Theories as Types

Dennis Müller[1(✉)], Florian Rabe[1,2(✉)], and Michael Kohlhase[1(✉)]

[1] Computer Science, FAU Erlangen-Nürnberg, Erlangen, Germany
{dennis.mueller,michael.kohlhase}@fau.de, florian.rabe@gmail.com
[2] LRI, Université Paris Sud, Orsay, France

Abstract. Theories are an essential structuring principle that enable modularity, encapsulation, and reuse in formal libraries and programs (called classes there). Similar effects can be achieved by dependent record types. While the former form a separate language layer, the latter are a normal part of the type theory. This overlap in functionality can render different systems non-interoperable and lead to duplication of work.

We present a type-theoretic calculus and implementation of a variant of record types that for a wide class of formal languages naturally corresponds to theories. Moreover, we can now elegantly obtain a contravariant functor that reflects the theory level into the object level: for each theory we obtain the type of its models and for every theory morphism a function between the corresponding types. In particular this allows shallow – and thus structure-preserving – encodings of mathematical knowledge and program specifications while allowing the use of object-level features on models, e.g. equality and quantification.

1 Introduction

In the area of formal systems like type theories, logics, and specification and programming languages, various language features have been studied that allow for inheritance and modularity, e.g., theories, classes, contexts, and records. They all share the motivation of grouping a list of declarations into a new entity such as in $R = [\![x_1 : A_1, \ldots, x_n : A_n]\!]$. The basic intuition behind it is that R behaves like a product type whose values are of the form $\langle\!\langle x_1 : A_1 := a_1, \ldots, x_n : A_n := a_n \rangle\!\rangle$. Such constructs are indispensable already for elementary applications such as defining the algebraic structure of Semilattices (as in Fig. 1), which we will use as a running example.

$$\texttt{Semilattice} = \left\{ \begin{array}{ll} U & : \texttt{type} \\ \wedge & : U \to U \to U \\ \texttt{assoc} & : \vdash \forall x, y, z : U.\ (x \wedge y) \wedge z \doteq x \wedge (y \wedge z) \\ \texttt{commutative} : \ldots \\ \texttt{idempotent} \ : \ldots \end{array} \right\}$$

Fig. 1. A grouping of declarations for semilattices

Many systems support **stratified grouping** (where the language is divided into a lower level for the base language and a higher level that introduces the grouping constructs) or **integrated grouping** (where the grouping construct is one out of many type-forming operations without distinguished ontological status), or both. The

System	Name of feature	
	Stratified	Integrated
ML	Signature/module	Record
C++	Class	Class, struct
Java	Class	Class
Idris [Bra13]	Module	Record
Coq [Coq15]	Module	Record
HOL Light [Har96]	ML signatures	Records
Isabelle [Wen09]	Theory, locale	Record
Mizar [TB85]	Article	Structure
PVS [ORS92]	Theory	Record
OBJ [Gog+93]	Theory	
FoCaLiZe [Har+12]	Species	Record

names of the grouping constructs vary between systems, and we will call them **theories** and **records** in the sequel. An overview of some representative examples is given in the table on the right. For a discussion of these concepts and a comprehensive review of the related work we refer the reader to [MRK].

The two approaches have different advantages. Stratified grouping permits a separation of concerns between the core language and the module system. It also captures high-level structure well in a way that is easy to manage and discover in large libraries, closely related to the advantages of the little theories approach [FGT92]. But integrated grouping allows applying base language operations (such as quantification or tactics) to the grouping constructs. For this reason, the (relatively simple) stratified Coq module system is disregarded in favor of records in major developments such as [Mat].

Allowing both features can lead to a duplication of work where the same hierarchy is formalized once using theories and once using records. A compromise solution is common in object-oriented programming languages, where classes behave very much like stratified grouping but are at the same time normal types of the type system. We call this **internalizing** the higher level features. While combining advantages of stratified and integrated grouping, internalizing is a very heavyweight type system feature: stratified grouping does not change the type system at all, and integrated grouping can be easily added to or removed from a type system, but internalization adds a very complex type system feature from the get-go. It has not been applied much to logics and similar formal systems: the only example we are aware of is the FoCaLiZe [Har+12] system. A much weaker form of internalization is used in OBJ and related systems based on stratified grouping: here theories may be used as (and only as) the types of parameters of parametric theories. Most similarly to our approach, OCaml's first-class modules internalize the theory (called *module type* in OCaml) M as the type `module` M; contrary to both OO-languages and our approach, this kind of internalization is in addition and unrelated to integrated grouping.

In any case, because theories usually allow for advanced declarations like imports, definitions, and notations, as well as extra-logical declarations,

systematically internalizing theories requires a correspondingly expressive integrated grouping construct. Records with defined fields are comparatively rare; e.g., present in [Luo09] and OO-languages. Similarly, imports between record types and/or record terms are featured only sporadically, e.g., in Nuprl [Con+86], maybe even as an afterthought only.

Finally, we point out a closely related trade-off that is orthogonal to our development: even after choosing either a theory or a record to define grouping, many systems still offer a choice whether a declaration becomes a parameter or a field. See [SW11] for a discussion.

Contribution. We present the first formal system that systematically internalizes theories into record types. The central idea is to use an operator Mod that turns the theory T into the type $\mathrm{Mod}\,(T)$, which behaves like a record type. We take special care not to naively compute this record type, which would not scale well to the common situations where theories with hundreds of declarations or more are used. Instead, we introduce record types that allow for defined fields and merging so that $\mathrm{Mod}\,(T)$ preserves the structure of T.

Our approach combines the advantages of stratified and integrated grouping in a lightweight language feature that is orthogonal to and can be easily combined with other foundational language features. Concretely, it is realized as a module in the Mmt framework [Rab14], which allows for the modular design of foundational languages. By combining our new modules with existing ones, we obtain many formal systems with internalized theories. In particular, our typing rules conform to the abstractions of Mmt so that Mmt's type reconstruction [Rab17] is immediately applicable to our features. We showcase the potential in a case study based on this implementation, and which is interesting in its own right: A formal library of elementary mathematical concepts that systematically utilizes $\mathrm{Mod}\,(\cdot)$ throughout for algebraic structures, topological spaces etc.

Overview. We formulate our approach in the setting of a dependently-typed λ-calculus, which we recall in Sect. 2. This section also serves as a gentle primer for defining language features in Mmt. Section 3 introduces our notion of record types, based on which we introduce the model-operator in Sect. 4. Section 5 presents our implementation and a major case study on elementary mathematics. This paper is a shortened version of [MRK], which also contains all the proofs.

2 Preliminaries

We introduce the well-known dependently-typed lambda calculus as the starting point of our development. The **grammar** is given in Fig. 2. The only surprise here is that we allow optional definitions in contexts; this is a harmless convenience at this point but will be critical later on when we introduce records with defined fields. As usual, we write $T \to T'$ instead of $\prod_{x:T} T'$ when possible. We also write $T[x/T']$ for the usual capture-avoiding **substitution** of T' for x in T.

Mmt uses a bidirectional type system, i.e., we have two separate judgments for type *inference* and type *checking*. Similarly, we have two equality judgments:

$$\begin{array}{ll}
\Gamma ::= \cdot \mid \Gamma, x[: T][:= T] & \text{contexts} \\
T ::= x \mid \texttt{type} \mid \texttt{kind} & \text{variables and universes} \\
\quad \mid \ \prod_{x:T'} T \mid \lambda x : T'.T \mid T_1 T_2 & \text{dependent function types}
\end{array}$$

Fig. 2. Grammar for contexts and expressions

one for checking equality of two given terms and one for reducing a term to another one. Our **judgments** are given in Fig. 3.

Adding record types in Sect. 3 will introduce non-trivial **subtyping**, e.g., $[\![x : T, y : S]\!]$ is a subtype of $[\![x : T]\!]$.[1] Therefore, we already introduce a subtyping judgment here even though it is not needed for dependent function types yet. For our purposes, it is sufficient (and desirable) to consider subtyping to be an abbreviation: $\Gamma \vdash T_1 <: T_2$ iff for all t $\Gamma \vdash t \Leftarrow T_1$ implies $\Gamma \vdash t \Leftarrow T_2$.

Judgment	Intuition
$\vdash \Gamma \ \texttt{ctx}$	Γ is a well-formed context
$\Gamma \vdash t \Leftarrow T$	t checks against type/kind T.
$\Gamma \vdash t \Rightarrow T$	type/kind of term t is inferred to be T
$\Gamma \vdash t_1 \equiv t_2 : T$	t_1 and t_2 are equal at type T
$\Gamma \vdash t_1 \rightsquigarrow t_2$	t_1 computes to t_2
$\Gamma \vdash T_1 <: T_2$	T_1 is a subtype of T_2

Fig. 3. Judgments

The **pre/postconditions** of these judgments are as follows: $\Gamma \vdash t \Leftarrow T$ assumes that T is well-typed and implies that t is well-typed. $\Gamma \vdash t \Rightarrow T$ implies that both t and T are well-typed. $\Gamma \vdash t_1 \rightsquigarrow t_2$ implies that t_2 is well-typed iff t_1 is (which puts additional burden on computation rules that are called on not-yet-type-checked terms). Equality and subtyping are only used for expressions that are assumed to be well-typed, i.e., $\Gamma \vdash t_1 \equiv t_2 : T$ implies $\Gamma \vdash t_i \Leftarrow T$, and $\Gamma \vdash T_1 <: T_2$ implies that T_i is a type/kind.

Remark 1 (Horizontal Subtyping and Equality). The equality judgment could alternatively be formulated as an untyped equality $t \equiv t'$. That would require some technical changes to the rules but would usually not be a huge difference. In our case, however, the use of typed equality is critical.

For example, consider record values $r_1 = \langle a := 1, b := 1 \rangle$ and $r_2 = \langle a := 1, b := 2 \rangle$ as well as record types $R = [\![a : nat]\!]$ and $S = [\![a : nat, b : nat]\!]$. Due to horizontal subtyping, we have $S <: R$ and thus both $r_i \Leftarrow S$ and $r_i \Leftarrow R$. This has the advantage that the function $S \rightarrow R$ that throws away the field

---

[1] This is sometimes called *horizontal* subtyping. In that case, the straightforward covariance rule for record types is called *vertical* subtyping.

b becomes the identity operation. Now our equality at record types behaves accordingly and checks only for the equality of those fields required by the type. Thus, $r_1 \equiv r_2 : R$ is true whereas $r_1 \equiv r_2 : S$ is false, i.e., the equality of two terms may depend on the type at which they are compared. While seemingly dangerous, this makes sense intuitively: r_1 can be replaced with r_2 in any context that expects an object of type R because in such a context the field b, where r_1 and r_2 differ, is inaccessible.

Of course, this treatment of equality precludes downcasts: an operation that casts the equal terms $r_1 : R$ and $r_2 : R$ into the corresponding unequal terms of type S would be inconsistent. But such downcasts are still possible (and valuable) at the meta-level. For example, a tactic $GroupSimp(G, x)$ that simplifies terms x in a group G can check if G is commutative and in that case apply more simplification operations.

The full rules of a lambda calculus can be found in the long version [MRK]. We can now show that the usual variance rule for function types is derivable.

Theorem 1. *The following subtyping rule is derivable:*

$$\frac{\Gamma \vdash A <: A' \quad \Gamma, x : A \vdash B' <: B}{\Gamma \vdash \prod_{x:A'} B' <: \prod_{x:A} B}$$

Moreover, we can show that every well-typed term t has a **principal type** T in the sense that (i) $\Gamma \vdash t \Leftarrow T$ and (ii) whenever $\Gamma \vdash t \Leftarrow T'$, then also $\Gamma \vdash T <: T'$. The principal type is exactly the one inferred by our rules (see Theorem 2).

3   Record Types with Defined Fields

We now introduce record types as an additional module of our framework by extending the grammar and the rules. The basic intuition is that $[\![\Gamma]\!]$ and $\langle\!\langle\Gamma\rangle\!\rangle$ construct record types and terms. We call a context **fully typed** resp. **defined** if all fields have a type resp. a definition. In $[\![\Gamma]\!]$, Γ must be fully typed and may additionally contain defined fields. In $\langle\!\langle\Gamma\rangle\!\rangle$, Γ must be fully defined; the types are optional and usually omitted in practice.

Because we frequently need fully defined contexts, we introduce a notational convention for them: a context denoted by a *lower* case letters like γ is always fully defined. In contrast, a context denoted by an *upper* case letter like Γ may have any number of types or definitions.

3.1   Records

We extend our grammar as in Fig. 4, where the previously existing parts are grayed out.

$$\Gamma ::= \cdot \mid \Gamma, x[: T][:= T]$$
$$T ::= x \mid \mathtt{type} \mid \mathtt{kind}$$
$$\mid \; \textstyle\prod_{x:T'} T \mid \lambda x : T'.T \mid T_1 T_2$$
$$\mid \; [\![\Gamma]\!] \mid \langle\!\langle \Gamma \rangle\!\rangle \mid T.x \qquad\qquad \text{record types, terms, projections}$$

Fig. 4. Grammar for records

Remark 2 (Field Names and Substitution in Records). Note that we use the same identifiers for variables in contexts and fields in records. This allows reusing results about contexts when reasoning about and implementing records. In particular, it immediately makes our records dependent, i.e., both in a record type and — maybe surprisingly — in a record term every variable x may occur in subsequent fields. In some sense, this makes x bound in those fields. However, record types are critically different from Σ-types: we must be able to use x in record projections, i.e., x can *not* be subject to α-renaming.

As a consequence, capture-avoiding substitution is not always possible. This is a well-known problem that is usually remedied by allowing every record to declare a name for itself (e.g., the keyword `this` in many object-oriented languages), which is used to disambiguates between record fields and a variable in the surrounding context (or fields in a surrounding record). We gloss over this complication here and simply make substitution a partial function.

Before stating the rules, we introduce a few critical auxiliary definition:

Definition 1 (Substituting in a Record). We extend substitution $t[x/t']$ to records:

- $[\![x_1 : T_1, \ldots, x_n : T_n]\!]\, [y/t]$
    $$= \begin{cases} [\![x_1 : T_1[y/t], \ldots, x_{i-1} : T_{i-1}[y/t], x_i : T_i, \ldots, x_n : T_n]\!] & \text{if } y = x_i \\ [\![x_1 : (T_1[y/t]), \ldots, x_n : (T_n[y/t])]\!] & \text{else} \end{cases}$$
    if none of the x_i are free in t. Otherwise the substitution is undefined.
- $\langle\!\langle x_1 := t_1, \ldots, x_n := t_n \rangle\!\rangle\, [y/t] = \begin{cases} \langle\!\langle x_1 := t_1, \ldots, x_n := t_n \rangle\!\rangle & \text{if } y \in \{x_1, \ldots, x_n\} \\ \langle\!\langle x_1 := (t_1[y/t]), \ldots, x_n := (t_n[y/t]) \rangle\!\rangle & \text{else} \end{cases}$
    if none of the x_i are free in t. Otherwise the substitution is undefined.
- $(r.x)[y/t] = (r[y/t]).x$.

Definition 2 (Substituting with a Record). We write $t[r/\Delta]$ for the result of substituting any occurrence of a variable x declared in Δ with $r.x$

In the special case where $r = \langle\!\langle \delta \rangle\!\rangle$, we simply write $t[\delta]$ for $t[\langle\!\langle \delta \rangle\!\rangle / \delta]$, i.e., we have $t[x_1 := t_1, \ldots, x_n := t_n] = t[x_n/t_n]\ldots[x_1/t_1]$.

Our rules for records are given in Fig. 5. Their roles are systematically similar to the rules for functions: three inference rules for the three constructors followed by a type and an equality checking rule for record types and the (in this case: two) computation rules. We remark on a few subtleties below.

Formation:

$$\frac{\vdash \Gamma, \Delta\ \mathtt{ctx} \qquad \Delta\ \text{fully typed} \qquad \max \Delta \in \{\mathtt{type}, \mathtt{kind}\}}{\Gamma \vdash [\![\Delta]\!] \Rightarrow \max \Delta}$$

where $\max \Delta$ is the maximal universe of all undefined fields in Δ

Introduction:

$$\frac{\vdash \Gamma, \Delta\ \mathtt{ctx} \quad \delta\ \text{fully defined} \quad \Delta\ \text{like}\ \delta\ \text{but with all missing types inferred}}{\Gamma \vdash \langle \delta \rangle \Rightarrow [\![\Delta]\!]}$$

Elimination:

$$\frac{\Gamma \vdash r \Rightarrow [\![\Delta_1, x : T[:= t], \Delta_2]\!]}{\Gamma \vdash r.x \Rightarrow T[r/\Delta_1]}$$

Type checking:

$$\frac{\overbrace{\Gamma \vdash r.x \Leftarrow T[r/\Delta]}^{\text{For } x : T[:= t] \in \Delta} \quad \overbrace{\Gamma \vdash r.x \equiv t[r/\Delta] : T[r/\Delta]}^{\text{additionally for } x : T := t \in \Delta} \quad \Gamma \vdash r \Rightarrow R}{\Gamma \vdash r \Leftarrow [\![\Delta]\!]}$$

Equality checking (extensionality):

$$\frac{\overbrace{\Gamma \vdash r_1.x \equiv r_2.x : T[r_1/\Delta]}^{\text{For every } (x : T) \in \Delta}}{\Gamma \vdash r_1 \equiv r_2 : [\![\Delta]\!]}$$

Computation:

$$\frac{\delta = \delta_1, x[: T] := t, \delta_2 \quad \Gamma \vdash \langle \delta \rangle \Rightarrow R}{\Gamma \vdash \langle \delta \rangle.x \leadsto t[\delta_1]} \qquad \frac{\Gamma \vdash r \Rightarrow [\![\Delta_1, x : T := t, \Delta_2]\!]}{\Gamma \vdash r.x \leadsto t[r/\Delta_1]}$$

Fig. 5. Rules for records

The formation rule is partial in the sense that not every context defines a record type or kind. This is because the universe of a record type must be as high as the universe of any undefined field to avoid inconsistencies. For example, $\max(a : \mathtt{nat}) = \mathtt{type}$, $\max(a : \mathtt{type}) = \mathtt{kind}$ and $\max(a : \mathtt{kind})$ is not defined. If we switched to a countable hierarchy of universes (which we avoid for simplicity), we could turn every context into a record type.

The introduction rule infers the principal type of every record term. Because we allow record types with defined fields, this is the singleton type containing only that record term. This may seem awkward but does not present a problem in practice, where type checking is preferred over type inference anyway.

The elimination rule is straightforward, but it is worth noting that it is entirely parallel to the second computation rule.[2]

---

[2] Note that it does not matter how the fields of the record are split into Δ_1 and Δ_2 as long as Δ_1 contains all fields that the declaration of x depends on.

The type checking rule has a surprising premise that r must already be well-typed (against some type R). Semantically, this assumption is necessary because we only check the presence of the fields required by $[\![\Delta]\!]$ — without the extra assumption, typing errors in any additional fields that r might have could go undetected. In practice, we implement the rule with an optimization: If r is a variable or a function application, we can efficiently infer some type for it. Otherwise, if $r = \langle \delta \rangle$, some fields of δ have already been checked by the first premise, and we only need to check the remaining fields. The order of premises matters in this case: we want to first use type checking for all fields for which $[\![\Delta]\!]$ provides an expected type before resorting to type inference on the remaining fields.

In the equality checking rule, note that we only have to check equality at undefined fields — the other fields are guaranteed to be equal by the assumption that r_1 and r_2 have type $[\![\Delta]\!]$.

Like the type checking rule, the first computation rule needs the premise that r is well-typed to avoid reducing an ill-typed into a well-typed term. In practice, our framework implements computation with a boolean flag that tracks whether the term to be simplified can be assumed to be well-typed or not; in the former case, this assumption can be skipped.

The second computation rule looks up the definition of a field in the type of the record. Both computation rules can be seen uniformly as definition lookup rules — in the first case the definition is given in the record, in the second case in its type.

Example 1. Figure 6 shows a record type of **Semilattices** (actually, this is a kind because it contains a type field) analogous to the grouping in Fig. 1 (using the usual encoding of axioms via judgments-as-types and higher-order abstract syntax for first-order logic).

$$\texttt{Semilattice} := \left[\!\!\left[\begin{array}{ll} U & : \texttt{type} \\ \wedge & : U \to U \to U \\ \texttt{assoc} & : \vdash \forall x, y, z : U.\ (x \wedge y) \wedge z \doteq x \wedge (y \wedge z) \\ \texttt{commutative} & : \dots \\ \texttt{idempotent} & : \dots \end{array}\right]\!\!\right]$$

Fig. 6. The (record-)kind of semilattices

Then, given a record $r : \texttt{Semilattice}$, we can form the record projection $r.\wedge$, which has type $r.U \to r.U \to r.U$ and $r.\texttt{assoc}$ yields a proof that $r.\wedge$ is associative. The intersection on sets forms a semilattice so (assuming we have proofs $\cap - \texttt{assoc}$, $\cap - \texttt{comm}$, $\cap - \texttt{idem}$ with the corresponding types) we can give an instance of that type as

$$\texttt{interSL} : \texttt{Semilattice} := \left\langle U := \texttt{Set}, \wedge := \cap, \texttt{assoc} := \cap - \texttt{assoc}, \dots \right\rangle$$

Theorem 2 (Principal Types). *Our inference rules infer a principal type for each well-typed normal term.*

In analogy to function types, we can derive the subtyping properties of record types. We introduce context subsumption and then combine horizontal and vertical subtyping in a single statement.

Definition 3 (Context Subsumption). For two fully typed contexts Δ_i we write $\Gamma \vdash \Delta_1 \hookrightarrow \Delta_2$ iff for every declaration $x : T[:= t]$ in Δ_1 there is a declaration $x : T'[:= t']$ in Δ_2 such that

- $\Gamma \vdash T' <: T$ and
- if t is present, then so is t' and $\Gamma \vdash t \equiv t' : T$

Intuitively, $\Delta_1 \hookrightarrow \Delta_2$ means that everything of Δ_1 is also in Δ_2. That yields:

Theorem 3 (Record Subtyping). *The following rule is derivable:*

$$\frac{\Gamma \vdash \Delta_1 \hookrightarrow \Delta_2}{\Gamma \vdash [\![\Delta_2]\!] <: [\![\Delta_1]\!]}$$

3.2 Merging Records

We introduce an advanced operation on records, which proves critical for both convenience and performance: Theories can easily become very large containing hundreds or even thousands of declarations. If we want to treat theories as record types, we need to be able to build big records from smaller ones without exploding them into long lists. Therefore, we introduce an explicit merge operator $+$ on both record types and terms.

In the grammar, this is a single production for terms:

$$T ::= T + T$$

The intended meaning of $+$ is given by the following definition:

Definition 4 (Merging Contexts). Given a context Δ and a (not necessarily well-typed) context E, we define a partial function $\Delta \oplus E$ as follows:

- $\cdot \oplus E = E$
- If $\Delta = d, \Delta_0$ where d is a single declaration for a variable x:
    - if x is not declared in E: $(d, \Delta_0) \oplus E = d, (\Delta_0 \oplus E)$
    - if $E = E_0, e, E_1$ where e is a single declaration for a variable x:
        - ∗ if a variable in E_0 is also declared in Δ_0: $\Delta \oplus E$ is undefined,
        - ∗ if d and e have unequal types or unequal definitions: $\Delta \oplus E$ is undefined[3],

---

[3] It is possible and important in practice to also define $\Delta \oplus E$ when the types/definitions in d and e are provably equal. We omit that here for simplicity.

∗ otherwise, $(d, \Delta_0) \oplus (E_0, e, E_1) = E_0, m, (\Delta_0, E_1)$ where m arises by merging d and e.

Note that $\oplus$ is an asymmetric operator: While Δ must be well-typed (relative to some ambient context), E may refer to the names of Δ and is therefore not necessarily well-typed on its own.

We do not define the semantics of $+$ via inference and checking rules. Instead, we give equality rules that directly expand $+$ into $\oplus$ when possible:

$$\frac{\vdash \Gamma, (\Delta_1 \oplus \Delta_2) \ \texttt{ctx}}{\Gamma \vdash [\![\Delta_1]\!] + [\![\Delta_2]\!] \rightsquigarrow [\![\Delta_1 \oplus \Delta_2]\!]} \qquad \frac{\vdash \Gamma, (\delta_1 \oplus \delta_2) \ \texttt{ctx}}{\Gamma \vdash \langle\!\langle \delta_1 \rangle\!\rangle + \langle\!\langle \delta_2 \rangle\!\rangle \rightsquigarrow \langle\!\langle \delta_1 \oplus \delta_2 \rangle\!\rangle}$$

$$\frac{\vdash \Gamma, (\Delta \oplus \delta) \ \texttt{ctx}}{\Gamma \vdash [\![\Delta]\!] + \langle\!\langle \delta \rangle\!\rangle \rightsquigarrow \langle\!\langle \Delta \oplus \delta \rangle\!\rangle}$$

In implementations some straightforward optimizations are needed to verify the premises of these rules efficiently; we omit that here for simplicity. For example, merges of well-typed records with disjoint field names are always well-typed, but e.g., $[\![x : nat]\!] + [\![x : bool]\!]$ is not well-typed even though both arguments are.

In practice, we want to avoid using the computation rules for $+$ whenever possible. Therefore, we prove admissible rules (i.e., rules that can be added without changing the set of derivable judgments) that we use preferentially:

Theorem 4. *If R_1, R_2, and $R_1 + R_2$ are well-typed record types, then $R_1 + R_2$ is the greatest lower bound with respect to subtyping of R_1 and R_2. In particular, $\Gamma \vdash r \Leftarrow R_1 + R_2$ iff $\Gamma \vdash r \Leftarrow R_1$ and $\Gamma \vdash r \Leftarrow R_2$.*
If $\Gamma \vdash r_i \Leftarrow R_i$ and $r_1 + r_2$ is well-typed, then $\Gamma \vdash r_1 + r_2 \Leftarrow R_1 + R_2$.

Inspecting the type checking rule in Fig. 5, we see that a record r of type $[\![\Delta]\!]$ must repeat all defined fields of Δ. This makes sense conceptually but would be a major inconvenience in practice. The merging operator solves this problem elegantly as we see in the following example:

Example 2. Continuing our running example, we can now define a type of semilattices *with order* (and all associated axioms) as in Fig. 7.

```
SemilatticeOrder := Semilattice + ⟦ ≤    : U → U → U := λx, y : U. x ≐ x ∧ y
                                     refl :⊢ ∀a : U. a ≤ a := (proof)
                                     ... ⟧
interSLO := SemilatticeOrder + interSL
```

Fig. 7. Running example

Now the explicit merging in the type `SemilatticeOrder` allows the projection `interSLO.` $\leq$, which is equal to $\lambda x, y : (\texttt{interSLO}.U) . (x \doteq x(\texttt{interSLO}.\wedge)y)$ and `interSLO.refl` yields a proof that this order is reflexive – without needing to define the order or prove the axiom anew for the specific instance `interSL`.

4   Internalizing Theories

4.1   Preliminaries: Theories

We introduce a minimal definition of stratified theories and theory morphisms, which can be seen as a very simple fragment of the MMT language [RK13]. The **grammar** is given in Fig. 8, again graying out the previously introduced parts.

$$\begin{array}{ll}
\Theta ::= \cdot \mid \Theta,\, X = \{\Gamma\} \mid \Theta,\, X : X_1 \to X_2 = \{\Gamma\} & \text{theory level} \\
\Gamma ::= \cdot \mid \Gamma, x[:T][:=T] \mid \Gamma,\, \mathtt{include}\, X & \text{includes} \\
T ::= x \mid \mathtt{type} \mid \mathtt{kind} & \\
\quad \mid \ \prod_{x:T'} T \mid \lambda x : T'.T \mid T_1 T_2 & \\
\quad \mid \ [\![\Gamma]\!] \mid \langle\!\langle \Gamma \rangle\!\rangle \mid T.x \mid T_1 + T_2 &
\end{array}$$

Fig. 8. A simple stratified language

Each of the two levels has its own context: Firstly, the **theory level** context Θ introduces names X, which can be either theories $X = \{\Gamma\}$ or morphisms $X : P \to Q = \{\Gamma\}$, where P and Q are the names of previously defined theories. Secondly, the **expression level** context Γ is as before but may additionally contain includes $\mathtt{include}\, X$ of other theories resp. morphisms X. We call a context **flat** if it does not contain includes.

All **judgments** are as before except that they acquire a second context, e.g., the typing judgment now becomes $\Theta; \Gamma \vdash t \Leftarrow T$. With this modification, all rules for function and record types remain unchanged. However, we add the restriction that Γ in $[\![\Gamma]\!]$ and $\langle\!\langle \Gamma \rangle\!\rangle$ must be flat.

We omit the **rules** for theories and morphisms for brevity and only sketch their intuitions. We think of theories as named contexts and of morphisms as named substitutions between contexts. Includes allow forming both modularly by copying over the declarations of a previously named object. While theories may contain arbitrary declarations, morphisms are restricted: Let Θ contain $P = \{\Gamma\}$ and $Q = \{\Delta\}$. Then a morphism $V : P \to Q = \{\delta\}$ is well-typed if δ is fully defined (akin to record terms) and contains for each declaration $x : T$ of P a declaration $x = t$ where t may refer to all names declared in Q. V induces a homomorphic extension $\overline{V}$ that maps P-expressions to Q-expressions. The key property of morphisms is that, if V is well-typed, then $\Theta; P \vdash t \Leftarrow T$ implies $\Theta; Q \vdash \overline{V}(t) \Leftarrow \overline{V}(T)$ and accordingly for equality checking and subtyping. Thus, theory morphisms preserve judgments and (via propositions-as-types representations) truth. Moreover, it is straightforward to extend the above with identity and composition so that theories and morphisms form a category. We refer to [Rab14] for details.

4.2   Internalization

We can now add the internalization operator, for which everything so far was preparation. We add one production to the **grammar**:

$$T ::= \mathtt{Mod}\,(X)$$

The intended meaning of $\mathrm{Mod}\,(X)$ is that it turns a theory X into a record type and a morphism $X : P \to Q$ into a function $\mathrm{Mod}\,(Q) \to \mathrm{Mod}\,(P)$. For simplicity, we only state the rules for the case where all include declarations are at the beginning of theory/morphism:

$$\frac{P = \{\mathtt{include}\,P_1, \ldots, \mathtt{include}\,P_n, \Delta\} \text{ in } \Theta \quad \Delta \text{ flat} \quad \max P \text{ defined}}{\Theta; \Gamma \vdash \mathrm{Mod}\,(P) \rightsquigarrow \mathrm{Mod}\,(P_1) + \ldots + \mathrm{Mod}\,(P_n) + [\![\Delta]\!]}$$

$$\frac{V : P \to Q = \{\mathtt{include}\,V_1, \ldots, \mathtt{include}\,V_n, \delta\} \text{ in } \Theta \quad \delta \text{ flat}}{\Theta; \Gamma \vdash \mathrm{Mod}\,(V) \rightsquigarrow \lambda r : \mathrm{Mod}\,(Q) . \, \mathrm{Mod}\,(P) + (\mathrm{Mod}\,(V_1)\,r) + \ldots + (\mathrm{Mod}\,(V_n)\,r) + \langle\!\langle \delta[r] \rangle\!\rangle}$$

where we use the following abbreviations:

- In the rule for theories, $\max P$ is the biggest universe occurring in any declaration transitively included into P, i.e., $\max P = \max\{\max P_1, \ldots, \max P_n, \max \Delta\}$ (undefined if any argument is).
- In the rule for morphisms, $\delta[r]$ is the result of substituting in δ every reference to a declaration of x in Q with $r.x$.

In the rule for morphisms, the occurrence of $\mathrm{Mod}\,(P)$ may appear redundant; but it is critical to (i) make sure all defined declarations of P are part of the record and (ii) provide the expected types for checking the declarations in δ.

Example 3. Consider the theories in Fig. 9. Applying $\mathrm{Mod}\,(\cdot)$ to these theories yields exactly the record types of the same name introduced in Sect. 3 (Figs. 6 and 7), i.e., we have $\mathtt{interSL} \Leftarrow \mathrm{Mod}\,(\mathtt{Semilattice})$ and $\mathtt{interSLO} \Leftarrow \mathrm{Mod}\,(\mathtt{SemilatticeOrder})$. In particularly, Mod preserves the modular structure of the theory.

$$\mathtt{theory\ Semilattice} = \begin{cases} U & : \mathtt{type} \\ \wedge & : U \to U \to U \\ \mathtt{assoc} & : \vdash \forall x, y, z : U. \, (x \wedge y) \wedge z \doteq x \wedge (y \wedge z) \\ \mathtt{commutative} : \ldots \\ \mathtt{idempotent} \ : \ldots \end{cases}$$

$$\mathtt{theory\ SemilatticeOrder} = \begin{cases} \mathtt{include\ Semilattice} \\ \mathtt{order} \ : U \to U \to U := \lambda x, y : U. \, x \doteq x \wedge y \\ \mathtt{refl} \ :\vdash \forall a : U. \, a \leq a := \text{ (proof)} \\ \ldots \end{cases}$$

Fig. 9. A theory of semilattices

The basic properties of $\mathrm{Mod}\,(X)$ are collected in the following theorem:

Theorem 5 (Functoriality). $\mathrm{Mod}\,(\cdot)$ *is a monotonic contravariant functor from the category of theories and morphisms ordered by inclusion to the category of types (of any universe) and functions ordered by subtyping. In particular,*

– *if P is a theory in Θ and $\max P \in \{\mathtt{type},\mathtt{kind}\}$, then $\Theta; \Gamma \vdash \mathtt{Mod}\,(P) \Leftarrow \max P$*
– *if $V : P \to Q$ is a theory morphism in $\Theta; \Gamma \vdash \mathtt{Mod}\,(V) \Leftarrow \mathtt{Mod}\,(Q) \to \mathtt{Mod}\,(P)$*
– *if P is transitively included into Q, then $\Theta; \Gamma \vdash \mathtt{Mod}\,(Q) <: \mathtt{Mod}\,(P)$.*

An immediate advantage of $\mathtt{Mod}\,(\cdot)$ is that we can now use the expression level to define expression-like theory level operations. As an example, we consider the **intersection** $P \cap P'$ of two theories, i.e., the theory that includes all theories included by both P and P'. Instead of defining it at the theory level, which would begin a slippery slope of adding more and more theory level operations, we can simply build it at the expression level:

$$P \cap P' := \mathtt{Mod}\,(Q_1) + \ldots + \mathtt{Mod}\,(Q_n)$$

where the Q_i are all theories included into both P and P'.[4]

Note that the computation rules for $\mathtt{Mod}$ are efficient in the sense that the structure of the theory level is preserved. In particular, we do not flatten theories and morphisms into flat contexts, which would be a huge blow-up for big theories.[5]

However, efficiently *creating* the internalization is not enough. $\mathtt{Mod}\,(X)$ is defined via $+$, which is itself only an abbreviation whose expansion amounts to flattening. Therefore, we establish admissible rules that allow *working with internalizations* efficiently, i.e., without computing the expansion of $+$:

Theorem 6. *Fix well-typed Θ, Γ and $P = \{\mathtt{include}\,P_1, \ldots, \mathtt{include}\,P_n, \Delta\}$ in Θ. Then the following rules are admissible:*

$$\frac{\overbrace{\Theta; \Gamma \vdash r \Leftarrow \mathtt{Mod}\,(P_i)}^{1 \le i \le n} \quad \overbrace{\Theta; \Gamma \vdash r.x \Leftarrow T[r/P]}^{x:T \in \Delta} \quad \overbrace{\Theta; \Gamma \vdash r.x \equiv t[r/P] : T[r/P]}^{x:T:=t \in \Delta} \quad \Gamma \vdash r \Rightarrow R}{\Theta; \Gamma \vdash r \Leftarrow \mathtt{Mod}\,(P)}$$

$$\frac{\overbrace{\Theta; \Gamma \vdash r_i \Leftarrow \mathtt{Mod}\,(P_i)}^{1 \le i \le n} \quad \overbrace{\Theta; \Gamma \vdash r_i \equiv r_j : P_i \cap P_j}^{1 \le i,j \le n} \quad \Theta; \Gamma \vdash \langle \delta \rangle\,[r/P] \Leftarrow [\![\Delta]\!] \quad \Gamma \vdash r \Rightarrow R}{\Theta; \Gamma \vdash \underbrace{\mathtt{Mod}\,(P) + r_1 + \ldots + r_n + \langle \delta \rangle}_{=:r} \Rightarrow \mathtt{Mod}\,(P)}$$

where $[r/P]$ abbreviates the substitution that replaces every x declared in a theory transitively-included into P with $r.x$.[6]

The first rule in Theorem 6 uses the modular structure of P to check r at type $\mathtt{Mod}\,(P)$. If r is of the form $\langle \delta \rangle$, this is no faster than flattening $\mathtt{Mod}\,(P)$

---

[4] Note that because $P \cap P'$ depends on the syntactic structure of P and P', it only approximates the least upper bound of $\mathtt{Mod}\,(P)$ and $\mathtt{Mod}\,(P')$ and is not stable under, e.g., flattening of P and P'. But it can still be very useful in certain situations.

[5] The computation of $\max P$ may look like it requires flattening. But it is easy to compute and cache its value for every named theory.

[6] In practice, these substitutions are easy to implement without flattening r because we can cache for every theory which theories it includes and which names it declares.

all the way. But in the typical case where r is also formed modularly using a similar structure as P, this can be much faster. The second rule performs the corresponding type inference for an element of $\mathrm{Mod}\,(P)$ that is formed following the modular structure of P. In both cases, the last premise is again only needed to make sure that r does not contain ill-typed fields not required by $\mathrm{Mod}\,(P)$. Also note that if we think of $\mathrm{Mod}\,(P)$ as a colimit and of elements of $\mathrm{Mod}\,(P)$ as morphisms out of P, then the second rule corresponds to the construction of the universal morphisms out of the colimit.

Example 4. We continue Example 3 and assume we have already checked `interSL` $\Leftarrow \mathrm{Mod}\,($`Semilattice`$)$ (*).

   We want to check `interSL` $+ \langle\delta\rangle \Leftarrow \mathrm{Mod}\,($`SemilatticeOrder`$)$. Applying the first rule of Theorem 6 reduces this to multiple premises, the first one of which is (*) and can thus be discharged without inspecting `interSL`.

   Example 4 is still somewhat artificial because the involved theories are so small. But the effect pays off enormously on larger theories.

5   Implementation and Case Study

We have implemented a variant of the record types and the $\mathrm{Mod}\,(\cdot)$-operator described here in the MMT-system (as part of [LFX]). They are used extensively in the *Math-in-the-Middle* archive (MitM), which forms an integral part in the OpenDreamKit [Deh+16] and MaMoRed [Koh+17] projects. In particular the formalizations of algebra and topology are systematically built on top of the concepts presented in this paper.

$$
\texttt{theory Ring} = \left\{
\begin{array}{ll}
U & : \texttt{type} \\
+ & : U \to U \to U \\
\cdot & : U \to U \to U \\
\texttt{assoc\_plus} & : \vdash \forall x, y, z : U.\ (x + y) + z \doteq x + (y + z) \\
\texttt{commutative\_plus} & : \dots \\
\dots &
\end{array}
\right\}
$$

$$
\texttt{theory Field} = \left\{
\begin{array}{l}
\texttt{include} \qquad \texttt{Ring} \\
\texttt{inverses\_times} : \vdash \forall x : U.\ x \neq 0 \Rightarrow \exists y.\ x \cdot y \doteq 1 \\
\dots
\end{array}
\right\}
$$

$$
\texttt{theory Module}(R : \mathrm{Mod}\,(\texttt{Ring})) = \left\{
\begin{array}{l}
\texttt{include} \qquad \texttt{AbelianGroup} \\
\texttt{scalar\_mult} : R.U \to U \to U \\
\dots
\end{array}
\right\}
$$

$$
\texttt{theory VSpace}(F : \mathrm{Mod}\,(\texttt{Field})) = \left\{
\begin{array}{l}
\texttt{include Module}(F) \\
\dots
\end{array}
\right\}
$$

Fig. 10. Theories for R-modules and vector spaces

The archive sources can be found at [Mit], and its contents can be inspected and browsed online at https://mmt.mathhub.info under MitM/smglom. Note that the $\text{Mod}\,(\cdot)$ operator is called ModelsOf here.

For a particularly interesting example that occurs in MitM, consider the theories for modules and vector spaces (over some ring/field) given in Fig. 10, which elegantly follow informal mathematical practice. Going beyond the syntax introduced so far, these use *parametric* theories. Our implementation extends Mod to parametric theories as well, namely in such a way that $\text{Mod}\,(\text{Module})$: $\prod_{R:\text{Mod}(\text{Ring})} \text{Mod}\,(\text{Module}(R))$ and correspondingly for fields. Thus, we obtain

$$\text{Mod}\,(\text{VSpace}) = \lambda F : \text{Mod}\,(\text{Field})\,.((\text{Mod}\,(\text{Module})\ F) + \ldots)$$

and, e.g., $\text{Mod}\,(\text{VSpace})\ \mathbb{R} <: \text{Mod}\,(\text{Module})\ \mathbb{R}$. Because of type-level parameters, this requires some kind of parametric polymorphism in the type system. For our approach, the shallow polymorphism module that is available in MMT is sufficient.

6   Conclusion

We have presented a formal system that allows to systematically combine the advantages of stratified and integrated grouping mechanisms found in type theories, logics, and specification/programming languages. Concretely, our system allows internalizing theories into record types in a way that preserves their defined fields and modular structure.

Our MitM case study shows that theory internalization is an important feature of any foundation; especially if it interfaces to differing mathematical software systems. Our experiments have also shown that (predicate) subtyping makes internalization even stronger in practice. But type-inference in the combined system induces non-trivial trade-offs; which we leave to future work.

References

[Bra13] Brady, E.: Idris, a general-purpose dependently typed programming language: design and implementation. J. Funct. Program. **23**(5), 552–593 (2013)

[Con+86] Constable, R., et al.: Implementing Mathematics with the Nuprl Development System. Prentice-Hall, Upper Saddle River (1986)

[Deh+16] Dehaye, P.-O., et al.: Interoperability in the OpenDreamKit project: the math-in-the-middle approach. In: Kohlhase, M., Johansson, M., Miller, B., de Moura, L., Tompa, F. (eds.) CICM 2016. LNCS (LNAI), vol. 9791, pp. 117–131. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-42547-4_9. https://github.com/OpenDreamKit/OpenDreamKit/blob/master/WP6/CICM2016/published.pdf

[FGT92] Farmer, W., Guttman, J., Thayer, F.: Little theories. In: Kapur, D. (ed.) Conference on Automated Deduction, pp. 467–581 (1992)

[Gog+93] Goguen, J., Winkler, T., Meseguer, J., Futatsugi, K., Jouannaud, J.: Introducing OBJ. In: Goguen, J., Coleman, D., Gallimore, R. (eds.) Applications of Algebraic Specification Using OBJ. Cambridge (1993)

[Har+12] Hardin, T., et al.: The FoCaLiZe Essential (2012). http://focalize.inria.fr/

[Har96] Harrison, J.: HOL light: a tutorial introduction. In: Srivas, M., Camilleri, A. (eds.) FMCAD 1996. LNCS, vol. 1166, pp. 265–269. Springer, Heidelberg (1996). https://doi.org/10.1007/BFb0031814

[Koh+17] Kohlhase, M., Koprucki, T., Müller, D., Tabelow, K.: Mathematical models as research data via flexiformal theory graphs. In: Geuvers, H., England, M., Hasan, O., Rabe, F., Teschke, O. (eds.) CICM 2017. LNCS (LNAI), vol. 10383, pp. 224–238. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-62075-6_16

[LFX] MathHub MMT/LFX Git Repository. http://gl.mathhub.info/MMT/LFX. Accessed 15 May 2015

[Luo09] Luo, Z.: Manifest fields and module mechanisms in intensional type theory. In: Berardi, S., Damiani, F., de'Liguoro, U. (eds.) TYPES 2008. LNCS, vol. 5497, pp. 237–255. Springer, Heidelberg (2009). https://doi.org/10.1007/978-3-642-02444-3_15

[Mat] Mathematical Components. http://www.msr-inria.fr/projects/mathematical-components-2/

[Mit] MitM/SMGLoM. https://gl.mathhub.info/MitM/smglom. Accessed 01 Feb 2018

[MRK] Müller, D., Rabe, F., Kohlhase, M.: Theories as Types. http://kwarc.info/kohlhase/submit/tatreport.pdf

[ORS92] Owre, S., Rushby, J.M., Shankar, N.: PVS: a prototype verification system. In: Kapur, D. (ed.) CADE 1992. LNCS, vol. 607, pp. 748–752. Springer, Heidelberg (1992). https://doi.org/10.1007/3-540-55602-8_217

[Rab14] Rabe, F.: How to identify, translate, and combine logics? J. Log. Comput. (2014). https://doi.org/10.1093/logcom/exu079

[Rab17] Rabe, F.: A modular type reconstruction algorithm. ACM Trans. Comput. Log. (2017). Accepted pending minor revision: https://kwarc.info/people/frabe/Research/rabe_recon_17.pdf

[RK13] Rabe, F., Kohlhase, M.: A scalable module system. Inf. Comput. **230**(1), 1–54 (2013)

[SW11] Spitters, B., van der Weegen, E.: Type classes for mathematics in type theory. CoRR abs/1102.1323 (2011). arXiv:1102.1323

[TB85] Trybulec, A., Blair, H.: Computer assisted reasoning with MIZAR. In: Joshi, A. (ed.) Proceedings of the 9th International Joint Conference on Artificial Intelligence, pp. 26–28. Morgan Kaufmann (1985)

[Wen09] Wenzel, M.: The Isabelle/Isar Reference Manual, 3 December 2009. http://isabelle.in.tum.de/documentation.html

[Coq15] Coq Development Team: The Coq Proof Assistant: Reference Manual. Technical report, INRIA (2015)