

Dependent Type Systems as Macros

STEPHEN CHANG, Northeastern University, USA

MICHAEL BALLANTYNE, Northeastern University, USA

MARCELA POFFALD, Unaffiliated, USA

WILLIAM J. BOWMAN, Northeastern University, USA

Increasingly, programmers want the power of dependent types, yet significant expertise is still required to write realistic dependently-typed programs. In response, domain-specific languages (DSLs) attempting to tame dependent types have proliferated, adding notation and tools tailored to a problem domain. This only shifts the problem, however, since implementing such languages requires at least as much expertise as using them.

We show how to lower the burden for implementing dependently-typed languages and DSLs, using a classic approach to DSL implementation not typically associated with typed languages: macros. By leveraging a macro system, programmers may reuse all of a host language's infrastructure when implementing a new, dependently-typed language or DSL, reducing the overall effort. We also extend the `TURNSTILE` language, a meta-DSL for implementing typed DSLs using syntax resembling "pen and paper" models, with support for dependent types. Using macros simplifies not only the initial language implementation, but also the addition of extensions like notation or tactic languages—all but required features for dependently-typed languages.

We evaluate our approach by building three languages in different parts of the design space: first, we present a video-editing DSL with a Dependent ML-like type system, demonstrating that our approach accommodates "lightweight" dependent types; second, we gradually extend MLTT to the Calculus of Inductive Constructions, demonstrating that our approach is modular, and scales to "heavyweight" dependent type systems; finally, we describe `CUR`, a prototype proof assistant with a design similar to `Coq`, which supports new notation and an extensible tactic language, demonstrating that our approach scales to realistic dependently-typed languages.

1 INTRODUCTION

Programmers are increasingly wanting and using dependent types. For example, Haskell has embraced type-level computation [Weirich et al. 2017], Rust is considering adding Π types [rus 2017], and new dependently typed languages such as `F*` have leveraged domain-specific languages (DSLs) to verify software such as Firefox's TLS [Beurdouche et al. 2017; Zinzindohoué et al. 2017].

Despite this progress, implementing and using dependent types remains complicated, and thus not all programmers are ready for them. At one end of the spectrum, language designers debate about the "right" amount of dependent types. For example, determining the ideal "power-to-weight" ratio has slowed adoption in Haskell [Yorgey et al. 2012] and has led to repeated rewrites of Rust's dependent type RFCs [rus 2016]. At the other end, proof assistants that ignore "weight" in favor of "power" must layer on companion DSLs (e.g., a tactic language) to help programmers use the language [Brady and Hammond 2006; Christiansen 2014; Devriese and Piessens 2013; Ebner et al. 2017; Gonthier and Mahboubi 2010; Gonthier et al. 2011; Krebbers et al. 2017; Malecha and Bengtson 2016; Pientka 2008; Stampoulis and Shao 2010; Ziliani et al. 2013].

Ideally, language designers or even users would simply construct a new DSL for each problem domain, choosing how much "power" to wield on a case-by-case basis. Indeed, DSLs have been used effectively to tame dependent types [Barthe et al. 2009; Chlipala 2011; Chlipala et al. 2017; Zinzindohoué et al. 2017], but so far, they have not been simple to build.

Authors' addresses: Stephen Chang, Northeastern University, USA, stchang@ccs.neu.edu; Michael Ballantyne, Northeastern University, USA, mballantyne@ccs.neu.edu; Marcela Poffald, Unaffiliated, USA, mpoffald@gmail.com; William J. Bowman, Northeastern University, USA, wjb@williamjbowman.com.

2018. 2475-1421/2018/1-ART1 \$15.00

<https://doi.org/>

50 We aim to change that, by showing how to use *macros* to create dependently-typed DSLs.
51 Procedural macros, in the style of LISP and its descendants, simplify the construction of DSLs [Fowler
52 and Parsons 2010] by reusing much of the host language infrastructure such as parsing, elaboration,
53 namespace management, and compilation. We show how this approach reduces the complexity of
54 dependent types for both language implementors—because the DSL may reuse the infrastructure of
55 the host language—and users—because the complexity of the type system may be exactly tailored
56 to a specific problem domain. Better yet, any languages created with this approach may leverage
57 the macro system to implement extensions to the core language such as new notation and tools for
58 automatically constructing proofs and programs.

59 The macro-based approach to building DSLs has not historically included typed languages,
60 but recently Chang et al. [2017] introduced the technique of “type systems as macros”, showing
61 how programmers may use macros to create typed DSLs as well. Specifically, they show that
62 with a contemporary macro system as found in Racket [Flatt and PLT 2010]—a LISP and Scheme
63 descendant—programmers may create typed DSLs simply by embedding type rule logic directly into
64 the macro definitions. This macro-based approach improves on the traditional approach of creating
65 typed DSLs—where domain-specific types are encoded into an existing host type system—because
66 it does not constrain DSL creators to the limitations of any particular type system. Instead, DSL
67 implementers have the flexibility to create the right type system for their domain. Finally, macros
68 are naturally expressed as local, modular transformations, and implementing type rules with macros
69 results in a type checker that naturally matches the modularity of its mathematical specification.
70 This is demonstrated in Chang et al. [2017]’s TURNSTILE, a meta-DSL that allows implementing
71 typed DSLs using a judgement-like syntax resembling what programmers would find in a textbook.

72 We extend the “type systems as macros” approach—and the TURNSTILE language—to support
73 creating dependently-typed languages. This is a major technical challenge, as dependent types
74 break many of the assumptions implicit in macro systems, and in the previous design of TURNSTILE.
75 For example, run-time and expansion time are distinct phases for a macro system, but there can be
76 no such distinction for a dependently-typed language, which may evaluate expressions while type
77 checking. There are new design challenges as well, *e.g.*, a framework for building dependently-
78 typed languages should support common type notation such as *telescopes*, *i.e.*, nested binding
79 environments [de Bruijn 1991; McBride 2000]. Specifically, we make the following contributions.

- 81 • We extend TURNSTILE with a new API for defining types that is syntactically concise, yet
82 robust enough to implement a range of constructs from base types, to binding forms like Π
83 types, to indexed inductive type families.
- 84 • We show how to leverage macros and macro expansion to perform the work of type-level re-
85 duction in an extensible manner, and also add TURNSTILE constructs that allow implementing
86 these reduction rules with familiar on-paper syntax.
- 87 • A key source of complexity in implementing dependent types is handling dependent binding
88 structure, *e.g.*, manipulating *telescopes*. For example, checking such binding types requires
89 interleaving checking with adding new environment bindings, and instantiating them requires
90 a folding substitution operation. We extend TURNSTILE’s pattern language to support these
91 operations, allowing us to express features with complex binding structure (such as indexed
92 inductive type families) using a concise, intuitive notation.
- 93 • We evaluate our approach by constructing three example languages.
94 (1) We present a video-editing DSL with a Dependent ML-like type system that statically
95 enforces guarantees about the lengths of videos, tracks, and playlists, demonstrating that
96 our approach allows tailoring dependent types into a more “lightweight” flavor.
97

- (2) We gradually build up a core calculus, culminating in the Calculus of Inductive Constructions (CIC) [Pfenning and Paulin-Mohring 1989], demonstrating that our macros-based approach is modular, extensible, and supports “heavyweight” dependent type systems.
- (3) To demonstrate that our approach supports creating realistic languages, we present CUR, a prototype proof assistant whose core—the impredicative CIC—resembles that of Coq, but requires only a few dozen lines of code in our extended TURNSTILE. With macros, we easily extend core Cur with features such as syntactic sugar and a tactic language. Using the latter, we worked through several chapters of “Logical Foundations” [Pierce et al. 2018], demonstrating that the tactic system is sophisticated enough to support Coq-style proofs.

2 CREATING MACRO-BASED DSLS WITH RACKET: PRIMER

This section introduces building languages—typed and untyped—with RACKET’s macro system.¹

```

112 #lang racket bool-lang
113 (provide true false and or not #%app ⊃ (rename truth-table λ))
114
115 (define-m (truth-table (xid ...) [argbool ... = resbool] ...)
116   #:with (dnf-clause-fn ...) (λ (x ...) (and res ((bool->lit arg) x) ...))
117   (λ (x ...) (or (dnf-clause-fn x ...) ...)))
118
119 (define ⊃ (truth-table (x y) [false false = true]
120                       [true false = false]
121                       [false true = true]
122                       [true true = true]))
123
124 (define-m bool->lit [(_ true) (λ (x) x)] [(_ false) not])

```

Fig. 1. A basic (untyped) Boolean-logic DSL created with Racket.

2.1 An Untyped DSL

A Racket language is defined by the exports of a module. Figure 1 presents the `bool-lang` module, an example language of Boolean logic, which we use to introduce notation used in the paper,² and DSL creation with macros. Key to defining languages as macros are the abilities to:

- (1) *reuse* host language (Racket) features for their own language; e.g., `bool-lang` reuses `true`, `false`, `and`, `or`, `not`, as well as the function application form `#%app`;
- (2) *add* functions and forms; e.g., `bool-lang` defines and exports the implication function `⊃`;
- (3) *interpose* on primitive forms, such as functions and application, using syntactic hooks such as `#%app` and `λ`, e.g., `bool-lang` redefines `λ` by the `truth-table` macro;
- (4) *exclude* features from the host language; e.g., `bool-lang` does not include first-class functions, numbers, or lists, but only the explicitly exported features.

To program with `bool-lang`, programmers use the `#lang` directive:

¹Appendix B lists the macro system features we use in more detail, and discusses other languages with the same features.

²To more clearly communicate concepts, we sacrifice code precision by stylizing code with abbreviations or non-syntactic elements like color and subscripts. For example, `define-m` is shorthand for a Racket macro defined via `define-syntax` and the `syntax-parse` pattern matching construct [Culpepper and Felleisen 2010]. Examples may not run as presented, but full implementations of all examples are available at <https://www.github.com/stchang/macrotypes>, <https://www.github.com/wilbowma/cur>. We summarize our style conventions in Appendix A.

```

148 #lang bool-lang bool-prog
149 (⊃ true false) ; result: false
150 ((λ (x) [true = false] [false = true]) (⊃ false true)) ; result: false
151 (+ 1 2) ; ERR: unsupported

```

152 In `BOOL-LANG`'s implementation, `truth-table` is a macro, *i.e.*, it is defined³ with `define-m`, that
 153 converts a table of boolean values into a function implementing an equivalent formula in disjunctive
 154 normal form (DNF). Macros consume and produce a *syntax object*, an AST data structure that
 155 combines a tree of symbols with context information like source location and binding structure. A
 156 macro typically pattern-matches on its input, using a *syntax pattern* (green in this paper) whose
 157 shape dictates how the macro is invoked. (Note that the name of the macro being invoked is included
 158 as part of the input for the macro, so all initial syntax patterns begin with a pattern representing
 159 the macro's own name. For example, `truth-table` is part of the syntax pattern in the definition
 160 of `truth-table`.) Most identifiers in a syntax pattern are bound as *pattern variables*, which are
 161 associated with corresponding pieces of syntax supplied by the programmer when they invoke the
 162 macro. The ellipses pattern `...` means “zero or more of the preceding pattern”.

163 `BOOL-LANG` invokes `truth-table` to define `⊃`, where the pattern `(x ...)` in `truth-table`'s input
 164 pattern matches syntax object `(x y)`, representing the arguments expected by `⊃`. A superscript *syntax*
 165 *class* may adorn a pattern variable, which refines what syntax matches that pattern variable. For
 166 example `truth-table`'s input parameters are tagged with `id`, so they only match identifiers. Further,
 167 its body consists of rows of literal boolean values representing the inputs and output, separated
 168 by `=` (a bolded pattern symbol, e.g., `=`, denotes an exact value that must be matched). A `#:with`
 169 keyword introduces additional pattern variables by matching on the syntax object computed by the
 170 second position after the keyword. For example, `truth-table` uses `#:with` to define `dnf-clause-fn`
 171 pattern variables, representing the “and” clauses in its “or of ands” DNF output.

172 Pattern variables are used in *syntax templates* (blue in the this paper). A template replaces
 173 references to pattern variables with their corresponding syntax object values. Ellipses that followed
 174 a variable in a syntax pattern must also accompany references to that variable in the syntax template.
 175 Macros frequently use syntax templates to construct their outputs, e.g., `truth-table`'s output is a
 176 template that references the `dnf-clause-fn` pattern variables.

177 Finally, the meaning of any non-pattern-variable identifiers in a syntax template is taken from
 178 the context of the macro definition. For example, the syntax template constructing `dnf-clause-fn`
 179 references Racket's `λ` and `and`, as well as a local macro `bool->lit`, which converts a boolean value
 180 into DNF formula literal. The `bool->lit` macro uses an alternate macro definition syntax with
 181 multiple clauses, whose patterns are tried in order. Observe that each pattern still includes the
 182 name of the macro in its first position, but our example ignores it using the `_` pattern.
 183

184 2.2 A Typed DSL

185 **Figure 2 (left)** presents `TYPED-LANG`, which adds arithmetic to `BOOL-LANG`, and a type system to
 186 ensure that operations receive the right values. It is created with the “type systems as macros”
 187 technique [Chang et al. 2017] and uses the same four “DSL tools” from [Section 2.1](#). Specifically,
 188 `TYPED-LANG` interposes on `λ` and `#%app` with two new macros, `typed-λ` and `typed-app`, respectively,
 189 replacing the `#%app` and (truth table) `λ` from `BOOL-LANG`. These macros use the computation and
 190 pattern matching performed by `#:with` to implement basic type checking. Since macros embody
 191 local transformations, however, successfully checking types in this manner requires additional
 192 coordination between macros, to communicate type information. Solving this coordination problem
 193

194
 195 ³We underline names being defined.
 196

```

197 #lang racket typed-lang #lang turnstile typed-lang
198 (provide [typed-λ λ] [typed-app #%app] (provide [typed-λ λ] [typed-app #%app]
199       [typed+ +] [typed-and and]))       [typed+ +] [typed-and and])
200 (define-m (typed-app f e) (define-tyrule (typed-app f e) >>
201   #:with [ $\bar{f}$  ( $\rightarrow \tau_{in} \tau_{out}$ )] (synth f)      [⊢ f >>  $\bar{f} \Rightarrow (\rightarrow \tau_{in} \tau_{out})$ ]
202   #:with  $\bar{e}$  (check e  $\tau_{in}$ )                       [⊢ e >>  $\bar{e} \Leftarrow \tau_{in}$ ]
203   (assign (hash-table f  $\bar{e}$ )  $\tau_{out}$ ))              -----
204                                           [⊢ (hash-table f  $\bar{e}$ )  $\Rightarrow \tau_{out}$ ]
205
206 (define-m (typed-λ [ $x^{id} : \tau_{in}$ ] e) (define-tyrule (typed-λ [ $x^{id} : \tau_{in}$ ] e) >>
207   #:with [ $\bar{x} \bar{e} \tau_{out}$ ] (synth e #:ctx [x :  $\tau_{in}$ ]) [[x >>  $\bar{x} : \tau_{in}$ ] ⊢ e >>  $\bar{e} \Rightarrow \tau_{out}$ ]
208   (assign (λ  $\bar{x} \bar{e}$ ) ( $\rightarrow \tau_{in} \tau_{out}$ )))      -----
209                                           [⊢ (λ  $\bar{x} \bar{e}$ )  $\Rightarrow (\rightarrow \tau_{in} \tau_{out})$ ]
210
211 (define-m (typed-and e1 e2) (define-tyrule (typed-and e1 e2) >>
212   #:with  $\bar{e}_1$  (check e1 Bool) [⊢ e1 >>  $\bar{e}_1 \Leftarrow \text{Bool}$ ]
213   #:with  $\bar{e}_2$  (check e2 Bool) [⊢ e2 >>  $\bar{e}_2 \Leftarrow \text{Bool}$ ]
214   (assign (and  $\bar{e}_1 \bar{e}_2$ ) Bool))              -----
215                                           [⊢ (and  $\bar{e}_1 \bar{e}_2$ )  $\Rightarrow \text{Bool}$ ]
216
217 (define-m typed+ (assign + ( $\rightarrow \text{Int Int Int}$ ))) (define-primop typed+ + : ( $\rightarrow \text{Int Int Int}$ ))

```

Fig. 2. (Part of) a typed extension of BOOL-LANG, (left) using Racket, and (right) using TURNSTILE.

is the essence of “type systems as macros”. Specifically, the synth, check, and assign metafunctions (Figure 3) implement such a communication protocol between type-checking macros.

The typed-app macro first uses synth to compute the type of function term f , which must match the pattern ($\rightarrow \tau_{in} \tau_{out}$). The synth function produces a second result \bar{f} representing an elaborated version of f . Because our type checker is embedded in macro definitions, type checking is interleaved with macro expansion, and synth necessarily expands f . To avoid redundant expansions, synth returns the elaborated \bar{f} so that typed-app may produce an elaborated term that includes \bar{f} .⁴ This turns out to be an effective and concise way to implement type checkers, since type systems often require an elaboration pass anyway, e.g., for type erasure.

The typed-app macro’s second premise uses check to ensure that argument e has type τ_{in} . Similar to synth, check expands its argument and returns the elaborated \bar{e} . Finally, typed-app constructs output term (hash-table $\bar{f} \bar{e}$), which uses the (untyped) host language hash-table, and “assigns” it type τ_{out} . This call to assign is the crucial step that communicates type information between macros, by attaching type information to the syntax objects, which other type checking macros understand.

In typed-λ, synth computes the type of body e in a context where x has type τ_{in} . (The function is passed the type environment via a named keyword argument #:ctx.) Here synth returns the elaborated \bar{e} , as well as the binder \bar{x} for references in \bar{e} . The latter is required to construct the output term in a hygienic macro system, i.e., one that tracks and enforces proper binding structure in all syntax objects. In other words, programmers may not create binding terms using any arbitrary identifier with the same name; instead a proper binder must carry the correct program context information, added via expansion (see Flatt [2016] for more details). Thus only \bar{x} may close over \bar{e} because they were expanded with the same context. It turns out that this knowledge of the program’s binding structure is extremely useful for implementing type systems and many type

⁴We use overlines to denote pattern variables bound to fully elaborated syntax

operations such as substitution and alpha equivalence. Finally, typed- λ uses `assign` to associate the elaborated syntax $(\lambda \bar{x} \bar{e})$ with its type $(\rightarrow \tau_{in} \tau_{out})$.

Figure 2 (left) includes a few other “type rules”: `typed-and` checks that its arguments are `Bool`, and `typed+` is a function with type $(\rightarrow \text{Int Int Int})$. In the latter case, `typed+` is an *identifier macro* that does not require any arguments to invoke it. Thus, `typed-app` handles type checking application of `typed+`. For now, we assume that types are literal pieces of syntax as in Figure 2 (left), e.g., `Bool` and `Int`. Section 3.1 presents a more thorough treatment of defining types.

```

253 (define (assign e  $\tau$ ) (attach e 'type  $\tau$ ))          (define (check e  $\tau$  #:ctx [ctx ()])
254 (define (synth e #:ctx [ctx ()])                    #:with [ $\bar{x}\bar{s}$   $\bar{e}$   $\tau_e$ ] (synth e  $\tau$  #:ctx ctx)
255   #:with [ $\bar{x}\bar{s}$   $\bar{e}$ ] (local-expand (letSTX ctx e))    (if ( $\tau = \tau_e \tau$ ) ( $\bar{x}\bar{s}$   $\bar{e}$ )
256     ( $\bar{x}\bar{s}$   $\bar{e}$  (detach  $\bar{e}$  'type)))                  (err "type mismatch"))))
257

```

Fig. 3. “Type systems as macros” core API.

Figure 3 shows the implementations of `synth`, `check`, and `assign`. They require only a few lower-level operations on syntax, demonstrating that the entire type checker is implemented “as macros”. Specifically, the functions rely on two features: (1) a `local-expand` function that initiates macro expansion on a syntax object, which allows invoking “type checking” macros on a subterm; and (2) a way of associating additional information (types) with syntax objects; we use *syntax properties* which, via `attach` and `detach`, to associate key-value pairs to syntax objects.

Individually, `assign` attaches a type to a term, at key `'type`. The `synth` function consumes an expression `e` and an optional environment `ctx`—which has shape $([x : \tau] \dots)$ —and invokes the macro expander via `local-expand` to type check `e`. To implement the type environment, it wraps `e` with `letSTX`, which allows defining local macros. In other words, `letSTX` defines new typed macros such that untyped variable references in \bar{e} are themselves macro invocations that return the desired type information, effectively using the macro environment to implement the type environment. Finally, `synth` returns a triple consisting of the expanded context variables, the expanded term \bar{e} , and its type. The coloring of `synth`’s output denotes a *quasiquoted* syntax template, i.e., the syntax object is constructed with references to pattern variables $\bar{x}\bar{s}$ and \bar{e} , and a call to the (meta)function `detach`. The `check` function first invokes `synth` on term `e`, checks that the actual and expected type match using type-equality function $\tau =$, and returns the expanded \bar{e} if successful. For this paper, we assume $\tau =$ (not shown) is syntactic equality up to alpha-equivalence; it is straightforward to implement since syntax objects are aware of the program’s binding structure.

2.3 A DSL for Typed DSLs

Chang et al. [2017] observed that Figure 2 (left)’s macros closely correspond to algorithmic specifications. Thus, they created `TURNSTILE`, a meta-DSL that allows writing type-checking macros using a type judgement-like syntax, as seen in Figure 2 (right). Specifically, `TURNSTILE` uses two relations which correspond to “`synth`” and “`check`” bidirectional type checking judgments [Pierce and Turner 1998], interleaved with elaboration. They are implemented with `synth` and `check` from Figure 3, respectively. The judgement $[ctx \vdash e \gg \bar{e} \Rightarrow \tau]$ says that, in environment `ctx`, `e` elaborates (\gg) to \bar{e} and synthesizes (\Rightarrow) type τ —i.e., τ is an output. Observe how the syntax pattern and syntax templates on the left and right of Figure 2 remain the same. The `check` judgment $[ctx \vdash e \gg \bar{e} \Leftarrow \tau]$ specifies that, in environment `ctx`, `e` elaborates to \bar{e} and checks against (\Leftarrow) type τ —i.e., τ is an input. Bindings are added to the type environment by writing them to the left of \vdash , as in typed- λ , but only *new* variables must be written. Since `TURNSTILE` reuses the macro environment as the type environment, existing bindings are automatically propagated by lexical scope. Figure 2 (right) uses `define-tyrule`, which has a few usage variations, “`synth`” (L) and “`check`” (R):

```

295 (define-tyrule input-pattern >> (define-tyrule input-pattern <=< input-type >>
296   premises ...                   premises ...
297   -----                         -----
298   [⊢ output-template ⇒ type])    [⊢ output-template])

```

Figure 2 (right) implements “synth” rules, which fire when a term matches `input-pattern`. If the premises—a series of synth and check judgements—holds, the macro produces the output specified by `output-template`, with `type` attached. Observe that a textbook would typically write these rules with the entire conclusion `[⊢ input-pattern >> output-template ⇒ type]` at the bottom, but TURNSTILE shifts the conclusion’s input (pattern) to the top, as in a typical macro definition. With a “check” rule, `input-type` is also an input and is written next to `input-pattern`, above the premises. Thus, a “check” type rule fires when: a term matches `input-pattern`, the term’s type may be inferred from its context, and that type matches pattern `input-type`. TURNSTILE automatically switches from “check” to “synth” rules when no corresponding “check” rules exists.

3 LIGHTWEIGHT DEPENDENT TYPES, FOR VIDEO

While Chang et al. [2017] implement a variety of languages with TURNSTILE, they can not handle dependent types since they assume an explicit phase distinction, *i.e.*, that terms and types are distinct. We show that maintaining this distinction is no longer possible when implementing dependent types, and how to improve TURNSTILE to cover this deficit. We do this in the context of an example, TYPED VIDEO, a DSL with *indexed types*—“lightweight” dependent types in the style of Dependent ML [Xi 2007]—implemented “as macros”. With indexed types, we can lift some terms (the index language) to the type level to express simple predicates about those terms. While Andersen et al. [2017] do briefly describe a few type rules, our work is the first to explain the underlying details required to implement indexed types as macros, such as the implementation of types and type-level computation. We focus on the new techniques required to express indexed types as macros, using TYPED VIDEO as an example, not on the use or implementation of TYPED VIDEO itself.

TYPED VIDEO is a typed version of Andersen et al. [2017]’s VIDEO language, a DSL for editing movies that has been used to create the video proceedings of several workshops, *e.g.*, OPLSS.⁵ TYPED VIDEO uses indexed types in order to statically rule out errors that arise when creating and combining video streams. A VIDEO program manipulates *producers*—streams of data such as audio, video, or some combination thereof—cutting, splicing, and mixing them together into a final product. Since producers ultimately represent physical data on disk, it’s possible to crash a program (usually during rendering) by accidentally using more data than exists. To prevent this, TYPED VIDEO assigns producer values a `Producer` type, indexed by its length. This is an ideal type system for VIDEO since programmers are already required to provide the length of many expressions.

Below is a function that combines audio, video, and slides to create a conference talk video.

```

332 #lang typed/video
333 (define (mk-conf-talk [n : Int] [aud : (Producer n)] [vid : (Producer n)]
334   [slid : (Producer n)]) #:when (> n 3) -> (Producer (+ n 9))
335   (playlist (img "conf-logo.png" #:len 9)
336     (fade #:len 3)
337     (overlay aud vid sli)))

```

The `mk-conf-talk` function consumes an integer length `n` and audio, video, and slides producers with types `(Producer n)`, meaning they must be at least `n` frames long. The function combines its inputs with `overlay`, and further adds a logo that fades into the main content. The function specifies

⁵<https://lang.video/community.html>

an additional constraint ($> n\ 3$), to ensure that the inputs contain enough data to perform the fade transition. Finally, the output type specifies that the input is extended by 9 frames, to account for the added logo. The function is assigned the type:

```
(→vid [n : Int] [aud : (Producer n)] [vid : (Producer n)] [sli : (Producer n)]
  (Prod (+ n 9)) #:when (> n 3))
```

In this binding variant of a function type, each argument is named and the type of each argument may reference the names preceding it. We demonstrate how to scale the “type systems as macros” approach to support this binding structure, and extend the TURNSTILE implementation with the *telescope* notation used on-paper to manipulate types like this [de Bruijn 1991; McBride 2000].

3.1 Defining Types

The first challenge is how one can define a typing rule for the function type just presented. This subsection describes how to extend the key typing judgments when encoded in macros (introduced in Section 2), and demonstrates our implementation of this approach via a new TURNSTILE API for defining types, including dependent types.

To type check *types themselves*, the obvious approach is to define *types as macros*, not just terms. We can then specify the semantics for type construction in the same judgments we used in Section 2. Figure 4 shows type rules for (single-arity) function types \rightarrow and \rightarrow_{vid} .

<pre>(struct → (in out)) (define-tyrule (→ τ_{in} τ_{out}) >> [⊢ τ_{in} >> $\bar{\tau}_{in} \Leftarrow \text{Type}$] [⊢ τ_{out} >> $\bar{\tau}_{out} \Leftarrow \text{Type}$] ----- [⊢ (#%app → $\bar{\tau}_{in}$ $\bar{\tau}_{out}$) ⇒ Type])</pre>	<pre>(struct →_{vid} (in out)) (define-tyrule (→_{vid} [x : τ_{in}] τ_{out}) >> [⊢ τ_{in} >> $\bar{\tau}_{in} \Leftarrow \text{Type}$] [[x >> $\bar{x} : \bar{\tau}_{in}$] ⊢ τ_{out} >> $\bar{\tau}_{out} \Leftarrow \text{Type}$] ----- [⊢ (#%app →_{vid} $\bar{\tau}_{in}$ (λ (\bar{x}) $\bar{\tau}_{out}$)) ⇒ Type])</pre>
--	---

Fig. 4. Some type rules (macro definitions) for single-arity function types, (left) \rightarrow , (right) \rightarrow_{vid}

In Figure 4 (left), the rule for a standard function type \rightarrow checks that its input and output have type `Type`, a type of types.⁶ But what should be the output of the rule? In other words, what is the “runtime” representation of a type? For typed lambda, we used the underlying host language’s lambda representation, but there is no analogous construct for types. Thus, for types, we have more freedom in what to put in the rule’s output. Any representation, however, has three criteria. The first two, which are simple to understand, are: 1) it should uniquely identify the type and 2) it should store the arguments to the type constructor. In Figure 4 (left), we use a named record \rightarrow , declared with `struct`, to represent the \rightarrow type. Thus, the output of the \rightarrow macro is a *syntax object* of an application of \rightarrow to its arguments.

The third criteria requires thinking about binding. In Figure 4 (right), the rule for \rightarrow_{vid} resembles that of \rightarrow , but differs in that: a), the type’s input has a name x ; b) the output τ_{out} is checked in the context of x because it may reference x ; and c) for the type’s representation in the rule’s conclusion, a *lambda wraps and binds* references to \bar{x} in $\bar{\tau}_{\text{out}}$. The last difference reveals the third criteria for a type’s internal representation: it must comply with hygiene. Syntax objects must have a valid binding structure at all times, or they are rejected during macro expansion. This last criteria is key to getting boilerplate operations—such as substitution, alpha-equivalence, and environment management—for free.

⁶This paper omits discussing the implementation of `Type`, which is not interesting, due to space. Note that the CUR language in Section 5 supports a proper type universe hierarchy, as found in languages like Coq.

3.2 Type Checking Telescopes

The above tells us how to represent dependent types as macros, but dependent types and their binding structure also complicate defining macros by pattern matching. Suppose we want to change our function type rules to accommodate multiple arguments. In Figure 5 (left), the plain function type may simply use the ellipses pattern, which effectively “maps” over the τ_{in} s. For dependent types, as in for \rightarrow_{vid} , this “map” operation is incorrect, and results in the wrong binding structure. For example, the \rightarrow_{wrong} rule in Figure 5 (right) tries to use the same ellipses pattern as on the left, but this type checks each argument’s type τ_{in} in a type environment with *every* argument x bound, *including itself*. On the left, without dependent types, this is not a problem since types cannot reference the term variables from the same type annotation.

<pre> 393 (struct $\overrightarrow{\quad}$ (in out)) 394 (define-tyrule ($\rightarrow \tau_{in} \dots \tau_{out}$) \gg 395 [$\vdash \tau_{in} \gg \overline{\tau_{in}} \leftarrow \text{Type}$] ... 396 [$\vdash \tau_{out} \gg \overline{\tau_{out}} \leftarrow \text{Type}$] 397 ----- 398 [$\vdash (\#\text{app } \overrightarrow{\quad} (\overline{\tau_{in}} \dots) \overline{\tau_{out}}) \Rightarrow \text{Type}$]) 399 400 401 402 403 404 405 406 407 408 409 410 411 412 413 414 </pre>	<pre> 404 (struct $\overrightarrow{\quad}_{vid}$ (types)) 405 (define-tyrule ($\rightarrow_{wrong} [x : \tau_{in}] \dots \tau_{out}$) \gg 406 [[$x \gg \overline{x} : \overline{\tau_{in}}$] ... $\vdash [\tau_{in} \gg \overline{\tau_{in}} \leftarrow \text{Type}] \dots$] 407 [[$x \gg \overline{x} : \overline{\tau_{in}}$] ... $\vdash \tau_{out} \gg \overline{\tau_{out}} \leftarrow \text{Type}$] 408 ----- 409 [$\vdash (\#\text{app } \overrightarrow{\quad}_{vid} (\lambda (\overline{x} \dots) \overline{\tau_{in}} \dots \overline{\tau_{out}})) \Rightarrow \text{Type}$]) 410 411 (define-tyrule ($\rightarrow_{vid} [x : \tau_{in}] \dots \tau_{out}$) 412 [[$x \gg \overline{x} : \tau_{in} \gg \overline{\tau_{in}} \leftarrow \text{Type}$] ... $\vdash \tau_{out} \gg \overline{\tau_{out}} \leftarrow \text{Type}$] 413 ----- 414 [$\vdash (\#\text{app } \overrightarrow{\quad}_{vid} (\lambda (\overline{x} \dots) \overline{\tau_{in}} \dots \overline{\tau_{out}})) \Rightarrow \text{Type}$]) </pre>
---	--

Fig. 5. Some type rules for multi-arity function types, (left) \rightarrow , (right) \rightarrow_{vid}

Instead, to use familiar macro notation to implement dependent types, we require that ellipses express a “fold” operation for recursively applying macro expansion (*i.e.*, type checking) to express the proper binding structure. Since environments themselves contain type annotations, this fold operation must interleave binding and checking. In Figure 6, we present such a fold operation which is part of the “dependent type systems as macros” core API (*i.e.*, an extension to the core API discussed in Section 2 needed to support dependent types). This new function consumes a name x , a target to check τ , an expected type $\kappa_{expected}$ for τ , and a previous context, and it checks that τ has type $\kappa_{expected}$ *while* adding x and τ to create the next context. This new context is returned along with expanded versions of x and τ .

```

426 (define (folding-check x  $\tau$   $\kappa_{expected}$  #:ctx [ctxprev ()])
427   #:with [ctxnew  $\overline{x}$   $\overline{\tau}$ ] (synth x #:ctx ([x  $\tau$ ] ctxprev))
428   #:with  $\kappa$  (detach  $\overline{\tau}$  'type)
429   (if ( $\tau = \kappa \kappa_{expected}$ ) (ctxnew  $\overline{x}$   $\overline{\tau}$ ) (err "type mismatch"))
430
431
432
433
434
435
436
437
438
439
440
441

```

Fig. 6. A folding variant of the check API function from Figure 3.

In our extension to TURNSTILE, we interpose on the ellipses to use folding-check instead of check when appropriate. In Figure 5 (right), we give the corrected definition of \rightarrow_{vid} using the new TURNSTILE syntax $[x \gg \overline{x} : \tau_{in} \gg \overline{\tau_{in}} \leftarrow \text{Type}] \dots$, which checks each τ_{in} , but also names it so that subsequent type checking invoked by the ellipses may reference the argument. Since the new syntax both checks and binds, subsuming what is typically on the left *and* right side of the (\vdash), programmers may use it on *either side*, *e.g.*, the following is equivalent to the definition in Figure 5:

```

438 (define-tyrule ( $\rightarrow_{vid} [x : \tau_{in}] \dots \tau_{out}$ )
439   [ $\vdash [x \gg \overline{x} : \tau_{in} \gg \overline{\tau_{in}} \leftarrow \text{Type}] \dots [\tau_{out} \gg \overline{\tau_{out}} \leftarrow \text{Type}]$ ]
440   -----
441

```

[\vdash ($\#app \xrightarrow{\text{vid}} (\lambda (\bar{x} \dots) \bar{\tau}_{in} \dots \bar{\tau}_{out})) \Rightarrow \text{Type}$)]

Note that a single lambda wrapping all the types in the macro’s output is sufficient due to hygiene. There will be no capture so long as each type was expanded in the appropriate context.

3.3 Macros for Pattern Matching

In general, the programmer does not need to know the underlying “run-time” type representation, and would prefer to simply pattern match on the *surface* syntax of the type instead of the “run-time” representation produced by macro expansion. In TURNSTILE, we can define “pattern” macros for each type, as in Figure 7. This macro is used exclusively in pattern positions and it matches on, but hides, a type’s internal representation.⁷ While this feature is not strictly necessary, it relieves some notational burden for programmers implementing “dependent type systems as macros.”

```
(define-m ( $\sim \xrightarrow{\text{vid}}$  [ $x : \tau_{in}$ ] ...  $\tau_{out}$ ) ( $\#app \xrightarrow{\text{vid}}$  ( $\lambda (x \dots) \tau_{in} \dots \tau_{out}$ )))
```

Fig. 7. Pattern matching macro for the \rightarrow_{vid} type.

3.4 Putting It All Together

We’ve now seen all the components necessary to define dependent types as macros: a `struct` record declaration for the internal representation, a `define-tyrule` implementing the rule for type construction and elaborating to the `struct`, and one or more pattern macros. As a convenience, we add a new construct to TURNSTILE, `define-type`, that automatically generates the boilerplate and allows language implementors to simply write down the rule for well-formed dependent types.

In Figure 8, we show how to implement (a simplification of) the typing rules for TYPED VIDEO using the approach we’ve presented, and our extensions to TURNSTILE. We see that integer terms may be lifted to the type level via the `Producer` type constructor. The new lambda rule resembles the rule for \rightarrow_{vid} from Figure 4, except a lambda has the \rightarrow_{vid} type. Similarly, the application rule requires that the type of its operator matches a \rightarrow_{vid} type, and that its argument has the type of the \rightarrow_{vid} type’s inputs.

The lambda rule is a multi-clause `define-tyrule`, analogous to multi-clause macros, because we only wish to allow lifting of integer terms to the type level. Notice that the first clause identifies `Int` cases with the $\sim \text{Int}$ pattern macro generated by `define-type`. When the parameter does not have integer type, type checking falls through to the second clause, where the output \rightarrow_{vid} type is constructed with a fresh dummy name, so it may not be referenced in subsequent types. As expected, in the output of the `appvid` type rule, we substitute references to the binder in $\bar{\tau}_{out}$ with the argument from the application.

3.5 Type-Level Computation

Since types may contain integer expressions, we must add type-level computation to normalize the types thus the integer constraints. We present two approaches to type-level computation “as macros”: a simple approach here, and a more modular and extensible approach in Section 4.

To enable the first approach, we again modify the “dependent types as macros” core API. First, we add a new interposition point in the `assign` metafunction, so the new definition is the following.

```
(define (assign e  $\tau$ ) (attach e 'type ( $\tau\text{-eval}$   $\tau$ )))
```

The interposition point `$\tau\text{-eval}$` enables customization of type normalization. Since `assign` is implicitly called in the conclusion of every `define-tyrule`, interposing on `$\tau\text{-eval}$` allows us to inject

⁷The pattern macros could have the same name as its analogous type, but to better distinguish pattern positions we (and TURNSTILE) follow Racket’s convention of prefixing pattern macros with \sim . See Figure 8 for a usage of a pattern macro.

```

491 #lang turnstile
492 (provide (rename [λvid λ] [appvid #%app])
493 (define-type Int : Type)
494 (define-type Producer : Int -> Type)
495 (define-type →vid #:binders ([X : Type]) : Type)
496 (define-tyrule λvid
497 [(λ [x : τin] e) >> ; Int case
498  [τin >> ~Int <= Type]
499  [[x >> x̄ : Int] ⊢ e >> ē ⇒ τout]
500 -----
501  [τ (λ (x̄) ē) ⇒ (→vid [x̄ : Int] τout)]])
502 [(λ [x : τin] e) >>
503  [τin >> τ̄in <= Type]
504  [[x >> x̄ : τ̄in] ⊢ e >> ē ⇒ τout]
505 -----
506  [τ (λ (x̄) ē) ⇒ (→vid [dummy : τ̄in] τout)]])
507 (define-tyrule (appvid f e) >>
508  [τ f >> f̄ ⇒ (→vid [x̄ : τ̄in] τ̄out)]
509  [τ e >> ē <= τ̄in]
510 -----
511  [τ (#app f̄ ē) ⇒ (subst ē x̄ τ̄out)]])

```

typed/video

Fig. 8. TYPED VIDEO type definitions, lambda, and function application rules.

```

513 #lang turnstile
514 (define-τ-eval
515  [nint n] [bbool b]
516  [(+ n m) #:with n*int (τ-eval n) #:with m*int (τ-eval m) (+ n*.val m*.val)]
517  [(< n m) #:with n*int (τ-eval n) #:with m*int (τ-eval m) (< n*.val m*.val)]
518  [(Producer n) (Producer (τ-eval n))]
519  [other other])

```

typed/video

Fig. 9. Excerpt of type-level evaluation in the TYPED VIDEO language.

the needed behavior. By default, τ -eval just expands a type; for TYPED VIDEO, we implement an interpreter for the index language. Figure 9 shows a (simplified version of) this function. We use $\text{define-}\tau\text{-eval}$ to redefine the τ -eval function used by other type rules. The definition is a series of pattern-body clauses. When τ -eval is called with a type τ , the first clause whose pattern matches τ is used. The first two clauses match literal values. The third clause matches on addition. This clause first recursively calls τ -eval on the arguments. If evaluating those terms produce syntactic literal numbers, then the actual arithmetic operation is performed. This fourth case is similar. If the input to τ -eval is a `Producer`, then its index is evaluated, otherwise the type is left unchanged.

4 A DEPENDENTLY-TYPED CALCULUS

The approach to type-level computation for the TYPED VIDEO language in Section 3 suffices when the index language is simple. It does not scale well when, for example, we want to define new reduction rules that can be used both for run-time and during type checking, as is common in type theory. This section presents a more general, extensible approach to adding type-level computation via macros where types and terms may mix.

Again, we do this in the context of an example language. We start with essentially the Calculus of Constructions (CC) [Coquand and Huet 1988], which features “heavyweight”, also called full-spectrum, dependent types, in which there is no distinction between terms and types. We gradually extend our initial implementation with type schemas, ala Martin-Löf Type Theory [Martin-Löf 1975], and finally extend to the Calculus of Inductive Constructions [Pfenning and Paulin-Mohring 1989]. This demonstrates that our approach scales to the same calculi used in contemporary proof assistants. Each extension is entirely modular: it does not require modifying any prior code, and only defines new macros. This demonstrates key features of the “dependent type systems as macros” approach: modularity and extensibility.

We start by upgrading Figure 2’s simply-typed language into CC by:

- (1) changing the \rightarrow type into a Π type, whose output type can refer to its input type;
- (2) modifying the lambda and application rules to introduce and eliminate the Π type; and
- (3) implementing reduction rules for type-level computation.

We first present the key concepts as they apply to macro systems in general, and then the new TURNSTILE abstractions that support on-paper notation.

4.1 Defining Type-Level Reductions

Figure 10 presents DEP-LANG, a dependent calculus with Π types, *i.e.*, dependently typed functions. The new lambda rule introduces the Π type and the function application rule eliminates it. The key difference from TYPED VIDEO’s calculus is in the conclusion of the function application rule.

```

561 #lang turnstile
562 (provide  $\Pi$  (rename [ $\lambda_{\text{dep}}$   $\lambda$ ] [appdep #%app])
563 (define-type  $\Pi$  #:binders ([X : Type]) -> Type)
564 (define-tyrule (appdep f e) >>
565   [ $\vdash f \gg \bar{f} \Rightarrow (\sim\Pi [\bar{X} : \bar{\tau}_{\text{in}}] \bar{\tau}_{\text{out}})$ ]
566   [ $\vdash e \gg \bar{e} \Leftarrow \bar{\tau}_{\text{in}}$ ]
567   -----
568   [ $\vdash (\beta \bar{f} \bar{e}) \Rightarrow (\uparrow/\vee 1 (\text{subst } \bar{e} \bar{X} \bar{\tau}_{\text{out}}))$ ])
569
570 (define-tyrule ( $\lambda_{\text{dep}}$  [x :  $\tau$ ] e) >>
571   [ $\vdash \tau \gg \bar{\tau} \Leftarrow \text{Type}$ ]
572   [[x >>  $\bar{x} : \bar{\tau} \vdash e \gg \bar{e} \Rightarrow \bar{\tau}_{\text{out}}$ ]
573   -----
574   [ $\vdash (\lambda (\bar{x}) \bar{e}) \Rightarrow (\Pi [\bar{x} : \bar{\tau}] \bar{\tau}_{\text{out}})$ ])

```

Fig. 10. A dependently-typed lambda calculus.

In app_{dep}’s output type, we replace the Π type binder \bar{x} with the argument of the application \bar{e} . To support arbitrary run-time terms in the type system, the type is also wrapped with a “reflect” operation $\uparrow/\vee 1$ that will be explained in detail shortly. To implement the reduction rule for Π , the output “term” (which is also a type) is wrapped with a β macro which implements the reduction rule for Π , enabling evaluation during type checking.

Figure 11 defines the β -reduction rule as a macro. The β macro expands its head expression and matches on that result using an explicit syntax-parse syntax pattern matcher. If the expanded head matches a λ (first case), occurrences of the λ parameter x in the body are replaced with the argument e . Performing the reduction, however, may create additional redexes, for example if the argument itself is a function. To further reduce these new redexes in the contractum, we need to “reflect” references to run-time representations #%app back to β , and the $\uparrow/\vee 1$ function performs this operation. Otherwise (second case), the result is an unreduced run-time #%app term.

While Figure 11 conceptually captures our approach to type-level computation via macros, this obvious implementation is not extensible since the reflection operation would need to know about all possible reduction rules in advance. Instead, we add a new core API function \uparrow for reflection, defined in Figure 12, which is extensible via annotation on syntax objects. Instead of just replacing

```

589 (define-m ( $\beta$  f e)
590   (syntax-parse (local-expand f)
591     [( $\lambda$  (x) body) ( $\uparrow$ /v1 (subst e x body))])
592     [ $\bar{f}$  (#%app  $\bar{f}$  e)])])
593
594 (define ( $\uparrow$ /v1 e) (subst  $\beta$  #%app e)) ; fn mapping #%app back to  $\beta$ ; not extensible

```

Fig. 11. Beta reduction rule, implemented as a plain macro.

598 `;%app`, it traverses a piece of syntax and checks for the syntax property `'reflected-name`. If the
599 property exists, the value associated with the key is used as the reflected name. The `mk-reflected`
600 function expects a `placeholder`, an identifier to use as the run-time representation of the elimination
601 form, and attaches a given identifier for use as the `'reflected-name` property. For example, the
602 `placeholder` for application is `;%app`, and the reflected id will be the `β` macro.

```

603
604 (define ( $\uparrow$  e) ; extensible fn mapping terms to surface stx
605   (syntax-parse e
606     [placeholderid (detach placeholder 'reflected-name)]
607     [(e ...) (( $\uparrow$  e) ...)]
608     [else e]))
609
610 (define (mk-reflected placeholder reflected-id)
611   (attach placeholder 'reflected-name reflected-id))
612
613 ; example: ((mk-reflected #%app  $\beta$ ) ( $\lambda$  (x) x) ( $\lambda$  (x) x)) = (%app ( $\lambda$  (x) x) ( $\lambda$  (x) x))
614 ; example: ( $\uparrow$  (mk-reflected #%app  $\beta$ )) =  $\beta$ 
615 ; example: ( $\uparrow$  ((mk-reflected #%app  $\beta$ ) ( $\lambda$  (x) x) ( $\lambda$  (x) x))) = ( $\beta$  ( $\lambda$  (x) x) ( $\lambda$  (x) x))
616 ; example: (local-expand ( $\uparrow$  ((mk-reflected #%app  $\beta$ ) ( $\lambda$  (x) x) ( $\lambda$  (x) x)))) = ( $\lambda$  (x) x)

```

Fig. 12. Core API for reflection

619 This API, while small involves a fundamental change in how macro expansion proceeds. Typically,
620 we think of the work flow with macros as: 1) macro expansion, 2) runtime evaluation. In the previous
621 “type systems as macros” work, this changes to: 1) macro expansion + type checking (interleaved)
622 2) runtime. Macro expansion is interleaved with type checking, since the type system is defined
623 in macros. This requires communicating types between macros, which happens through syntax
624 properties on syntax objects. With our extension to the “types systems as macros” API, this changes
625 again. Instead, we have: 1) macro expansion + type checking + evaluation (interleaved), 2) runtime
626 evaluation. All are mutually recursive and we must coordinate information between each stage.
627 Now we must communicate how a reduced term corresponds to a type (*i.e.*, macro), and therefore,
628 to its own type-level reduction semantics; this is the role of the reflection API above. We describe
629 this interleaved semantics and the requirements it imposes on macros systems in [Appendix B](#).

630 For our TURNSTILE implementation, we add abstractions to avoid the boilerplate in the pattern
631 described above. [Figure 13](#) defines `define-red`, a macro-defining macro. Given a redex and contractum,
632 it generates macros like `β` in [Figure 11](#), where the generated macro automatically handles reflecting
633 the contractum with `\uparrow` . In essence, multiple `define-red` declarations cooperate with each other
634 through the API in [Figure 12](#).

635 The definition of `define-red` is a multi-clause macro, which itself generates a macro representing
636 a reduction rule, such as `β` , but automatically inserts the `\uparrow` and `mk-reflected` calls as required. The

637

```

638 (define-m define-red ; TURNSTILE form for defining reduction rules
639   [(define-red red-name redex ~> contractum) ; single-redex case
640    (define-red red-name [redex ~> contractum])] ; rewrite to match second case
641   ; multi-redex case
642   [(define-red red-name [(placeholder redex-hd redex-rst ...) ~> contractum] ...)
643    ⌈(define-m (red-name hd arg ...)
644      (syntax-parse ((local-expand hd) arg ...)
645        [(redex-hd redex-rst ...) (↑ contractum)] ...)
646      ⌋(e ...) ((mk-reflected placeholder red-name) e ...))]⌋)

```

Fig. 13. TURNSTILE form for defining reduction rules.

```

648 (define-red  $\beta$  (#%app ( $\lambda$  (x) body) e) ~> (subst e x body))
649

```

Fig. 14. Beta reduction rule, implemented with define-red.

652

653

654 first case is short-hand for easy reduction rules, and recursively calls `define-red` to invoke the second
655 case. The second case accommodates multiple redexes and contractums. (The output of the second
656 clause is marked with `⌈` instead of the usual blue text color, to avoid obscuring nested patterns and
657 templates in the generated macro.) The generated reduction macro, `red-name`, behaves essentially
658 like a generalized version of β in Figure 11. More specifically, `red-name` expands the head, and if
659 it matches the supplied redex, rewrites it to the specified contractum. Otherwise, it expands to a
660 term represented by the placeholder but marked with reflection property *reflected-name*. If further
661 evaluation (*i.e.*, macro-based reductions) causes the term to become redex, then `↑` ensures that
662 `red-name` is invoked again to reduce the redex. Figure 14 shows the definition of the β -reduction rule
663 implemented with TURNSTILE’s `define-red`, where the redex is a macro pattern from the underlying
664 macro system and the contractum rewrites components of that pattern. This definition concisely
665 matches how such a rule would be written in a textbook.

666

667

4.2 A Little Sugar

668

669

670

671

672

673

674

675

676

677

678

679

680

681

682

683

684

685

686

The DEP-LANG language from Figure 10 and Figure 14 is equivalent to the Calculus of Constructions (CC) [Coquand and Huet 1988]. There are many tutorials on implementing dependent types, and they typically end here, but it is very difficult to actually program or prove with CC. Fortunately, a fundamental feature of the “(dependent) type system as macros” approach is that DSLs gain extensibility via macros, for free. We demonstrate this for dependent types in Figure 15, which presents DEP-LANG/SUGAR, a library that adds syntactic sugar for DEP-LANG using macros—local, modular elaboration passes. In contrast, a typical implementation of a dependently-typed language would add a whole-program elaboration pass on top of the core. We subsequently use our DEP-LANG/SUGAR library to add additional type schemas.

We define automatically-currying, multiple-argument versions of Π , λ , and function application. We may also define \rightarrow and \forall as shorthands for Π , where the former generates an arbitrary name, and the latter inserts implicit Type annotations. These new variations are exported with the same name as their single-arity versions, using the interposition feature of the macro system to interpose on the definitions from DEP-LANG. Thus new DEP-LANG programs importing this library will automatically use the new sugary forms.

The last macro in our library, `define-data-constructor`, wraps `define-type` with the just-defined λ/c to allow partial application of its constructor. Like `define-type`, `define-data-constructor` supports syntax for declaring data structures with either named or unnamed arguments.

```

687 #lang dep-lang
688 (provide → ∀ (rename [λ/c λ] [app/c #%app] [Π/c Π]))
689 (define-m Π/c
690   [(_ e) e]
691   [(_ x . rst) (Π x (Π/c . rst))])
692 (define-m (→ τin ... τout) (Π [TMP : τin] ... τout) ; TMP fresh
693 (define-m (∀ X ... τ) (Π [X : Type] ... τ))
694 (define-m λ/c
695   [(_ e) e]
696   [(_ x . rst) (λ x (λ/c . rst))])
697 (define-m app/c
698   [(_ e) e]
699   [(_ f e . rst) (app/c (=%app f e) . rst)])
700 (define-m define-data-constructor
701   [(define-data-constructor name : τ ... → τout)
702    (define-data-constructor name : [TMP : τ] ... → τout) ; TMP fresh
703    [(define-data-constructor name : [x : τ] ... → τout)
704     ↖(define-type name : [x : τ] ... -> τout) ↗
705     (define-m name (λ/c [x : τ] ... (name x ...)))
706     ↙(define-m ~name ~name) ↘]

```

dep-lang/sugar

Fig. 15. A DEP-LANG library that adds some syntactic sugar, e.g., currying.

4.3 A Library of Natural Numbers

Technically, we could Church-encode all our programs and proofs, but this is somewhat impractical. Luckily, we have already developed all tools to extend DEP-LANG with new datatypes. Figure 16 extends DEP-LANG with a natural number library, using a type schema in the style of Martin-Löf Type Theory [Martin-Löf 1975]. Each type schema defines a type, an introduction rule, and an elimination rule; we ignore equivalence rules for this presentation. Specifically, we define Nat using the TURNSTILE define-type form. We use the define-data-constructor variant of define-type (from Figure 15) to define the standard introduction rules, Z and S, corresponding to “zero” and “successor”. The elimination form, elim^{Nat} , corresponds to a fold over the datatype. Following the terminology of McBride [2000], the form $(\text{elim}^{\text{Nat}} n P m_z m_s)$ takes *target* to eliminate n , a *motive* P that describes the return type of this form, and one *method* for each case of natural numbers: m_z when n is zero and m_s when n is the successor of a number. Method m_z must have type $(P Z)$, i.e., the motive applied to zero, while m_s must have type $(\Pi [k : \text{Nat}] (\rightarrow (P k) (P (S k))))$, which mirrors an induction proof: for any k , given a proof of $(P k)$, we show $(P (S k))$.

Using define-red, we can define reduction rules for elim^{Nat} , one each for Z and S, as succinctly as in a textbook. Observe that the pattern macros $\sim Z$ and $\sim S$, defined by define-type, are useful when specifying the reduction. For convenience, the DEP-LANG/NAT library also extends $\#\text{datum}$, an interposition point for interpreting literal data. With new-datum, users of the DEP-LANG/NAT library can write numeric literals in place of the more cumbersome Z and S constructors. The last new-datum clause falls back to the current $\#\text{datum}$, making this library compatible with other literal data. We can even support diamond extensions by importing *two* existing versions of $\#\text{datum}$ (under different names) and using them in separate new-datum clauses and, of course, writing some macros to automate such boilerplate.

```

736 #lang dep-lang
737 (provide Nat Z S elimNat (rename [new-datum #%datum]) +)
738
739 (define-type Nat : Type)
740 (define-data-constructor Z : Nat)
741 (define-data-constructor S : Nat → Nat)
742
743 (define-tyrule (elimNat n P mz ms) >>
744   [⊢ n >> n̄ ← Nat] ; target
745   [⊢ P >> P̄ ← (→ Nat Type)] ; prop / motive
746   [⊢ mz >> m̄z ← (P̄ Z)] ; method for Z
747   [⊢ ms >> m̄s ← (Π [k : Nat] (→ (P̄ k) (P̄ (S k))))] ; method for S
748   -----
749   [⊢ (evalNat n̄ P̄ m̄z m̄s) ⇒ (P̄ b̄)]
750
751 (define-red evalNat
752   [(elimNat ~Z P mz ms) ~> mz]
753   [(elimNat (~S k) P mz ms) ~> (ms k (evalNat k P mz ms))]
754
755 (define-m new-datum
756   [(new-datum nnat) #:when (zero? n) Z]
757   [(new-datum nnat) (S (new-datum (- n 1)))]
758   [(new-datum x) (new-datum x)])
759
760 (define + ; implements n + m
761   (λ [n : Nat]
762     (elimNat n
763       (λ [m : Nat] (→ Nat Nat))
764       (λ [m : Nat] m)
765       (λ [n-1 : Nat] [ih : (→ Nat Nat)] (λ [m : Nat] (S (ih m)))))))

```

Fig. 16. A DEP-LANG library for natural numbers.

4.4 An Equality Type Library, and Applying Telescopes

Figure 17 shows an implementation of the standard equality, or identity, type. The $\text{elim}^=$ rule resembles elim^{Nat} from Figure 16: for any motive P such that $(P\ a)$ holds, eliminating a proof that $a = b$ allows concluding that $(P\ b)$ holds.

The implementation of $\text{elim}^=$ demonstrates the implicit support for *telescopes* we’ve added to TURNSTILE to simplify implementing dependent types. The arguments are named and subsequent arguments may reference previous names. Note that *checking* a telescope and *applying* a constructor with telescoping arguments, which involves substitution, are two distinct operations. Section 3.2 presented our implementation of abstractions for the former; the rest of this subsection addresses the latter with a novel, pattern-based substitution technique for instantiating types in a telescope.

Figure 18 shows the relevant parts of `define-type`, which generates a `define-tyrule` that uses this technique. `define-type` first validates the $\kappa_{\text{in}} \dots \kappa_{\text{out}}$ annotations supplied by the programmer, using the new `folding-check` from Section 3.2. The conclusion ($a >$ variant accommodates emitting definitions) produces the type rule for constructing name types.

The key is the reuse of the $\bar{A} \dots$ pattern variables from the premises of the `define-type` as the pattern variables of the generated `define-tyrule`. When the name type constructor is called, \bar{A} is bound to the arguments supplied to that constructor. Since \bar{A} binds references in $\bar{\kappa}_{\text{in}} \dots$, use of $\bar{\kappa}_{\text{in}} \dots$ in

dep-lang/nat


```

785 #lang dep-lang
786 (provide = refl elim=)
787
788 (define-type ≡ : [A : Type] [a : A] [b : A] → Type)
789 (define-data-constructor refl : [A : Type] [e : A] → (= A e e))
790
791 (define-tyrule (elim= t P pt w peq) >>
792   [⊢ t >>  $\bar{t} \Rightarrow A$ ]
793   [⊢ P >>  $\bar{P} \Leftarrow (\rightarrow A \text{Type})$ ]
794   [⊢ pt >>  $\bar{pt} \Leftarrow (\bar{P} \bar{t})$ ]
795   [⊢ w >>  $\bar{w} \Leftarrow A$ ]
796   [⊢ peq >>  $\bar{peq} \Leftarrow (= A \bar{t} \bar{w})$ ]
797   -----
798   [⊢  $\bar{pt} \Rightarrow (\bar{P} \bar{w})$ ])

```

dep-lang/eq

Fig. 17. A DEP-LANG library for the equality type.

the output syntax template *automatically* instantiates any \bar{A} s in $\bar{\kappa}_{in}$ with the concrete arguments to the name type constructor, which is the desired behavior. In other words, we hijack substitutions that the macro system already performs with pattern variables in templates to instantiate type variables. The technique is safe, *i.e.*, no variables are captured, thanks to hygiene.

```

807 (define-tyrule (define-type name : [Aid :  $\kappa_{in}$ ] ... →  $\kappa_{out}$ ) >>
808   [[A >>  $\bar{A} : \kappa_{in} \gg \bar{\kappa}_{in} \Leftarrow \text{Type}$ ] ... ⊢  $\kappa_{out} \gg \bar{\kappa}_{out} \Leftarrow \text{Type}$ ]
809   -----
810   [⊢ (define-tyrule (name  $\bar{A}$  ...) >>⌈
811     [⊢  $\bar{A} \gg \bar{A} \Leftarrow \bar{\kappa}_{in}$ ] ...
812     -----
813     [⊢ (name  $\bar{A}$  ...) ⇒  $\bar{\kappa}_{out}$ ])
814     ⌋ ; rest of the macro elided ⌋])

```

Fig. 18. (Part of) the implementation of define-type.

With numbers and equality, we can now write a simple example proof in DEP-LANG. Figure 19 proves the additive identity of the natural numbers. The left identity is simple since + is defined by recursion on its first argument and (+ 0 n) trivially reduces to n. The right identity requires an argument by induction, since (+ n 0) cannot evaluate until we know more about n.

4.5 Indexed Inductive Type Families

TURNSTILE easily supports type schemas, but type schemas are a modification to the trusted core. Realistic proof assistants instead support safe extension through inductively-defined type families Dybjer [1994]. Inductive types can be implemented in one set of general-purpose rules that are proven sound, and then users can declare new inductive types without extending the trusted core. Adding indexed inductive type families is straightforward using the constructs we have already presented.

Figure 21 presents define-datatype, which enables defining inductive types. Our version is based on Brady's presentation of TT [Brady 2005]. The complete implementation requires 18 lines of

```

834 #lang dep-lang dep-lang-prog
835 (require dep-lang/nat dep-lang/eq)
836 (ann (λ [n : Nat] (refl Nat n)) : (Π [n : Nat] (= Nat (+ 0 n) n)))
837 (ann (λ [n : Nat]
838     (elimNat n
839     (λ [m : Nat] (= (+ m 0) m))
840     (refl Nat 0)
841     (λ [k : Nat] [p : (= (+ k 0) k)]
842     (elim= (+ k 0)
843     (λ x (= Nat (S (+ k 0)) (S x)))
844     (refl Nat (S (+ k 0)))
845     k
846     p))))
847 : (Π [n : Nat] (= Nat (+ n 0) n)))
848

```

Fig. 19. An example DEP-LANG program: proving additive identity.

code.⁸ It makes DEP-LANG equivalent to the Calculus of Inductive Constructions, *i.e.*, the core of the Coq proof assistant, demonstrating that “dependent type system as macros” scales to expressive type theories, and supports on-paper notation even for advanced typing rules like inductive types.

$$\frac{A : \star \quad n : \mathbb{N}}{\text{Vec } A \ n : \star} \quad \text{where} \quad \frac{}{\text{nil} : \text{Vec } A \ 0} \quad \frac{k : \mathbb{N} \quad x : A \quad xs : \text{Vec } A \ k}{\text{cons } x \ xs : \text{Vec } A \ (S \ k)}$$

```

858
859 #lang dep-lang list-prog
860 (define-datatype Vec [A : Type] : [i : Nat] → Type
861   [nil : (Vec A 0)]
862   [cons : [k : Nat] [x : A] [xs : (Vec A k)] → (Vec A (S k))])
863

```

Fig. 20. Indexed list data definitions, (Top) by hand, and (Bottom) in DEP-LANG.

To help our explanation of `define-datatype`, we begin with a concrete example. Figure 20 shows two length-indexed list data definitions, the first using a natural deduction style as commonly written in the literature (*e.g.*, [McBride and McKinna 2004]) and the second as written with `define-datatype` in DEP-LANG, which is based on Coq’s notation. The main source of complexity compared to previous type definitions is that indexed inductive type families distinguish between *parameters* (the A in the figure), and *indices* (the i in the figure). Briefly, parameters are invariant across the definition while indices may vary. The key is that in both the formal notation and the code, the rules for the data constructors reference the parameter A that is bound in the type definition, while the index argument is specific to each rule. It turns out that this invariance of parameters can be used to simplify the implementation of `define-datatype`, as the following prose explains.

At a high-level, `define-datatype` is “just” a macro that produces four output definitions:

- (1) a `define-type` type definition;
- (2) `define-data-constructor` data constructor definitions;
- (3) a `define-tyrule` elimination rule; and
- (4) a `define-red` reduction rule.

⁸Admittedly, we elide positivity checking for simplicity.

```

883 #lang turnstile dep-lang
884 (define-tyrule (define-datatype TY [A : τA] ... : [i : τi] ... → τ
885                [C : [i|x : τin] ... → τout] ...) >>
886     [[A >>  $\bar{A}$  : τA >>  $\bar{\tau}_A$  <= Type] ... [[i >>  $\bar{i}$  : τi >>  $\bar{\tau}_i$  <= Type] ... ⊢ τ >>  $\bar{\tau}$  <= Type]
887     [[i|x >>  $\bar{i}|x$  : τin >>  $\bar{\tau}_{in}$  <= Type] ... ⊢ τout >>  $\bar{\tau}_{out}$  <= Type]...]
888     #:with (TY _ ...  $\bar{\tau}_{outi}$  ...) (τout ...)
889     #:with (( $\bar{i}_{rec}$  ...  $\bar{x}_{rec}$ ) ...) (find-recur (([i|x  $\bar{\tau}_{in}$ ] ...) ...) TY)
890     -----
891     [> Γ(define-type TY : [A : τA] ... [i : τi] ... → τ) ; define the type and constructors ⊢
892         (define-data-constructor  $\bar{C}$  : [A : τA] ... [i|x : τin] ... → τout) ...
893
894     (define-tyrule (elimTY v P m ...) ; eliminator for TY
895         [⊢ v >>  $\bar{v}$  ⇒ (~TY  $\bar{A}$  ...  $\bar{i}_{inferred}$  ...)] ; target
896         [⊢ P >>  $\bar{P}$  <= (Π [ $\bar{i}$  :  $\bar{\tau}_i$ ] ... (→ (TY  $\bar{A}$  ...  $\bar{i}$  ...) Type))] ; motive
897         [⊢ m >>  $\bar{m}$  <= (Π [ $\bar{i}|x$  :  $\bar{\tau}_{in}$ ] ...
898             (→ ( $\bar{P}$   $\bar{i}_{rec}$  ...  $\bar{x}_{rec}$ ) ... ( $\bar{P}$   $\bar{\tau}_{outi}$  ... (C  $\bar{A}$  ...  $\bar{i}|x$  ...))))] ...
899     -----
900     [⊢ (evalTY  $\bar{v}$   $\bar{P}$   $\bar{m}$  ...) ⇒ ( $\bar{P}$   $\bar{i}_{inferred}$  ...  $\bar{v}$ )]
901
902     (define-red evalTY ; define reduction rule
903     L [[(elimTY (~C A ...  $\bar{i}|x$  ...) P m ...) ~> (m  $\bar{i}|x$  ... (evalTY  $\bar{x}_{rec}$  P m ...) ...) ] ⊢])
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     -----
    
```

Fig. 21. Implementation of define-datatype makes DEP-LANG equivalent to CIC.

But the details are dense so we go line-by-line.

- (define-tyrule (define-datatype TY [A : τ_A] ... : [i : τ_i] ... → τ
[C : [i|x : τ_{in}] ... → τ_{out}] ...)

This defines a new type checking macro named `define-datatype`. The first part of the inputs consists of the name of a new type `TY`, its parameter names `A ...`, the types of those parameters `τA ...`, index names `i ...`, and the types of those indices `τi ...`. Together, `[A : τA] ... [i : τi] ...` is a telescope, where each `τA ... τi ...` may reference the names that come before it. The result type of the type constructor `TY` itself is the type `τ`. The second line of the input specifies the constructors for `TY`, `C ...`. The type of the constructors is described by the telescopes `[i|x : τin] ...` (where `i|x` is a literal identifier). Finally, a fully-applied constructor has type of shape `τout`, which we refine later. The key is that the `A` binders range over the entire declaration, *i.e.*, the `τin` and `τout` types may also reference `A`, which is not true of the `i` binders.

- $$[[A \gg \bar{A} : \tau_A \gg \bar{\tau}_A \leftarrow \text{Type}] \dots [[i \gg \bar{i} : \tau_i \gg \bar{\tau}_i \leftarrow \text{Type}] \dots \vdash \tau \gg \bar{\tau} \leftarrow \text{Type}]$$

$$[[i|x \gg \bar{i}|x : \tau_{in} \gg \bar{\tau}_{in} \leftarrow \text{Type}] \dots \vdash \tau_{out} \gg \bar{\tau}_{out} \leftarrow \text{Type}] \dots]$$

These premises validate that the types supplied by the programmer in a `define-datatype` declaration have type `Type`. It uses the new folding `TURNSTILE` syntax introduced in [Section 3.2](#), but with a new twist. Since the `A` ranges over the entire definition, we are essentially checking *nested* telescopes; `TURNSTILE` supports this, using the slightly altered syntax, compared to [Section 3.2](#), above.

- #:with (TY _ ... τ_{outi} ...) (τ_{out} ...)

This extracts the index arguments in the data constructor output types and binds them to the `τouti...` pattern, which is later used to check the eliminator methods. For example, `τouti...` would correspond to `θ` and `(S k)` in [Figure 20](#)'s definitions.

932 • #:with (((i_{rec} ... x_{rec}) ...) ...) (find-recur (((i|x τ_{in}] ...) ...) TY)

933 This line finds the recursive arguments for each constructor C. The function, find-recur (elided),
934 returns arguments x_{rec} ∈ i|x ... whose type is equal to TY. In addition, find-recur finds the indices
935 i_{rec} ... ⊆ i|x ... that the type of x_{rec} references.

936 • (define-type TY : [A : τ_A] ... [i : τ_i] ... → τ)

937 This defines type TY, using parameters and indices given to define-datatype.

938 • (define-data-constructor C : [A : τ_A] ... [i|x : τ_{in}] ... → τ_{out}) ...

939 This line defines data constructors C ..., with the type specified in the input to define-datatype.
940 It uses the define-data-constructor form, from DEP-LANG/SUGAR, so that the constructors may be
941 partially applied. Note that parameters [A : τ_A] ... are added to each constructor declaration.

942 • (define-tyrule (elim^{TY} v P m ...))

943 This implements the type rule for the eliminator named elim^{TY}, which has three kinds of inputs:
944 a target v, a motive P, and methods m ..., one for each C ...

945 • [⊢ v ≫ \bar{v} ⇒ (~TY \bar{A} ... i_{inferred} ...)]

946 The target v must have type TY, with parameters bound to \bar{A} ... and indices to i_{inferred} ... This
947 reuse of pattern variables \bar{A} ... (from the premises to define-datatype) is instance of the pattern-
948 based type instantiation technique introduced in Section 4.4. Within this elimination rule, any
949 other pattern variables from define-datatype's input with references to \bar{A} ..., e.g., $\bar{\tau}_A$, $\bar{\tau}_i$, or $\bar{\tau}_{in}$, will
950 automatically be instantiated with v's parameters by the macro system. Note that we *do not use*
951 this technique for the indices. Instead we bind new pattern variables i_{inferred} ...

952 • [⊢ P ≫ \bar{P} ⇐ (Π [\bar{i} : $\bar{\tau}_i$] ... (→ (TY \bar{A} ... \bar{i} ...) Type))]

953 The motive P is a function that consumes indices \bar{i} ... and a value with type TY at those indices,
954 and returns a type that is the result of the elimination. Here $\bar{\tau}_i$... are the types of indexes specified
955 in the input to define-datatype, but automatically instantiated with the inferred concrete parameters
956 of the target v. The \bar{A} ... passed to TY are those same parameters.

957 • [⊢ m ≫ \bar{m} ⇐ (Π [$\bar{i}|x$: $\bar{\tau}_{in}$] ...
958 (→ (P \bar{i}_{rec} ... \bar{x}_{rec}) ... (P $\bar{\tau}_{out}$... (C \bar{A} ... $\bar{i}|x$...)))] ...

959 A call to the eliminator must include one method for each constructor C ... Each method m
960 consumes the constructor inputs [$\bar{i}|x$: $\bar{\tau}_{in}$] ..., as specified in the input to define-datatype, and an
961 argument for each recursive argument \bar{x}_{rec} . These latter arguments represent recursive applications
962 of the eliminators, so have types specified by the motive P, i.e., (P \bar{i}_{rec} ... \bar{x}_{rec}) ... The type
963 (P $\bar{\tau}_{out}$... (C \bar{A} ... $\bar{i}|x$...)) of each method's result is also determined by the motive. The $\bar{\tau}_{out}$...
964 comes from the constructor output types specified in the input to define-datatype.

965 • [⊢ (eval^{TY} \bar{v} \bar{P} \bar{m} ...) ⇒ (P i_{inferred} ... \bar{v})]

966 The eliminator output calls reduction rule eval^{TY} to reduce redexes where \bar{v} is a fully-applied
967 constructor. Its type is determined by the motive applied to the indices of the target \bar{v} and \bar{v} itself.

968 • (define-red eval^{TY}
969 [(elim^{TY} (~C A ... i|x ...) P m ...) ~> (m i|x ... (eval^{TY} x_{rec} P m ...) ...)] ...)

970 The last definition produced by a define-datatype declaration is a reduction rule consisting of
971 a series of redexes, one for each constructor C ... The rule states that elimination of a fully-
972 applied constructor C reduces to application of the method for that constructor, where the recursive
973 arguments to the method are additional invocations of the eliminator on the recursive constructor
974 arguments. Observe how the macro system's pattern language naturally associates each C with its
975 method m, again leading to concise definition that matches what language designers write on paper.
976 For comparison, see the specification of inductive type families from Brady [2005].

5 CUR: A PROOF ASSISTANT AS MACROS

To demonstrate that our approach scales to a realistic language, we implemented a prototype proof assistant called CUR. Proof assistants based on dependently-typed languages are unusual in that they directly implement a formal calculus—typically intended as only a theoretical model of computation—as their core language, thus ensuring a small, consistent, trusted computing base. CUR’s core is the DEP-LANG dependent calculus presented in Section 4. To make programming and proving in a calculus practical, proof assistants typically build separate layers of features and DSLs, such as unification for generating annotations, notational support to generating definitions, or tactic systems for constructing proofs. By building CUR with macros from the beginning, we already have a framework in which we, and CUR users, can easily build such DSLs. Further, *all* new DSLs are integrated into the language, instead of as third-party preprocessors.

This section presents two such DSLs: OLLY—a DSL for modeling programming languages inspired by Ott [Sewell et al. 2007]—and NTAC—a tactic language for scripting proofs. These DSLs elaborate to core CUR during macro expansion, but before type checking, and thus we are able to extend the functionality of our language yet keep the trusted base small. We demonstrate how these DSLs simplify formal development by allowing users to express programs and proofs using familiar notation, rather than the syntax of dependent type theory.

5.1 Olly

OLLY is an Ott-like [Sewell et al. 2007] DSL for modeling programming languages in CUR. Specifically, programmers may write BNF notation or inference rule notation to specify language syntax and relations, respectively, and OLLY automatically generates the inductive type definitions to represent them. Both notations support extracting the models to \LaTeX and Coq, in addition to using the models directly in Cur. Unlike Ott, however, which is an external tool chain, Olly is a user-written library for CUR. As such, it can take advantage of the existing elaboration framework, and is integrated into the standard CUR development environment and language.

Figure 22 shows how one may define the syntax of a simply-typed λ -calculus using OLLY. This language includes booleans, unit, pairs, and functions. The definition uses standard BNF notation, with optional annotations of the form `#:bind <var>` to specify a binding position. Note that the `let` form eliminates pairs in this language, and thus binds two names.

```

1013 #lang cur
1014 (require cur/olly)
1015 (define-language stlc #:vars (x)
1016 #:output-coq "stlc.v" #:output-latex "stlc.tex"
1017 val (v) ::= true false unit
1018 type (A B) ::= boolty unitty (-> A B) (* A A)
1019 term (e) ::= x v (lambda (#:bind x : A) e) (app e e)
1020 (cons e e) (let (#:bind x #:bind x) = e in e))

```

olly-prog

Fig. 22. An STLC example using OLLY, a notation extension for CUR.

The first argument, `stlc`, is the language name. The next three are optional arguments: `#:vars` specifies meta-variables for variables in the syntax; `#:output-coq` specifies a Coq output file; and `#:output-latex` specifies a file for a \LaTeX rendering of the BNF grammar. After the optional arguments, an arbitrary number of non-terminal definitions are specified.

1030 The `define-language` form generates an inductive type definition for each non-terminal. It uses
 1031 the language name and the non-terminal names to generate the inductive type and the constructors.
 1032 For example, below is the definition generated for the `term` non-terminal.

```
1033 (define-datatype stlc-term : Type
1034   (Var->stlc-term : Var → stlc-term)
1035   (stlc-val->stlc-term : stlc-value → stlc-term)
1036   (stlc-lambda : stlc-type stlc-term → stlc-term)
1037   (stlc-app : stlc-term stlc-term → stlc-term)
1038   (stlc-cons : stlc-term stlc-term → stlc-term)
1039   (stlc-let : stlc-term stlc-term → stlc-term))
```

1040 It has a constructor for each kind of term. In addition, a conversion constructor is produced
 1041 for references to other non-terminals, e.g., `Var->stlc-term`. Internally, `define-language` uses an
 1042 intermediate Racket data structure to represent the grammar, which may then be converted to `CUR`,
 1043 `Coq`, `LaTeX`, and other outputs. Since extensions are supported linguistically, programmers may use
 1044 `OLLY` forms alongside normal `CUR` code, rather than switch to an external tool.

1045 `OLLY` demonstrates how “dependent types as macros” supports domain-specific modeling—here,
 1046 the domain is programming language theory. By starting from macros, the proof assistant is
 1047 extensible with domain-specific support *by default*, rather than as an after thought. We can tailor
 1048 all aspects of the proof assistant, from the object theory to the syntax, to our domain.

1049

1050 5.2 A Tactic Language

1051 Tactic systems are a popular addition to proof assistants to enable interactive, command-based
 1052 construction of proof terms; some proof assistants even feature multiple tactic languages. A tactic
 1053 system also exercises many interesting features of the elaboration system—they require pre-type-
 1054 checking-time general purpose computation, traversal and pattern matching of object language
 1055 terms, interesting data structures in the elaboration system for manipulating proof states, an API
 1056 to the object language in order to type check and evaluate the terms while constructing proofs,
 1057 interactivity, and syntactic integration into the language. A macro system, such as Racket’s, provides
 1058 all but the API to the object language, but by developing the type system as macros, we get this
 1059 meta-language API to the object language by construction.

1060 For a simple example, in [Figure 23](#), we present a hypothetical library of propositional logic, along
 1061 with a tactic, `tauto`, that automatically finds a term that proves the given type. The tactic is simply
 1062 a macro that traverses and pattern matches on terms. It uses the pattern combinators generated
 1063 by `define-datatype`, e.g., `~True`, along with the backtracking inherent in the matching algorithm, to
 1064 concisely specify the proof search. For more complex tactics, however, we require slightly more
 1065 than plain macros, e.g., to help maintain proof state and track intermediate theorems.

1066 Thus, we create `NTAC`, a tactic language for `CUR`. The `NTAC` tactic system builds on the basic idea
 1067 of tactics as macros, but uses a zipper data structure to allow navigation of the proof term and
 1068 to track the program context. Rather than plain macros, `NTAC` tactics are host-language (Racket)
 1069 functions over this zipper that are executed by the macro system during elaboration. While the
 1070 details of `NTAC`’s navigation and construction of proof terms are not particularly novel as far as
 1071 tactic systems go, the use of macros allowed us to easily develop the tactic system and integrate it
 1072 into `CUR`.

1073 To demonstrate how `NTAC` integrates into `CUR`, here is an `NTAC` proof for a trivial theorem.

```
1074 #lang cur
1075 (require cur/ntac)
1076 (ntac (forall (A : Type) (a : A) A)
1077   (by-intros A a))
```

tactic-eg

1078

```

1079 (define-type False : Type)
1080
1081 (define-datatype True : Type
1082   [I : True])
1083
1084 (define-datatype And [X : Type] [Y : Type] → Type
1085   [conj : [x : X] [y : Y] → (And X Y)])
1086
1087 (define-datatype Or [X : Type] [Y : Type] → Type
1088   [or-introL : [x : X] → (Or X Y)]
1089   [or-introR : [y : Y] → (Or X Y)])
1090
1091
1092
1093
1094
1095
1096
1097
1098
1099
1100
1101
1102
1103
1104
1105
1106
1107
1108
1109
1110
1111
1112
1113
1114
1115
1116
1117
1118
1119
1120
1121
1122
1123
1124
1125
1126
1127

```

```

(define-m tauto
  [(tauto ~True) I]
  [(tauto (~And X Y))
   #:with x (tauto #'X)
   #:with y (tauto #'Y)
   (conj X Y x y)]
  [(tauto (~Or X Y))
   #:with x (tauto #'X)
   (or-introL X Y x)]
  [(tauto (~Or X Y))
   #:with y (tauto #'Y)
   (or-introR X Y y)]
  [(~fail "no proof") _])

```

Fig. 23. A library for type-level propositions.

```

1093 (by-assumption))

```

The `ntac` form builds an expression given an initial goal, *e.g.*, the polymorphic identity type, and a tactic script. It is similar to Coq’s `Goal`, which introduces an anonymous goal that can be solved using an `Ltac` script. Unlike `Goal`, however, `ntac` produces an expression, meaning it can be used in any expression position in CUR, not just as a top-level command. This is the natural design for `ntac`, since macros are naturally extensions to the *expression* language. This example uses the `by-intros` tactic, which takes arguments representing names to bind as assumptions in the local proof environment. Then we conclude the proof with `by-assumption`, which takes no arguments and searches the local environment for a term that matches the current goal. We can create a `define-theorem` definition to assign a name to an NTAC script, so it may be used to help with another proof:

```

1105 #lang cur
1106 (define-theorem id (forall (A : Type) (a : A) A)
1107   (by-intros A a)
1108   (by-assumption))

```

tactic-eg

A definition (`define-theorem name goal script ...`) is essentially syntactic sugar for (`define name (ntac goal script ...)`).⁹ Implementationwise, NTAC builds a proof tree data structure in the metalanguage—*i.e.*, Racket—that can represent partial CUR terms, *e.g.*, terms with holes. The tree is a single *goal* node, representing a completely unknown term of some type. Each tactic is an operation that manipulates this tree, usually by changing the current goal node to a larger subtree that represents a partial term with a new subgoal. Once there are no goals left, the tree is translated into a CUR term. The `ntac` form then, simply applies tactics to this tree:

```

1116 #lang cur
1117 (define-m (ntac goal tactic ...) (ntac-interp goal (list tactic ...)))
1118

```

ntac

The `ntac` form calls `ntac-interp`, which constructs an initial proof tree from goal and runs each tactic, in order, on the tree. If the resulting tree contains unsolved goals, it raises an error; otherwise, it converts the tree to a CUR term. As an example, we define the `intro` tactic below.

```

1122 #lang cur
1123 (define-tactic (intro name ctx prooftree)
1124   (define goal (get-current-goal prooftree))
1125   (ntac-match goal

```

ntac

⁹In fact, it is somewhat more complicated to support resugaring and rewriting.

```

1128 [(~forall (x : A) B)
1129 (make-apply-node
1130 goal
1131 (make-ctxt-node
1132 (λ (old-ctxt) (dict-set old-ctxt name A))
1133 (make-new-goal (subst name x B)))
1134 (λ (body-proof) (λ (name : A) body-proof)))]))

```

1135 This tactic introduces a new variable, name when the goal has the shape `(forall (x : A) B)`.
 1136 The tactic extracts the current goal from the proof tree and pattern matches on it. When the
 1137 goal matches, we construct a new proof tree using `make-apply-node`. This node describes how to
 1138 construct a term of type `(forall (x : A) B)`, if it is provided a term body of type B. It then creates a
 1139 new subgoal from the type B (with references to `x` in B replaced with `name`) with `make-new-goal`. This
 1140 new goal is wrapped with a `make-ctxt-node`, which adds `name` bound to A in the environment. When
 1141 this proof tree is complete, the `ntac-interp` function will apply the Racket function on the last
 1142 line, `(λ (body-proof) (λ (name : A) body-proof))`, to the completed subtree that has replaced
 1143 `(make-new-goal B)`. Observe how the coloring denotes the use of quasiquotation to build up the
 1144 term. More specifically, the outer lambda, produces the inner *syntax object* lambda, except it embeds
 1145 the variable name as its parameter, and `body-proof` as the body.

1146 By using a flexible macro system as the basis for our tactic system, we can even equip user-defined
 1147 tactics with features like interactivity, as shown in Figure 24 (left). Specifically, the interactive
 1148 tactic uses the `print` tactic to print the proof state, then starts a read-eval-print-loop (REPL).
 1149 Figure 24 (right) shows an example interactive session. The REPL repeatedly reads in a command
 1150 and runs it via `run-tactic`; when it sees `quit`, it returns the proof tree.

```

1151
1152 (define-tactic (interactive pt)          -----
1153 (print pt)                             (forall (A : Type) (forall (a : A) A))
1154 (match (read-syntax)                    > (by-intro A)
1155 [(quit) pt]                             A : Type
1156 [tactic                                  -----
1157   (interactive (run-tactic pt tactic))]) (forall (a : A) A)
1158                                         ....
1159 (ntac (forall (A : Type) (a : A) A)      > by-assumption
1160   interactive)                          Proof complete.
1161                                         > (quit) ; => < procedure >

```

1161 Fig. 24. (Left) Implementation and use of interactivity tactic; (Right) An interactive proof session.

1162 Our macros-based approach makes it simple to develop a prototype proof assistant capable of
 1163 realistic proofs. To demonstrate this, we used CUR and NTAC to implement the exercises for several
 1164 chapters of the *Software Foundations* curriculum [Pierce et al. 2018], totaling several thousand lines
 1165 of proof scripts.¹⁰ Table 1 presents a list of tactics available in NTAC. The rewrite tactics, where
 1166 most of the work lies, support two versions of the equality type: Coq’s default Paulin-Mohring
 1167 equality, and Martin-Löf’s original version. When applied to quantified hypotheses, these tactics
 1168 will try to automatically instantiate the theorems with a basic search over the current proof state.

```

1170      assert      intros      assumption      simpl      obvious      destruct
1171      induction  reflexivity  interactive  rewriteL  rewriteR      print

```

1172 Table 1. List of tactics available in NTAC.

1175 ¹⁰<https://www.github.com/stchang/macrotypes>, <https://www.github.com/wilbowma/cur>

6 FUTURE WORK

A interesting next step is to experiment with typed tactic DSLs, à la Mtac [Ziliani et al. 2013]’s design. We conjecture that use of Racket’s #lang framework and TURNSTILE will make it straightforward to do so. We also plan to experiment with automation and integration with other tools, e.g., by calling out to solvers or even using the foreign-function interface during expansion.

Resugaring is another direction for future work. Since type checking is interleaved with macro expansion, some effort is required to prevent abstraction leaks that could expose users to elaborated syntax. For example CUR and NTAC use resugaring during interactive proof sessions. The current resugaring approach is rather ad-hoc, however, but recent advances [Pombrio and Krishnamurthi 2018] could help improve this part of the language, and apply generally to macro-based approaches. Another solution could be to stage expansion to avoid the need for resugaring at all. We are experimenting with “stop lists”, i.e., finer-grained knobs for controlling expansion, so that we can maintain the benefits of type checking with macros, but do not expand beyond abstractions that the user cares about during the process.

7 RELATED WORK

Much has been written about **implementing basic dependent types** [Altenkirch et al. 2010; Augustsson 2007; Bauer 2012; Löh et al. 2010; Weirich 2014]. All of these tutorials, however, start from scratch and typically stop short of a practical language. For example, most manually deal with type environments and rely on deBruijn indices for α -equality. Further, they do not include practical features such as user-defined inductive datatypes, and they are not easily extensible with sugar, interactivity, or other companion DSLs that programmers typically need to use with their dependently-typed language. In contrast, we show how our macros-based approach enables both rapid creation of a core dependently-typed language, and scales to a realistic full-spectrum proof assistant with user-defined inductive datatypes and extensible notation.

Extending proof assistants is an active area of research. For example, some dependently-typed languages have explored adding metaprogramming [Brady and Hammond 2006; Christiansen 2014; Devriese and Piessens 2013; Ebner et al. 2017] capabilities. This feature, however, typically requires extending the core language. Other languages like Coq often require writing extensions in a less integrated manner, e.g, programming plugins with OCaml and then linking it with other language binaries. With our approach, we use the metaprogramming facilities inherited from the host language, and thus get to write extensions in a more linguistically supported manner.

One of the most common extensions created by dependently-typed programmers, using many clever methods, is **new tactic languages** [Gonthier and Mahboubi 2010; Gonthier et al. 2011; Krebbers et al. 2017; Malecha and Bengtson 2016]. This suggests that (1) the ability to create domain-specific tactic languages is critical, and (2) that linguistic support for creation of such DSLs would be well received. While we have yet to conduct a thorough comparison of all tactic languages and their implementations, we conjecture that our macros-based approach could accommodate many of them in a convenient manner. For example, there has been recent exploration of typed tactic languages Beluga [Pientka 2008], Mtac [Ziliani et al. 2013], and VeriML [Stampoulis and Shao 2010]. We conjecture that it would be straightforward to add a typed tactic language to CUR using our macros-based approach. This could be done either by utilizing TURNSTILE, or using CUR’s reflection API to use CUR as it’s own meta-language, following the approach of Lean [Ebner et al. 2017] or Typed Template Coq [Anand et al. 2018].

1226 8 CONCLUSION

1227 To fully leverage the power of dependent types, programmers should be able to quickly develop
1228 their own dependently-typed DSLs with just the right expressiveness for their domain. Further,
1229 these DSLs should be easily extensible with any new notation or companion DSLs that might be
1230 required to make the language practical for realistic programming. We have demonstrated that a
1231 macros-based approach to building dependently-typed DSLs satisfies this criteria. For future work,
1232 we hope to leverage the rapid prototyping benefit of our approach to experiment with new type
1233 theory features like extensions for parametricity modalities and homotopy type theory, and to
1234 leverage the extensibility to further explore other domain-specific dependent type applications.
1235

1236

1237

1238

1239

1240

1241

1242

1243

1244

1245

1246

1247

1248

1249

1250

1251

1252

1253

1254

1255

1256

1257

1258

1259

1260

1261

1262

1263

1264

1265

1266

1267

1268

1269

1270

1271

1272

1273

1274

REFERENCES

- 1275
1276 2016. RFC: Const-dependent Type System. (June 2016). <https://github.com/rust-lang/rfcs/pull/1657>
- 1277 2017. RFC: The pi type trilogy. (Feb. 2017). <https://github.com/rust-lang/rfcs/issues/1930>
- 1278 Thorsten Altenkirch, Nils Anders Danielsson, Andres Löf, and Nicolas Oury. 2010. $\Pi\Sigma$: Dependent Types Without the Sugar. In *Proceedings of the 10th International Conference on Functional and Logic Programming (FLOPS'10)*. 40–55.
- 1279 Abhishek Anand, Simon Boulier, Cyril Cohen, Matthieu Sozeau, and Nicolas Tabareau. 2018. Towards Certified Meta-Programming with Typed Template-Coq. (2018). www.irif.fr/~sozeau/research/publications/drafts/Towards_Certified_Meta-Programming_with_Typed_Template-Coq.pdf
- 1280 Leif Andersen, Stephen Chang, and Matthias Felleisen. 2017. Super 8 Languages for Making Movies (Functional Pearl). *Proc. ACM Program. Lang.* 1, ICFP, Article 30 (Aug. 2017), 29 pages. <https://doi.org/10.1145/3110274>
- 1283 Lennart Augustsson. 2007. Simpler, Easier! (2007). <http://augustss.blogspot.ru/2007/10/simpler-easier-in-recent-paper-simply.html>
- 1284 Gilles Barthe, Benjamin Grégoire, and Santiago Zanella-Béguelin. 2009. Formal Certification of Code-based Cryptographic Proofs. In *Symposium on Principles of Programming Languages (POPL)*. <https://doi.org/10.1145/1480881.1480894>
- 1286 Andrej Bauer. 2012. How to Implement Dependent Type Theory. (2012). <http://math.andrej.com/2012/11/08/how-to-implement-dependent-type-theory-i/>
- 1288 Benjamin Beurdouche, Franziskus Kiefer, and Tim Taubert. 2017. Verified cryptography for Firefox 57. (July 2017). <https://blog.mozilla.org/security/2017/09/13/verified-cryptography-firefox-57/>
- 1289 Edwin Brady and Kevin Hammond. 2006. Dependently Typed MetaProgramming. In *7th Symposium on Trends in Functional Programming*.
- 1292 Edwin C. Brady. 2005. *Practical Implementation of a Dependently Typed Functional Programming Language*. Ph.D. Dissertation. University of Durham.
- 1293 Stephen Chang, Alex Knauth, and Ben Greenman. 2017. Type Systems As Macros. In *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages*. 694–705.
- 1294 Adam Chlipala. 2011. Mostly-automated Verification of Low-level Programs in Computational Separation Logic. In *Proceedings of the 32nd ACM SIGPLAN Conference on Programming Language Design and Implementation*. 234–245.
- 1296 Adam Chlipala, Benjamin Delaware, Samuel Duchovni, Jason Gross, Clément Pit-Claudel, Sorawit Suriyakarn, Peng Wang, and Katherine Ye. 2017. The End of History? Using a Proof Assistant to Replace Language Design with Library Design. In *Summit on Advances in Programming Languages (SNAPL)*. <https://doi.org/10.4230/LIPICs.CVIT.2016.23>
- 1299 David Raymond Christiansen. 2014. Type-Directed Elaboration of Quasiquotations: A High-Level Syntax for Low-Level Reflection. In *of the 26nd 2014 International Symposium on Implementation and Application of Functional Languages (IFL 2014)*. ACM, 1.
- 1300 Thierry Coquand and Gérard P. Huet. 1988. The Calculus of Constructions. *Inf. Comput.* 76, 2/3 (1988), 95–120. [https://doi.org/10.1016/0890-5401\(88\)90005-3](https://doi.org/10.1016/0890-5401(88)90005-3)
- 1303 Ryan Culpepper and Matthias Felleisen. 2010. Fortifying macros. In *Proceeding of the 15th ACM SIGPLAN International Conference on Functional Programming*. 235–246.
- 1304 N.G. de Bruijn. 1991. Telescopic Mappings in Typed Lambda-Calculus. *Information and Computation* 91, 2 (1991), 189–204.
- 1306 Dominique Devriese and Frank Piessens. 2013. Typed Syntactic Meta-programming. In *of the 18th ACM SIGPLAN International Conference on Functional Programming (ICFP 2013)*. 73–86.
- 1307 Peter Dybjer. 1994. Inductive families. *Formal Aspects of Computing* 6, 4 (01 Jul 1994), 440–465.
- 1308 Gabriel Ebner, Sebastian Ullrich, Jared Roesch, Jeremy Avigad, and Leonardo de Moura. 2017. A metaprogramming framework for formal verification. *Proceedings of the ACM on Programming Languages (PACMPL)* 1, ICFP (2017), 34:1–34:29. <https://doi.org/10.1145/3110278>
- 1310 Matthew Flatt. 2016. Binding As Sets of Scopes. In *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. 705–717.
- 1312 Matthew Flatt, Ryan Culpepper, David Darais, and Robert Bruce Findler. 2012. Macros That Work Together: Compile-time Bindings, Partial Expansion, and Definition Contexts. 22, 2 (March 2012), 181–216. <https://doi.org/10.1017/S0956796812000093>
- 1314 Matthew Flatt and PLT. 2010. *Reference: Racket*. Technical Report PLT-TR-2010-1. PLT Design Inc. <http://racket-lang.org/tr1/>
- 1316 Martin Fowler and Rebecca Parsons. 2010. *Domain-Specific Languages*. Addison-Wesley.
- 1317 Georges Gonthier and Assia Mahboubi. 2010. An introduction to small scale reflection in Coq. *Journal of Formalized Reasoning* 3, 2 (2010), 95–152. <https://hal.inria.fr/inria-00515548>
- 1318 Georges Gonthier, Beta Ziliani, Aleksandar Nanevski, and Derek Dreyer. 2011. How to Make Ad Hoc Proof Automation Less Ad Hoc. In *Proceedings of the 16th ACM SIGPLAN International Conference on Functional Programming (ICFP '11)*. ACM, New York, NY, USA, 163–175. <https://doi.org/10.1145/2034773.2034798>
- 1320 Robbert Krebbers, Amin Timany, and Lars Birkedal. 2017. Interactive Proofs in Higher-order Concurrent Separation Logic. In *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages (POPL 2017)*. ACM, New
- 1322
1323

- 1324 York, NY, USA, 205–217. <https://doi.org/10.1145/3009837.3009855>
- 1325 Andres Löh, Conor McBride, and Wouter Swierstra. 2010. A Tutorial Implementation of a Dependently Typed Lambda
1326 Calculus. *Fundam. Inform.* 102, 2 (2010), 177–207.
- 1327 Gregory Malecha and Jesper Bengtson. 2016. *Programming Languages and Systems: 25th European Symposium on Program-*
1328 *ming, ESOP 2016, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2016, Eindhoven,*
1329 *The Netherlands, April 2–8, 2016, Proceedings.* Springer Berlin Heidelberg, Berlin, Heidelberg, Chapter Extensible and
Efficient Automation Through Reflective Tactics, 532–559. https://doi.org/10.1007/978-3-662-49498-1_21
- 1330 Per Martin-Löf. 1975. An intuitionistic theory of types: Predicative part. *Studies in Logic and the Foundations of Mathematics*
1331 80 (1975), 73–118.
- 1332 Conor McBride. 2000. *Dependently Typed Functional Programs and Their Proofs.* Ph.D. Dissertation. University of Edinburgh,
1333 UK. <http://hdl.handle.net/1842/374>
- 1334 Conor McBride and James McKinna. 2004. The View from the Left. *J. Funct. Program.* 14, 1 (2004), 69–111. <https://doi.org/10.1017/s0956796803004829>
- 1335 Frank Pfenning and Christine Paulin-Mohring. 1989. Inductively Defined Types in the Calculus of Constructions. In
1336 *Mathematical Foundations of Programming Semantics, 5th International Conference, Tulane University, New Orleans,*
1337 *Louisiana, USA, March 29 - April 1, 1989, Proceedings (Lecture Notes in Computer Science),* Michael G. Main, Austin Melton,
1338 Michael W. Mislove, and David A. Schmidt (Eds.), Vol. 442. Springer, 209–228. <https://doi.org/10.1007/BFb0040259>
- 1339 Brigitte Pientka. 2008. A Type-theoretic Foundation for Programming with Higher-order Abstract Syntax and First-
1340 class Substitutions. In *Proceedings of the 35th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming*
Languages (POPL '08). 371–382.
- 1341 Benjamin C. Pierce, Arthur Azevedo de Amorim, Chris Casinghino, Marco Gaboardi, Michael Greenberg, Cătălin Hrițcu,
1342 Vilhelm Sjöberg, and Brent Yorgey. 2018. *Logical Foundations.* Electronic textbook.
- 1343 Benjamin C. Pierce and David N. Turner. 1998. Local Type Inference. In *Proceedings of the 25th ACM SIGPLAN-SIGACT*
Symposium on Principles of Programming Languages. 252–265.
- 1344 Justin Pombrio and Shriram Krishnamurthi. 2018. Inferring Type Rules for Syntactic Sugar. In *Proceedings of the 39th ACM*
1345 *SIGPLAN Conference on Programming Language Design and Implementation.* 812–825.
- 1346 Peter Sewell, Francesco Zappa Nardelli, Scott Owens, Gilles Peskine, Thomas Ridge, Susmit Sarkar, and Rok Strniša. 2007. Ott:
1347 Effective Tool Support for the Working Semanticist. In *of the 12th ACM SIGPLAN International Conference on Functional*
1348 *Programming (ICFP 2007).* ACM, New York, NY, USA, 1–12. <https://doi.org/10.1145/1291151.1291155>
- 1349 Antonis Stampoulis and Zhong Shao. 2010. VeriML: Typed Computation of Logical Terms Inside a Language with Effects.
1350 In *of the 15th ACM SIGPLAN International Conference on Functional Programming (ICFP 2010).* 333–344.
- 1351 Stephanie Weirich. 2014. Pi Forall: notes from OPLSS. (2014). <https://github.com/sweirich/pi-forall>
- 1352 Stephanie Weirich, Antoine Voizard, Pedro Henrique Azevedo de Amorim, and Richard A. Eisenberg. 2017. A Specification
1353 for Dependent Types in Haskell. *Proceedings of the ACM on Programming Languages (PACMPL)* 1, ICFP (Aug. 2017).
<https://doi.org/10.1145/3110275>
- 1354 Hongwei Xi. 2007. Dependent ML: An approach to practical programming with dependent types. *Journal of Functional*
1355 *Programming* 17, 2 (2007), 215–286. <https://doi.org/10.1017/S0956796806006216>
- 1356 Brent A. Yorgey, Stephanie Weirich, Julien Cretin, Simon L. Peyton Jones, Dimitrios Vytiniotis, and José Pedro Magalhães.
1357 2012. Giving Haskell a promotion. In *Types in Language Design and Implementation (TLDI).* <https://doi.org/10.1145/2103786.2103795>
- 1358 Beta Ziliani, Derek Dreyer, Neelakantan R Krishnaswami, Aleksandar Nanevski, and Viktor Vafeiadis. 2013. Mtac: A Monad
1359 for Typed Tactic Programming in Coq. In *Proceedings of the 18th ACM SIGPLAN International Conference on Functional*
Programming (ICFP 2013). ACM, New York, NY, USA, 87–100. <https://doi.org/10.1145/2500365.2500579>
- 1360 Jean Karim Zinzindohoué, Karthikeyan Bhargavan, Jonathan Protzenko, and Benjamin Beurdouche. 2017. HACL*: A Verified
1361 Modern Cryptographic Library. In *Conference on Computer and Communications Security, (CCS).* <https://doi.org/10.1145/3133956.3134043>

1363 A STYLE AND GLOSSARY

1364 This section summarizes the various style choices used in the paper, and also presents some macro
1365 terminology that may help with reading the paper.

1366 A.1 Macros Glossary

- 1367 • A *syntax object* is the Racket AST representation. It’s a tree of symbols, accompanied by
1368 context information such as source locations, the program’s binding structure, and even
1369 arbitrary user-specified metadata.

- *Syntax patterns* deconstruct syntax objects, binding *pattern variables* to different parts of a syntax object, which are themselves syntax objects.
- *Syntax templates* construct syntax objects. They may reference pattern variables, whose corresponding syntax object gets embedded into the constructed syntax object.
- *Syntax properties* are key-value pairs associated with syntax object nodes. We use syntax properties to propagate type and other meta information about a piece of syntax.

A.2 Style

In the paper, to help readability, we often stylize code with colors or abbreviations. This section summarizes a few of our choices.

- `syntax pattern` positions, which deconstruct syntax objects, are highlighted with green. A macro definition's input is frequently a pattern position.
- `syntax template` positions, which construct syntax objects, are in blue. A macro definition's output is frequently a syntax template position. When there are nested positions, e.g., for a macro defining macro, we might instead enclose the syntax template with `□`, so that the coloring of any nested pattern and syntax template positions are not obscured.
- The name being defined (e.g., a macro, typerule, function, etc.) is always underlined.
- In a pattern, literal symbols to match are marked with bold (latex `pmb`).
- Pattern variables representing elaborated, i.e., expanded, syntax objects are marked with an overline.
- *Syntax classes*, which additionally constrain the shape of pattern variables, are written with a superscript.

B MACRO SYSTEM FEATURES

This section, beginning with table 2, summarizes the macro system features used in the paper, and their availability in other macro systems. While, to our knowledge, Racket is the only language that combines all the features needed for “type systems as macros”, many other popular languages are rapidly adopting the same features in their macro systems.

B.1 Procedural macros

Procedural macros are syntax transformations defined in a general-purpose language supporting arbitrary computations. They are essential to allow arbitrary type-checking logic during expansion.

B.2 Quasiquote and syntax pattern matching

With quasiquote, macros construct their expansion using syntax matching the textual form of the language, plus escapes for inserting computed elements. Similarly, macros using syntax pattern matching deconstruct the AST of macro invocations using syntax matching the textual form of the language. We use pattern matching and quasiquote to give type rule macros relatively readable syntax, even without the Turnstile DSL layer.

B.3 Extensible pattern matching

Turnstile types expand into a common internal representation which enables simple implementations of type equality and substitution. Pattern matching this internal representation is verbose. We use Racket's *pattern expanders*, e.g., in section 3.3, to abstract away the internal representation and create simple pattern forms for each type constructor.

	Racket	Lisp	Clojure	Scala	Rust	Julia	Elixir	Crystal
1422								
1423	X	X	X	X		X	X	
1424	X	X	X	X	X	X	X	X
1425	X			X	X			
1426	X			X ₁				
1427	X				X	X	X	
1428	X	X ₂	X ₃	X ₄				
1429	X	X	X					X
1430	X	X	~ ₅					
1431	X	X	~ ₅					
1432	X							
1433	(1) Referred to as “extractor macros”							
1434	(2) Referred to as “property lists”							
1435	(3) Referred to as “metadata”							
1436	(4) Referred to as “attachments”							
1437	(5) Clojure does not have symbol macros or expansion in the environment of local macro							
1438	definitions, but they can be added as a library. See https://github.com/clojure/tools.macro .							

Table 2. Macro system features used in the paper, and their availability in other macro systems.

1440

1441

1442

1443

1444

B.4 Automatic hygiene

1445

1446

1447

1448

1449

1450

1451

1452

B.5 Syntax properties

1453

1454

1455

1456

1457

1458

B.6 Macro-defining macros and identifier macros

1459

1460

1461

1462

1463

1464

1465

B.7 Local expansion

1466

1467

1468

1469

1470

Macro systems with local expansion allow macros to request the expansion of subexpressions, including in the environment of local variable and macro bindings. We use local expansion to typecheck subexpressions and access their inferred type while typechecking the parent expression. Racket’s approach to local expansion is discussed in Flatt et al. [2012].

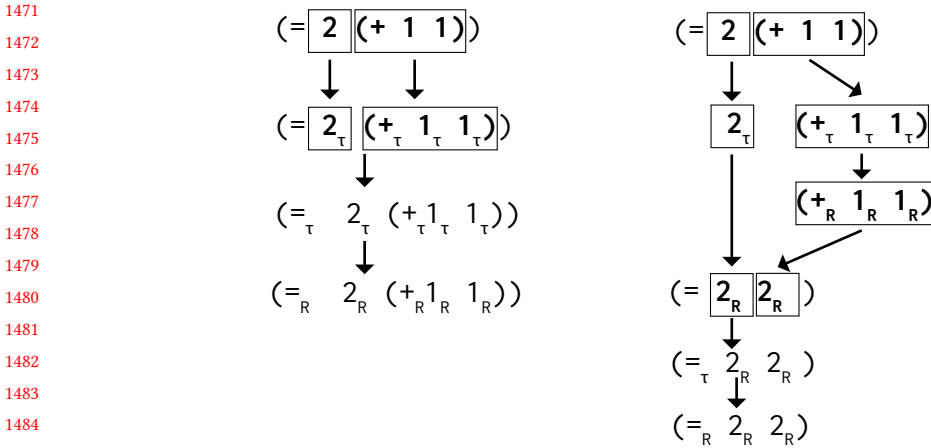


Fig. 25. Example of the interleaving

B.8 Interposition points

Racket’s expander automatically inserts hooks at various points, e.g., the `#:app` macro at each function application and the `#:datum` macro at each literal datum (like a number), to allow customizing the behavior of these and other constructs such as modules and the REPL. By redefining these macros, types as macros can add typechecking to the expansion of these syntactic elements.

B.9 Interleaving Semantics

Combined, the above features allow us to interleave macro expansion, type checking, and evaluation, and to communicate information across each stage effectively. In Figure 25 (left), we give an example of a term as it proceeds through the interleaved macro expansion and type checking process, while in Figure 25 (right), we show an example of the interleaving of macro expansion, type checking, and evaluation. Specifically, we show the expansion of an equality type $(= 2 (+ 1 1))$. Without type level reduction, this elaborates each subexpression into the type-annotated version (denoted by the τ subscript), before generating the fully elaborated run-time representations (denoted by the R subscript). The type-annotated versions represent the output of type-rule macros, while the run-time representations represent the output of the reduction-rule macros. Without dependent types, the type-annotated and run-time representations are the same. However, once we support dependent types and reduction rules as macros, they are different and require the reflection process described in Section 4.

Notice that, on the right of the figure with reduction during type-checking, we end up with a run-time subterm that must be interleaved with type-annotated terms. Supporting this, particularly when any term (such as a function defined in another module) can be evaluated at expansion time, is the key challenge in the type systems as macros approach, which we solve using many of the above macro system features.