

# Closure Converting the Universes

?? and STEFAN MONNIER, DIRO - Université de Montréal

Type preserving closure conversion of languages with dependent types has proved difficult. It took until 2018 to get the first solution to the problem, and that solution relies on language constructs custom-made for the purpose. A first part of the problem is the fact that closure conversion inevitably requires some form of impredicativity since a function can have free variables that belong to a higher universe than itself, but none of the existing forms of impredicativity (other than those known to be inconsistent) satisfy the needs of closure conversion. A second part is that closure conversion exposes internal details of functions, and those details affect the definitional equality of (converted) functions, hence breaking type preservation.

In this paper we propose to solve the first problem with the use of a new form of impredicativity, and the second with a more restrictive notion of function equality. This allows us to define a closure conversion that relies only on generic language constructs, yet handles a dependently typed language with a tower of universes.

CCS Concepts: • **Theory of computation** → **Type theory**; *Higher order logic*; Logic and verification; • **Software and its engineering** → *Compilers*; Functional languages.

Additional Key Words and Phrases: Closure conversion, Dependent types, Universe polymorphism, Impredicativity, Function equality

## 1 INTRODUCTION

Closure conversion is a core part of the implementation of a functional programming language. Preserving the full type information across this compilation stage is nowadays common for traditional functional languages, but not so for dependently typed languages, where it stayed an open problem until recently and where the existing solutions are not fully satisfactory yet.

Preserving the types across the various stages of the compiler is particularly important for dependently typed programming languages where the embedded proofs in the source code apply only to ... the source code, whereas for the purpose of formal verification of software we would like to be able to guarantee that those proofs also apply to the compiled code, and that the various pieces that were compiled separately can indeed be linked without breaking any of their invariants.

This tends to get harder the further you progress in the compiler pipeline. And dependent types make it a lot more difficult since runtime terms can appear in the types: as the compilation phases modify those runtime terms, the typing information tends to be affected in profound ways.

In the case of closure conversion, there are three main hurdles: the first is the need to explain to the type system that the closure object we pass at runtime to the closed code indeed contains those precise values that were held in the free variables when we constructed the closure. The second is the fact that closure objects can contain not only other closures of the same type, but also those closures's types, and hence their own type, hence requiring a form of impredicativity. The third is the fact that after closure conversion, the variables captured by a closure are now in full view, which tends to completely change the notion of equality between functions and hence equality between types.

[Bowman and Ahmed \[2018\]](#) provided a first, and only, solution to preserve the types across the closure conversion phase of a dependently typed language, but that solution still has two main shortcomings: first, it handles only the Calculus of Constructions, which is only a subset of most

---

Authors' address: ??, ??@iro.umontreal.ca; Stefan Monnier, monnier@iro.umontreal.ca, DIRO - Université de Montréal, Montréal, Canada.

---

2023. 2475-1421/2023/1-ART \$15.00

<https://doi.org/>

dependently typed languages used nowadays, so to be usable for an actual system it would typically need to be extended to cover inductive types and a tower of universes. Second, it relies on a custom construct to represent closures in the output language. While this is a very sensible pragmatic choice, it does beg the questions: what is so special about closures that we do not know how to represent them explicitly without resorting to a custom construct? What extra features would a language need in order to be able to accommodate closure converted code?

In this article, we show an alternate path which makes different tradeoffs in order to find a way to preserve the types across the closure conversion phase while sticking to generic language constructs like dependent tuples and existential packages. Admittedly, our solution still retains a few unusual characteristics, but gets closer to the ideal, and gives a partial answer to the previous questions.

Our contributions are the following:

- A type-preserving closure conversion algorithm for a dependently typed language with a tower of universes.
- A novel solution to the problem of aligning the definition of function equivalence in the code before and after closure conversion.
- A new form of universe polymorphism to handle the need for impredicativity in the closure conversion algorithm.
- The use of equality proofs instead of the translucent types used by [Minamide et al. \[1996\]](#). This is admittedly already folklore at this point, and was also sketched by [Bowman and Ahmed \[2018\]](#), but we are not aware of any previous work that shows the actual details.

We present the basic problem of type-preserving closure conversion in Section 2. In Section 3 we present each of the three difficulties specific to closure conversion of dependently typed languages, along with how we solved them. In Section 4 we show more formally the input and output languages we use and the closure conversion algorithm. In Section 5 we extend the input language to match the output language. In Section 6 we discuss some of our design decisions.

## 2 BACKGROUND

Let's consider a predicative Pure Type System (PTS) with a tower of universes as our input language:

$$\begin{array}{ll}
 (\text{levels}) & \ell ::= 1 \mid S \ell \\
 (\text{sorts}) & s ::= \mathcal{U}_\ell \\
 (\text{exps}) & e, \tau ::= x \mid s \mid (e : \tau) \\
 & \quad \mid (x : \tau_1) \rightarrow \tau_2 \mid e_1 e_2 \mid \lambda x. e
 \end{array}$$

Here  $\mathcal{U}_\ell$  denotes the universe of level  $\ell$  and we use  $(x : \tau_1) \rightarrow \tau_2$  to denote the type of (dependent) functions, which we will shorten to  $\tau_1 \rightarrow \tau_2$  when  $x$  is not used in  $\tau_2$ .  $(e : \tau)$  is a type annotation to help the bidirectional type checking, so that  $\lambda x. e$  does not need a type annotation. 1 is the base universe level and  $S \ell$  returns the successor of  $\ell$ . Our universe levels start counting at 1 rather than 0 for reasons that will become clear later.

Closure conversion needs to reify closures as data-structures usually represented using tuples, so our output language will additionally require some kind of tuples:

$$\begin{array}{ll}
 (\text{telescopes}) & \Gamma ::= \bullet \mid \Gamma, x : \tau \\
 (\text{exps}) & e, \tau ::= \dots \\
 & \quad \mid \langle \Gamma \rangle \mid \langle e_1, \dots, e_n \rangle \mid e.i
 \end{array}$$

$\langle \Gamma \rangle$  is the type constructor for (dependent) tuples, where  $\Gamma$  lists the types of the fields and where the type of later fields can refer to values of earlier fields;  $\langle e_1, \dots, e_n \rangle$  is the tuple constructor; and  $e.i$  returns the  $i^{\text{th}}$  field of the tuple  $e$ . For convenience we will use a bit of syntactic sugar and write

let  $\langle x_1, \dots, x_n \rangle = e$  in  $e'$  to mean  $e'[e.1/x_1, \dots, e.n/x_n]$  and  $\lambda\langle \vec{x} \rangle.e$  to mean  $\lambda y.\text{let } \langle \vec{x} \rangle = y \text{ in } e$  where  $\vec{x}$  stands for  $x_1, \dots, x_n$ .

Disregarding types for now, the closure conversion of a function like  $f = \lambda x.x + y + z$  may look like the following:

$$\llbracket f \rrbracket = \langle \langle y, z \rangle, \lambda\langle x_e, x \rangle.\text{let } \langle y, z \rangle = x_e \text{ in } x + y + z \rangle$$

I.e. a pair whose first element holds the “environment”, i.e. the values of the variables captured by the closure ( $y$  and  $z$ ), and whose second element holds the “code”, i.e. a closed function. The code in turn expects as argument a pair  $\langle x_e, x \rangle$  whose second element ( $x$ ) is the actual argument to the function, and whose first element ( $x_e$ ) should be the environment, holding the values of the captured variables, i.e. the first element of the closure.

Accordingly, after closure conversion, a call like  $f \ 42$  would turn into:

$$\llbracket f \ 42 \rrbracket = \text{let } \langle f_e, f_c \rangle = f \text{ in } f_c \ \langle f_e, 42 \rangle$$

If we build the closure naïvely like we did above, its type would look like the following:

$$\langle \text{env} : \langle y : \text{Int}, z : \text{Int} \rangle, \text{code} : \langle x_e : \langle y : \text{Int}, z : \text{Int} \rangle, x : \text{Int} \rangle \rightarrow \text{Int} \rangle$$

But the type of this pair representing a function whose original type was  $\text{Int} \rightarrow \text{Int}$  now exposes the number and types of the captured variables. This implies that after closure conversion, two functions which originally had the same type can end up being represented by data structures of incompatible types, thus breaking the type preservation property. For this reason, closures are usually given an existential type that hides the type of the inner tuple representing the environment. Since our tuples are dependently typed, we can do that by prepending to our closures a third field that holds this “hidden” type:

$$\begin{aligned} \llbracket f \rrbracket = & \langle \langle y : \text{Int}, z : \text{Int} \rangle, \\ & \langle y, z \rangle, \\ & \lambda\langle x_e, x \rangle.\text{let } \langle y, z \rangle = x_e \text{ in } x + y + z \rangle \\ & : \langle t \quad : \mathcal{U}, \\ & \quad \text{env} : t, \\ & \quad \text{code} : \langle x_e : t, x : \text{Int} \rangle \rightarrow \text{Int} \rangle \end{aligned}$$

Now the type does not expose the shape of the captured environment, so two different functions that had the same type before conversion will still have the same type after conversion even if they capture a different number of variables or variables of different types.

This approach works well for System-F [Morrisett et al. 1998], but when we try to apply it to a dependently-typed language there are 3 problems that come up:

- (1) Some of the captured variables may also appear in the *type* of the function. In that case, our closure will be ill-typed because the type checker cannot see that the values extracted from  $x_e$  are the same as the ones that were captured.
- (2) The closure will tend to belong to too high a universe compared to the original function, because it contains the types of the captured variables.
- (3) The conversion does not preserve equivalence of terms. For example, when  $y$  is equal to 7, the above closure will look very different from the closure generated for the equivalent function  $\lambda x.x + 7 + z$ . Since terms can appear in types, this means that types may also fail to be equivalent after closure conversion.

All three problems need to be solved if we want the closure conversion of properly typed code to still be properly typed.

### 3 OUR APPROACH

We present here in more detail the three mentioned problems that afflict type preserving closure conversion in the specific case of a dependently typed language, and we present the solution we used to solve each one.

#### 3.1 Taming dependencies

The first problem we face was identified by [Minamide et al. \[1996\]](#) already: if some of the free variables over which we close a function can appear in its type, then the simple existential encoding fails. For example, say we have the following primitive:

$$\text{makevec} \quad : \quad (t : \mathcal{U}) \rightarrow (len : \text{Nat}) \rightarrow t \rightarrow \text{Vec } t \text{ len}$$

and we want to perform closure conversion on the following function:

$$\lambda x. \text{makevec } \alpha \ n \ x \quad : \quad \alpha \rightarrow \text{Vec } \alpha \ n$$

where  $\alpha$  and  $n$  are its two free variables. The encoding shown before would give us:

$$\begin{aligned} & \llbracket \lambda x. \text{makevec } \alpha \ n \ x \rrbracket \\ & = \langle \langle \alpha : \mathcal{U}, n : \text{Nat} \rangle, \\ & \quad \langle \alpha, n \rangle, \\ & \quad \lambda \langle x_e, x \rangle. \text{let } \langle \alpha', n' \rangle = x_e \text{ in makevec } \alpha' \ n' \ x \rangle \\ & : \langle t \quad : \mathcal{U}, \\ & \quad env \quad : t, \\ & \quad code : \langle x_e : t, x : \alpha \rangle \rightarrow \text{Vec } \alpha \ n \rangle \end{aligned}$$

But this code is ill-typed: in  $\text{makevec } \alpha' \ n' \ x$ , we tell  $\text{makevec}$  that we will provide an element of type  $\alpha'$  but then pass it  $x$  which has type  $\alpha$ . Also, even if we were generous enough to accept the argument  $x$ , the return type would not match its expected type because  $\text{makevec } \alpha' \ n' \ x$  returns a value of type  $\text{Vec } \alpha' \ n'$  rather than  $\text{Vec } \alpha \ n$ .

In the case of a language like System-F, [Morrisett et al. \[1998\]](#) showed that you can circumvent the problem by not closing over type variables, which are the only variables that can appear in the type in such a language, and since types can be erased we don't really need to close over them. In the above example, maybe  $\alpha$  would not be used at run-time and we could then leave it as a free variable, but that is not an option for  $n$  since that argument is needed at run time to determine the size of the returned vector.

Arguably the only value that  $x_e$  above can take is  $\langle \alpha, n \rangle$  and thus  $\alpha'$  is always equal to  $\alpha$  and  $n'$  is always equal to  $n$ . You might even prove it via parametricity. Nevertheless, while we may know this, the type system does not. Worse, there is simply no way to write a closed function of the above type because in order to return something of type  $\text{Vec } \alpha \ n$  it would *have to* refer to  $\alpha$  and  $n$ , defeating the purpose of the closure conversion. For this reason, if we want the code to match its expected type, we need to change the type so as to make it more obvious that we will always receive in  $x_e$  the exact value stored in the *env* field of the tuple. [Minamide et al. \[1996\]](#) did this using a feature they call “translucent types”, and we could also solve it also using some form of singleton types, but in languages with dependent types, the more natural solution is to use an

197 equality proof:

```

198      $\llbracket \lambda x. \text{makevec } \alpha \ n \ x \rrbracket$ 
199     =
200      $\langle \langle \alpha : \mathcal{U}, n : \text{Nat} \rangle,$ 
201        $\langle \alpha, n \rangle,$ 
202        $\lambda \langle x_e, x, p \rangle. \text{let } \langle \alpha', n' \rangle = x_e \text{ in}$ 
203          $\text{let } x' = \text{cast } (eq\_comm \ p) \ (\lambda x'_e. x'_e.1) \ x \text{ in}$ 
204          $\text{let } res = \text{makevec } \alpha' \ n' \ x' \text{ in}$ 
205          $\text{cast } p \ (\lambda x'_e. \text{Vec } (x'_e.1) \ (x'_e.2)) \ res \rangle$ 
206
207     :  $\langle t \quad : \mathcal{U},$ 
208        $env \ : \ t,$ 
209        $code \ : \ \langle x_e : t, x : \alpha, p : (x_e = env) \rangle \rightarrow \text{Vec } \alpha \ n \rangle$ 
210
211

```

211 On the last line we see that the *code* now takes an additional argument  $p$  holding a proof that  $x_e$  is  
 212 the same as  $env$ . In the corresponding code, we see that this proof object is first passed to *cast* in  
 213 order to turn the input argument  $x$  of type  $\alpha$  into  $x'$  of type  $x_e.1$  (which is also known here as  $\alpha'$ ),  
 214 and then used a second time at the end to convert the result from  $\text{Vec } (x_e.1) \ (x_e.2)$  (also known as  
 215  $\text{Vec } \alpha' \ n'$ ) to  $\text{Vec } \alpha \ n$ .

216 The proof  $p$  allows us to convert back and forth between the external types which refer to the  
 217 surrounding variables and the internal types which refer only to variables local to the function.

218 There is one wrinkle remaining here: the example function we convert is not dependently typed,  
 219 so the function  $\lambda x'_e. \text{Vec } (x'_e.1) \ (x'_e.2)$  we pass to the second *cast* to describe the return type does  
 220 not need to refer to the argument  $x$ . In the general case, the return type may refer to the argument  
 221  $x$ . But this function can't refer to  $x$  for two reasons: first, because it would then not be closed, but  
 222 more importantly because the  $x$  it needs would be the actual  $x$  on one side of the equality and  $x'$   
 223 on the other side. The usual solution to this problem is to merge both *casts* into one as follows:

```

224
225     let  $f' = \lambda x'. \text{makevec } \alpha' \ n' \ x' \text{ in}$ 
226     let  $f = \text{cast } p \ (\lambda x'_e. (x'_e.1) \rightarrow \text{Vec } (x'_e.1) \ (x'_e.2)) \ f' \text{ in}$ 
227      $f \ x$ 
228

```

229 But we cannot use this solution either because, again,  $f'$  is not a closed function. This is a variant  
 230 of the *convoy pattern* [Chlipala 2013], which always relies crucially on non-closed functions, and  
 231 it is important to make sure our target language supports it. For this reason, our target language  
 232 replaces this *cast* operation with a *letcast* which includes the core element of the convoy pattern. It  
 233 takes the following form:

```

234
235     letcast  $[e_=: e_m] \ x = e_1 \text{ in } e_2 \ \sim \ (\text{cast } e_=: e_m \ (\lambda x. e_2)) \ e_1$ 
236

```

237 The  $e_=:$  is the proof of equality,  $e_m$  is the *motive* which describes how to use the equality to change  
 238 the types. This *letcast* operation is arguably the only primitive in our target language that is most  
 239 visibly adjusted to accommodate the specific needs of a language after closure conversion.

### 241 3.2 Taming universes

242 In the previous section, we just used  $\mathcal{U}$  as the universe of types, but in the context of a language  
 243 with a tower of universes, we need to qualify it with the corresponding level.

Let us consider the source function  $g = \lambda x.1 + f(x - 1)$ , of type  $\text{Int} \rightarrow \text{Int}$ . Its type after closure conversion becomes:

$$\begin{aligned} & \llbracket \text{Int} \rightarrow \text{Int} \rrbracket \\ & = \langle t \quad : \mathcal{U}_\ell, \\ & \quad \text{env} : t, \\ & \quad \text{code} : \langle x_e : t, x : \text{Int}, p : (x_e = \text{env}) \rangle \rightarrow \text{Int} \rangle \end{aligned}$$

Where the  $\ell$  subscript in  $\mathcal{U}_\ell$  is the universe level inhabited by the captured environment. This tuple type inhabits the universe  $\mathcal{U}_{(S \ell)}$ . But when building the closure for  $g = \lambda x.1 + f(x - 1)$ , the captured environment contains  $f$  which is also of type  $\text{Int} \rightarrow \text{Int}$  and whose type after closure conversion will thus also be the tuple type above. So we would need to fit into the  $t$  field of the tuple above a tuple type belonging to  $\mathcal{U}_{(S \ell)}$ , which is clearly too large to fit into the  $\mathcal{U}_\ell$  type of this field.

More specifically, we fundamentally need here some form of impredicativity such that our closure's existential quantification can quantify over a universe which includes its own. One solution is to use a language like  $\lambda^*$  that collapses all the universes into a single  $\mathcal{U}$  that belongs to itself, but those languages are known to be inconsistent [Hurkens 1995]. All forms of impredicativity known to be consistent are too weak to accommodate our needs. Bowman and Ahmed [2018] were the first to provide a solution to this problem by circumventing it and introducing a custom-made type construct for closures instead of relying on existential quantification.

While their solution only accommodates the Calculus of Constructions, it can likely be adapted to a language with a tower of universes. Yet, we would prefer a solution that does not rely on such custom constructs. So, instead we go back to the tuple type above and see another problem with it: the universe level  $\ell$  needed for the field  $t$  of the tuple depends on the types of captured variables, and hence exposes details of the captured variables. So again, we face the problem that the closure conversion of two different closures of the same original type may end up having different types depending on the set of captured variables, thus breaking again our dear type preservation property.

So we apply the same existential quantification trick, but this time quantifying over the universe level:

$$\begin{aligned} & \llbracket \text{Int} \rightarrow \text{Int} \rrbracket \\ & = \exists l. \langle t \quad : \mathcal{U}_l, \\ & \quad \text{env} : t, \\ & \quad \text{code} : \langle x_e : t, x : \text{Int}, p : (x_e = \text{env}) \rangle \rightarrow \text{Int} \rangle \end{aligned}$$

We don't store the universe level  $l$  in the tuple but use a separate  $\exists l.\tau$  construct instead for two reasons: first, because manipulating universe levels as first class values is fraught with danger, and second because it lets us make sure that universe levels can be erased, so we do not need to close over them, saving us from a lot of extra complications such as the need to manipulate proofs of equality between universe levels.

With this extra existential quantification, we recover the property that the converted type of a function is always the same regardless of its free variables. But there still remains the question of the universe to which this type should belong. Since it can hold values from arbitrary universe levels, a predicative type theory such as Agda would put such a type in a special universe level  $\omega$  beyond all other levels and over which  $\exists l.\tau$  cannot quantify. This of course would not satisfy our impredicative needs.

A naïve impredicative choice would put this type in the bottom universe instead, but this would immediately lead to an inconsistent type theory because we could then use dummy  $\exists l.\tau$  wrappers to bring any type down to the bottom universe, making the language equivalent to  $\lambda^*$ .

295 So instead we put  $\exists l.\tau$  in the universe  $\mathcal{U}_{\ell[0/l]}$  where  $\ell$  is the universe level of  $\tau$ . For the above  
 296  $\llbracket \text{Int} \rightarrow \text{Int} \rrbracket$ , it results in the universe  $\mathcal{U}_l$  since the tuple type itself belong to universe  $\mathcal{U}_{(S\ l)}$ .

297 Whether this choice is sound is currently unknown: impredicativity is notoriously tricky and  
 298 this particular form of impredicativity has not been investigated to any significant extent. The  
 299 proportion of impredicative systems which have been found to be inconsistent suggests that the  
 300 odds are not in our favor. This said, we have not been able to adapt known paradoxes like that of  
 301 Hurkens [1995] to this system, although it might be just a reflection of our inexperience. Our hope  
 302 is that even if it proves unsound, there might still be a more restrictive version of it which is sound  
 303 and which at the same time covers the very specific use we make of it: since the  $l$  of an object of  
 304 type  $\exists l.\tau$  is actually erased, we cannot extract  $l$  out of it, nor can we extract from it any type that  
 305 belongs to  $\mathcal{U}_l$  (such as the field  $t$  of our tuple), nor for that matter any value whose type belongs to  
 306  $\mathcal{U}_l$  (such as the field  $env$ ) or contains a element that belongs to  $\mathcal{U}_l$  (such as the field  $code$ ), so really,  
 307 the only thing we can do with those closure objects is to call them. In this sense, they really are  
 308 strictly equivalent to the closures they represent from our source language, which is known to be  
 309 consistent and does not involve any impredicativity.

### 3.3 Taming function equality

312 The final remaining problem is the preservation of equality for functions: for the closure conversion  
 313 to preserve types, we also need to make sure that closure conversion preserves equality between  
 314 types, i.e. if  $\tau_1 \simeq \tau_2$  then  $\llbracket \tau_1 \rrbracket \simeq \llbracket \tau_2 \rrbracket$ .

315 But this is not the case: in our source language  $\lambda x.x + 7$  is equivalent to let  $y = 7$  in  $\lambda x.x + y$   
 316 because let  $y = 7$  in  $\lambda x.x + y$  can be reduced to  $\lambda x.x + 7$ . But after closure conversion, these two  
 317 functions are not equivalent any more. The first will be a closure capturing an empty environment:

$$\begin{aligned} & \llbracket \lambda x.x + 7 \rrbracket \\ & = \langle 1, \langle \langle \bullet \rangle, \langle \rangle, \lambda \langle x_e, x, p \rangle . x + 7 \rangle \rangle \end{aligned}$$

321 While the second will be a closure capturing an environment containing the value of  $y$ :

$$\begin{aligned} & \llbracket \text{let } y = 7 \text{ in } \lambda x.x + y \rrbracket \\ & \text{let } y = 7 \text{ in } \langle 0, \langle \langle y : \text{Int} \rangle, \langle y \rangle, \lambda \langle x_e, x, p \rangle . x + x_e.1 \rangle \rangle \\ & \rightsquigarrow \\ & \langle 1, \langle \langle y : \text{Int} \rangle, \langle 7 \rangle, \lambda \langle x_e, x, p \rangle . x + x_e.1 \rangle \rangle \end{aligned}$$

327 These two closures are clearly different and it seems very difficult to adjust our language's reduction  
 328 rules so as to allow them to treat those two objects as equivalent. The custom-made closure construct  
 329 used by Bowman and Ahmed [2018] to circumvent the problem of impredicativity saves them again  
 330 here, since it allows them to provide a specific  $\eta$ -equivalence rule for those objects. But with our  
 331 use of tuples, our hands are tied.

332 We could try to replace the tuple field  $t$  with an actual existential quantification (so as to make  
 333 this  $t$  erasable), and then equip this existential quantification with an appropriate  $\eta$ -equivalence  
 334 rule, as done in [Bowman 2018], but instead we decided to attack this problem at the other end:  
 335 given the fact that we won't consider those functions equivalent after closure conversion, we  
 336 change the source language's notion of function equality such that they are also considered as  
 337 non-equivalent in the source language.

338 Substitution is not the only operation that threatens the preservation of function equivalence.  
 339 Another case is  $\beta$ -reductions in the body of a function. For example  $\lambda x.x$  and  $\lambda x.((\lambda y.x) z)$  are  
 340 traditionally considered as equivalent in type theories, but after closure conversion the former  
 341 will be a closure capturing no variables while the latter captures a variable  $z$ , so they will not be  
 342 considered equivalent any more.

So we need a new notion of function equality which does not treat those examples as equivalent, yet still allows functions to be equal to themselves. We do that in two parts: first, functions are considered equivalent only if their body is identical (rather than equivalent), basically preventing “reduction under  $\lambda$ ”; second we prevent substitutions from entering the body of  $\lambda$ -expressions until the moment they are called, so as to mimic in the source language what happens after closure conversion.

More specifically, we use a limited form of explicit substitutions [Abadi et al. 1990] and change the syntax of functions to  $\lambda^\sigma x.e$  where  $\sigma$  is a pending substitution (and where we denote the identity substitution by an empty  $\sigma$ ). For example, we have:

$$\text{let } y = 7 \text{ in } \lambda x.x + y \simeq \lambda^{7/y} x.x + y$$

While  $\lambda x.x + 7$  and  $\lambda^{7/y} x.x + y$  behave the same, they are not equivalent any more according to the equivalence relation  $\simeq$ .

Experience teaches us that confluence of explicit substitutions is not a given [Curien et al. 1996], so we need to define substitutions with care. There are two issues:

- **Membership:** We can omit the substitution of a variable which is not in the set of free variables. But the set of free variables of a term can be affected by  $\beta$ -reductions, so the sequencing of those operations can affect the end result.
- **Ordering:** we need to impose some ordering on the elements of a pending substitution to avoid situations where depending on the sequencing of reductions we can end up with either  $\lambda^{7/y,3/z} x.x + y + z$  or  $\lambda^{3/z,7/y} x.x + y + z$ .

While the set of free variables can be different for equivalent terms because of  $\beta$ -reductions, this does not affect  $\lambda$ -expressions of the source language now that we have decided to disallow reductions under  $\lambda$ . So we are free to omit from the pending substitutions those variables which do not appear free in the body of the function without the risk of losing confluence, as long as we make this choice in a deterministic way. We basically have two choices: either never omit them, or always omit them. Omitting them gives us a slightly stronger notion of equivalence but can break the preservation of equivalence of our closure conversion if the closure conversion captures more variables than necessary.

There are also basically two ways to enforce a stable ordering: either place the variables in the pending substitution in the order of their corresponding scope, so a variable cannot “overtake” another, or use an arbitrary ordering defined by the algorithm that collects the set (well: list) of free variables. This latter choice is applicable only if we omit those variables which do not appear freely in the body of the function.

## 4 CLOSURE CONVERTING THE UNIVERSES

In this section we define the source and target languages for our closure conversion as well as the algorithm itself.

### 4.1 Source language

The source language we intend to convert has the following syntax:

$$\begin{array}{ll}
 \text{(levels)} & \ell ::= 1 \mid S \ell \\
 \text{(sorts)} & s ::= \mathcal{U}_\ell \\
 \text{(substitutions)} & \sigma ::= \mid \sigma, e/x \\
 \text{(terms)} & e, \tau ::= x \mid s \mid (e : \tau) \\
 & \quad \mid (x : \tau_1) \rightarrow \tau_2 \mid e_1 e_2 \mid \lambda^\sigma x.e
 \end{array}$$



$\boxed{\Gamma \vdash e \Rightarrow \tau}$  and  $\boxed{\Gamma \vdash e \Leftarrow \tau}$ :  $e$  has type  $\tau$  in context  $\Gamma$ :

$$\begin{array}{c}
\Gamma \vdash \mathcal{U}_\ell \Rightarrow \mathcal{U}_{(s \ell)} \quad \frac{\Gamma(x) = \tau}{\Gamma \vdash x \Rightarrow \tau} \text{(VAR)} \quad \frac{\Gamma \vdash e \Leftarrow \tau}{\Gamma \vdash e \Leftarrow (\tau : s)} \text{(STRIP)} \\
\frac{\Gamma \vdash \tau \Rightarrow s \quad \Gamma \vdash e \Leftarrow \tau}{\Gamma \vdash (e : \tau) \Rightarrow \tau} \text{(ANN)} \quad \frac{\Gamma \vdash e \Rightarrow \tau_1 \quad \tau_1 \simeq \tau_2}{\Gamma \vdash e \Leftarrow \tau_2} \text{(CONV)} \\
\frac{\Gamma \vdash \tau_1 \Rightarrow \mathcal{U}_{\ell_1} \quad \Gamma, x : \tau_1 \vdash \tau_2 \Rightarrow \mathcal{U}_{\ell_2}}{\Gamma \vdash (x : \tau_1) \rightarrow \tau_2 \Rightarrow \mathcal{U}_{\max(\ell_1, \ell_2)}} \text{(PI)} \quad \frac{\Gamma \vdash e_1 \Rightarrow (x : \tau_1) \rightarrow \tau_2 \quad \Gamma \vdash e_2 \Leftarrow \tau_1}{\Gamma \vdash e_1 e_2 \Rightarrow \tau_2[(e_2 : \tau_1)/x]} \text{(APP)} \\
\frac{\Gamma', x : \tau'_1 \vdash e \Leftarrow \tau'_2 \quad \Gamma \vdash \sigma : \Gamma' \quad \tau_1 = \tau'_1[\sigma] \quad \tau_2 = \tau'_2[\sigma]}{\Gamma \vdash \lambda^{\sigma} x. e \Leftarrow ((x : \tau_1) \rightarrow \tau_2)} \text{(LAM)} \\
\frac{\Gamma \vdash \sigma : \Gamma' \quad \Gamma \vdash e \Rightarrow \tau[\sigma]}{\Gamma \vdash \sigma, e/x : \Gamma', x : \tau}
\end{array}$$

$\boxed{e[\sigma]}$ : substitution of  $\sigma$  applied to  $e$ :

$$\begin{array}{l}
x[\sigma] = \begin{cases} \sigma(x) & \text{if } x \in \text{Dom}(\sigma) \\ x & \text{otherwise} \end{cases} \\
s[\sigma] = s \\
(e : \tau)[\sigma] = (e[\sigma] : \tau[\sigma]) \\
((x : \tau_1) \rightarrow \tau_2)[\sigma] = (x : \tau_1[\sigma]) \rightarrow (\tau_2[\sigma]) \\
(e_1 e_2)[\sigma] = (e_1[\sigma]) (e_2[\sigma]) \\
(\lambda^{\sigma_1} y. e)[\sigma] = \lambda^{\sigma_1} y. e \\
(\lambda^{\sigma_1} y. e)[\sigma_2, e_x/x] = \begin{cases} (\lambda^{\sigma_1} y. e)[\sigma_2] & \text{if } x \notin \text{fv}(\lambda y. e : (y : \tau_1) \rightarrow \tau_2) \\ (\lambda^{e_x/\sigma_1, \sigma_1[e_x/x]} y. e)[\sigma_2] & \text{otherwise} \end{cases}
\end{array}$$

$\boxed{e_1 \simeq e_2}$ :  $e_1$  is equivalent to  $e_2$ :

$$(e : \tau) \simeq e \quad (\lambda^{\sigma} x. e_1) e_2 \simeq e_1[\sigma, e_2/x] \quad \frac{e \simeq e'}{E[e] \simeq E[e']} \text{(CONG)}$$

$$\begin{array}{l}
\text{(subst contexts)} \quad \Sigma ::= \Sigma, e/x \mid \sigma, E/x \\
\text{(eval contexts)} \quad E ::= \bullet \mid (E : \tau) \mid (e : E) \\
\quad \mid (x : E) \rightarrow \tau \mid (x : \tau) \rightarrow E \mid E e \mid e E \mid \lambda^{\Sigma} x. e
\end{array}$$

Fig. 1. Typing rules of our source language.

This is the same language as shown at the beginning of Section 2, except for the  $\sigma$  annotation that we added on the  $\lambda$ -expression. The typing rules for this language are shown in Figure 1. We show them in the style of bidirectional rules, so the main judgment is split between  $\Gamma \vdash e \Rightarrow \tau$  and  $\Gamma \vdash e \Leftarrow \tau$ , but other than this detail, the main unusual part of those rules comes from the  $\sigma$  annotation on  $\lambda$ -expressions.

442	(levels)	$\ell$	$::=$	$l \mid 1 \mid S \ell \mid \ell_1 \sqcup \ell_2$
443	(sorts)	$s$	$::=$	$\mathcal{U}_\ell$
444	(ctx)	$\Gamma$	$::=$	$\bullet \mid \Gamma, x : \tau$
445	(substitutions)	$\sigma$	$::=$	$\mid \sigma, e/x$
446	(terms)	$e, \tau$	$::=$	$x \mid s \mid (e : \tau)$
447				$\mid (x : \tau_1) \rightarrow \tau_2 \mid e_1 e_2 \mid \lambda^\sigma x. e$
448				$\mid \langle \Gamma \rangle \mid \langle e_1, \dots, e_n \rangle \mid e.i$
449				$\mid \text{Eq } e_1 e_2 \mid \text{refl} \mid \text{letcast}[e_-, e_m] x = e_1 \text{ in } e_2$
450				$\mid \forall l. \tau \mid \Lambda l. e \mid e[\ell]$
451				$\mid \exists l. \tau \mid \langle \ell, e \rangle \mid \text{open } \langle l, x \rangle = e_1 \text{ in } e_2$
452				
453				
454				
455				

Fig. 2. Syntax of the target language

456 A simpler version of the typing rule for  $\lambda^\sigma x. e$  could just propagate the substitution  $\sigma$  into the  
 457 body  $e$  and check that  $\Gamma, x : \tau_1 \vdash e[\sigma] \Leftarrow \tau_2$ . Sadly, this would allow  $e$  and  $\sigma$  to be ill-typed when  
 458 considered on their own. So instead the LAM rule enforces that both  $e$  and  $\sigma$  are properly typed. The  
 459 downside is that when  $\sigma$  is not the identity substitution this rule does not have enough information  
 460 to be decidable. In practice, what this means is that  $\lambda^\sigma x. e$  needs to be annotated with its  $\Gamma'$ ,  $\tau'_1$ , and  
 461  $\tau'_2$  when  $\sigma$  is a non-identity substitution, but those should normally not occur in actual source code,  
 462 they are only generated internally during reductions.

463 The figure includes the definition of the substitution operation  $e[\sigma]$ . While we write it  $e[\sigma]$  for  
 464 convenience, it is actually defined on the typing derivation of  $e$  because it requires more typing  
 465 information than immediately available from the term. This operation shows how a substitution  $\sigma_2$   
 466 applied to a term  $\lambda^{\sigma_1} y. e$  is accumulated into the  $\sigma$  annotation instead of continuing its way more  
 467 eagerly into the body of the function. In that computation we filter the substitution to only keep  
 468 those variables that are relevant (according to the set of free variables returned by  $\text{fv}$ ), and we keep  
 469 them in the order they appear in the context.

470 Finally the figure also includes a definition of the equivalence relation  $e_1 \simeq e_2$ , where we use  
 471 evaluation contexts  $E$  (and substitution evaluation contexts  $\Sigma$ ) to express the set of congruence  
 472 rules. In type theories, we often omit this part because the primitive rules are presumed to be  
 473 applicable anywhere, but here we need to clarify that they cannot be used within the body of a  
 474 function, so the most important element here is the fact that  $E$  does *not* include a case of the form  
 475  $\lambda^\sigma x. E$ .

476 In all other respects, this is a fairly conventional dependently typed  $\lambda$ -calculus with a predicative  
 477 tower of universes and without universe subsumption.

## 479 4.2 Closure conversion

480 Our target language is a superset of our source language. Its syntax is given in Figure 2. The  
 481 quantification over universe levels adds to our levels  $\ell$  two new cases, one for level variables  $l$ ,  
 482 and another for the maximum of two levels  $\ell_1 \sqcup \ell_2$ . The sorts, contexts, and substitutions stay  
 483 unchanged. But we add many new elements to the terms:

- 484 •  $\langle \Gamma \rangle$  is the type of dependent tuples where  $\Gamma$  describes the fields and their types;  $\langle e_1, \dots, e_n \rangle$   
 485 is the corresponding term constructor; and  $e.i$  is the elimination form which projects the  $i^{\text{th}}$   
 486 field out of  $e$ .
- 487 •  $\text{Eq } e_1 e_2$  is the type of proofs that  $e_1$  and  $e_2$  are equal;  $\text{refl}$  is the corresponding constructor  
 488 of the proof by reflexivity; and  $\text{letcast}[e_-, e_m] x = e_1 \text{ in } e_2$  is the elimination form which  
 489

$$\begin{aligned}
491 \quad \llbracket x \rrbracket &= x \\
492 \quad \llbracket s \rrbracket &= s \\
493 \quad \llbracket (e : \tau) \rrbracket &= (\llbracket e \rrbracket : \llbracket \tau \rrbracket) \\
494 \quad \llbracket (x : \tau_1) \rightarrow \tau_2 \rrbracket &= \exists l. \langle t \quad : \mathcal{U}_l, \\
495 &\quad env : t, \\
496 &\quad code : \langle x_e : t, x : \llbracket \tau_1 \rrbracket, p : (x_e = env) \rangle \rightarrow \llbracket \tau_2 \rrbracket \rangle \\
497 \quad \llbracket e_1 e_2 \rrbracket &= ((open \langle l, c \rangle = \llbracket e_1 \rrbracket \text{ in } c.3 \langle c.2, \llbracket e_2 \rrbracket, refl \rangle) : \llbracket \tau \rrbracket) \\
498 &\quad \text{where } \Gamma \vdash e_1 e_2 \Rightarrow \tau \\
499 \quad \llbracket \lambda^\sigma x. e \rrbracket &= \langle \ell, \langle \llbracket \Gamma' \rrbracket \rangle, \langle \llbracket \vec{x}[\sigma] \rrbracket \rangle, \lambda \langle x_e, x', x_- \rangle. body \rangle \\
500 &\quad \text{where } \Gamma, x : \tau_x \vdash e \Leftarrow \tau_r \\
501 &\quad \vec{x}_f = fv(\lambda x. e : (x : \tau_x) \rightarrow \tau_r) \\
502 &\quad \Gamma' = \Gamma |_{\vec{x}_f} = \vec{x} : \vec{\tau} \\
503 &\quad \Gamma \vdash \langle \Gamma' \rangle \Rightarrow \mathcal{U}_\ell \\
504 &\quad f_m = \lambda \langle \vec{x} \rangle. (x : \llbracket \tau_x \rrbracket) \rightarrow \llbracket \tau_r \rrbracket \\
505 &\quad body = \text{let } \langle \vec{x} \rangle = x_e \text{ in} \\
506 &\quad \quad \text{letcast}[x_-, f_m] x = x' \text{ in } \llbracket e \rrbracket \\
507 \\
508 \\
509 \\
510 \\
511 \\
512 \\
513 \\
514 \\
515 \\
516 \\
517 \\
518 \\
519 \\
520
\end{aligned}$$

Fig. 3. The closure conversion itself

takes a proof  $e_ = \text{Eq } e_i e_o$ , a motive  $e_m$ , and returns  $\text{let } x = e_1 \text{ in } e_2$  except that  $e_2$  is typed in a context where the  $e_o$  are replaced by  $e_i$ . Other than this type-twist, it is expected to be compiled to the same code as  $\text{let } x = e_1 \text{ in } e_2$ .

- $\forall l. \tau$  is the type of code that is polymorphic over universe level  $l$ ;  $\Lambda l. e$  is the corresponding introduction form; and  $e[\ell]$  the corresponding elimination form which instantiates the level variable with level  $\ell$ .
- $\exists l. \tau$  is the existential type that quantifies over universe level  $l$ ;  $\langle \ell, e \rangle$  is the corresponding introduction form; and  $\text{open } \langle l, x \rangle = e_1 \text{ in } e_2$  is its elimination form.

We can now show the actual closure conversion itself, denoted  $\llbracket \cdot \rrbracket$ , which is in Figure 3. An important detail to note in that figure is an abuse of notation: while we write  $\llbracket e \rrbracket$ , the conversion algorithm does not take a mere term  $e$  as argument but it really operates on a typing derivation of  $e$  because it needs more type information than is readily provided in  $e$  itself. We use this notational abuse in the hope to make the code more readable. In contrast,  $\llbracket \cdot \rrbracket$  does return a mere term and not a full typing derivation. While we like to think of it as a conversion from intrinsically typed terms to intrinsically typed terms, we prefer to return a mere term so that we can separately state and prove that it does indeed preserve typing.

The first three cases are of no interest. The first interesting case is the one for  $(x : \tau_1) \rightarrow \tau_2$  where we state the type of a closure to be fundamentally a 4-tuple made of (in reverse order) a closed function we denote as *code*, a captured environment *env*, its type  $t$ , and its universe level  $l$ .

The case for  $e_1 e_2$  takes such a closure object and calls it: it uses *open* to get to the universe level  $l$  and the rest of the closure  $c$ . At that point  $c.3$  holds the actual code and  $c.2$  is the captured environment. The proof that  $x_e = env$  is trivially *refl* since at this point both  $x_e$  and *env* have been replaced by  $c.2$ .

The more intricate case is for  $\lambda^\sigma x. e$ : there we actually build the closure object, made of its code  $\lambda \langle x_e, x', x_- \rangle. body$ , its captured environment  $\langle \llbracket \vec{x}[\sigma] \rrbracket \rangle$ , its type  $\langle \llbracket \Gamma' \rrbracket \rangle$ , and its universe level  $\ell$ . By  $\llbracket \vec{x}[\sigma] \rrbracket$  and  $\llbracket \Gamma' \rrbracket$  we mean to apply  $\llbracket \cdot \rrbracket$  and  $\cdot[\sigma]$  pointwise to, respectively,  $\vec{x}$  and  $\Gamma'$ . To build

(level contexts)  $L ::= \bullet \mid S L \mid L \sqcup \ell \mid \ell \sqcup L$   
 (eval contexts)  $E ::= \dots \mid \mathcal{U}_L$

$$\frac{\Gamma \vdash \tau_1 \Rightarrow \mathcal{U}_{\ell_1} \quad \Gamma, x : \tau_1 \vdash \tau_2 \Rightarrow \mathcal{U}_{\ell_2}}{\Gamma \vdash (x : \tau_1) \rightarrow \tau_2 \Rightarrow \mathcal{U}_{(\ell_1 \sqcup \ell_2)}} \text{(PI)}$$

$$0 \sqcup \ell \simeq \ell \quad \ell \sqcup 0 \simeq \ell \quad (S \ell_1) \sqcup (S \ell_2) \simeq S (\ell_1 \sqcup \ell_2)$$

Fig. 4. Differences with source language for the core target language.

$$\Gamma \vdash \langle \bullet \rangle \Rightarrow \mathcal{U}_1 \quad \frac{\Gamma \vdash \langle \Gamma' \rangle \Rightarrow \mathcal{U}_{\ell_1} \quad \Gamma, \Gamma' \vdash \tau \Rightarrow \mathcal{U}_{\ell_2}}{\Gamma \vdash \langle \Gamma', x : \tau \rangle \Rightarrow \mathcal{U}_{\ell_1 \sqcup \ell_2}} \text{(\Sigma)}$$

$$\Gamma \vdash \langle \vec{e} \rangle \Leftarrow \langle \Gamma' \rangle \quad \Gamma' = \vec{x} : \vec{\tau} \quad \Gamma \vdash e \Leftarrow \tau[\vec{e}/\vec{x}] \text{(TUP)}$$

$$\frac{\Gamma \vdash e \Rightarrow \langle \Gamma' \rangle \quad \Gamma' = \vec{x} : \vec{\tau}}{\Gamma \vdash e.i \Rightarrow \tau_i[e.1/x_1, \dots, e.(i-1)/x_{i-1}]} \text{(PROJ)}$$

$$\langle \vec{e} \rangle.i \simeq e_i$$

Fig. 5. Typing rules for dependent tuples.

those, we first need to extract the type information of  $e$  (provided by the typing derivation we receive as argument), then we compute the *set* of free variables  $\vec{x}_f$  and from that we compute  $\Gamma'$  to get the *list* of free variables, placed in the same order they appear in the context so as to preserve any type dependencies between them. We then compute the type of  $\langle \Gamma' \rangle$  to find the universe level  $\ell$  in which it lives.  $f_m$  is the motive of the letcast we use later, which describes how to redirect variable references between  $x_e$  and  $env$ . Finally the *body* just unpacks the environment into the corresponding variables and evaluates  $\llbracket e \rrbracket$ , while carefully adjusting the types of  $x$  and of the result to redirect references from the types to variables captured by the function.

### 4.3 Target language

We have already shown the syntax of the target language, but we present here its actual definition in the form of its typing and conversion rules.

Since our target language is a superset of our source language, the core elements are the same and basically share the same rules, except for changes to the universe levels. Figure 4 shows the parts of the rules that changed, fundamentally due to the fact that the  $\max(\ell_1, \ell_2)$  computation that used to be performed at the metalevel is now internalized as  $\ell_1 \sqcup \ell_2$ , which in turn requires new conversion rules to define the semantics of this operation.

Figure 5 shows the rules governing the dependent tuples, which are completely standard. The only significant complexity is in the rules PROJ and TUP which need to take into account the possible

$$\begin{array}{c}
589 \\
590 \\
591 \\
592 \\
593 \\
594 \\
595 \\
596 \\
597 \\
598 \\
599 \\
600 \\
601 \\
602
\end{array}$$

$$\frac{\Gamma \vdash e_1 \Rightarrow \tau \quad \Gamma \vdash e_2 \Leftarrow \tau \quad \Gamma \vdash \tau \Rightarrow s}{\Gamma \vdash \text{Eq } e_1 e_2 \Rightarrow s} \text{ (EQ)} \qquad \frac{e_1 \simeq e_2}{\Gamma \vdash \text{refl} \Leftarrow \text{Eq } e_1 e_2} \text{ (REFL)}$$

$$\frac{e_m e_1 = (x:\tau_{a1}) \rightarrow \tau_{r1} \quad e_m e_2 = (x:\tau_{a2}) \rightarrow \tau_{r2} \quad \Gamma \vdash e_m \Leftarrow (x:\tau) \rightarrow s \quad \Gamma \vdash e_a : \tau_{a2} \quad \Gamma, x:\tau_{a1} \vdash e_r : \tau_{r1}}{\Gamma \vdash \text{letcast}[e_=: e_m] x = e_a \text{ in } e_r \Rightarrow \tau_{r2}[e_a/x]}$$

$$\text{letcast}[\text{refl}, e_m] x = e_a \text{ in } e_r \simeq e_r[e_a/x]$$

Fig. 6. Typing rules for the identity type.

$$\begin{array}{c}
603 \\
604 \\
605 \\
606 \\
607 \\
608 \\
609 \\
610 \\
611 \\
612 \\
613 \\
614 \\
615 \\
616 \\
617 \\
618
\end{array}$$

$$\frac{\Gamma \vdash \tau \Rightarrow \mathcal{U}_\ell}{\Gamma \vdash \forall l. \tau \Rightarrow \mathcal{U}_{\ell[0/l]}} \text{ (U-}\forall\text{)} \qquad \frac{\Gamma \vdash e \Leftarrow \tau}{\Gamma \vdash \Lambda l. e \Leftarrow \forall l. \tau} \text{ (U-}\Lambda\text{)} \qquad \frac{\Gamma \vdash e \Rightarrow \tau}{\Gamma \vdash e[\ell] \Rightarrow \tau[\ell/l]} \text{ (U-APP)}$$

$$\frac{\Gamma \vdash \tau \Rightarrow \mathcal{U}_\ell}{\Gamma \vdash \exists l. \tau \Rightarrow \mathcal{U}_{\ell[0/l]}} \text{ (U-}\exists\text{)} \qquad \frac{\Gamma \vdash e \Rightarrow \tau}{\Gamma \vdash \langle \ell, e \rangle \Rightarrow \exists l. \tau} \text{ (U-PACK)}$$

$$\frac{\Gamma \vdash e_1 \Rightarrow \exists l. \tau_1 \quad l' \text{ is fresh} \quad \Gamma, x:\tau_1[l'/l] \vdash e_2 \Leftarrow \tau}{\Gamma \vdash \text{open } \langle l, x \rangle = e_1 \text{ in } e_2 \Leftarrow \tau} \text{ (U-OPEN)}$$

$$(\Lambda l. e)[\ell] \simeq e[\ell/l] \qquad \text{open } \langle l, x \rangle = \langle \ell, e_1 \rangle \text{ in } e_2 \simeq e_2[\ell/l][e_1/x]$$

Fig. 7. Typing rules for universe polymorphism.

dependencies between the fields, which requires applying substitutions to replace references to previous fields with those fields' values when returning the type of fields.

Figure 6 shows the rules governing the equality (also called identity) type. The EQ and REFL rules are standard, and while the letcast primitive is non-standard due to the need to accommodate the constraint that we want it to be able to write our code such that all  $\lambda$ -expressions are closed, it is more verbose than complex: its rule simply encodes what a standard elimination rule like  $J$  would provide when combined with the convoy pattern.

Figure 7 shows the rules that govern universe polymorphism and the associated existential quantification of universe level. Note that our closure conversion algorithm uses only the  $\exists$  quantification on universe levels, but we include the  $\forall$  quantification anyway since it is very closely tied and corresponds to a more common feature in proof assistants.

Rules U- $\Lambda$ , U-APP, E-PACK, and U-OPEN for the introduction and elimination forms of  $\forall$  and  $\exists$  as well as the corresponding conversion rules are reasonably standard in the sense that they are almost identical to the corresponding rules of System-F, except for the fact that they quantify over universe levels rather than over types. Indeed, we even chose a presentation where we do not keep track of the set of universe level variables in scope, as is done in some presentations of System-F. This was done to avoid having to carry around another context in all the typing rules, so as to make the rules easier on the eyes, but it is otherwise of no particular significance.

An important detail to note here is that the substitutions  $e[\ell/l]$  are different from the substitutions  $e[\sigma]$ : universe level variables are treated completely normally with the usual substitution rules, which means that they propagate into the body of  $\lambda$  expressions unlike substitutions for term variables  $x$ .

The novelty here is rather in the rules  $U\text{-}\forall$  and  $U\text{-}\exists$  which determine the universe level in which those universe-polymorphic terms live. These are the rules which introduce impredicativity in our language.

*4.3.1 The base universe question.* A careful comparison of the syntax of our rules shows that our treatment of the base universe is not quite standard: while our language syntax uses 1 as the base universe level and our unit type  $\langle \bullet \rangle$  indeed lives in this universe, our universe rules in Figure 4 as well as rules  $U\text{-}\forall$  and  $U\text{-}\exists$  use 0 as the base universe level instead. These are not inconsistencies. Indeed, we use 0 as a kind of basement, or arguably more like a *crawl space* because no type lives in this universe and the source code cannot access it. It is only used internally in  $U\text{-}\forall$  and  $U\text{-}\exists$ .

In a predicative PTS with a tower of universes, the base universe is always empty, at least until one explicitly extends the language by adding elements to it, e.g. by adding a unit type. For the same reason, our  $U\text{-}\forall$  and  $U\text{-}\exists$  rules can never place a type in universe  $\mathcal{U}_0$ : the stratification of our universes means that if  $l$  appears somewhere in  $\tau$ ,  $\tau$  itself necessarily lives in a universe level strictly higher than  $l$ , so  $\ell[0/l]$  cannot be 0. This makes universe  $\mathcal{U}_0$  annoyingly special. For example it means our encoding of closures can only represent closures in universes  $\mathcal{U}_1$  or higher. So we decided to make all type constructors “equal” by keeping  $\mathcal{U}_0$  empty, moving the unit type to  $\mathcal{U}_1$  and preventing the source code from referring to  $\mathcal{U}_0$ . The only use for universe level 0 in our system is to make it possible for  $U\text{-}\forall$  and  $U\text{-}\exists$  to place their types in the “real” bottom universe  $\mathcal{U}_1$ .

## 4.4 Properties

Our closure conversion algorithm is designed like a syntactic model, following the approach advocated in [?].

Note: since our conversion takes an actual typing derivation as input, we cannot really manipulate “naked” terms, and thus all the properties and proofs below require that we first redefine our conversion relation  $e_1 \simeq e_2$  to be a typed conversion relation  $\Gamma \vdash e_1 \simeq e_2 : \tau$ . This is a boring and mechanical process which result in significantly more verbose rules, so we preferred to present the untyped rules which focus on the important elements. Sadly, this also forces the three lemmas below to be mutually recursive.

LEMMA 4.1 (SUBSTITUTION LEMMA).

If  $e_1$  and  $e_2$  are expressions in our source language and  $\Gamma, x : \tau_2 \vdash e_1 \Leftarrow \tau_1$  and  $\Gamma \vdash e_2 \Leftarrow \tau_2$  then  $\llbracket e_1[e_2/x] \rrbracket \simeq \llbracket e_1 \rrbracket \llbracket \llbracket e_2 \rrbracket / x \rrbracket$ .

PROOF. By induction on  $e_1$ . □

LEMMA 4.2 (COMPUTATIONAL SOUNDNESS).

If  $e_1$  and  $e_2$  are two expressions in our source language, and  $e_1 \simeq e_2$ , and  $\Gamma \vdash e_1 \Leftarrow \tau$  and  $\Gamma \vdash e_2 \Leftarrow \tau$ , then  $\llbracket e_1 \rrbracket \simeq \llbracket e_2 \rrbracket$ .

PROOF. By induction on the derivation of  $e_1 \simeq e_2$ . The only non-trivial case is the  $\beta$ -reduction rule, where it is easy to see that the various introduction forms used by the encoding of the  $\lambda$ -expression are all canceled by the corresponding elimination forms of the encoding of the application. □

LEMMA 4.3 (TYPING SOUNDNESS).

If  $e$  is an expression in our source language and  $\Gamma \vdash e \Leftarrow \tau$ , then  $\llbracket \Gamma \rrbracket \vdash \llbracket e \rrbracket \Leftarrow \llbracket \tau \rrbracket$ .

687 PROOF. This first requires proving that the premises imply that  $\Gamma \vdash \tau \Rightarrow s$ , without which one  
 688 cannot talk about  $\llbracket \tau \rrbracket$ . This is a trivial side lemma proved by induction on the typing derivation.  
 689 The rest of the proof is by induction on the typing derivation. It is rather tedious and follows the  
 690 same general structure as the usual proof of type preservation of closure conversion, such as found  
 691 in [?].  $\square$

692 LEMMA 4.4 (CONSERVATION OF FREEDOM).

693 *If  $e$  is an expression in our source language and  $\Gamma \vdash e \Leftarrow \tau$ , then  $\text{fv}(e : \tau) = \text{fv}(\llbracket e \rrbracket : \llbracket \tau \rrbracket)$ .*

695 PROOF. By induction on the typing derivation.  $\square$

696 LEMMA 4.5 (CLOSEDNESS).

697 *If  $e$  is an expression in our source language and  $\Gamma \vdash e \Leftarrow \tau$ , then for all  $\lambda$ -expressions  $\lambda^\sigma x.e$  in  $\llbracket e \rrbracket$   
 698 we have that  $\text{fv}(\lambda^\sigma x.e) = \emptyset$ .*

699 PROOF. By induction on the typing derivation. We can see that all the  $\lambda^\sigma x.e$  that can occur in  
 700 the converted code come from the conversion of a  $\lambda^\sigma x.e$  in the input. Furthermore, we can see that  
 701 for such an input we generate two  $\lambda$ -expressions, both of which are closed by construction, thanks  
 702 to the conservation of freedom lemma.  $\square$

## 703 5 CLOSURE CONVERTING A BIGGER LANGUAGE

704 Our source language was purposefully very limited, so that we could focus on the important details,  
 705 but of course we want to be able to scale this to a more realistic language. Luckily, this source  
 706 language also hit the most problematic spots of closure conversion, so it is fairly straightforward to  
 707 extend our result to a more general source language.

708 To get started, we can extend our source language with (dependent) tuples, using the same syntax  
 709 and rules as we used in our target language. Extending the conversion function to handle these  
 710 constructs is simply:

$$\begin{aligned} 711 \llbracket \langle \Gamma \rangle \rrbracket &= \langle \llbracket \Gamma \rrbracket \rangle \\ 712 \llbracket \langle \vec{e} \rangle \rrbracket &= \langle \llbracket \vec{e} \rrbracket \rangle \\ 713 \llbracket \langle e.i \rangle \rrbracket &= \llbracket e \rrbracket . i \end{aligned}$$

714 Next step, we can extend our source language with the same  $\exists$  and  $\forall$  quantification over universe  
 715 levels as we have in our target language. Universe levels are second class citizens which can be  
 716 erased, just like types in System-F, so our closures do not need to close over universe level variables,  
 717 which means we can use the same simple approach as was used in [Morrisett et al. 1998]:

$$\begin{aligned} 718 \llbracket \forall l. \tau \rrbracket &= \forall l. \llbracket \tau \rrbracket \\ 719 \llbracket \Lambda l. e \rrbracket &= \Lambda l. \llbracket e \rrbracket \\ 720 \llbracket e[\ell] \rrbracket &= \llbracket e \rrbracket [\ell] \\ 721 \llbracket \exists l. \tau \rrbracket &= \exists l. \llbracket \tau \rrbracket \\ 722 \llbracket \langle \ell, e \rangle \rrbracket &= \langle \ell, \llbracket e \rrbracket \rangle \\ 723 \llbracket \text{open } \langle l, x \rangle = e_1 \text{ in } e_2 \rrbracket &= \text{open } \langle l, x \rangle = \llbracket e_1 \rrbracket \text{ in } \llbracket e_2 \rrbracket \end{aligned}$$

724 Not shown here: in order for this to work correctly, one has to be careful to define  $\text{fv}(e)$  such that  
 725 it only considers term variables  $x$  and ignores universe level variables  $l$ .

726 Finally, we can try to extend our source language to be the same as our target language by adding  
 727 the remaining equality type constructs. But this hits a minor hurdle: when trying to convert letcast  
 728 expressions, we find that we cannot do that because it only accepts a  $\lambda$ -expression for its  $e_m$  motive  
 729 and we can neither guarantee that this  $\lambda$ -expression is always closed nor convert it into something  
 730 closed when it is not.

$$\begin{array}{c}
\Gamma \vdash e_1 \Rightarrow \tau \quad \Gamma \vdash e_2 \Rightarrow \tau' \quad \Gamma \vdash e_c \Rightarrow \tau_c \\
\Gamma \vdash e_m \Leftarrow (x : \langle \tau_c, \tau' \rangle) \rightarrow s \quad e_m \langle e_c, e_1 \rangle = (x : \tau_{a1}) \rightarrow \tau_{r1} \quad e_m \langle e_c, e_2 \rangle = (x : \tau_{a2}) \rightarrow \tau_{r2} \\
\Gamma \vdash e_a : \tau_{a2} \quad \Gamma, x : \tau_{a1} \vdash e_r : \tau_{r1} \\
\hline
\Gamma \vdash \text{letcast}[e_1, e_c, e_m] x = e_a \text{ in } e_r \Rightarrow \tau_{r2}[e_a/x] \\
\text{letcast}[\text{refl}, e_c, e_m] x = e_a \text{ in } e_r \approx e_r[e_a/x]
\end{array}$$

Fig. 8. Refinement of letcast with  $e_c$ 

To solve this problem, we take a page from traditional C programming conventions where similar problems are solved by making higher-order functions take an additional `void*` argument that they just pass on to their first-class function argument. Concretely we add an  $e_c$  argument to `letcast` which is simply passed on to the  $e_m$  function.

$$\text{letcast}[e_1, e_c, e_m] x = e_1 \text{ in } e_2 \sim (\text{cast } e_1 = (\lambda y. e_m \langle e_c, y \rangle) (\lambda x. e_2)) e_1$$

Figure 8 shows how to adjust the typing and equivalence rules for the new construct. It is easy to adjust our closure conversion to use this new `letcast`: since do not need  $e_c$  when converting a  $\lambda$ -expression, we can just pass a dummy unit type argument for it.

With this wrinkle adjusted, we can now define the closure conversion of the equality type constructs as well:

$$\begin{array}{l}
\llbracket \text{Eq } e_1 \ e_2 \rrbracket = \text{Eq } \llbracket e_1 \rrbracket \ \llbracket e_2 \rrbracket \\
\llbracket \text{refl} \rrbracket = \text{refl} \\
\llbracket \text{letcast}[e_1, e_c, e_m] x = e_1 \text{ in } e_2 \rrbracket = \text{letcast}[\llbracket e_1 \rrbracket, e'_c, e'_m] x = \llbracket e_1 \rrbracket \text{ in } \llbracket e_2 \rrbracket \\
\text{where } \vec{x}_f = \text{fv}(e_m) \\
e'_c = \langle \vec{x}_f, \llbracket e_c \rrbracket \rangle \\
e'_m = \lambda \langle \langle \vec{x}_f, x_c \rangle, x_m \rangle. \llbracket e_m \langle x_c, x_m \rangle \rrbracket
\end{array}$$

With this extension, our closure conversion algorithm accepts the same input language as its output language.

## 6 DISCUSSION

Our aim was to cover a fairly realistic source language and to use a target language that is as “generic” as we can. Here we discuss several design decisions as well as the limits of our current work.

### 6.1 Source language features

While our source language does not cover all of the features of a real language like Idris or Agda, we do cover a significant part. The main missing functionality would be things like inductive types, coinductive types, a Prop universe, erasable arguments, and linearity.

Adding support for inductive types should not be problematic: we already support dependent tuples and equality types, so fundamentally all that is missing is sum types and recursive types. We do not foresee any particular difficulty here, as long as the usual syntactic checks termination checks are sufficiently refined not to be confused by the layers of tuples and `letcast` added by the closure conversion. The same should hold for coinductive types.

There will probably need to be a need for some adjustments to some constructs, like we had to do with `letcast`. E.g. Coq’s elimination principle for inductive types also takes a motive to describe the dependent output type and that annotation has to be a function rather than a closure, just like the  $e_m$  of `letcast`. This could be solved in the same way we did for `letcast`.



785 In contrast, it is very much unclear how an impredicative Prop universe would interact with the  
786 rest of this language. It is also unclear how the impredicativity already provided by our language  
787 compares to that of Prop.

788 Adding support for erasable arguments and linearity seems to fall in-between. It is not imme-  
789 diately obvious, but it might be feasible: adding those features to the language itself should not  
790 pose any specific difficulty; the main difficulty would be performing closure conversion without  
791 changing the erasability/linearity of variables.

## 793 6.2 Efficiency

794 Performing closure conversion correctly is a good first step, but generating efficient code is also  
795 important. The current closure conversion is not fully satisfactory in this regard and solving some  
796 of those issues may not be straightforward.

797 The main issue with our conversion is that our closures are represented as triples of the code,  
798 the environment, and its type (the universe level can be assumed to be erased by later phases).  
799 We may be able erase also the type of the environment, leaving us with a pair of the code and  
800 the environment. While this is the “official” definition of a closure, in practice closures are more  
801 often implemented by merging the inner tuple representing the environment with the outer pair,  
802 resulting in a tuple that contains the code in one field and the free variables in the other fields.  
803 Worse: it also the code receives as argument not just the environment, but the whole closure (since  
804 they do not exist separately any more), which requires makes the code’s type recursive. It is unclear  
805 how to allow this kind of recursion without breaking the logic’s soundness.

806 Another related aspect is that it is common to place the code in the first field and the captured  
807 variables afterwards, whereas our encoding requires fundamentally the environment to come first  
808 because the type of the code must refer to the environment. Dependent type system always assume  
809 a kind of left-to-right ordering of dependencies, but sometimes practical concerns may require  
810 fields to be layed out differently. At the lower level these ordering issues should not matter, so it  
811 would be good to find a way to encode dependencies separately from the ordering, but how to  
812 allow that without breaking the logic is again unclear.

## 814 6.3 Function equality

815 One of the most unusual aspect of our language is the way it defines function equality. This clearly  
816 weakens the definitional equality of our logic, and to make matters worse it makes our language  
817 significantly more complex.

818 An naive counter argument might point out that the definitional equality between functions in  
819 languages like Agda and Coq are already too weak anyway, so it is very commonplace to supplement  
820 with the functional extensionality axiom. Sadly, this breaks our closure conversion: this axiom  
821 applies to  $\lambda$ -expressions and not to the corresponding closure objects after conversion, for the very  
822 same reason that pushed us to weaken our definitional equality.

823 Clearly, this is the most problematic part of our approach. We hope to revise it in a future work by  
824 strengthening our target language with the use of quotient types so that it can faithfully represent  
825 the full definitional equality of source  $\lambda$ -expression after they’re encoded as closures. This should  
826 also be amenable to an extension that handles the functional extensionality axiom.

827 There might still be a potential other use for our functions’ weak definitional equality: this notion  
828 of equality seems to be compatible with a notion of runtime equality testing, although it would  
829 impose a few restrictions on the kinds of optimizations that the compiler can safely apply (for  
830 example, the  $\eta$  rule is not valid in general). More specifically it could be used to provide a precise  
831 semantics for runtime equality tests between functions, thus solving in a different way the problem  
832

834 that arguably lead the SML designers to introduce the unsatisfactory notion of eqtype [?], and  
 835 arguably motivated the invention of type classes.

836

#### 837 6.4 Universe polymorphism

838 While our minimal source language is known to be sound, the logical consistency of our target  
 839 language is a big question mark because of its reliance on a novel notion of impredicativity  
 840 introduced by more aggressive rules for universe polymorphism.

841 In the specific way the  $\exists$  quantification is used here, the rules proposed seem eminently rea-  
 842 sonable: they place the closure objects right in the exact same universe level that they original  
 843  $\lambda$ -expression occupied in the source code. Sadly, that does not guarantee that they are sound in  
 844 general.

845 Our intuition as for why they *may* not be completely crazy is that the second-class status of  
 846 universe levels makes universe polymorphic definitions enjoy a strong form of parametricity. Agda’s  
 847 position says that  $\forall l. \tau$  can be modeled as a set theoretic function which for every level  $l$  returns  
 848 the corresponding  $\tau$ , so this function is clearly very large since it includes all the possible  $\tau$  one can  
 849 get for all the possible levels with which we can instantiate it. For this reason, if  $\Gamma \vdash \tau : \mathcal{U}_l$  Agda  
 850 places  $\forall l. \tau$  in the universe  $\text{sup}_l \ell$  which they represent as  $\omega$ . Our typing rules basically take the  
 851 opposite position, considering that the type  $\forall l. \tau$  is arguably smaller than any given instantiation of  
 852  $\tau$  since it only holds those functions which can be used at any universe level (just like the type  
 853  $\forall t. t \rightarrow t$  is so small that it only contains a single element), so it places it in the universe  $\text{inf}_l \ell$ , hence  
 854  $\ell[0/l]$ .

855 We do not know that this intuition holds water, sadly. But we can’t see how to type closure  
 856 converted code of a source language with a tower of universes without using something similar to  
 857 the rules we propose: in a sense, the rules we use arise naturally in our encoding. Of course, there  
 858 is always the emergency escape hatch of resorting to custom-made constructs like the one used  
 859 in [Bowman and Ahmed 2018].

860

#### 861 ACKNOWLEDGMENTS

862 This work was supported by the Natural Sciences and Engineering Research Council of Canada grant  
 863 N° 298311/2012 and RGPIN-2018-06225 as well as by the Fonds de Recherche du Québec - Nature  
 864 et Technologie grant N° 253521. Any opinions, findings, and conclusions or recommendations  
 865 expressed in this material are those of the author and do not necessarily reflect the views of the  
 866 NSERC or the FRQNT.

867

#### 868 REFERENCES

- 869 Martin Abadi, Luca Cardelli, Pierre-Louis Curien, and Jean-Jacques Lévy. 1990. Explicit Substitutions. In *Symposium on*  
 870 *Principles of Programming Languages*. ACM Press, 31–46.
- 871 Bowman. 2018. *Compiling with Dependent Types*. Ph. D. Dissertation. Northeastern University. <https://williamjbowman.com/resources/wjb-dissertation.pdf>
- 872 William J. Bowman and Amal Ahmed. 2018. Typed closure conversion for the calculus of constructions. In *Programming*  
 873 *Languages Design and Implementation*. 797–811. <https://doi.org/10.1145/3192366.3192372>
- 874 Adam Chlipala. 2013. *Certified Programming with Dependent Types*. MIT Press.
- 875 Pierre-Louis Curien, Thérèse Hardin, and Jean-Jacques Lévy. 1996. Confluence properties of Weak and Strong Calculi of  
 876 Explicit Substitutions. *Journal of the ACM* 43, 2 (1996), 362–397.
- 877 Antonius Hurkens. 1995. A simplification of Girard’s paradox. In *International conference on Typed Lambda Calculi and*  
 878 *Applications*. 266–278. <https://doi.org/10.1007/BFb0014058>
- 879 Yasuhiko Minamide, Greg Morrisett, and Robert Harper. 1996. Typed Closure Conversion. In *Symposium on Principles of*  
 880 *Programming Languages*. ACM Press, 271–283. <https://doi.org/10.1145/237721.237791>
- 881 Stefan Monnier and Nathaniel Bos. 2019. Is Impredicativity Implicitly Implicit?. In *Types for Proofs and Programs (Leibniz*  
 882 *International Proceedings in Informatics (LIPIcs))*. 9:1–9:19. <https://doi.org/10.4230/LIPIcs.TYPES.2019.9>

883 Stefan Monnier and David Haguenauer. 2010. Singleton types here, Singleton types there, Singleton types everywhere. In  
884 *Programming Languages meets Program Verification*. ACM Press, 1–8. <https://doi.org/10.1145/1707790.1707792>

885 Greg Morrisett, David Walker, Karl Crary, and Neal Glew. 1998. From System F to Typed Assembly Language. In *Symposium*  
886 *on Principles of Programming Languages*. 85–97. <https://doi.org/10.1145/319301.319345>

887 Zhong Shao, Bratin Saha, Valery Trifonov, and Nikolaos Papaspyrou. 2002. A Type System for Certified Binaries. In  
888 *Symposium on Principles of Programming Languages*. 217–232. <https://doi.org/10.1145/503272.503293>

889 Vladimir Voevodsky. 2012. A universe polymorphic type system. [https://www.math.ias.edu/~dgrayson/Voevodsky-old-](https://www.math.ias.edu/~dgrayson/Voevodsky-old-files/files/files-annotated/Dropbox/Unfinished_papers/Type_systems/UPTS_current/Universe_polymorphic_type_sytem.pdf)  
890 [files/files/files-annotated/Dropbox/Unfinished\\_papers/Type\\_systems/UPTS\\_current/Universe\\_polymorphic\\_type\\_](https://www.math.ias.edu/~dgrayson/Voevodsky-old-files/files/files-annotated/Dropbox/Unfinished_papers/Type_systems/UPTS_current/Universe_polymorphic_type_sytem.pdf)  
891 [sytem.pdf](https://www.math.ias.edu/~dgrayson/Voevodsky-old-files/files/files-annotated/Dropbox/Unfinished_papers/Type_systems/UPTS_current/Universe_polymorphic_type_sytem.pdf)

892

893

894

895

896

897

898

899

900

901

902

903

904

905

906

907

908

909

910

911

912

913

914

915

916

917

918

919

920

921

922

923

924

925

926

927

928

929

930

931