e' \tau} \quad (\text{C-Elim}) \quad \frac{C; \Gamma \vdash_{GADT} e : \tau_1 \sim \tau_2 \Rightarrow \eta \rightsquigarrow e' \quad C \vdash_{\text{co}} \gamma : \tau_1 \sim \tau_2}{C; \Gamma \vdash_{GADT} e : \eta \rightsquigarrow e' \gamma} \\
\boxed{C; \Gamma \vdash_{GADT} p \rightarrow e : \pi \rightarrow \pi \rightsquigarrow p' \rightarrow e'} \\
(\text{Alt}) \quad \frac{K :: \forall \bar{a}, \bar{b}. \overline{\tau'} \Rightarrow \bar{\tau} \rightarrow T \bar{a} \quad \bar{a} \cap \bar{b} = \emptyset \quad \text{fv}(\bar{\tau}, \bar{\tau}', \bar{\tau}'') = \text{fv}(\bar{a}, \bar{b}) \quad \theta = \overline{[v/a]}}{C, c : \theta(\tau') \sim \theta(\tau''); \Gamma, x : \theta(\tau) \vdash_{GADT} e : \tau' \rightsquigarrow e' \quad \bar{c} \text{ fresh}} \\
C; \Gamma \vdash_{GADT} K \bar{x} \rightarrow e : T \bar{v} \rightarrow \tau' \rightsquigarrow K(\bar{b} : *) (c : \theta(\tau') \sim \theta(\tau'')) (x : \theta(\tau)) \rightarrow e'
\end{array}$$

Figure 5: Type-Directed GADT to F_C Translation (interesting cases)

types, Rules (C-Intro) and (C-Elim), relate to the standard type-class translation [17], but where class constraints induce value abstraction and application, equality constraints induce type abstraction and application.

The translation of pattern clauses in Rule (Case) is as expected. We replace each GADT constructor by an appropriate F_C constructor which additionally carries coercion types representing the GADT type equalities. We assume that source patterns are already flat.

Rule (Eq) applies the cast construct to coerce types. For this, we need a coercion γ witnessing the equality of the two types, and we simply re-use the F_C judgement $\Gamma \vdash_{\text{co}} \gamma : \tau_1 \sim \tau_2$ from Figure 2. In this context, γ is an “output” of the judgement, a coercion whose syntactic structure describes the proof of $\tau_1 \sim \tau_2$. In other words, $C \vdash_{\text{co}} \gamma : \tau_1 \sim \tau_2$ represents the GADT condition that the equality context “ C implies $\tau_1 \sim \tau_2$ ”.

Finding a γ is decidable, using an algorithm inspired by the unification algorithm [23]. The key observation is that the statement “ C implies $\tau_1 \sim \tau_2$ ” holds if $\theta(\tau_1) = \theta(\tau_2)$ where θ is the most general unifier of C . W.l.o.g., we neglect the case that C has no unifier, i.e. C is unsatisfiable. Program parts which make use of unsatisfiable constraints effectively represent dead-code.

Roughly, the type coercion construction procedure proceeds as follows. Given the assumption set C and our goal $\tau_1 \sim \tau_2$ we perform the following calculations:

Step 1 : We normalise the constraints $C = \overline{c : \tau' \sim \tau''}$ to the solved form $\overline{\gamma : a \sim v}$ where $a_i < a_{i+1}$ and $\text{fv}(\bar{a}) \cap \text{fv}(\bar{v}) = \emptyset$ by decomposing with Rule (Right) (we neglect higher-kinded types for simplicity) and applying Rule (Sym) and (Trans). We assume some suitable ordering among variables with $<$ and disallow recursive types.

Step 2 : Normalise $c' : \tau_1 \sim \tau_2$ where c' is fresh to the solved form $\overline{\gamma' : a' \sim v'}$ where $a'_j < a'_{j+1}$.

Step 3 : Match the resulting equations from Step 2 against equations from Step 1.

Step 4 : We obtain γ by reversing the normalisation steps in Step 2.

Failure in any of the steps implies that $C \vdash_{\text{co}} \gamma : \tau_1 \sim \tau_2$ does not hold for any γ . A constraint-based formulation of the above algorithm is given in [39].

To illustrate the algorithm, let’s consider $C = \{c_1 : [a] \sim [b], c_2 : b = c\}$ and $c_3 : [a] \sim [c]$, with $a < b < c$.

Step 1: Normalising C yields $\{\text{right } c_1 : a \sim b, c_2 : b = c\}$ in an intermediate step. We apply rule (Trans) to obtain the solved form $\{\text{right } c_1 \circ c_2 : a \sim c, c_2 : b = c\}$

Step 2: Normalising $c_3 : [a] \sim [c]$ yields $(\text{right } c_3) : a \sim c$.

Step 3: We can match $\text{right } c_3 : a \sim c$ against $(\text{right } c_1 \circ c_2) : a \sim c$.

Step 4: Reversing the normalisation steps in Step 2 yields $c_3 = [\text{right } c_1 \circ c_2]$, as $\vdash_{\text{co}} [] : [] \sim []$.

The following result can be straightforwardly proven by induction over the derivation.

LEMMA 1 (Type Preservation). *Let $C; \emptyset \vdash_{GADT} e : t \rightsquigarrow e'$. Then, $C \vdash_e e' : t$.*

In §3.5, we saw that only consistent F_C programs are sound. It is not hard to show that this the case for GADT F_C programs, as GADT programs only make use of syntactic (a.k.a. Herbrand) type equality, and so, require no type functions at all.

THEOREM 4 (GADT Consistency). *If $\text{dom}(\Gamma)$ contains no type variables or coercion constants, and $\Gamma \vdash_{\text{co}} \gamma : \sigma_1 \sim \sigma_2$, then $\sigma_1 = \sigma_2$ (i.e. the two are syntactically identical).*

The proof is by induction on the structure of γ . Consistency is an immediate corollary of Theorem 4. Hence, all GADT F_C programs are sound. From the Erasure Soundness Corollary 2, we can immediately conclude that the semantics of GADT programs remains unchanged (where e° is e after type erasure).

LEMMA 2. *Let $\emptyset; \emptyset \vdash_{GADT} e : t \rightsquigarrow e'$. Then, $e' \rightsquigarrow^* v$ iff $e^\circ \rightsquigarrow^* v$ where v is some base value, e.g. integer constants.*

5. Translating Associated Types

In §2.2, we claimed that F_C permits a more direct and more general type-preserving translation of associated types than the translation to plain System F described in [6]. In fact, the translation of associated types to F_C is almost embarrassingly simple, especially given the translation of GADTs to F_C from §4. In the following, we outline the additions required to the standard translation of type classes to System F [17] to support associated types.

5.1 Translating expressions

To translate expressions, we need to add three rules to the standard system of [17], namely Rules (Eq), (C-Intro), and (C-Elim) from

Figure 5 of the GADT translation. Rule (Eq) permits casting expression with types including associated types to equal types where the associated types have been replaced by their definition. Strictly speaking, the Rules (C-Intro) and (C-Elim) are used in a more general setting during associated type translation than during GADT translation. Firstly, the set C contains not only equalities, but both equality and class constraints. Secondly, in the GADT translation only GADT data constructors carry equality constraints, whereas in the associated type translation, any function can carry equality constraints.

5.2 Translating class predicates

In the standard translation, predicates are translated to dictionaries by a judgement $C \Vdash_D D \tau \rightsquigarrow \nu$. In the presence of associated types, we have to handle the case where the type argument to a predicate contains an associated type. For example, given the class

```
class Collects c where
  type Elem c      -- associated type synonym
  empty  :: c
  insert :: Elem c -> c -> c
  toList :: c -> [Elem c]
```

we might want to define

```
sumColl :: (Collects c, Num (Elem c))
  => c -> Elem c
sumColl c = sum (toList c)
```

which sums the elements of a collection, provided these elements are members of the Num class; i.e., provided we have Num (Elem c). Here we have an associated type as a parameter to a class constraint. Wherever the function `sumColl` is used, we will have to check the constraint Num (Elem c), which will require a cast of the resulting dictionary if `c` is instantiated. We achieve this by adding the following rule:

$$\text{(Subst)} \quad \frac{C \Vdash_D D \tau_1 \rightsquigarrow w \quad C \vdash_{\text{TY}} \gamma : D \tau_1 = D \tau_2}{C \Vdash_D D \tau_2 \rightsquigarrow w \blacktriangleright \gamma}$$

It permits to replace type class arguments by equal types, where the coercion γ witnessing the equality is used to adapt the type of the dictionary w , which in turn witnesses the type class instance. Interestingly, we need this rule also for the translation as soon as we admit qualified constructor signatures in GADT declarations.

5.3 Translating declarations

Strictly speaking, we also have to extend the translation rules for class and instance definitions, as these can now declare and define associated types. However, the extension is so small that we omit the formal rules for space reasons. In summary, each declaration of an associated type in a type class turns into the declaration of a type function in F_C , and each definition of an associated type in an instance turns into an equality **axiom** in F_C . We have seen examples of this in §2.2.

5.4 Observations

In the translation of associated types, it becomes clear why F_C includes coercions over type constructors of higher kind. Consider the following class of monads with references:

```
class Monad m => RefMonad m where
  type Ref m :: * -> *
  newRef  :: a -> m (Ref m a)
  readRef :: Ref m a -> m a
  writeRef :: Ref m a -> a -> m ()
```

This class may be instantiated for the IO monad and the ST monad. The associated type Ref is of higher-kind, which implies that the coercions generated from its definitions will also be higher kinded.

The translation of associated types to plain System F imposes two restrictions on the formation of well-formed programs [5, §5.1], namely (1) that equality constraints for an n parameter type class must have type variables as the first n arguments to its associated types and (2) that class method signatures cannot constrain type class parameters. Both constraints can be lifted in the translation to F_C .

5.5 Guaranteeing consistency for associated types

How do we know that the axioms generated by the source-program associated types and their instance declarations are consistent? The answer is simple. The source-language type system for associated types only makes sense if the instance declarations obey certain constraints, such as non-overlap [6]. Under those conditions, it is easy to guarantee that the axioms derived from the source program are consistent. In this section we briefly sketch why this is the case.

The axiom generated by an instance declaration for an associated type has the form¹ $C : (\forall \bar{a} : \star. S \sigma_1) \sim (\forall \bar{a} : \star. \sigma_2)$, where (a) σ_1 does not refer to any type function, (b) $f\nu(\sigma_1) = \bar{a}$, and (c) $f\nu(\sigma_2) \subseteq \bar{a}$. This is an entirely natural condition and can also be found in [5]. We call an axiom of this form a *rewrite axiom*, and a set of such axioms defines a rewrite system among types.

Now, the source language rules ensure that this rewrite system is *confluent* and *terminating*, using the standard meaning of these terms [2]. We write $\sigma_1 \downarrow \sigma_2$ to mean that σ_1 can be rewritten to σ_2 by zero or more steps, where σ_2 is a normal form. Then we prove that each type has a canonical normal form:

THEOREM 5 (Canonical Normal Forms). *Let Γ be well-formed, terminating and confluent. Then, $\Gamma \vdash_{\text{CO}} \gamma : \sigma_1 \sim \sigma_2$ iff $\sigma_1 \downarrow \sigma'_1$ and $\sigma_2 \downarrow \sigma'_2$ such that $\sigma'_1 = \sigma'_2$.*

Using this result we can decide type equality via a canonical normal form test, and thereby prove consistency:

COROLLARY 3 (AT Consistency). *If Γ contains only rewrite axioms that together form a terminating and confluent rewrite system, then Γ is consistent.*

For example, assume $\Gamma \vdash_{\text{CO}} \gamma : T_1 \overline{\sigma_1} \sim T_2 \overline{\sigma_2}$. Then, we find $T_1 \overline{\sigma_1} \downarrow \sigma'_1$ and $T_2 \overline{\sigma_2} \downarrow \sigma'_2$ such that $\sigma'_1 = \sigma'_2$. None of the rewrite rules affect T_1 or T_2 . Hence, σ'_1 must have the shape $T_1 \overline{\sigma''_1}$ and σ'_2 the shape $T_2 \overline{\sigma''_2}$. Immediately, we find that $T_1 = T_2$ and we are done.

We can state similar results for type functions resulting from functional dependencies. Again, the canonical normal form property is the key to obtain consistency. While sufficient the canonical normal form property is not a necessary condition. Consider the non-confluent but consistent environment $\Gamma = \{c_1 : S_1 [Int] \sim S_2, c_2 : (\forall a : \star. S_1 [a]) \sim (\forall a : \star. [S_1 a])\}$. We find that $\Gamma \vdash_{\text{CO}} \gamma : S_1 [Int] \sim S_2$. But there exists $S_1 [Int] \downarrow [S_1 Int]$ and $S_2 \downarrow S_2$ where $[S_1 Int] \neq S_2$. Similar observations can be made for ill-formed, consistent environments.

6. Related Work

System F with GADTs. Xi et al. [43] introduced the explicitly typed calculus $\lambda_{2,G\mu}$ together with a translation from an implicitly typed source language supporting GADTs. Their calculus has the typing rules for GADTs built in, just like Pottier & Régis-Gianas's MLGI [34]. This is the approach that GHC initially took. F_C is the result of a search for an alternative.

Encoding GADTs in plain System F and F_ω . There are several previous works [3, 9, 30, 42, 7, 39] which attempt an encoding

¹For simplicity, we here assume unary associated types that do not range over higher-kinded types.

of GADTs in plain System F with (boxed) existential types. We believe that these primitive encoding schemes are not practical and often non-trivial to achieve. We discuss this in more detail in [37].

An encoding of a restricted subset of GADT programs in plain System F_{ω} can be found in [33], but this encoding only works for limited patterns of recursion.

Intentional type analysis and beyond. Harper and Morrisett’s visionary paper on intensional type analysis [20] introduced the calculus λ_i^{ML} , which was already sufficiently expressive for a large range of GADT programs, although GADTs only became popular later. Subsequently, Crary and Weirch’s language LX [12] generalised the approach significantly by enabling the analysis of source language types in the intermediate language and by providing a type erasure semantics, among other things. LX’s type analysis is sufficiently powerful to express *closed* type functions which must be primitive recursive. This is related, but different to $F_C(X)$, where type functions are *open* and need not be terminating (see [37]).

Trifonov et al. [41] generalised λ_i^{ML} in a different direction than LX, such that they arrived at a fully reflexive calculus; i.e., one that can analyse the type of any runtime value of the calculus. In particular, they can analyse types of higher kind, an ability that was also crucial in the design of $F_C(X)$. However, Trifonov et al.’s work corresponds to λ_i^{ML} and LX in that it applies to *closed*, primitive-recursive type functions.

Calculi with explicit proofs. Licata & Harper [25] introduced the calculus LH to represent programs in the style of Dependent ML. LH’s type terms include lambdas, and its definitional equality therefore includes a beta rule, whereas F_C ’s definitional equality is simpler, being purely syntactic. LH’s propositional equality enables explicit proofs of type equality, much as in $F_C(X)$. These explicit proofs are the basis for the definition of *retyping functions* that play a similar role to our cast expressions. In contrast, F_C ’s propositional equality lacks some of LH’s equalities, namely those including certain forms of inductive proofs as well as type equalities whose retypings have a computational effect. The price for LH’s added expressiveness is that retypings — even if they amount to the identity on values — can incur non-trivial runtime costs and (together with LH types) cannot be erased without meta-level proofs that assert that particular forms of retypings are guaranteed to be identity functions.

Another significant difference is that in LH, as in LX, type functions are *closed* and must be *primitive recursive*; whereas in $F_C(X)$, they are open and need not be terminating. These properties are very important in our intended applications, as we argued in Section 3.4. Finally, $F_C(X)$ admits optimising transformations that are not valid in LH, as we discussed in Section 3.8.

Shao et al.’s impressive work [35] illustrates how to integrate an entire proof system into typed intermediate and assembly languages, such that program transformations preserve proofs. Their type language TL resembles the calculus of inductive constructions (CIC) and, among other things, can express retypings witnessed by explicit proofs of equality [35, Section 4.4], not unlike LH. TL is much more expressive and complex than $F_C(X)$ and, like LH, does not support open type functions.

Coercion-based subtyping. Mitchell [29] introduced the idea of inserting coercions during type inference for an ML-like languages. However, Mitchell’s coercion are not identities, but perform coercions between different numeric types and so forth. A more recent proposal of the same idea was presented by Kießling and Luo [22]. Subsequently, Mitchell [28] also studied coercions that are operationally identities to model type refinement for type inference in systems that go beyond Hindley/Milner.

Much closer to F_C is the work by Breazu-Tannen et al. [4] who add a notion of coercions to System F to translate languages featuring inheritance polymorphism. In contrast to F_C , their coercions model a subsumption relationship, and hence are not symmetric. Moreover, their coercions are values, not types. Nevertheless, they introduce coercion combinators, as we do, but they don’t consider decomposition, which is crucial to translating GADTs. The focus of their paper is the translation of an extended version of Cardelli & Wegner’s Fun, and in particular, the coherence properties of that translation.

Similarly, Crary [11] introduces a coercion calculus for inclusive subtyping. It shares the distinction between plain values and coercion values with our system, but does not require quantification over coercions, nor does it consider decomposition.

Intuitionistic type theory, dependent types, and theorem provers. The ideas from Mitchell’s work [29, 28] have also been transferred to dependently typed calculi as they are used in theorem provers; e.g., based on the Calculus of Constructions [8]. Generally, our coercion terms are a simple instance of the proof terms of logical frameworks, such as LF [18], or generally the evidence in intuitionistic type theory [26]. This connection indicates several directions for extending the presented system in the direction of more powerful dependently typed languages, such as Epigram [27].

Translucency and singleton kinds. In the work on ML-style module systems, type equalities are represented as singleton kinds, which are essential to model translucent signatures [14]. Recent work [15] demonstrated that such a module calculus can represent a wide range of type class programs including associated types. Hence, there is clearly an overlap with $F_C(X)$ equality axioms, which we use to represent associated types. Nevertheless, the current formulation of modular type classes covers only a subset of the type class programs supported by Haskell systems, such as GHC. We leave a detailed comparison of the two approaches to future work.

7. Conclusions and further work

We showed that explicit evidence for type equalities is a convenient mechanism for the type-preserving translation of GADTs, associative types, and functional dependencies. We implemented $F_C(X)$ in its full glory in GHC, a widely used, state-of-the-art, highly optimising Haskell compiler. At the same time, we re-implemented GHC’s support for `newtypes` and GADTs to work as outlined in §2 and added support for associated (data) types. Consequently, this implementation instantiates the decision procedure for consistency, “X”, to a combination of that described in Section 4 and 5. The F_C -version of GHC is now *the* main development version of GHC and supports our claim that $F_C(X)$ is a practical choice for a production system.

An interesting avenue for future work is to find good source language features to expose more of the power of F_C to programmers.

Acknowledgements

We thank James Cheney, Karl Crary, Roman Leshchinskiy, Dan Licata, Conor McBride, Benjamin Pierce, François Pottier, Robert Harper, and referees for ICFP’06, POPL’07 and TLDI’07 for their helpful comments on previous versions of this paper. We are grateful to Stephanie Weirich for interesting discussions during the genesis of F_C , to Lennart Augustsson for a discussion on encoding associated types in System F, and to Andreas Abel for pointing out the connection to [33].

References

- [1] M. Abadi, L. Cardelli, and P.-L. Curien. Formal parametric polymorphism. In *Proc. of POPL ’93*, pages 157–170, New York,

- NY, USA, 1993. ACM Press.
- [2] F. Baader and T. Nipkow. *Term rewriting and all that*. Cambridge University Press, 1999.
- [3] A. I. Baars and S. D. Swierstra. Typing dynamic typing. In *Proc. of ICF'02*, pages 157–166. ACM Press, 2002.
- [4] V. Breazu-Tannen, T. Coquand, C. Gunter, and A. Scedrov. Inheritance as implicit coercion. *Information and Computation*, 93(1):172–221, July 1991.
- [5] M. M. T. Chakravarty, G. Keller, and S. Peyton Jones. Associated type synonyms. In *Proc. of ICFP '05*, pages 241–253, New York, NY, USA, 2005. ACM Press.
- [6] M. M. T. Chakravarty, G. Keller, S. Peyton Jones, and S. Marlow. Associated types with class. In *Proc. of POPL '05*, pages 1–13. ACM Press, 2005.
- [7] C. Chen, D. Zhu, and H. Xi. Implementing cut elimination: A case study of simulating dependent types in Haskell. In *Proc. of PADL'04*, volume 3057 of *LNCS*, pages 239–254. Springer-Verlag, 2004.
- [8] G. Chen. Coercive subtyping for the calculus of constructions. In *Proc. of POPL'03*, pages 150–159, New York, NY, USA, 2003. ACM Press.
- [9] J. Cheney and R. Hinze. A lightweight implementation of generics and dynamics. In *Proc. of Haskell Workshop '02*, pages 90–104. ACM Press, 2002.
- [10] J. Cheney and R. Hinze. First-class phantom types. TR 1901, Cornell University, 2003.
- [11] K. Cray. Typed compilation of inclusive subtyping. In *Proc. of ICFP'00*, pages 68–81, New York, NY, USA, 2000. ACM Press.
- [12] K. Cray and S. Weirich. Flexible type analysis. In *Proc. of ICFP'99*, pages 233–248. ACM Press, 1999.
- [13] N. Dershowitz and J.-P. Jouannaud. *Handbook of Theoretical Computer Science (Volume B: Formal Models and Semantics)*, chapter 6: Rewrite Systems. Elsevier Science Publishers, 1990.
- [14] D. Dreyer, K. Cray, and R. Harper. A type system for higher-order modules. In *Proc. of POPL'03*, pages 236–249, New York, NY, USA, 2003. ACM Press.
- [15] D. Dreyer, R. Harper, and M. M. T. Chakravarty. Modular type classes. In *Proc. of POPL '07*. ACM Press, 2007. To appear.
- [16] G. J. Duck, S. Peyton Jones, P. J. Stuckey, and M. Sulzmann. Sound and decidable type inference for functional dependencies. In *Proc. of (ESOP'04)*, number 2986 in *LNCS*, pages 49–63. Springer-Verlag, 2004.
- [17] C. V. Hall, K. Hammond, S. L. Peyton Jones, and P. L. Wadler. Type classes in Haskell. *ACM Trans. Program. Lang. Syst.*, 18(2):109–138, 1996.
- [18] R. Harper, F. Honsell, and G. Plotkin. A Framework for Defining Logics. In *Proc. of LICS'87*, pages 194–204. IEEE Computer Society Press, 1987.
- [19] R. Harper and J. C. Mitchell. On the type structure of Standard ML. *ACM Transactions on Programming Languages and Systems*, 15(2):211–252, 1993.
- [20] R. Harper and G. Morrisett. Compiling polymorphism using intensional type analysis. In *Proc. of POPL'95*, pages 130–141. ACM Press, 1995.
- [21] M. P. Jones. Type classes with functional dependencies. In *Proceedings of the 9th European Symposium on Programming (ESOP 2000)*, number 1782 in *Lecture Notes in Computer Science*. Springer-Verlag, 2000.
- [22] R. Kießling and Z. Luo. Coercions in Hindley-Milner systems. In *Types for Proofs and Programs: Third International Workshop, TYPES 2003*, number 3085 in *LNCS*, pages 259–275, 2004.
- [23] J. Lassez, M. Maher, and K. Marriott. Unification revisited. In *Foundations of Deductive Databases and Logic Programming*. Morgan Kaufman, 1987.
- [24] K. Läufer and M. Odersky. Polymorphic type inference and abstract data types. *ACM Transactions on Programming Languages and Systems*, 16(5):1411–1430, 1994.
- [25] D. R. Licata and R. Harper. A formulation of Dependent ML with explicit equality proofs. Technical Report CMU-CS-05-178, Carnegie Mellon University, Dec. 2005.
- [26] P. Martin-Löf. *Intuitionistic Type Theory*. Bibliopolis-Napoli, 1984.
- [27] C. McBride and J. McKinna. The view from the left. *Journal of Functional Programming*, 14(1):69–111, 2004.
- [28] J. Mitchell. Polymorphic type inference and containment. In *Logical Foundations of Functional Programming*, pages 153–193. Addison-Wesley, 1990.
- [29] J. C. Mitchell. Coercion and type inference. In *Proc of POPL'84*, pages 175–185. ACM Press, 1984.
- [30] E. Pasalic. *The Role of Type Equality in Meta-Programming*. PhD thesis, Oregon Health & Science University, OGI School of Science & Engineering, September 2004.
- [31] S. Peyton Jones, D. Vytiniotis, S. Weirich, and G. Washburn. Simple unification-based type inference for GADTs. In *Proc. of ICFP'06*, pages 50–61. ACM Press, 2006.
- [32] B. Pierce. *Types and programming languages*. MIT Press, 2002.
- [33] B. Pierce, S. Dietzen, and S. Michaylov. Programming in higher-order typed lambda-calculi. Technical report, Carnegie Mellon University, 1989.
- [34] F. Pottier and Y. Régis-Gianas. Stratified type inference for generalized algebraic data types. In *Proceedings of the 33rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 232–244. ACM Press, 2006.
- [35] Z. Shao, V. Trifonov, B. Saha, and N. Pappaspyrou. A type system for certified binaries. *ACM Trans. Program. Lang. Syst.*, 27(1):1–45, 2005.
- [36] T. Sheard. Languages of the future. In *OOPSLA '04: Companion to the 19th Annual ACM SIGPLAN Conference on Object-Oriented Programming Systems, Languages, and Applications*, pages 116–119, New York, NY, USA, 2004. ACM Press.
- [37] M. Sulzmann, M. M. T. Chakravarty, S. Peyton Jones, and K. Donnelly. System F with type equality coercions. Technical Report UNSW-CSE-TR-0624, The University of New South Wales, 2006. <ftp://ftp.cse.unsw.edu.au/pub/doc/papers/UNSW/0624.pdf>.
- [38] M. Sulzmann, T. Schrijvers, and P. J. Stuckey. Type inference for GADTs via Herbrand constraint abduction. <http://www.comp.nus.edu.sg/~sulzmann>, July 2006.
- [39] M. Sulzmann and M. Wang. A systematic translation of guarded recursive data types to existential types. Technical Report TR22/04, The National University of Singapore, 2004.
- [40] M. Sulzmann, J. Wazny, and P. J. Stuckey. A framework for extended algebraic data types. In *Proc. of FLOPS'06*, volume 3945 of *LNCS*, pages 47–64. Springer-Verlag, 2006.
- [41] V. Trifonov, B. Saha, and Z. Shao. Fully reflexive intensional type analysis. In *Proc. of ICFP'00*, pages 82–93, New York, NY, USA, 2000. ACM Press.
- [42] S. Weirich. Type-safe cast (functional pearl). In *Proc. of ICFP'00*, pages 58–67. ACM Press, 2000.
- [43] H. Xi, C. Chen, and G. Chen. Guarded recursive datatype constructors. In *Proc. of POPL '03*, pages 224–235, New York, NY, USA, 2003. ACM Press.