Cubical Agda: A Dependently Typed Programming Language with Univalence and Higher Inductive Types

ANDREA VEZZOSI, IT University of Copenhagen, Denmark
ANDERS MÖRTBERG, Stockholm University, Sweden and Carnegie Mellon University, USA
ANDREAS ABEL, Chalmers and Gothenburg University, Sweden

Proof assistants based on dependent type theory provide expressive languages for both programming and proving within the same system. However, all of the major implementations lack powerful extensionality principles for reasoning about equality, such as function and propositional extensionality. These principles are typically added axiomatically which disrupts the constructive properties of these systems. Cubical type theory provides a solution by giving computational meaning to Homotopy Type Theory and Univalent Foundations, in particular to the univalence axiom and higher inductive types. This paper describes an extension of the dependently typed functional programming language Agda with cubical primitives, making it into a full-blown proof assistant with native support for univalence and a general schema of higher inductive types. These new primitives make function and propositional extensionality as well as quotient types directly definable with computational content. Additionally, thanks also to copatterns, bisimilarity is equivalent to equality for coinductive types. This extends Agda with support for a wide range of extensionality principles, without sacrificing type checking and constructivity.

CCS Concepts: • **Theory of computation → Type theory**; • **Software and its engineering → Functional languages**; **Patterns**.

Additional Key Words and Phrases: Cubical Type Theory, Univalence, Higher Inductive Types, Dependent Pattern Matching

ACM Reference Format:
Andrea Vezzosi, Anders Mörtberg, and Andreas Abel. 2019. Cubical Agda: A Dependently Typed Programming Language with Univalence and Higher Inductive Types. *Proc. ACM Program. Lang.* 3, ICFP, Article 87 (August 2019), 29 pages. https://doi.org/10.1145/3341691

1 INTRODUCTION

A core idea in programming and mathematics is *abstraction*: the exact details of how an object is represented should not affect its abstract properties. In other words, the implementation details should not matter. This is exactly what the principle of univalence captures by extending the equality on the universe of types to incorporate equivalent types.[1] This then gives a form of abstraction,

[1]For the sake of this introduction, "equivalent" may be read as "isomorphic". In Homotopy Type Theory (HoTT), *isomorphism* coincides with equivalence for *sets* (in the sense of HoTT). Equivalence for *types* in general is a refinement of the concept of isomorphism.

Authors' addresses: Andrea Vezzosi, Computer Science, IT University of Copenhagen, Rued Langgaards Vej 7, Copenhagen, 2300, Denmark, avez@itu.dk; Anders Mörtberg, Department of Mathematics, Stockholm University, Kräftriket Hus 6, Stockholm, 10691, Sweden, anders.mortberg@math.su.se, School of Computer Science, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA, 15213, USA, mortberg@cmu.edu; Andreas Abel, andreas.abel@gu.se, Chalmers and Gothenburg University, Department of Computer Science and Engineering, Rännvägen 6b, Gothenburg, 41296, Sweden.

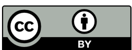

This work is licensed under a Creative Commons Attribution 4.0 International License.
© 2019 Copyright held by the owner/author(s).
2475-1421/2019/8-ART87
https://doi.org/10.1145/3341691

or invariance up to equivalence, in the sense that equivalent types will share the same structures and properties. The fact that equality is *proof relevant* in dependent type theory is the key to enabling this; the data of an equality proof can store the equivalence and transporting along this equality should then apply the function underlying the equivalence. In particular, this allows programs and properties to be transported between equivalent types, hereby increasing modularity and decreasing code duplication. A concrete example are the equivalent representations of natural numbers in unary and binary format. In a univalent system it is possible develop theory about natural numbers using the unary representation, but compute using the binary representation, and as the two representations are equivalent they share the same properties.

The principle of univalence is the major new addition in Homotopy Type Theory and Univalent Foundations (HoTT/UF) [Univalent Foundations Program 2013]. However, these new type theoretic foundations add univalence as an *axiom* which disrupts the good constructive properties of type theory. In particular, if we transport addition on binary numbers to the unary representation we will not be able to compute with it as the system would not know how to reduce the univalence axiom. Cubical type theory [Cohen et al. 2018] addresses this problem by introducing a novel representation of equality proofs and thereby providing computational content to univalence. This makes it possible to constructively transport programs and properties between equivalent types. This representation of equality proofs has many other useful consequences, in particular function and propositional extensionality (pointwise equal functions and logically equivalent propositions are equal), and the equivalence between bisimilarity and equality for coinductive types [Vezzosi 2017].

Dependently typed functional languages such as `Agda` [2018], `Coq` [2019], `Idris` [Brady 2013], and `Lean` [de Moura et al. 2015], provide rich and expressive environments supporting both programming and proving within the same language. However, the extensionality principles mentioned above are not available out of the box and need to be assumed as axioms just as in HoTT/UF. Unsurprisingly, this suffers from the same drawbacks as it compromises the computational behavior of programs that use these axioms, and even make subsequent proofs more complicated.

So far, cubical type theory has been developed with the help of a prototype `Haskell` implementation called `cubicaltt` [Cohen et al. 2015], but it has not been integrated into one of the main dependently typed functional languages. Recently, an effort was made, using `Coq`, to obtain effective transport for restricted uses of the univalence axiom [Tabareau et al. 2018], because, as the authors mention, *"it is not yet clear how to extend [proof assistants] to handle univalence internally"*.

This paper achieves this, and more, by making `Agda` into a cubical programming language with native support for univalence and higher inductives types (HITs). We call this extension `Cubical Agda` as it incorporates and extends cubical type theory. In addition to providing a fully constructive univalence theorem, `Cubical Agda` extends the theory by allowing proofs of equality by copatterns, HITs as in Coquand et al. [2018] with nested pattern matching, and interval and partial pre-types. This paper aims to provide a formal account of the extensions to the language of `Agda` and its type-checking algorithm needed to accommodate the new features. In particular, as it required the most care, we will dedicate a large portion of this paper to the handling of pattern matching.

Contributions: The main contribution of this paper is the implementation of `Cubical Agda`; a fully fledged proof assistant with constructive support for univalence and HITs. This makes a variety of extensionality principles provable and we show how these can be used for programming and proving in Sect. 2. We explain how `Agda` was extended to support cubical type theory (Sect. 3),

in particular we describe how some primitive notions of cubical type theory are internalized as pre-types, like the interval (Sect. 3.1), partial elements and cubical subtypes (Sect. 3.3). The technical contributions are:

- We extend cubical type theory by records and coinductive types (Sect. 3.2.2).
- We add support for a general schema of HITs, and extend the powerful dependent pattern-matching of Agda to also support pattern matching on HITs (Sect. 4).
- We describe an optimization to the algorithm for transport in Glue types, (Sect. 5), which gives a simpler proof of the univalence theorem compared to Cohen et al. [2018] (Sect. 5.2).

Using the optimization to transport for Glue types we discuss an improved canonicity theorem for cubical type theory with HITs (Sect. 6). The paper finishes with some concluding remarks and an overview of related and future work (Sect. 7).

Acknowledgements: The authors are grateful to the anonymous reviewer's for their helpful comments on earlier versions of the paper. The second author is also grateful for the feedback from everyone in the Agda learning group at CMU and especially to Loïc Pujet for porting the HIT integers from `cubicaltt` and Zesen Qian for the formalization of set quotients. We would also like to thank Martín Escardó for encouraging us to develop the `agda/cubical` library and everyone who has contributed to it since then.

2 PROGRAMMING AND PROVING IN CUBICAL AGDA

In this section, we show some examples of how the new cubical features in Agda enable interesting and useful ways for both programming and proving in dependent type theory. No expert knowledge of HoTT/UF is assumed. By using univalence and other ideas from HoTT/UF we can:

1. Transfer programs and proofs between equivalent types. (Sect. 2.1.1)
2. Prove properties for proof oriented datatypes using computation oriented ones. (Sect. 2.1.2)
3. Treat bisimilar elements of coinductive types as equal. (Sect. 2.2)
4. Define and reason about quotient types. (Sect. 2.3)
5. Represent topological spaces as datatypes and reason about them synthetically. (Sect. 2.4)

The examples are taken from the open-source library `agda/cubical` hosted at https://github.com/agda/cubical.

2.1 Unary and Binary Numbers

An example of two equivalent types that are well-suited for different tasks are unary and binary numbers. The unary representation is useful for proving because of the direct induction principle and the binary representation is much better for computation as it is exponentially more compact. By utilizing computational univalence we can transfer results between the two representations in a convenient way, this lets us get the best of both worlds without having to duplicate results.

The unary numbers, ℕ, are built into Agda and are inductively generated by the constructors zero and suc (successor). We encode binary numbers as:

```
data Bin : Set where
  bin0   : Bin
  binPos : Pos → Bin

data Pos : Set where
  pos1 : Pos
  x0   : Pos → Pos
  x1   : Pos → Pos
```

```
six : Bin
six = binPos (x0 (x1 (pos1)))
```

A binary number is hence either zero (bin0) or a positive binary number represented as a list of zeroes and ones with no trailing zeroes. Least significant bits come first (little-endian format), thus, the number 6 is binary 011. This way every number has a unique binary representation and it is straightforward to write maps to and from the unary representation (Bin→ℕ and ℕ→Bin) with proofs that they cancel (ℕ→Bin→ℕ and Bin→ℕ→Bin). This means that the two types ℕ and Bin are isomorphic which implies that they are *equivalent*, in the sense of having *"contractible fibers"* following the terminology of Voevodsky [2015] and Univalent Foundations Program [2013]. Spelled out, a map $f : A \to B$ is an equivalence if the preimage of any point in B is a singleton type.

Given types A and B we write $A \simeq B$ for the type of equivalences between them, the univalence theorem[2] then states that

$$(A \equiv B) \simeq (A \simeq B).$$

In particular there is a function ua : $A \simeq B \longrightarrow A \equiv B$, sending a proof that two types are equivalent to an equality between the types. We use ua to turn the equivalence of ℕ and Bin into an equality.

```
ℕ≃Bin : ℕ ≃ Bin
ℕ≃Bin = isoToEquiv (iso ℕ→Bin Bin→ℕ Bin→ℕ→Bin ℕ→Bin→ℕ)

ℕ≡Bin : ℕ ≡ Bin
ℕ≡Bin = ua ℕ≃Bin
```

In fact, the equality in ℕ≡Bin is not really a normal type-theoretic equality à la Martin-Löf (in the sense of being inductively generated from constructor refl for reflexivity), but rather a *path* equality. The core idea of HoTT/UF is the close correspondence between proof-relevant equality, as in type theory, and paths, as in topology. The idea that equality correspond to paths is taken very literally in cubical type theory; by adding a primitive interval type I, paths in a type A can be represented as functions $\mathsf{I} \longrightarrow A$. Iterating these function types lets us represent squares, cubes, and hypercubes; making the type theory *cubical*.

The interval I has two distinguished endpoints i0 and i1. Since paths are functions, we introduce them using λ-abstraction and eliminate them using function application; by applying a path to i0 we get its left endpoint and by applying it to i1 we get the right one. We often want to specify the endpoints of a path (or, more generally, the boundary of a cube) in its type; in `Cubical Agda`, there is a special primitive for this:

```
PathP : (A : I ⟶ Set ℓ) ⟶ A i0 ⟶ A i1 ⟶ Set ℓ
```

we introduce these paths by lambda abstractions like so, $\lambda i \longrightarrow t$: PathP A $t[\mathsf{i0} / i]$ $t[\mathsf{i1} / i]$, provided that $t : A\ i$ for $i : \mathsf{I}$. Consequently, we can apply p : PathP A a_0 a_1 to an $r : \mathsf{I}$ to obtain $p\ r : A\ r$. Also, no matter how f is given, we have that p i0 reduces to a_0 and p i1 reduces to a_1.

The PathP types should be thought of as heterogeneous equalities since the two endpoints are in different types; this is similar to the dependent paths in HoTT [Univalent Foundations Program 2013, Sect. 6.2]. We can define homogeneous non-dependent path equality in terms of PathP as follows:

```
_≡_ : {A : Set ℓ} ⟶ A ⟶ A ⟶ Set ℓ
_≡_ {A = A} x y = PathP (λ _ ⟶ A) x y
```

In the previous definition, the syntax $\{A = A\}$ tells `Agda` to bind the hidden argument A of type _≡_ (first A) to a variable A (second A) that can be used on the right hand side. Note also that some

[2] As the univalence "axiom" is provable in `Cubical Agda` we refer to it as the *univalence theorem*.

definitions are polymorphic in universe level ℓ which is implicitly universally quantified. Further, from now we will omit the explicit quantification over $\{A : \mathsf{Set}\ \ell\}$.

Viewing equalities as functions out of the interval allows us to reason elegantly about equality; for instance, the constant path represents a proof of reflexivity.

```
refl : {x : A} → x ≡ x
refl {x = x} = λ i → x
```

We can also directly apply a function to a path in order to prove that dependent functions respect path equality, as shown in the definition of cong below. Simply by computation cong satisfies some new definitional equalities compared to the corresponding definition for the inductive equality type à la Martin-Löf [1975]. For instance, cong is functorial by definition: we can prove congId and congComp by plain reflexivity (refl).

```
cong : ∀ {B : A → Set ℓ} (f : (a : A) → B a) {x y} (p : x ≡ y) → PathP (λ i → B (p i)) (f x) (f y)
cong f p i = f (p i)

congId : ∀ {x y : A} (p : x ≡ y) → cong (λ a → a) p ≡ p
congId p = refl

congComp : ∀ (f : A → B) (g : B → C) {x y} (p : x ≡ y) → cong (g ∘ f) p ≡ cong g (cong f p)
congComp f g p = refl
```

Path types also let us prove new things that are not provable in standard Agda with Martin-Löf propositional equality. For example, function extensionality, stating that pointwise equal functions are equal themselves, has an extremely simple proof:

```
funExt : {f g : A → B} → ((x : A) → f x ≡ g x) → f ≡ g
funExt p i x = p x i
```

The proof of function extensionality for dependent and n-ary functions is equally direct. Since funExt is a *definable* notion in `Cubical Agda`, it is, in contrast to Martin-Löf Type Theory, not an axiom. This means that it has computational content: it simply swaps the arguments to p.

The facts that paths can be manipulated as functions and that we have heterogeneous path types makes many equality proofs simpler compared to the corresponding proofs in standard `Agda` or HoTT. For instance, the equality of second projections of dependent pair types is a simple instance of cong:

```
Σ-eq₂ : ∀ {A : Set} {B : A → Set} {p q : Σ[ x ∈ A ] (B x)} → (e : p ≡ q) →
          PathP (λ i → B (e i .fst)) (p .snd) (q .snd)
Σ-eq₂ = cong snd
```

The corresponding result in regular type theory has to be stated with the homogeneous notion of equality using transport, making equality in dependent pair types notoriously difficult to work with.

2.1.1 Univalent Transport. One of the key properties of type theoretic equality is *transport*.

```
transport : A ≡ B → A → B
transport p a = transp (λ i → p i) i0 a
```

This is defined using another primitive of `Cubical Agda` called transp. It is a generalization of the regular transport principle which lets us specify where the transport is the identity function. In particular, when the second argument to transp is i1 it will reduce to a, which let us prove that transp $A\ r\ a$ is always path equal to a (cf. addp later). A consequence is that whenever we have an

equivalence $e : A \simeq B$ we have that transport (λ i ⟶ F (ua e i)) is an equivalence as well. This ability to lift equivalences through arbitrary type operators F is an easily overlooked benefit of a language with computational univalence.

The substitution principle is obtained as an instance of transport.

```
subst : (B : A ⟶ Set ℓ) {x y : A} ⟶ x ≡ y ⟶ B x ⟶ B y
subst B p b = transport (λ i ⟶ B (p i)) b
```

Function subst invokes transport with a proof of $B\ x \equiv B\ y$; this proof $\lambda\ i \longrightarrow B\ (p\ i)$ is an inlining of cong, stating that families B respect equality p.

After this digression about path types, let us revisit the ℕ≡Bin path. Using transport, we can transfer zero along ℕ≡Bin, and since univalence is a theorem with computational content in `Cubical Agda`, this will reduce to bin0. In contrast, if we were working in a system with axiomatic univalence, we could still define ℕ≡Bin, but the transport of zero along that equality would be *stuck*; the system would not know how to transport with the ua constant.

Having computational univalence lets us do a lot more than just transporting constructors. We can for example transport the *addition* function from binary to unary numbers in order to get a more efficient addition function on unary numbers. However, as we do not want to prove properties about the more complex binary addition function, we instead transport the unary one with its properties to binary.

```
_+Bin_ : Bin ⟶ Bin ⟶ Bin
_+Bin_ = transport (λ i ⟶ ℕ≡Bin i ⟶ ℕ≡Bin i ⟶ ℕ≡Bin i) _+_
```

In this case, the path that we transport along is between the function types ℕ ⟶ ℕ ⟶ ℕ and Bin ⟶ Bin ⟶ Bin. This way, we obtain an addition function on binary numbers and the fact that ua has computational content lets us run it:

```
_ : (binPos (x0 (x0 pos1))) +Bin (binPos pos1) ≡ binPos (x1 (x0 pos1))
_ = refl
```

In order to reduce the left hand side `Cubical Agda` will convert all of the arguments to the unary representation, add them using _+_ and then convert the result back to binary. This is of course a very inefficient way of adding binary numbers. However, we could now prove that the efficient addition function on binary numbers is equal to _+Bin_ using function extensionality and as the functions are equal they will share the same properties. The main reason for defining _+Bin_ as above is that it lets us transport results about the unary addition function. For example, we transport the proof of associativity as follows:

```
addp   : PathP (λ i ⟶ ℕ≡Bin i ⟶ ℕ≡Bin i ⟶ ℕ≡Bin i) _+_ _+Bin_
addp i = transp (λ j ⟶ ℕ≡Bin (i ∧ j) ⟶ ℕ≡Bin (i ∧ j) ⟶ ℕ≡Bin (i ∧ j)) (~ i) _+_

+Bin-assoc : (m n o : Bin) ⟶ m +Bin (n +Bin o) ≡ (m +Bin n) +Bin o
+Bin-assoc
  = transport (λ i ⟶ (m n o : ℕ≡Bin i) ⟶ addp i m (addp i n o) ≡ addp i (addp i m n) o)
              +-assoc
```

In addp we utilize the interval operators *minimum* (_∧_) and *reversal* (~_); further, `Cubical Agda` features the *maximum* operator (_∨_). The intuition is that elements of I correspond to points in the real unit interval $[0, 1]$. The _∧_ and _∨_ operations take the minimum and maximum of i, j : I while the reversal operation computes $1 - i$. These operations satisfy the laws of a *De Morgan algebra*. This means, for one, that the min/max operations form a bounded distributive lattice, with i0 and i1 as bottom and top elements. Further, the reversal is a De Morgan involution, so, for

instance, $\sim (i \wedge j) = \sim i \vee \sim j$. Note that this is still *not* a Boolean algebra, since $i \wedge \sim i = \mathsf{i0}$ and $i \vee \sim i = \mathsf{i1}$ are not valid for points of the unit interval, except for the endpoints.

Let us assert the well-typedness of addp: When i is i0, the first argument to transp in addp is constantly ℕ ⟶ ℕ ⟶ ℕ, since i0 ∧ i is then just i0 and ℕ≡Bin i0 reduces to ℕ. The second argument ~ i becomes i1 so that the left endpoint of the path is _+_, exploiting that transp (...) is the identity function when applied to i1. On the other hand, when i is i1, then addp i reduces to transp (λ j ⟶ ℕ≡Bin j⟶ℕ≡Bin j⟶ℕ≡Bin j) i0 _+_ which is exactly the definition of _+Bin_. This establishes that addp indeed constitutes a path from _+_ to _+Bin_. Note that the path type of addp is heterogeneous as the two addition functions have different types.

The desired result is then obtained by transporting the proof that unary addition is associative along a path from

$$(m\ n\ o : \mathbb{N}) \longrightarrow m + (n + o) \equiv (m + n) + o$$

to

$$(m\ n\ o : \mathsf{Bin}) \longrightarrow m \mathbin{+\mathsf{Bin}} (n \mathbin{+\mathsf{Bin}} o) \equiv (m \mathbin{+\mathsf{Bin}} n) \mathbin{+\mathsf{Bin}} o\,.$$

Another very useful consequence of being able to transport proofs between types is that we can prove some result by computation for binary numbers and then transport the proof to the unary representation where the computation might have been infeasible.

2.1.2 Univalent Program and Data Refinements. Sometimes concrete computations are necessary in proofs. For example one could imagine a situation where one needs to check an equality between two terms that are expensive to compute like:

$$2^{20} \cdot 2^{10} = 2^{5} \cdot 2^{15} \cdot 2^{10}\,.$$

As this is part of a proof it is likely that one is using a unary representation of natural numbers, however this makes it impossible to verify the above equation by computation. One way to resolve this would be to redo the formalization using binary numbers, but this could involve a complete rewrite of the formalization. Another alternative would be to use algebraic manipulations to prove the above equality manually, however sometimes this is not feasible as the computation might be very complicated.

This kind of issue can be resolved by what we call univalent *program and data refinements* following Cohen et al. [2013]. As binary numbers are equivalent to unary numbers we can prove the property for binary numbers by computation and then transport the proof to unary numbers. To do this we define a "doubling structure" in which we can express the above equation, and instantiate with unary and binary numbers. We omit the concrete definitions, but as expected the doubling function is of linear complexity for unary numbers and constant for binary.

```
record Double (A : Set) : Set where
  constructor doubleStruct
  field
    double : A ⟶ A
    elt : A

Doubleℕ : Double ℕ
Doubleℕ = doubleStruct doubleℕ 1024

DoubleBin : Double Bin
DoubleBin = doubleStruct doubleBin bin1024
```

The equality between binary and unary numbers lifts to an equality of doubling structures. We omit the details of this definition as they are quite technical, however the whole definition is only 7 lines of code so it is not particularly difficult to write once one is familiar with all of the features of `Cubical Agda`. Note that, unlike for sucℤ, we are using the inverse of the above path from unary to binary numbers as we are going the opposite direction compared to the previous examples.

```
DoubleBin≡Doubleℕ : PathP (λ i → Double (Bin≡ℕ i)) DoubleBin Doubleℕ
```

We can now formulate the equation that we originally wanted to verify. We put this in a record and use copattern matching when proving it for DoubleBin so that Agda won't eagerly try to unfold the unary instance of it (more about this in the next section) when we transport the proof over to unary numbers along the equality of doubling structures.

```
doubles : {A : Set} (D : Double A) → ℕ → A → A
doubles D n x = iter n (double D) x

record propDouble {A : Set} (D : Double A) : Set where
  field
    proof : doubles D 20 (elt D) ≡ doubles D 5 (doubles D 15 (elt D))

propDoubleBin : propDouble DoubleBin
proof propDoubleBin = refl

propDoubleℕ : propDouble Doubleℕ
propDoubleℕ = transport (λ i → propDouble (DoubleBin≡Doubleℕ i)) propDoubleBin
```

The fact that equivalences of types lift to equivalences of structures is called the *structure identity principle* in HoTT/UF [Univalent Foundations Program 2013, Sect. 9.8]. This has been formalized in Agda for algebraic structures and isomorphisms in Coquand and Danielsson [2013]. Combining this with univalence lets us lift equalities of types to equalities of structures on these types.

2.2 Univalence for Coinductive Types

Coinductive types allow the direct manipulation of infinite structures without breaking the consistency of the language. However, in their treatment in Coq and Agda, reasoning about them was impeded by the inability to prove two elements equal whenever they have the same unfolding, rather than when they are the same by definition [McBride 2009]. `Cubical Agda` solves this by exploiting the interaction between path and projection copatterns [Abel et al. 2013].

The prototypical example of a coinductive type is infinite streams, which can be declared in Agda as a coinductive record type with two fields: .head and .tail. A function returning a stream is then defined by explaining how it computes when applied to the projections. For example, here we define mapS which applies a function to every element of a stream.

```
record Stream (A : Set) : Set where
  coinductive; constructor _,_
  field
    head : A
    tail  : Stream A

mapS : (A → B) → Stream A → Stream B
mapS f xs .head = f (xs .head)
mapS f xs .tail  = mapS f (xs .tail)
```

The result is that mapS f xs by itself will not unfold further and the termination checker is happy to accept this definition as productive as it always reaches a weak head normal form in finite time when applied to projections. As shown in the following proof of the identity law for mapS, Cubical Agda extends the notion of productivity by allowing the same recursion pattern also for paths between streams.

```
mapS-id : (xs : Stream A) → mapS (λ x → x) xs ≡ xs
mapS-id xs i .head = xs .head
mapS-id xs i .tail  = mapS-id (xs .tail) i
```

To define a path between (mapS ($\lambda\, x \to x$) xs) and xs we introduce an interval variable i and then are left to define (mapS-id xs i) of type Stream A, so we can proceed by copatterns and corecursion.

To convince ourselves that mapS-id defines the required path, we note that if mapS-id xs is supposed to be a path between mapS ($\lambda\, x \to x$) xs and xs, then the type of ($\lambda\, i \to$ mapS-id $xs\, i$ .head) should be mapS ($\lambda\, x \to x$) xs .head ≡ xs .head. This type in turn reduces to xs .head ≡ xs .head, by definition of mapS, and so the constant path suffices. A similar reasoning applies to the tail case, this time using the tail clause of mapS to realize that we need a path between mapS ($\lambda\, x \to x$) (xs.tail) and xs.tail, which we provide with a corecursive call. We give a systematic description of how we compute such *boundary* constraints from the left hand sides of clauses in Sect. 4.

More generally, we can define bisimilarity as a coinductive record and show that two bisimilar streams are equal.

```
record _≈_ (xs ys : Stream A) : Set where
  coinductive
  field
    ≈head : xs .head ≡ ys .head
    ≈tail  : xs .tail  ≈ ys .tail

bisim : ∀ {xs ys : Stream A} → xs ≈ ys → xs ≡ ys
bisim xs≈ys i .head = xs≈ys .≈head i
bisim xs≈ys i .tail  = bisim (xs≈ys .≈tail) i
```

Finally we note that bisim is actually an equivalence, and so equality of streams is indeed bisimilarity.

```
path≡bisim : ∀ {xs ys : Stream A} → (xs ≡ ys) ≡ (xs ≈ ys)
```

The agda/cubical library contains the complete proof, see https://github.com/agda/cubical/Cubical/Codata/Stream/Properties.agda, as well as a proof of the universal property of indexed M-types [Ahrens et al. 2015], see https://github.com/agda/cubical/Cubical/M-types.agda

2.3 Quotient Types as Higher Inductive Types

Another major new addition in HoTT are *higher inductive types* (HITs). These are datatypes in which we can specify "higher" constructors representing non-trivial paths of the type (representing identifications of elements), in addition to the normal "point" constructors. These types enable many interesting constructions in type theory, in particular quotient types.

The addition of HITs in systems like Agda or Coq is usually done by postulating their existence, however this suffer from the same issues, in terms of computation, as postulating the univalence axiom. Cubical Agda extends the datatype declarations of Agda to also support a general schema of HITs so that it is not necessary to postulate their existence axiomatically.

In this section, we illustrate how we can use HITs to define quotient types in `Cubical Agda`. The first example of a quotient type is a very simple encoding of the integers—while this example might seem rather trivial it will help us showcase quite a few interesting possibilities of working with HITs. The second example is a general formulation of quotient types and set quotients.

2.3.1 Integers as a HIT. The integers are often represented as $\mathbb{N} + \mathbb{N}$, however this has the drawback that there are two zeroes (inl 0 and inr 0). This is usually resolved by shifting one of them by 1 (so that for example inl 0 represents -1, etc.), however this can easily lead to confusion and off-by-one errors. A better solution is to identify the two zeroes. This can be achieved with the following HIT.

```
data ℤ : Set where
  pos : (n : ℕ) → ℤ
  neg : (n : ℕ) → ℤ
  posneg : pos 0 ≡ neg 0

sucℤ : ℤ → ℤ
sucℤ (pos n)       = pos (suc n)
sucℤ (neg zero)    = pos 1
sucℤ (neg (suc n)) = neg n
sucℤ (posneg i)    = pos 1
```

This type is similar to $\mathbb{N} + \mathbb{N}$, except that there is also a *path* constructor posneg which identifies the two zeroes. For an element i of the interval type, posneg i is an integer which reduces to pos 0 in case $i = \mathsf{i0}$ and neg 0 in case $i = \mathsf{i1}$. These so-called *boundary conditions* of posneg have to be respected by any function on $\mathbb{Z}$. For example, the successor function on $\mathbb{Z}$ can be written as above. The final case maps the path constructor constantly to pos 1 which is accepted by `Cubical Agda` as the following equations hold definitionally:

$$\mathsf{suc}\mathbb{Z}\ (\mathsf{pos}\ 0) = \mathsf{suc}\mathbb{Z}\ (\mathsf{neg}\ 0) = \mathsf{pos}\ 1.$$

It is direct to define an inverse to sucℤ (i.e., the predecessor function) and hence get an equivalence from $\mathbb{Z}$ to $\mathbb{Z}$ which, combined with ua, gives a non-trivial path from $\mathbb{Z}$ to $\mathbb{Z}$. Transporting along this path applies the successor function.

```
sucPathℤ : ℤ ≡ ℤ
sucPathℤ = isoToPath (iso sucℤ predℤ sucPredℤ predSucℤ)
```

We can also define addition and prove that addition with a fixed number is an equivalence, however this takes a bit of work as we need to define subtraction and prove that they are inverse. Using univalence we can take a shortcut and define an alternative addition function so that addition with a fixed number is automatically an equivalence. Consider the following path equality that composes sucPathℤ with itself n times.

```
addEq : ℕ → ℤ ≡ ℤ
addEq zero    = refl
addEq (suc n) = addEq n · sucPathℤ
```

The · operation is binary composition of paths which will be discussed in detail in Sect. 3.4. Similarly we can define a path composing the predecessor path with itself n times. By transporting along these paths we get an addition function.

```
addℤ : ℤ → ℤ → ℤ
addℤ m (pos n)      = transport (addEq n) m
addℤ m (neg n)      = transport (subEq n) m
addℤ m (posneg _) = m
```

By using that transporting along a path is an equivalence we get that addition by a fixed number is an equivalence.

```
isEquivAddℤ : (m : ℤ) → isEquiv (λ n → addℤ n m)
isEquivAddℤ (pos n)      = isEquivTransport (addEq n)
isEquivAddℤ (neg n)      = isEquivTransport (subEq n)
isEquivAddℤ (posneg i) = isEquivTransport refl
```

2.3.2 *General Quotient Types and Set Quotients.* Quotient types have so far been a possibility in Agda and other dependently typed programming languages only axiomatically. Here we show how to define them as a HIT in Cubical Agda.

A first attempt is the following definition.

```
data _/_ (A : Set ℓ) (R : A → A → Set ℓ) : Set ℓ where
  [_] : A → A / R
  eq  : (a b : A) → R a b → [ a ] ≡ [ b ]
```

This type has a constructor for mapping elements of A to the quotient by R and an equality identifying the image of each pair of related elements. However this is not exactly what we want as the resulting quotient type might have a too complex notion of equality. For example, if we use this construction to quotient Unit by the total relation, then we will get a type with a point [tt] and an identification of this point with itself. We would expect this to be equivalent to Unit, however it is in fact equivalent to the HIT circle that we will discuss in Sect. 2.4.1. As Unit and the circle do not satisfy the same properties, they are not equivalent; the loop space of Unit, here the type of paths from the only element to itself, is contractible while the loop space of the circle is $\mathbb{Z}$ [Univalent Foundations Program 2013].

We get the expected notion of quotients if we switch to *set* quotients. We add another higher constructor that eliminates all of the higher-dimensional structure from the quotient type, in other words, we *set truncate* the type.

```
data _/_ (A : Set ℓ) (R : A → A → Set ℓ) : Set ℓ where
  [_]   : A → A / R
  eq    : (a b : A) → (r : R a b) → [ a ] ≡ [ b ]
  trunc : (x y : A / R) → (p q : x ≡ y) → p ≡ q
```

This makes the quotient into a *recursive* HIT as the trunc constructor quantifies over elements of the type that we are constructing. It forces the quotient to be a set in the sense of satisfying the *uniqueness of identity proofs* (UIP) principle, in other words, any two proofs of equality of members of A / R are equal. Thanks to trunc, we can prove the universal property of set quotients:

```
setQuotUniversal : {A B : Set ℓ} {R : A → A → Set ℓ} → isSet B →
                   (A / R → B) ≃ (Σ[ f ∈ (A → B) ] ((a b : A) → R a b → f a ≡ f b))
```

This says that maps out of the set quotient is the same as maps sending related elements to equal elements in the quotient (assuming that the image satisfies UIP). If we furthermore assume that R is an equivalence relation then the set quotients are effective, in the sense that if [a] ≡ [b] then

also $R\ a\ b$. As an interesting application, we could for example define the positive fractions as a quotient of ℕ × ℕ by relating (n_1, d_1) and (n_2, d_2) if $(n_1 \cdot (1 + d_2) \equiv n_2 \cdot (1 + d_1))$.

2.4 Synthetic Homotopy Theory in Cubical Agda

One of the main applications of HITs in HoTT is the ability to reason *synthetically* about topological spaces inside type theory. This means that we can define topological spaces (like spheres, tori, etc.) as datatypes and reason about them using functional programming. The semantic justification for this is the standard model in Kan simplicial sets, a combinatorial representation of topological spaces [Kapulkin and Lumsdaine 2012; Lumsdaine and Shulman 2017]. In this section, we illustrate how we can do synthetic homotopy theory in `Cubical Agda` by proving that the torus is equivalent to two circles.

2.4.1 The Torus and Two Circles. We can define the circle and the torus as the following higher inductive types.

```
data S¹ : Set where
  base : S¹
  loop : base ≡ base

data Torus : Set where
  point  : Torus
  line1  : point ≡ point
  line2  : point ≡ point
  square : PathP (λ i → line1 i ≡ line1 i) line2 line2
```

The idea is that the circle, S¹, is generated by a base point and a non-trivial path constructor loop connecting base to itself. The Torus on the other hand also has a base point with two non-trivial path constructors connecting it to itself and a square relating the two paths. This square can be illustrated by:

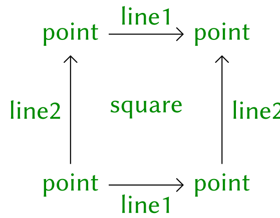

The idea is that the square constructor identifies line2 with itself over an identification of line1 with itself. This has the effect of identifying the opposite sides of the square, making it into a torus (imagine the square being a sheet of soft paper that one folds so that the opposite sides match).

As we demonstrate in the following, a torus is equivalent to the product of two circles.

```
t2c : Torus → S¹ × S¹
t2c point        = (base   , base)
t2c (line1 i)    = (loop i , base)
t2c (line2 j)    = (base   , loop j)
t2c (square i j) = (loop i , loop j)

c2t : S¹ × S¹ → Torus
c2t (base , base)     = point
c2t (loop i , base)   = line1 i
c2t (base , loop j)   = line2 j
c2t (loop i , loop j) = square i j
```

The functions back and forth are directly definable by pattern-matching. As a consequence, proving that they are mutually inverse is trivial, and we get an equality between the two types.

```
c2t-t2c : (t : Torus) → c2t (t2c t) ≡ t
c2t-t2c point          = refl
c2t-t2c (line1 _)      = refl
c2t-t2c (line2 _)      = refl
c2t-t2c (square _ _)   = refl

t2c-c2t : (p : S¹ × S¹) → t2c (c2t p) ≡ p
t2c-c2t (base   , base)   = refl
t2c-c2t (base   , loop _) = refl
t2c-c2t (loop _ , base)   = refl
t2c-c2t (loop _ , loop _) = refl

Torus≡S¹×S¹ : Torus ≡ S¹ × S¹
Torus≡S¹×S¹ = isoToPath (iso t2c c2t t2c-c2t c2t-t2c)
```

This is a rather elementary result in topology. However, it had a surprisingly non-trivial proof in HoTT because of the lack of definitional computation for higher constructors [Licata and Brunerie 2015; Sojakova 2016]. With the additional definitional computation rules of Cubical Agda this proof is now almost entirely trivial.

2.4.2 Further Synthetic Homotopy Theory in Cubical Agda. The agda/cubical library contains several further results from synthetic homotopy theory. For instance, we have a direct proof that the fundamental group of the circle is $\mathbb{Z}$, inspired by Licata and Shulman [2013]. Combined with the above characterization of the torus it proves that the fundamental group of the Torus is $\mathbb{Z} \times \mathbb{Z}$. The fact that univalence and HITs compute in Cubical Agda lets us then compute winding numbers of iterated loops around the circle and torus.

The library also features more substantial results: a proof that $\mathbb{S}^3$, i.e., the four dimensional sphere, is equivalent to the join of two circles, and a proof that the total space of the Hopf fibration is $\mathbb{S}^3$. We also ave a definition of the "Brunerie number": a number $n \in \mathbb{Z}$ such that $\pi_4(\mathbb{S}^3) \simeq \mathbb{Z}/n\mathbb{Z}$.[3] However, despite considerable efforts we have not been able to reduce n to a normal form yet (even though the absolute value of the expected result is just 2 as proved by Brunerie [2016]).

3 MAKING AGDA CUBICAL

In the remainder of the paper we will describe how Agda was extended to become cubical. The key additions to Agda are:

(1) The interval and path types (Sect. 3.1).
(2) Generalized transport, transp (Sect. 3.2).
(3) Partial elements (Sect. 3.3).
(4) Homogeneous composition, hcomp (Sect. 3.4).
(5) Higher inductive types (Sect. 4).
(6) Glue types (Sect. 5).

We have already discussed the first two points in some detail in the examples, however, the implementation of the transp operation is especially interesting as it is what makes Cubical Agda compute. This operation is defined by cases on the type formers of Agda; a contribution of this paper is the extension of these to types that are present in Agda but not covered by Cohen et al. [2018], namely, record and coinductive types. Furthermore, the way the hcomp operation works in Cubical Agda differs from Coquand et al. [2018] in a subtle way which enables us to optimize the transp operation for Glue types. These types are what allows us to give computational content to univalence. By optimizing how their composition operation compute we obtain simpler and more

[3]For details see https://github.com/agda/cubical/blob/master/Cubical/Experiments/Brunerie.agda#L268.

efficient proofs of univalence. As discussed in Sect. 6, this also has metatheoretical consequences for the canonicity theorem for HITs.

3.1 The Interval and Path Types

The first thing `Cubical Agda` add is an interval type I. Then, we add the PathP types that behave like function types out of the interval, but with fixed endpoints. Note that the interval in `Cubical Agda` is not inductively defined, so we cannot pattern match on it. This follows from the intuition that path types are *continuous* functions from the interval into a space so that they cannot provide arbitrarily different results for i0 and i1.

3.2 Generalized Transport

The next key thing that `Cubical Agda` adds is the generalized transport operation.

$$\mathsf{transp} : (A : \mathsf{I} \to \mathsf{Set}\,\ell) \to \mathsf{I} \to A\,\mathsf{i0} \to A\,\mathsf{i1}$$

Given a type line $A : \mathsf{I} \to \mathsf{Set}\,\ell$ and an element at end A i0, the transp operation gives an element at A i1, the other end of the line. This is generalized compared to regular transport in the sense that transp lets us specify where it should behave as the identity function. In particular there is an additional side condition to be satisfied for transp $A\ r\ a$ to type check, which is that A should be a constant function whenever the constraint $r = \mathsf{i1}$ is satisfied. When r is equal to i1 the transp function will compute as the identity function,

$$\mathsf{transp}\ A\ \mathsf{i1}\ a = a.$$

and this would not be sound in general if A was allowed to be a more complex line that is non-constant when $r = \mathsf{i1}$. In case $r = \mathsf{i0}$ there is nothing to check, thus, transp A i0 a is well-formed for any A, as in the definition of transport $A\ a$.

Internally, the transp operation computes differently for each of the type formers of Agda. We will show how this works in the special case of transport, but the general transp operation is not much more complicated. For a detailed type theoretic presentation of these definitions see Huber [2017]. This formulation of the computation rules for cubical type theory is based on a variation of the comp operation of Cohen et al. [2018] that was introduced in Coquand et al. [2018] in order to support HITs.

3.2.1 Function Types. Given two type lines $A\ B : \mathsf{I} \to \mathsf{Set}$ we seek to transport a function $f : A\ \mathsf{i0} \to B\ \mathsf{i0}$ to a function $A\ \mathsf{i1} \to B\ \mathsf{i1}$. To this end, we compose backward transport $A\ \mathsf{i1} \to A\ \mathsf{i0}$ along A, function f, and forward transport $B\ \mathsf{i0} \to B\ \mathsf{i1}$ along B.

```
transportFun : (A B : I → Set) → (A i0 → B i0) → (A i1 → B i1)
transportFun A B f = transport (λ i → B i) ∘ f ∘ transport (λ i → A (~ i))
```

By evaluating transport (λ i → A i → B i) f we can see that the definition of transportFun A B f is definitionally the same as the internal definition for how transp computes in `Cubical Agda`.

```
transportFunEq : (A B : I → Set) → (f : A i0 → B i0) →
                   transportFun A B f ≡ transport (λ i → A i → B i) f
transportFunEq A B f = refl
```

The definition for dependent functions is very similar, except that some extra work is required to correct the type in the outer transport. This definition clarifies why we need to consider transp and not just the simpler transport operation.

```
transportPi : (A : I → Set) (B : (i : I) → A i → Set)
            → ((x : A i0) → B i0 x)
```

$\longrightarrow ((x : A\ \mathsf{i1}) \longrightarrow B\ \mathsf{i1}\ x)$
$\mathsf{transportPi}\ A\ B\ f = \lambda\ (x : A\ \mathsf{i1}) \longrightarrow \mathsf{transport}\ (\lambda\ j \longrightarrow B\ j\ (\mathsf{transp}\ (\lambda\ i \longrightarrow A\ (j \vee \sim i))\ j\ x))$
$(f\ (\mathsf{transport}\ (\lambda\ i \longrightarrow A\ (\sim i))\ x))$

If we would have used the same definition as for non-dependent functions the outer transport would have been ill-typed. The reason is that f has a dependent type, meaning that $f\ x'$ has type $B\ x'$ for $x' := \mathsf{transport}\ (\lambda\ i \longrightarrow A\ (\sim i))\ x$. The first argument of the outer transport must hence be a line between $B\ \mathsf{i0}\ x'$ and $B\ \mathsf{i1}\ x$. This line is constructed by abstracting over j and considering $B\ j\ (\mathsf{transp}\ (\lambda\ i \longrightarrow A\ (j \vee \sim i))\ j\ x)$. When j is i0 this is indeed $B\ \mathsf{i0}\ x'$ and when j is i1 this is $B\ \mathsf{i1}\ x$ by virtue of transp being the identity function when applied to i1.

3.2.2 Records and Coinductive Types. Since record types are a generalization of Σ-types, the transport operation is computed pointwise, i.e. independently for every field. The only subtlety is when the record is dependent, in which case a similar type correction has to be done as for dependent functions.

As coinductive types in Agda are just record types these are handled the same way. In this case, however, we have to consider the issue of productivity, which is taken care of by how transport for record types unfolds only when projected from. Analogously to the Stream example from Sect. 2.2, we have that $\mathsf{transport}\ (\lambda i \to \mathsf{Stream}\ A)\ xs$ will not reduce further, while if we apply .tail to it we get $\mathsf{transport}\ (\lambda i \to \mathsf{Stream}\ A)\ (xs\ .\mathsf{tail})$.[4] Such controlled unfolding generally leads to smaller normal forms so `Cubical Agda` adopts it for record types in general. The same kind of controlled unfolding is also implemented for other "negative" types like function and path types.

3.2.3 Datatypes. The transport operation for inductive datatypes without parameters, for instance, the natural numbers, is trivial as they cannot vary along the interval.

$$\mathsf{transport}\ (\lambda i \to \mathbb{N})\ x = x$$

This would not work for inductive types with parameters like the disjoint union $A + B$ for which the transport operation would need to reduce to the transport operation in A or B depending on the argument.

$\mathsf{transportSum} : (A\ B : \mathsf{I} \longrightarrow \mathsf{Set}) \longrightarrow A\ \mathsf{i0} + B\ \mathsf{i0} \longrightarrow A\ \mathsf{i1} + B\ \mathsf{i1}$
$\mathsf{transportSum}\ A\ B\ (\mathsf{inl}\ x) = \mathsf{inl}\ (\mathsf{transport}\ (\lambda\ i \longrightarrow A\ i)\ x)$
$\mathsf{transportSum}\ A\ B\ (\mathsf{inr}\ x) = \mathsf{inr}\ (\mathsf{transport}\ (\lambda\ i \longrightarrow B\ i)\ x)$

The transp operation for HITs is a bit more involved. It also computes by cases on the argument, but for the higher constructors some extra care has to be taken. In particular, complicated parameterized HITs, like pushouts, require additional endpoint corrections [Coquand et al. 2018, Sect. 3.3.5].

3.2.4 Path Types. For path types we will need a new operation to provide the computation rules for transport as we need some way to record the endpoints of the path after transporting it. Indeed, consider the following naïve definition:

$\mathsf{transportPath} : (A : \mathsf{I} \to \mathsf{Set})\ (x\ y : (i : \mathsf{I}) \to A\ i) \to x\ \mathsf{i0} \equiv y\ \mathsf{i0} \to x\ \mathsf{i1} \equiv y\ \mathsf{i1}$
$\mathsf{transportPath}\ A\ x\ y\ p = \lambda\ i \to \mathsf{transport}\ (\lambda\ j \to A\ j)\ (p\ i)$

This might look plausible as a definition, but the resulting path does not have the correct boundary. When i is i0, for instance, the left boundary is $\mathsf{transport}\ (\lambda\ j \longrightarrow A\ j)\ (x\ \mathsf{i0})$ and not just $x\ \mathsf{i1}$. Note that these elements are equal up to a path (using $\lambda\ k \longrightarrow \mathsf{transp}\ (\lambda\ j \longrightarrow A\ (j \vee k))\ k\ (x\ k)$), so what we

[4]Productivity for the case of dependent records then relies on the type of later fields only being able to depend on earlier fields.

need is a way to compose the result with this path in order to correct the endpoints. To do this we introduce the homogeneous composition operation (hcomp) that generalizes binary composition of paths to n-ary composition of higher dimensional cubes.

3.3 Partial Elements

In order to describe the homogeneous composition operation we need to be able to write partially specified n-dimensional cubes, i.e., cubes where some faces are missing. Given an element of the interval r : I there is a new primitive predicate IsOne r which represents the constraint r = i1. This comes with a proof 1=1 that i1 is in fact equal to i1, i.e., 1=1 : IsOne i1. The type IsOne (i ∨ ~ i) corresponds to the formula (i = i1) ∨ (i = i0) which represents the two endpoints of the line specified by i, so by considering formulas made out of more variables we can specify the boundary of cubes. The type IsOne r is also proof-irrelevant, meaning that any two of its elements are definitionally equal.

Building on IsOne we have extended `Cubical Agda` with partial cubical types, written Partial r A. The idea is that Partial r A is the type of cubes in A that are only defined when IsOne r holds.[5] Concretely, Partial r A is a special version of the function space IsOne r → A with a more extensional equality: two of its elements are considered judgmentally equal if they represent the same subcube of A. Concretely they are equal whenever they reduce to equal terms for all the possible assignment of variables that make r equal to i1.

Elements of these partial cubical types are introduced using pattern matching lambdas. For this purpose `Cubical Agda` supports a new form of patterns, here (i = i0) and (i = i1), that specify the cases where IsOne (i ∨ ~ i) is true. Similarly to pattern matching on an inductive family, some variables from the context might get refined, in this case i, even if otherwise we would not be able to pattern match on them.

```
partialBool : ∀ i → Partial (i ∨ ~ i) Bool
partialBool i = λ { (i = i0) → true ; (i = i1) → false }
```

The term partialBool should be thought of as a boolean with different values when i is i0 and when i is i1. This is hence just the endpoints of a line and there is no way to connect them, since true is not equal to false. The pattern-matching cases must match the interval expression in the type (under the image of IsOne) and if there are overlapping cases then they must agree up to definitional equality. Furthermore, IsOne i0 is actually absurd and lets us define an empty partial element (also known as an *"empty system"* [Cohen et al. 2018, Sect. 4.2]).

```
empty : Partial i0 A
empty = λ { () }
```

`Cubical Agda` also has cubical subtypes as in Cohen et al. [2018]; given A : Set ℓ, r : I and u : Partial r A we can form the type A [$r \mapsto u$]. A term v of this type is a term of type A that is definitionally equal to u when IsOne r is satisfied.[6] Any term u : A can be seen as a term of type A [$r \mapsto u$] that agrees with itself when IsOne r.

```
inS : {r : I} (a : A) → A [ r ↦ (λ _ → a) ]
```

We can also forget that a partial element agrees with u when IsOne r holds:

```
outS : {r : I} {u : Partial r A} → A [ r ↦ u ] → A
```

[5] Partial is somewhat analogous to constrained set $P \Rightarrow A = \{a \in A \mid P\}$ where P = IsOne r, only that the proof of P matters.

[6] In the set-theoretic analogy, A [$r \mapsto u$] $= \{a \in A \mid \text{if } r \text{ then } (a = u)\} \subseteq A$, given $u \in (r \Rightarrow A)$. We have $a \in A$ [$r \mapsto \lambda_ \to a$] always for $a \in A$.

We have that both outS (inS v) $= v$ and inS (outS v) $= v$ hold if well typed. Moreover, outS v will reduce to u 1=1 when $r = $ i1.

With all of this cubical infrastructure we can now describe the hcomp operation.

3.4 Homogeneous Composition

The homogeneous composition operation generalizes binary composition of paths so that we can compose multiple composable cubes.

```
hcomp : {r : I} (u : I → Partial r A) (u₀ : A [ r ↦ u i0 ]) → A
```

When calling hcomp u u_0, `Cubical Agda` makes sure that u_0 agrees with u i0 on r; this is specified in the type of u_0. The idea is that u_0 is the base and u specifies the sides of an open box where the side opposite of u_0 is missing. The hcomp operation then gives us the missing side opposite of u_0, which we refer to as the *lid* of the open box. For example binary composition of paths can be written as:

```
_·_ : {x y z : A} → x ≡ y → y ≡ z → x ≡ z
_·_ {x = x} p q i = hcomp (λ j → λ { (i = i0) → x ; (i = i1) → q j }) (inS (p i))
```

Pictorially we are given $p : x \equiv y$ and $q : y \equiv z$, and the composite of the two paths is obtained by computing the dashed lid at the top of the following square.

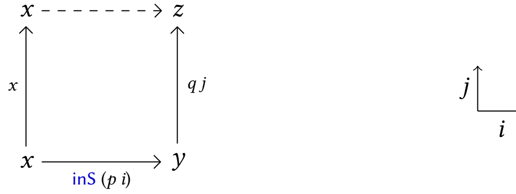

As we are constructing a path from x to z along i we have i : I in context and put inS (p i) as bottom. The direction j is abstracted in the first argument to hcomp and we use pattern-matching to specify the sides.

We can also define homogeneous filling of open boxes as

```
hfill : {r : I} (u : I → Partial r A) (u₀ : A [ r ↦ u i0 ]) → I → A
hfill {r = r} u u₀ i = hcomp (λ j → λ { (r = i1) → u (i ∧ j) 1=1 ; (i = i0) → outS u₀ }) (inS (outS u₀))
```

When i is i0 this is outS u_0 and when i is i1 it is hcomp (λ j → λ { (r = i1) → u j 1=1 } u_0 as the absurd faces (i0 = i1) gets filtered out. By the extensionality of partial elements this hence gives a line along i between outS u_0 and hcomp u u_0 which geometrically corresponds to the filling of an open box as it connects the base with the lid computed using hcomp. In the special case when q is refl the filler of the above square gives us a direct cubical proof that composing p with refl is p.

```
compPathRefl : {x y : A} (p : x ≡ y) → p · refl ≡ p
compPathRefl {x = x} {y = y} p j i = hfill (λ _ → λ { (i = i0) → x ; (i = i1) → y }) (inS (p i)) (~ j)
```

This way we can hence do even more equality reasoning by directly working with higher dimensional cubes.

By combining hcomp and transp we can define the heterogeneous composition operation of Cohen et al. [2018].

```
comp : (A : I → Set ℓ) {r : I} (u : (i : I) → Partial r (A i)) (u0 : A i0 [ r ↦ u i0 ]) → A i1
comp A {r = r} u u0 = hcomp (λ i → λ { (r = i1) → transp (λ j → A (i ∨ j)) i (u _ 1=1) })
                            (inS (transport (λ i → A i) (outS u0)))
```

With the comp operation we can then finally give the definition of transp for path types.

$$\text{transportPath} : (A : \mathbb{I} \longrightarrow \text{Set})\ (x\ y : (i : \mathbb{I}) \longrightarrow A\ i) \longrightarrow x\ \text{i0} \equiv y\ \text{i0} \longrightarrow x\ \text{i1} \equiv y\ \text{i1}$$
$$\text{transportPath}\ A\ x\ y\ p = \lambda\ i \longrightarrow \text{comp}\ A\ (\lambda\ j \longrightarrow \lambda\ \{\ (i = \text{i0}) \longrightarrow x\ j\ ;\ (i = \text{i1}) \longrightarrow y\ j\})\ (\text{inS}\ (p\ i))$$

The computation rules for hcomp are also defined by cases on the type formers of Agda, just like for transp. These are all quite direct to define and we refer the interested reader to Huber [2017] for details. We will note, however, that for HITs hcomp $(\lambda\, i \rightarrow \lambda\ \{\ (r = \text{i1}) \rightarrow u\})\ u_0$ only reduces to u[i1 / i] when $r = \text{i1}$, and is to be considered a canonical element otherwise. Therefore, functions defined by pattern matching on a HIT also have to make progress when provided an element built with hcomp. We will often refer to such an element as hcomp $r\ u\ u_0$ for ease of notation. The next section describes how this can be achieved, in the context of a core type theory for `Cubical Agda`.

4 HIGHER INDUCTIVE TYPES AND PATTERN MATCHING

Our main technical contribution is an elaboration algorithm for (co)pattern matching definitions in the presence of HITs and path applications. Following Cockx and Abel [2018] we formulate our algorithm as a translation from (co)pattern matching clauses to case trees. The main challenges are generating the computational behavior on hcomp elements of HITs and making sure clauses for path constructors agree with what the function does at the endpoints of the path.

4.1 Elaboration by Example

Here we illustrate how we can handle the case for hcomp when pattern matching using an example. Consider the function c2t from Sect. 2.4.1, it is defined by four clauses, which pattern match on a pair of elements of the circle. Recalling that hcomp $r\,u\,u_0$ is also a canonical element of the circle, we can see that we additionally have to cover the following cases:

$$\begin{array}{lll} \text{c2t}\ (\text{hcomp}\ r\,u\,u_0\ , y) & & = ?0 \\ \text{c2t}\ (\text{base} & , \text{hcomp}\ r\,u\,u_0) & = ?1 \\ \text{c2t}\ (\text{loop}\ i & , \text{hcomp}\ r\,u\,u_0) & = ?2 \end{array}$$

We can cover the first by setting $?0 :=$ hcomp $(\lambda\{j\ (r = \text{i1}) \rightarrow \text{c2t}\ (u\ j\ 1{=}1, y)\})\ (\text{c2t}\ (u_0, y))$ which not only produces an element of the right type, but also satisfies $?0 = \text{c2t}\ (u\ i\ 1{=}1, y)$ when $r = \text{i1}$, which is required to preserve the equality hcomp i1 $u\ u_0 = u\ i\ 1{=}1$. The case ?1 can be solved analogously, while ?2, since it matches on loop, has additional constraints: ?2 should be equal to c2t $(\text{base}, \text{hcomp}\ r\,u\,u_0)$ whenever $i = \text{i0}$ or $i = \text{i1}$. We can satisfy all of these constraints at once by including them in the composition, i.e. setting

$$?2 := \text{hcomp}\ (\lambda\ j \longrightarrow \lambda \left\{ \begin{array}{l} (r = \text{i1}) \longrightarrow \text{c2t}\ (\text{loop}\ i, u\ j\ 1{=}1) \\ (i = \text{i0}) \longrightarrow \text{c2t}\ (\text{base}, \text{hcomp}\ r\,u\,u_0) \\ (i = \text{i1}) \longrightarrow \text{c2t}\ (\text{base}, \text{hcomp}\ r\,u\,u_0) \end{array} \right\})\ (\text{c2t}\ (\text{loop}\ i, u_0))$$

where all the components of the partial element match up because they are all different specializations of c2t $(\text{loop}\ i, \text{hcomp}\ r\,u\,u_0)$ under the different boundary conditions. In the general case the return type of the function can depend on the HIT argument, so a heterogeneous composition will be necessary.

4.2 Syntax of the Core Type Theory

We recall some definitions from Cockx and Abel [2018], extended to allow for the new cubical primitives.

Expressions (Fig. 1) are given in spine-normal form, so that the head symbol of an application is easily accessible. Rather than adding the cubical types and operations described in Sect. 3 as new

x, i			variables
ℓ	$::=$	$n \mid \omega$	universe levels
A, B, u, v	$::=$	w	weak head normal form
	$\mid$	$\mathsf{f}\ \bar{e}$	defined function or primitive applied to eliminations
	$\mid$	$x\ \bar{e}$	variable applied to eliminations
	$\mid$	$\mathsf{c}\ \bar{e}_{\mathsf{c}}$	constructor applied to eliminations
W, w	$::=$	$(x : A) \to B$	dependent function type
	$\mid$	Set_ℓ	universe ℓ
	$\mid$	$\mathsf{D}\ \bar{u}$	datatype fully applied to parameters
	$\mid$	$\mathsf{R}\ \bar{u}$	record type fully applied to parameters
	$\mid$	$\lambda x. u$	lambda abstraction
e	$::=$	e_{c}	elimination for constructors
	$\mid$	$.\pi$	projection
e_{c}	$::=$	u	application
	$\mid$	$@_{u_0,u_1}\ v$	path application

Fig. 1. Syntax of Terms.

expression formers, we subsume them under $\mathsf{f}\ \bar{e}$. This happens also in the implementation of Agda, thanks to the pre-existing support for builtins and primitives.

We include a universe level ω for types like $\mathbb{I}$ which do not support transp or hcomp. Eliminations e include, beyond function application to u and projections $.\pi$, *path applications* $@_{u_0,u_1}\ v$. Path applications to interval element v are annotated by the endpoints u_0 and u_1 used for reduction, in case v becomes i0 or i1.

In contrast to ordinary applications, path applications of stuck terms can reduce; for instance, $x\ @_{u_0,u_1}\ \mathsf{i0}$ reduces to u_0. Thus, variable eliminations $x\ \bar{e}$ are not necessarily in weak head normal form. Thanks to HITs, this is not even the case for constructor applications; $\mathsf{c}\ \bar{e}$ might also reduce!

Binary application $\boxed{u\ e}$ is defined as a partial function on the syntax, by β reduction $(\lambda x.u)\ v = u[v\,/\,x]$ in case of abstractions, or by accumulating eliminations $(x\ \bar{e})\ e = x\ (\bar{e}, e)$, $(\mathsf{f}\ \bar{e})\ e = \mathsf{f}\ (\bar{e}, e)$, $(\mathsf{c}\ \bar{e})\ e_{\mathsf{c}} = \mathsf{c}\ (\bar{e}, e_{\mathsf{c}})$, and otherwise it is undefined.

Patterns are augmented with path application *copatterns*, also in the spine for constructors.

p	$::=$	x	variable pattern
	$\mid$	$\mathsf{c}\ \bar{q}_{\mathsf{c}}$	fully applied constructor pattern
	$\mid$	$\lfloor \mathsf{c} \rfloor\ \bar{q}_{\mathsf{c}}$	forced constructor pattern
	$\mid$	$\lfloor u \rfloor$	forced argument
	$\mid$	$\emptyset$	absurd pattern
q	$::=$	q_{c}	copattern for constructors
	$\mid$	$.\pi$	projection copattern
q_{c}	$::=$	p	application copattern
	$\mid$	$@_{u_0,u_1}\ i$	path application copattern

Note that we retain the boundary annotations of path applications even in patterns, since we convert patterns to expressions during type and coverage checking for case trees.

We write $\boxed{\mathrm{PV}(\bar{q})}$ for the set of variables appearing as variable patterns x or as i in a path application copattern. We will also often drop the subscript from e_{c} and q_{c}.

s	$::=$	$\ominus$	status: unchecked
	$\mid$	$\oplus$	status: checked
$decl^s$	$::=$	$\mathsf{data}\ \mathsf{D}\ \Delta : \mathsf{Set}_n\ \mathsf{where}\ \overline{con}$	datatype declaration
	$\mid$	$\mathsf{record}\ self : \mathsf{R}\ \Delta : \mathsf{Set}_n\ \mathsf{where}\ \overline{field}$	record declaration
	$\mid$	$\mathsf{definition}\ \mathsf{f} : A\ \mathsf{where}\ \overline{cls^s}$	function declaration
con	$::=$	$\mathsf{c}\ \Delta\ [\bar{i} \mid b]$	constructor declaration
b	$::=$	$\epsilon \mid (u_0, u_1)\ b$	boundary terms
$field$	$::=$	$\pi : A$	field declaration
$cls^\ominus$	$::=$	$\bar{q} \hookrightarrow rhs$	unchecked clause
$cls^\oplus$	$::=$	$\Delta \vdash \bar{q} \hookrightarrow u : B$	checked clause
rhs	$::=$	u	clause body: expression
	$\mid$	$\mathsf{impossible}$	empty body for absurd pattern
Σ	$::=$	$\overline{decl^\oplus}$	signature

Fig. 2. Declarations.

The theory is parameterized by a list of declarations Σ whose grammar is shown in Fig. 2. Declaration forms are due to Cockx and Abel [2018] except for datatype constructors $\mathsf{c}\ \Delta\ [\bar{i} \mid b]$. These take a telescope of arguments Δ, i.e. a list of variable typings $(x : A)$, but now also a *boundary* $[\bar{i} \mid b]$, which specifies the dimensions $\bar{i}$ and endpoints b of path constructors. For example, posneg from Sect. 2.3.1 would be specified by $\mathsf{posneg}\ \epsilon\ [i \mid (\mathsf{pos}\ 0, \mathsf{neg}\ 0)]$. We will write $\boxed{\Delta[\bar{i} \mid b] \to A}$ for the iterated function and path type defined by the following equations:

$$\begin{aligned}
[] \to A &= A \\
[i\ \bar{i} \mid (u_0, u_1)\ b] \to A &= \mathsf{PathP}\ (\lambda i.\ [\bar{i} \mid b] \to A)\ (\lambda \bar{i}.\ u_0)\ (\lambda \bar{i}.\ u_1) \\
(x : B)\Delta[\bar{i} \mid b] \to A &= (x : B) \to \Delta[\bar{i} \mid b] \to A
\end{aligned}$$

Further, we write $\hat{\Delta}[\bar{i} \mid b]$ for the appropriate sequence of function and path applications. A constructor $\mathsf{c}\ \Delta'\ [\bar{i} \mid b]$ for a datatype $\mathsf{D}\ \Delta$ will then have type $\Delta\Delta'[\bar{j} \mid b] \to \mathsf{D}\ \hat{\Delta}$.

The core definition of Cockx and Abel [2018] is the elaboration judgment $\Sigma; \Gamma \vdash P \mid \mathsf{f}\ \bar{q} := Q : C \rightsquigarrow \Sigma'$ which performs type and coverage checking for the user supplied clauses of f represented by the problem set P. It further computes the corresponding case tree Q and checked clauses $cls^\oplus$ in Σ'. Case trees Q are previously specified by a typing judgment $\Sigma; \Gamma \vdash \mathsf{f}\ \bar{q} := Q : C \rightsquigarrow \Sigma'$ which follows the same structure but takes the case tree as input. In the following, we only present the rules for the case tree typing judgement and refer to the supplemental material for the elaboration judgement.

4.3 Case Trees

Fig. 3 describes the case tree typing judgment $\boxed{\Sigma; \Gamma \vdash \mathsf{f}\ \bar{q} := Q : C \mid \Theta \rightsquigarrow \Sigma'}$. In our version, the judgment takes an extra input Θ which is a possibly empty list of *boundary assignments* α, which in turn are a list of assignments of interval variables to either $\mathsf{i0}$ or $\mathsf{i1}$. We denote with $\boxed{[\alpha]}$ the substitution implied by the equalities in α itself. The list Θ is used to keep track of which faces of

the current definition have accumulated some boundary constraints, due to the rules to introduce a path (CtIntroPath), a partial element (CtSplitPartial), or to pattern match on an higher inductive type (CtSplitConHIT).

CtDone A leaf of a case tree consists of a term v of the expected type C. Moreover v has to fulfill the boundary constraints on the faces specified by Θ: for every α_i we require that $\mathsf{f}\ \bar{q}$ and v agree when substituted with $[\alpha_i]$, i.e., when restricted to the face in question. Note that we impose the boundary constraints in the signature Σ where we have not added the clause $\mathsf{f}\ \bar{q} \hookrightarrow v$ yet, so they are non-trivial to satisfy. We use $\boxed{\Gamma_\alpha}$ to denote the context obtained by removing the variables in α from Γ and substituting their occurrences with the specified values.

CtIntroPath If the expected type C is a path type $\mathsf{PathP}\ B\ u_0\ u_1$ then we can extend the left hand side to $\mathsf{f}\ \bar{q}\ @_{u_0,u_1}\ i$. We also extend the list of boundary assignments to include the two faces $(i = \mathsf{i0})$ and $(i = \mathsf{i1})$, which will in CtDone ensure that Q produces an element that connects u_0 and u_1. To see this, note that u_0 is judgmentally equal to expression $\mathsf{f}\ \bar{q}\ @_{u_0,u_1}\ \mathsf{i0}$ and u_1 to $\mathsf{f}\ \bar{q}\ @_{u_0,u_1}\ \mathsf{i1}$, because of equality for path applications.

CtSplitPartial If the expected type is equal to $\mathsf{PartialP}\ r\ A$,[7] then we can proceed by splitting on the faces $\alpha_1, \ldots, \alpha_n$ as long as they together *cover* all the ways in which we can have $r = \mathsf{i1}$. This is ensured by the premise $r = \bigvee_i \bigwedge \alpha_i$, where $\boxed{\bigwedge \alpha}$ is defined by mapping $(i = \mathsf{i1})$ to i and $(i = \mathsf{i0})$ to $\sim i$, and combining the resulting elements with $\wedge$. For each face, the left hand side is first refined to $\mathsf{f}\ \bar{q}[\alpha_i]$ so that the variables in the copatterns $\bar{q}$ match their assignment, and then extended by $\lfloor \mathsf{1{=}1} \rfloor$ which is the right elimination for $\mathsf{PartialP}\ r\ A$ as we have $r = \mathsf{i1}$ in Γ_{α_i}. Additionally, since the faces α_i can have overlaps we need to make sure that the different case trees Q_i agree on the intersections, this is accomplished by extending Θ with the assignments of the previous cases. Finally when substituting into a list of assignments, as in $\boxed{\Theta[\alpha]}$, any trivial equalities get removed and any contradictory ones cause the whole assignment in which they appear to be removed.

CtSplitConHIT If the left-hand side $\mathsf{f}\ \bar{q}$ contains a variable x of a data type $\mathsf{D}\ \bar{v}$, then we can pattern match on x, building a $\mathsf{case}_x\{\ldots\}$ node that covers all the alternatives. In this rule we deal with the case in which $\mathsf{D}\ \bar{v}$ is an HIT, as at least one of the $\mathsf{c_i}$ has a non-empty boundary. For each of the constructors $\mathsf{c_i}$ we check the case tree Q_i. To do so we (1) refine $\mathsf{f}\ \bar{q}$ by replacing x with $\mathsf{c_i}$ fully applied to variable and path application copatterns according to its type; and (2) expand the list of assignments with both $(j = \mathsf{i0})$ and $(j = \mathsf{i1})$ for each j in $\bar{j}_i$, which is what $\boxed{\textsc{boundary}(\bar{j}_i)}$ denotes. Moreover we need to consider the case for $\mathsf{hcomp}\ r\ u\ u_0$: we check Q_{hc} by (1) replacing x by the pattern $\mathsf{hcomp}\ r\ u\ u_0$ in the left-hand side and (2) extending Θ with $(r = \mathsf{i1})$ since that is the face where $\mathsf{hcomp}\ r\ u\ u_0$ computes to $u\ \mathsf{i1}\ \mathsf{1{=}1}$.

The remaining rules do not directly involve any of the new features of `Cubical Agda`, so we ask the reader to refer back to the corresponding judgment in Cockx and Abel [2018].

4.3.1 Inferring Right-Hand Side of $\mathsf{hcomp}\ r\ u\ u_0$ *Cases.* In the examples given in Sect. 2 we have never given a clause for the hcomp constructor when pattern matching on an element of a higher inductive type. That is because `Cubical Agda` generates a suitable clause for us during elaboration. How to deal with the hcomp case was already explained in Cohen et al. [2018], but only for the respective induction principle, while in our case we have to deal with user clauses that include multiple-argument and nested pattern matching.

[7] PartialP is a dependent version of Partial where A is a partial element as well, i.e., $A : \mathsf{Partial}\ r\ \mathsf{Set}_n$.

$\boxed{\Sigma; \Gamma \vdash \mathsf{f}\ \bar{q} := Q : C \mid \Theta \rightsquigarrow \Sigma'}$ Presupposes: $\Sigma; \Gamma \vdash \mathsf{f}\ \lceil \bar{q} \rceil : C$ and $\mathrm{dom}(\Gamma) = \mathrm{PV}(\bar{q})$ and $\Theta = \alpha_1; \ldots; \alpha_n$ where $\alpha ::= \epsilon \mid (i = \mathsf{i0})\ \alpha \mid (i = \mathsf{i1})\ \alpha$ such that $\Sigma; \Gamma \vdash i : \mathbb{I}$.
Checks case tree Q and outputs an extension Σ' of Σ by the clauses represented by "$\mathsf{f}\ \bar{q} \hookrightarrow Q$".

$$\frac{\Sigma; \Gamma \vdash v : C \qquad \Theta = \alpha_1; \ldots; \alpha_n \qquad (\Sigma; \Gamma_{\alpha_i} \vdash \mathsf{f}\ \bar{q}[\alpha_i] = v[\alpha_i] : C[\alpha_i])_{i=1\ldots n}}{\Sigma; \Gamma \vdash \mathsf{f}\ \bar{q} := v : C \mid \Theta \rightsquigarrow \Sigma, (\mathsf{clause}\ \Gamma \vdash \mathsf{f}\ \bar{q} \hookrightarrow v : C)}\ \textsc{CtDone}$$

$$\frac{\Sigma; \Gamma \vdash C = (x : A) \to B : \mathsf{Set}_\ell \qquad \Sigma; \Gamma(x : A) \vdash \mathsf{f}\ \bar{q}\ x := Q : B \mid \Theta \rightsquigarrow \Sigma'}{\Sigma; \Gamma \vdash \mathsf{f}\ \bar{q} := \lambda x.\ Q : C \mid \Theta \rightsquigarrow \Sigma'}\ \textsc{CtIntro}$$

$$\frac{\Sigma; \Gamma \vdash C = \mathsf{PathP}\ B\ u_0\ u_1 : \mathsf{Set}_n \qquad \Theta' = (i = 0); (i = 1); \Theta \qquad \Sigma; \Gamma(i : \mathbb{I}) \vdash \mathsf{f}\ \bar{q}\ @_{u_0, u_1}\ i := Q : B\ i \mid \Theta' \rightsquigarrow \Sigma'}{\Sigma; \Gamma \vdash \mathsf{f}\ \bar{q} := \lambda i.\ Q : C \mid \Theta \rightsquigarrow \Sigma'}\ \textsc{CtIntroPath}$$

$$\frac{\Sigma_0; \Gamma \vdash C = \mathsf{PartialP}\ r\ A : \mathsf{Set}_\omega \qquad \Sigma_0; \Gamma \vdash r = \bigvee_i \bigwedge \alpha_i : \mathbb{I} \qquad (\Sigma_{i-1}; \Gamma_{\alpha_i} \vdash (\mathsf{f}\ \bar{q}[\alpha_i]\ \lfloor \mathsf{1{=}1} \rfloor) := Q_i : (A\ \mathsf{1{=}1}) \mid (\alpha_1; \ldots; \alpha_{i-1}; \Theta)[\alpha_i] \rightsquigarrow \Sigma_i)_{i=1\ldots n}}{\Sigma_0; \Gamma \vdash \mathsf{f}\ \bar{q} := \mathsf{split}\{\alpha_1 \mapsto Q_1; \ldots; \alpha_n \mapsto Q_n\} : C \mid \Theta \rightsquigarrow \Sigma_n}\ \textsc{CtSplitPartial}$$

$$\frac{\Sigma_0; \Gamma \vdash C = \mathsf{R}\ \bar{v} : \mathsf{Set}_n \qquad \mathsf{record}\ \mathit{self} : \mathsf{R}\ \Delta : \mathsf{Set}_n\ \mathsf{where}\ \overline{\pi_i : A_i} \in \Sigma_0 \qquad \sigma = [\bar{v} \,/\, \Delta, \mathsf{f}\ \lceil \bar{q} \rceil \,/\, \mathit{self}] \qquad (\Sigma_{i-1}; \Gamma \vdash \mathsf{f}\ \bar{q}\ .\pi_i := Q_i : A_i\sigma \mid \Theta \rightsquigarrow \Sigma_i)_{i=1\ldots n}}{\Sigma_0; \Gamma \vdash \mathsf{f}\ \bar{q} := \mathsf{record}\{\pi_1 \mapsto Q_1; \ldots; \pi_n \mapsto Q_n\} : C \mid \Theta \rightsquigarrow \Sigma_n}\ \textsc{CtCosplit}$$

$$\frac{\begin{array}{c}\Sigma_0; \Gamma_1 \vdash A = \mathsf{D}\ \bar{v} : \mathsf{Set}_n \qquad \mathsf{data}\ \mathsf{D}\ \Delta : \mathsf{Set}_n\ \mathsf{where}\ \overline{\mathsf{c}_i\ \Delta_i\ []} \in \Sigma_0 \\ (\Delta'_i = \Delta_i[\bar{v} \,/\, \Delta])_{i=1\ldots n} \qquad (\rho_i = \mathbb{1}_{\Gamma_1} \uplus [\mathsf{c}_i\ \hat{\Delta}'_i \,/\, x] \quad \rho'_i = \rho_i \uplus \mathbb{1}_{\Gamma_2})_{i=1\ldots n} \\ (\Sigma_{i-1}; \Gamma_1 \Delta'_i (\Gamma_2 \rho_i) \vdash \mathsf{f}\ \bar{q}\rho'_i := Q_i : C\rho'_i \mid \Theta \rightsquigarrow \Sigma_i)_{i=1\ldots n}\end{array}}{\Sigma_0; \Gamma_1 (x : A) \Gamma_2 \vdash \mathsf{f}\ \bar{q} := \mathsf{case}_x\{\mathsf{c}_1\ \hat{\Delta}'_1 \mapsto Q_1; \ldots; \mathsf{c}_n\ \hat{\Delta}'_n \mapsto Q_n\} : C \mid \Theta \rightsquigarrow \Sigma_n}\ \textsc{CtSplitCon}$$

$$\frac{\begin{array}{c}\Sigma_0; \Gamma_1 \vdash A = \mathsf{D}\ \bar{v} : \mathsf{Set}_n \qquad \Gamma = \Gamma_1 (x : A) \Gamma_2 \\ \mathsf{data}\ \mathsf{D}\ \Delta : \mathsf{Set}_n\ \mathsf{where}\ \overline{\mathsf{c}_i\ \Delta_i\ [\bar{j}_i \mid b_i]} \in \Sigma_0 \qquad \exists k.\ [\bar{j}_k \mid b_k] \neq [] \\ \left(\begin{array}{c}\Delta'_i = \Delta_i(\bar{j}_i : \mathbb{I})[\bar{v} \,/\, \Delta] \qquad \bar{q}_i = \hat{\Delta}_i[\bar{j}_i \mid b_i][\bar{v} \,/\, \Delta] \\ \rho_i = \mathbb{1}_{\Gamma_1} \uplus [\mathsf{c}_i\ \bar{q}_i \,/\, x] \qquad \rho'_i = \rho_i \uplus \mathbb{1}_{\Gamma_2} \\ \Theta_i = \textsc{boundary}(\bar{j}_i); \Theta \\ \Sigma_{i-1}; \Gamma_1 \Delta'_i (\Gamma_2 \rho_i) \vdash \mathsf{f}\ \bar{q}\rho'_i := Q_i : C\rho'_i \mid \Theta_i \rightsquigarrow \Sigma_i\end{array}\right)_{i=1\ldots n} \\ \Delta_{\mathsf{hc}} = (r : \mathbb{I})(u : \mathbb{I} \to \mathsf{Partial}\ r\ (\mathsf{D}\ \bar{v}))(u_0 : \mathsf{D}\ \bar{v}\ [\, r \mapsto u\ \mathsf{i0}\,]) \\ \rho_{\mathsf{hc}} = \mathbb{1}_{\Gamma_1} \uplus [\mathsf{hcomp}\ r\ u\ u_0 \,/\, x] \qquad \rho'_{\mathsf{hc}} = \rho_{\mathsf{hc}} \uplus \mathbb{1}_{\Gamma_2} \\ \Sigma_n; \Gamma_1 \Delta_{\mathsf{hc}} (\Gamma_2 \rho_{\mathsf{hc}}) \vdash \mathsf{f}\ \bar{q}\rho'_{\mathsf{hc}} := Q_{\mathsf{hc}} : C\rho'_{\mathsf{hc}} \mid (r = 1); \Theta \rightsquigarrow \Sigma_{n+1}\end{array}}{\Sigma_0; \Gamma \vdash \mathsf{f}\ \bar{q} := \mathsf{case}_x \left\{\begin{array}{l}\mathsf{c}_1\ \bar{q}_1 \mapsto Q_1; \ldots; \mathsf{c}_n\ \bar{q}_n \mapsto Q_n \\ \mathsf{hcomp}\ r\ u\ u_0 \mapsto Q_{\mathsf{hc}}\end{array}\right\} : C \mid \Theta \rightsquigarrow \Sigma_{n+1}}\ \textsc{CtSplitConHIT}$$

Fig. 3. The typing rules for case trees.

$$\boxed{\Sigma;\Gamma(x:\mathsf{D}\ \bar{v})\Delta \vdash \mathsf{f}\ \bar{q} : C \mid \Theta \Rightarrow \textsc{hc-rhs}(rhs)}$$

$$\frac{\begin{array}{l}\Theta = \alpha_1;\ldots;\alpha_n \quad \Delta_{\mathsf{hc}} = (r:\mathbb{I})(u:\mathbb{I}\to \mathsf{Partial}\ r\ (\mathsf{D}\ \bar{v}))(u_0 : \mathsf{D}\ \bar{v}\ [\ r\mapsto u\ \mathsf{i0}\ ]) \\ \Delta^i = \Delta[\mathsf{hfill}\ u\ u_0\ i\,/\,x] \\ \delta^i : \Delta^i = \textsc{transp-tel}\ (j.\,\Delta^{\sim j \vee i})\ i\ \hat{\Delta}^{\mathsf{i1}} \\ sys = \lambda i.\begin{cases}(r=\mathsf{i1}) \to \mathsf{f}\ \bar{q}[u\ i\ 1{=}1\,/\,x, \delta^i\,/\,\Delta^i] \\ \alpha_1 \qquad\quad \to \mathsf{f}\ \bar{q}[\mathsf{hfill}\ u\ u_0\ i\,/\,x, \delta^i\,/\,\Delta^i][\alpha_1] \\ \vdots \\ \alpha_n \qquad\quad \to \mathsf{f}\ \bar{q}[\mathsf{hfill}\ u\ u_0\ i\,/\,x, \delta^i\,/\,\Delta^i][\alpha_n]\end{cases} \\ rhs = \mathsf{comp}\ (\lambda i.\, C[\mathsf{hfill}\ u\ u_0\ i\,/\,x, \delta^i\,/\,\Delta^i])\ sys\ (\mathsf{f}\ \bar{q}[\mathsf{outS}\ u_0\,/\,x, \delta^{\mathsf{i0}}\,/\,\Delta]) \\ \text{Derivable typing:} \quad \Gamma\Delta_{\mathsf{hc}}\Delta[\mathsf{hcomp}\ u\ u_0\,/\,x] \vdash rhs : C[\mathsf{hcomp}\ u\ u_0\,/\,x]\end{array}}{\Sigma;\Gamma(x:\mathsf{D}\ \bar{v})\Delta \vdash \mathsf{f}\ \bar{q} : C \mid \Theta \Rightarrow \textsc{hc-rhs}(rhs)}$$

Fig. 4. Computing the right hand side of a hcomp match.

Fortunately in the context of the rule CtSplitConHIT we have the right information available to construct a term that would be suitable for Q_{hc}. This is accomplished in Fig. 4 by the only rule of the judgment $\Sigma;\Gamma(x:\mathsf{D}\ \bar{v})\Delta \vdash \mathsf{f}\ \bar{q} : C \mid \Theta \Rightarrow \textsc{hc-rhs}(rhs)$. If we had only to handle the case $\Sigma;(x:\mathsf{D}\ \bar{v}) \vdash \mathsf{f}\ x : C \mid \epsilon$ we could apply the strategy from Cohen et al. [2018] and set rhs to $\mathsf{comp}\ (\lambda i.\, C[\mathsf{hfill}\ u\ u_0\ i\,/\,x])\ (\lambda i.\{(r=\mathsf{i1}) \to \mathsf{f}\ (u\ i\ 1{=}1)\})\ (\mathsf{f}\ (\mathsf{outS}\ u_0))$, because that satisfies the expected constraint for the case $(r=\mathsf{i1})$. The extra complication comes from the variables in Δ, since their type might depend on x, and the extra constraints to satisfy because of Θ. Let us first look at the type argument for comp, since here $C[\mathsf{hfill}\ u\ u_0\ i\,/\,x]$ depends on $\Delta^i = \Delta[\mathsf{hfill}\ u\ u_0\ i\,/\,x]$ we need to transport the variables of type $\Delta[\mathsf{hcomp}\ u\ u_0\,/\,x]$, thankfully $\Delta^{\sim j \vee i}$ connects the two telescopes and we can use a version of transp generalized to telescopes, $\boxed{\textsc{transp-tel}}$, to construct the terms, δ^i, to substitute. A complication is that the telescope might contain types in the universe ω: $\textsc{transp-tel}$ will fail if those types actually mention j, otherwise that component of δ^i will be the corresponding variable from $\hat{\Delta}^i$. If $\textsc{transp-tel}$ succeeds we can then use δ^i to correct the calls to f as well, given that $\Delta^{\mathsf{i0}} = \Delta[x\,/\,u_0]$ and $\Delta^i[\mathsf{i1}\,/\,r] = \Delta[u\ i\ 1{=}1\,/\,x$ by the computation behavior of $\mathsf{hfill}\ u\ u_0\ i$. Finally whenever all the assignments in α_i are satisfied we want $rhs = \mathsf{f}\ \bar{q}[\mathsf{hcomp}\ u\ u_0\,/\,x]\ [\alpha_i]$, but that follows because the comp call will reduce to the corresponding case of sys with i replaced by i1 and $\mathsf{hfill}\ u\ u_0\ \mathsf{i1} = \mathsf{hcomp}\ u\ u_0$ and $[(\delta/^{\mathsf{i1}}\Delta^{\mathsf{i1}})]$ will just be the identity substitution because $\textsc{transp-tel}\ (j.\,\Delta^{\sim j \vee \mathsf{i1}})\ \mathsf{i1}\ \hat{\Delta}^{\mathsf{i1}}$ will compute to the last argument, like transp would.

During elaboration we can then run the algorithm expressed by this judgment and generate internally a clause for $\mathsf{hcomp}\ r\ u\ u_0$.

5 GLUE TYPES IN CUBICAL AGDA

Glue types are the key contribution of Cohen et al. [2018] for equipping the univalence principle with computational content. Given that a type in cubical type theory stands for a higher dimensional cube, Glue types let us construct a cube where some faces have been replaced by equivalent types. This is analogous to how hcomp lets us replace some faces of a cube by composing it with other cubes, however for Glue types we can compose with equivalences instead of paths. This implies the univalence principle and it is what lets us transport along paths built out of equivalences.

5.1 Glue Types and Univalence

As everything in Cubical Agda has to work up to higher dimensions the Glue types take a partial family of types A that are equivalent to the base type B. The idea is then that these types get glued onto B so that the equivalence data gets packaged up into a new datatype.

```
Glue : (B : Set ℓ) {r : I} → Partial r (Σ[ A ∈ Set ℓ ] (A ≃ B)) → Set ℓ
```

When r is i1 the type Glue B Ae reduces to Ae 1=1 .fst.

Using Glue types we can turn an equivalence of types into a path and hence define ua.

```
ua : {A B : Set ℓ} → A ≃ B → A ≡ B
ua {A = A} {B = B} e i = Glue B (λ { (i = i0) → (A , e) ; (i = i1) → (B , idEquiv B) })
```

The idea is that we glue A onto B when i is i0 using e and B onto itself when i is i1 using the identity equivalence. The term ua e is a path from A to B as the Glue type reduces when the face conditions are satisfied, so when i is i0 this reduces to A and when i is i1 it reduces to B. Pictorially we can describe ua e as the dashed line in:

$$\begin{array}{ccc} A & \dashrightarrow & B \\ e \downarrow \wr & & \wr \downarrow \text{idEquiv } B \\ B & \xrightarrow[B]{} & B \end{array}$$

The transp operation for Glue types is the most complicated part of the internals of Cubical Agda. The algorithm closely follows Huber [2017, Sect. 3.6], which is a variation of the original algorithm from Cohen et al. [2018, Sect. 6.2]. We will focus on the special case of transport (λ i → ua e i) a for simplicity. This will transport a from A to B by going through the three fully filled lines in the above picture.

Unfolding ua gives

```
transport (λ i → Glue B (λ { (i = i0) → (A , e) ; (i = i1) → (B , idEquiv B) }) a
```

By the boundary equations for Glue types we get that $a : A$ (as it is in the i = i0 face of the Glue type). The algorithm then applies the function of e (i.e., e .fst : $A \to B$) to a giving an element in B. As B is constant along i we could now be done, however for the general algorithm there is no reason for the base to be constant along i; it could for example be another Glue type! We must hence transport along (λ i → B) to get an element in the bottom-right B in the diagram. In order to go up to the top-right corner we then use the inverse of the identity equivalence.[8] Since this is the identity function we end up with:

```
transport (λ i → B) (e .fst a)
```

Using the same path as in the definition of transport for path types we can prove that this is equal to e .fst a up to a path:

```
uaβ : {A B : Set ℓ} (e : A ≃ B) (a : A) → transport (ua e) a ≡ e .fst a
uaβ {B = B} e a = λ i → transp (λ _ → B) i (e .fst a)
```

Transporting along the path that we get from applying ua to an equivalence is, thus, the same as applying the equivalence. This makes it possible to use the univalence axiom computationally in Cubical Agda: we can package up equivalences as paths, do equality reasoning using these paths, and in the end transport along the paths to compute with the equivalences. Furthermore,

[8] In general this might not be the identity function, thus, this step might actually do something.

the combination of ua and uaβ is sufficient to prove that ua is an equivalence which gives the full univalence theorem, i.e., an equivalence between paths and equivalences.[9]

$$\text{univalence} : \forall\{\ell\}\{A\ B : \text{Set}\ \ell\} \rightarrow (A \equiv B) \simeq (A \simeq B)$$

5.2 General Case of transp for Glue Types and the ghcomp Operation

While the special case of transp for Glue types above is quite simple the general case is a lot more complex. The reason is that the input might depend on many more variables than just i. When considering

$$\text{transport}\ (\lambda\ i \rightarrow \text{Glue}\ B\ (\lambda\ \{\ (r = \text{i1}) \rightarrow (A\ ,\ e)\ \})\ a$$

the interval element r might be quite complex and its disjunctive normal form might contain clauses that do not involve i. On these parts the transp function should compute like the transp function for A by the boundary rules for Glue types. This in turn means that additional corrections have to be made compared to the ua case. In Cohen et al. [2018] the part of r that does not mention i is written $\forall i.r$ (as this operation corresponds to universal quantification on the interval).[10]

One of the modifications we have to do in the general case of transp for Glue types is that the simple transport in B has to be a comp with suitable corrections for the $\forall i.r$ faces. While this is easily achieved it has some unfortunate consequences in the case of transporting along ua. In this particular case r is $i \vee \sim i$ so that $(\forall i.r) = \text{i0}$ as there is no part that does not mention i. This means that the comp correction will introduce an empty system which implies that our simple proof of uaβ does not work anymore. In order to fix this we have to extend the proof of uaβ with a suitable hfill in order to compensate for the additional empty system.

Luckily there is a simple trick in `Cubical Agda` that lets us adapt the correction to eliminate the empty system. The problem with the above sketched definition is that the comp does not reduce when $\forall i.r$ is i0, however if we add a clause mapping to the base for this case the issue with the empty system goes away. This relies on a subtle difference between the hcomp operation in `Cubical Agda` and the one in Coquand et al. [2018]. In the latter the boundary constraints were elements of the face lattice $\mathbb{F}$ generated by formal generators $(i = \text{i0})$ and $(i = \text{i1})$ subject to the relation $(i = \text{i0}) \wedge (i = \text{i1}) = \bot$. In `Cubical Agda` on the other hand the hcomp operation takes a family of partial elements that are specified by some $r : \mathsf{I}$. This means that we in `Cubical Agda` can add a face when $(r = \text{i0})$ which was not possible in Coquand et al. [2018] as there is no corresponding operation for $\mathbb{F}$.

The reason that $\mathbb{F}$ in Coquand et al. [2018] does not admit such an operation is that while every $\varphi : \mathbb{F}$ is expressible as $r = \text{i1}$ the choice of r is not unique. In particular for $\varphi = 0_{\mathbb{F}}$ we can choose either i0 or $i \wedge \neg i$ which would give different results when equated to i0. Using $r : \mathsf{I}$ to specify boundaries in `Cubical Agda` avoids the need to make such a choice, and in particular $(r = \text{i0})$ is represented by $\neg r$. It would be tempting to instead extend $\mathbb{F}$ with a negation operation, however that would allow us to represent new kinds of boundaries, like the open interval $(0, 1]$ as $\neg(i = \text{i0})$, and it is not clear how they would impact decidability of typechecking. Modifying hcomp and transp to take a $r : \mathsf{I}$ is semantically justified by the fact that it is not necessary for boundaries to be specified by a subobject of the subobject classifier Ω in the presheaf topos of cubical sets in order to obtain a model of univalent type theory.[11]

[9]For details see https://github.com/agda/cubical/blob/master/Cubical/Foundations/Univalence.agda#L63.

[10]Technically speaking the $\forall$ operation in Cohen et al. [2018] is not an operation on the interval, but rather on the face lattice $\mathbb{F}$. However it is direct to define an analogous operation on the interval and it is this one we use here.

[11]This generalization has been formally verified in `Agda` in https://github.com/mortberg/gen-cart/.

Inspired by Angiuli et al. [2017, Page 53] we call the homogeneous version of this operation *generalized homogeneous composition*, ghcomp. The heterogeneous version used above can be implemented by using ghcomp in the definition of comp. We can write the ghcomp operation as:

$$\text{ghcomp} : \{r : \mathbb{I}\}\ (u : \mathbb{I} \longrightarrow \text{Partial}\ r\ A)\ (u_0 : A\ [\ r \mapsto u\ \text{i0}\ ]) \longrightarrow A$$
$$\text{ghcomp}\ \{r = r\}\ u\ u_0 = \text{hcomp}\ (\lambda\ j \longrightarrow \lambda\ \{\ (r = \text{i1}) \longrightarrow u\ j\ 1{=}1\ ;\ (r = \text{i0}) \longrightarrow \text{outS}\ u_0\ \}\ )\ (\text{inS}\ (\text{outS}\ u_0))$$

By using this in all of the places where the ∀ correction has to be made in the general algorithm for transp for Glue we obtain a better algorithm which does not produce any new empty systems. This way the proof of uaβ can stay as simple as above and no additional corrections has to be made. This is an improvement compared to the algorithm in Cohen et al. [2018] (that is implemented in `cubicaltt`) which produced a surprisingly large number of empty systems even in simple cases.

6 METATHEORY OF CUBICAL TYPE THEORY AND CUBICAL AGDA

The original formulation of cubical type theory as in Cohen et al. [2018] has a model in Kan cubical sets with connections and reversals, that is, presheaves on a suitable cube category where types has structure corresponding to the comp operation. This model has been formally verified in both the NuPRL proof assistant [Bickford 2018] and using Agda as the internal language of the presheaf topos of cubical sets [Licata et al. 2018; Orton and Pitts 2016]. This hence provides semantic consistency proofs for the cubical type theory that `Cubical Agda` is based on. A syntactic consistency proof, using Tait's computability method, for this cubical type theory was given in Huber [2016] by defining an operational semantics and proving that any term of natural numbers type computes to a numeral.

The syntax and semantics of HITs in cubical type theory were studied in Coquand et al. [2018]. The canonicity proof has been shown to extend to the circle and propositional truncation in Huber [2016, Sect. 5]. One technical consequence of the way the system in Coquand et al. [2018] is designed is that there are closed terms of the circle in an empty context that are not base, for example hcomp ($\lambda\ i \longrightarrow$ empty) base. These degenerate elements were a serious problem in `cubicaltt` as they complicated both programming and proving, affecting the efficiency of the system.

These elements arose from the way comp reduces for Glue types in Cohen et al. [2018], but with the optimization discussed in Sect. 5.2 using ghcomp we can eliminate them. This requires us to impose a "validity" constraint on partial elements (following Angiuli et al. [2018, Def. 12]) which says that Partial r A is *valid* if it cannot become empty from a dimension substitution (a concrete condition is that r is a classical tautology). Validity combined with ghcomp eliminates all of the ways that a partial element can become empty in the system. As `Cubical Agda` implement the ghcomp optimization we expect it to be possible to prove a refinement of the canonicity theorem stating that the point constructors are the only elements of HITs in the empty context.

While the comp operation is complicated a recent result by Coquand et al. [2019] shows that for the Cohen et al. [2018] cubical type theory any implementation of the comp operation yield the same result for natural numbers up to a path. As `Cubical Agda` is based on this cubical type theory the result also applies, so even though the implementation of comp differs from the way comp was defined in Cohen et al. [2018] the result for closed terms of type natural numbers will be the same up to a path.

7 CONCLUSION

In this paper we presented `Cubical Agda`, an extension of Agda with features from cubical type theory. This brings to a proof assistant both a fully computational univalence principle and HITs. Moreover, induction on HITs and construction of paths are integrated into Agda's very expressive pattern matching, providing support for more idiomatic definitions than direct use of eliminators.

We expect that such a development environment will lead to more widespread use and experimentation not only of cubical type theory but also of HoTT/UF, in particular for programming applications.

7.1 Related Work

This work is based on the work on cubical type theory of Cohen et al. [2018] and Coquand et al. [2018] and the `cubicaltt` prototype implementation [Cohen et al. 2015]. However, that implementation did not have support for many of the features of a modern proof assistant (implicit arguments, type inference, powerful pattern-matching, etc.) so `Cubical Agda` can be seen as its successor.

The most closely related cubical proof assistant to `Cubical Agda` is **redtt** [The RedPRL Development Team 2018], which also supports computable univalence and HITs. It is based on a variation of cubical type theory called *cartesian* cubical type theory. This has models in cartesian cubical sets [Angiuli et al. 2019] and cartesian cubical computational type theory [Angiuli et al. 2018; Cavallo and Harper 2019]. The **redtt** system has been developed from scratch in order to be a proof assistant for cubical type theory and it has some features that are not in `Cubical Agda` yet, like pre-type universes and extension types inspired by Riehl and Shulman [2017].

The work of Tabareau et al. [2018] extends `Coq` with the ability to transport programs and properties along equivalences using what the authors call *univalent parametricity*. While this achieves some consequences of constructive univalence it does not provide computational content to the full univalence axiom, in particular to neither function nor propositional extensionality. There is also no support for HITs.

The computation rules for equality are also defined by cases on the type in *Observational Type Theory* (OTT) [Altenkirch and McBride 2006; Altenkirch et al. 2007]. This type theory also proves function and propositional extensionality without sacrificing typechecking and constructivity, however it satisfies UIP. Recently, the `XTT` type theory has been developed [Sterling et al. 2019] to reconstruct OTT's exact equality using cubical methods, satisfying UIP rather than univalence. Languages like `XTT` and OTT can be used as an extensional substrate for a two-level type theory [Annenkov et al. 2017; Voevodsky 2013], which would have both equality and path types.

Examples of ideas from HoTT/UF in computer science include [Angiuli et al. 2016] where the authors use univalence and HITs to model `Darcs` style patch theory. This work envisioned what could be done if these notions were computing, but at the time it was unknown how to make this happen. However, now that `Cubical Agda` supports this, it would be interesting to redo the examples as the implementation would now compute. Another example is `HoTTSQL` [Chu et al. 2017] which defines a formal `SQL` style language. The use of HoTT/UF is restricted to reasoning about cardinal numbers and it is not clear how much would be gained from doing this cubically.

7.2 Future Work

Interesting further directions would be to study meta-theoretical properties of cubical type theory, including a proof of decidability of type-checking and a complete correctness proof of the conversion checking algorithm with respect to a declarative specification of equality. We believe this can be done by extending the canonicity proof of Huber [2016] using ideas from Abel et al. [2017].

We would also like to extend `Cubical Agda` with more cubical features, like cubical extension types inspired by Riehl and Shulman [2017] and inductive families following Cavallo and Harper [2019]. An important open problem in the area of constructive synthetic homotopy theory is to compute the Brunerie number [Brunerie 2016] which so far has proved to be infeasible using `cubicaltt` and `Cubical Agda`. It would hence be interesting to study compilation and efficient closed term evaluators of cubical languages in order to be able to do this kind of computations.

REFERENCES

Andreas Abel, Joakim Öhman, and Andrea Vezzosi. 2017. Decidability of Conversion for Type Theory in Type Theory. *Proc. ACM Program. Lang.* 2, POPL, Article 23 (Dec. 2017), 29 pages. https://doi.org/10.1145/3158111

Andreas Abel, Brigitte Pientka, David Thibodeau, and Anton Setzer. 2013. Copatterns: Programming Infinite Structures by Observations. *SIGPLAN Not.* 48, 1 (Jan. 2013), 27–38. https://doi.org/10.1145/2480359.2429075

Agda development team. 2018. *Agda 2.5.4.2 documentation*. http://agda.readthedocs.io/en/v2.5.4.2/

Benedikt Ahrens, Paolo Capriotti, and Régis Spadotti. 2015. Non-wellfounded trees in Homotopy Type Theory. *CoRR* abs/1504.02949 (2015). http://arxiv.org/abs/1504.02949

Thorsten Altenkirch and Conor McBride. 2006. Towards Observational Type Theory. (2006). Unpublished draft.

Thorsten Altenkirch, Conor McBride, and Wouter Swierstra. 2007. Observational Equality, Now!. In *PLPV '07: Proceedings of the 2007 workshop on Programming languages meets program verification*. ACM, New York, NY, USA, 57–68.

Carlo Angiuli, Guillaume Brunerie, Thierry Coquand, Kuen-Bang Hou (Favonia), Robert Harper, and Daniel R. Licata. 2019. Syntax and Models of Cartesian Cubical Type Theory. (February 2019). https://github.com/dlicata335/cart-cube Preprint.

Carlo Angiuli, Kuen-Bang Hou (Favonia), and Robert Harper. 2017. Computational Higher Type Theory III: Univalent Universes and Exact Equality. (2017). Preprint arXiv:1712.01800v1.

Carlo Angiuli, Kuen-Bang Hou (Favonia), and Robert Harper. 2018. Cartesian Cubical Computational Type Theory: Constructive Reasoning with Paths and Equalities. In *27th EACSL Annual Conference on Computer Science Logic (CSL 2018) (Leibniz International Proceedings in Informatics (LIPIcs))*, Dan Ghica and Achim Jung (Eds.), Vol. 119. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 6:1–6:17. https://doi.org/10.4230/LIPIcs.CSL.2018.6

Carlo Angiuli, Edward Morehouse, Daniel R. Licata, and Robert Harper. 2016. Homotopical patch theory. *Journal of Functional Programming* 26 (2016). https://doi.org/10.1017/S0956796816000198

Danil Annenkov, Paolo Capriotti, and Nicolai Kraus. 2017. Two-Level Type Theory and Applications. (2017). https://arxiv.org/abs/1705.03307

Mark Bickford. 2018. Formalizing Category Theory and Presheaf Models of Type Theory in Nuprl. *CoRR* abs/1806.06114 (2018). arXiv:1806.06114 http://arxiv.org/abs/1806.06114

Edwin Brady. 2013. Idris, a general-purpose dependently typed programming language: Design and implementation. *Journal of Functional Programming* 23, 5 (2013), 552–593. https://doi.org/10.1017/S095679681300018X

Guillaume Brunerie. 2016. *On the homotopy groups of spheres in homotopy type theory*. Ph.D. Dissertation. Université de Nice.

Evan Cavallo and Robert Harper. 2019. Higher Inductive Types in Cubical Computational Type Theory. *Proc. ACM Program. Lang.* 3, POPL, Article 1 (Jan. 2019), 27 pages. https://doi.org/10.1145/3290314

Shumo Chu, Konstantin Weitz, Alvin Cheung, and Dan Suciu. 2017. HoTTSQL: Proving Query Rewrites with Univalent SQL Semantics. *SIGPLAN Not.* 52, 6 (June 2017), 510–524. https://doi.org/10.1145/3140587.3062348

Jesper Cockx and Andreas Abel. 2018. Elaborating Dependent (Co)Pattern Matching. *Proc. ACM Program. Lang.* 2, ICFP, Article 75 (July 2018), 30 pages. https://doi.org/10.1145/3236770

Cyril Cohen, Thierry Coquand, Simon Huber, and Anders Mörtberg. 2015. Cubicaltt. (2015). https://github.com/mortberg/cubicaltt.

Cyril Cohen, Thierry Coquand, Simon Huber, and Anders Mörtberg. 2018. Cubical Type Theory: A Constructive Interpretation of the Univalence Axiom. In *Types for Proofs and Programs (TYPES 2015) (LIPIcs)*, Vol. 69. 5:1–5:34.

Cyril Cohen, Maxime Dénès, and Anders Mörtberg. 2013. Refinements for Free!. In *Certified Programs and Proofs (Lecture Notes in Computer Science)*, Georges Gonthier and Michael Norrish (Eds.), Vol. 8307. Springer International Publishing, 147–162. https://doi.org/10.1007/978-3-319-03545-1_10

Thierry Coquand and Nils Anders Danielsson. 2013. Isomorphism is equality. *Indagationes Mathematicae* 24, 4 (2013), 1105–1120.

Thierry Coquand, Simon Huber, and Anders Mörtberg. 2018. On Higher Inductive Types in Cubical Type Theory. In *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science (LICS '18)*. ACM, New York, NY, USA, 255–264. https://doi.org/10.1145/3209108.3209197

Thierry Coquand, Simon Huber, and Christian Sattler. 2019. Homotopy canonicity for cubical type theory. (2019). Preprint available at http://www.cse.chalmers.se/ simonhu/papers/can.pdf.

Leonardo de Moura, Soonho Kong, Jeremy Avigad, Floris van Doorn, and Jakob von Raumer. 2015. The Lean Theorem Prover. In *Automated Deduction - CADE-25, 25th International Conference on Automated Deduction, Berlin, Germany, August 1-7, 2015, Proceedings*.

Simon Huber. 2016. Canonicity for Cubical Type Theory. (2016). Preprint arXiv:1607.04156.

Simon Huber. 2017. A Cubical Type Theory for Higher Inductive Types. (2017). http://www.cse.chalmers.se/~simonhu/misc/hcomp.pdf

Chris Kapulkin and Peter LeFanu Lumsdaine. 2012. The Simplicial Model of Univalent Foundations (after Voevodsky). (Nov. 2012). Preprint arXiv:1211.2851v4.

Daniel R. Licata and Guillaume Brunerie. 2015. A Cubical Approach to Synthetic Homotopy Theory. In *30th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS'15*. ACM, 92–103.

Daniel R. Licata, Ian Orton, Andrew M. Pitts, and Bas Spitters. 2018. Internal Universes in Models of Homotopy Type Theory. In *3rd International Conference on Formal Structures for Computation and Deduction, FSCD 2018, July 9-12, 2018, Oxford, UK (LIPIcs)*, Hélène Kirchner (Ed.), Vol. 108. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 22:1–22:17. https://doi.org/10.4230/LIPIcs.FSCD.2018.22

Daniel R. Licata and Michael Shulman. 2013. Calculating the Fundamental Group of the Circle in Homotopy Type Theory. In *Proceedings of the 2013 28th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS'13)*. 223–232. https://doi.org/10.1109/LICS.2013.28

Peter LeFanu Lumsdaine and Michael Shulman. 2017. Semantics of higher inductive types. (May 2017). Preprint arXiv:1705.07088.

Per Martin-Löf. 1975. An Intiutionistic Theory of Types: Predicative Part. In *Logic Colloquium '73*, H. E. Rose and J. Shepherdson (Eds.). North–Holland, Amsterdam, 73–118.

Conor McBride. 2009. Let's See How Things Unfold: Reconciling the Infinite with the Intensional. In *Proceedings of the 3rd International Conference on Algebra and Coalgebra in Computer Science (CALCO'09)*. Springer-Verlag, Berlin, Heidelberg, 113–126. http://dl.acm.org/citation.cfm?id=1812941.1812953

Ian Orton and Andrew M. Pitts. 2016. Axioms for Modelling Cubical Type Theory in a Topos. In *25th EACSL Annual Conference on Computer Science Logic (CSL 2016) (LIPIcs)*, Vol. 62. 24:1–24:19.

Emily Riehl and Michael Shulman. 2017. A type theory for synthetic ∞-categories. *Higher Structures* 1, 1 (2017), 147–224.

Kristina Sojakova. 2016. The Equivalence of the Torus and the Product of Two Circles in Homotopy Type Theory. *ACM Transactions on Computational Logic* 17, 4 (Nov. 2016), 29:1–29:19.

Jonathan Sterling, Carlo Angiuli, and Daniel Gratzer. 2019. Cubical Syntax for Reflection-Free Extensional Equality. (2019). http://www.jonmsterling.com/pdfs/xtt.pdf

Nicolas Tabareau, Éric Tanter, and Matthieu Sozeau. 2018. Equivalences for Free: Univalent Parametricity for Effective Transport. *Proc. ACM Program. Lang.* 2, ICFP, Article 92 (July 2018), 29 pages. https://doi.org/10.1145/3236787

The Coq Development Team. 2019. The Coq Proof Assistant, version 8.9.0. https://doi.org/10.5281/zenodo.2554024

The RedPRL Development Team. 2018. The **redtt** Proof Assistant. https://github.com/RedPRL/redtt/

The Univalent Foundations Program. 2013. *Homotopy Type Theory: Univalent Foundations of Mathematics.* Institute for Advanced Study.

Andrea Vezzosi. 2017. Streams for Cubical Type Theory. (2017). http://saizan.github.io/streams-ctt.pdf.

Vladimir Voevodsky. 2013. A simple type system with two identity types. (2013). https://www.math.ias.edu/vladimir/sites/math.ias.edu.vladimir/files/HTS.pdf

Vladimir Voevodsky. 2015. An experimental library of formalized Mathematics based on the univalent foundations. *Mathematical Structures in Computer Science* 25 (2015), 1278–1294.