

A constraint system for a SML type error slicer

Vincent Rahli J. B. Wells Fairouz Kamareddine

ULTRA group, Heriot-Watt University

Abstract

Existing compilers for many languages have confusing type error messages. *Type error slicing* (TES) helps the programmer by isolating the part of a program contributing to a type error, but unfortunately TES was initially done for a tiny toy language. Extending TES to a full programming language is extremely challenging, and for SML we needed a number of innovations and generalisations. Some issues would be faced for any language, and some are SML-specific but representative of the complexity of language-specific issues likely to be faced for other languages. We solve both kinds of issues and present a simple, general constraint system for providing type error slices for ill-typed programs. Our constraint system elegantly and efficiently handles features like the intricate *open* SML feature. We show how the simple clarity of type error slices can demystify language features known to confuse users.

We also provide in an appendix a case study on how to use our TES to help modifying user data types, and extend the core language presented in the main body of this report to handle more of the implementation of our system. These extensions allow handling local declarations, type declarations and some uses of signatures.

1. Introduction

Higher-order type inference. SML is a higher-order function-oriented imperative programming language. SML (and similar languages like OCaml, Haskell, etc.) has polymorphic types allowing considerable flexibility, and almost fully automatic type inference, which frees the programmer from writing explicit types. We say “almost fully” because some explicit types are required in SML, e.g., as part of datatype definitions, module types, and type annotations sometimes needed in special circumstances. Milner’s W algorithm [8] is the original type-checking algorithm of the functional core of ML (variables, abstractions, applications and polymorphic let-expressions). W implementations generally give error messages relative to the syntax tree node the algorithm was visiting when unification failed, and this is often unsatisfactory.

Moving the error spot. Following W, other algorithms try to get better locations by arranging that untypability will be discovered when visiting a different syntax tree node. For example, Lee and Yi proved that the folklore algorithm M [18] finds errors “earlier” than W and claimed that their combination “can generate strictly more informative type-error messages than either of the two algorithms alone can”. Similar claims are made for W’ [20] and UAE [30]. McAdam observes that W suffers a left-to-right bias and tries to eliminate it using “unification of substitutions”. Yang claims that UAE’s primary advantage is that it also eliminates this bias. However, all the algorithms mentioned above retain a left-to-right bias in handling of let-bindings and they all blame only one syntax tree node for each type error when in fact a node set is at fault.

When only one node is reported as the error site, this node is often far away from the actual programming error. The situation is made worse because which node is blamed depends on internal implementation details, i.e., the tree node traversal order and

which constraints are accumulated and solved at different times in the traversal. The confusion is worsened because these algorithms usually exhibit in error messages (1) an internal representation of the program subtree at the blamed location which often has been transformed substantially from what the programmer wrote, and (2) details of inferred types which were not written by the programmer and which are anyway erroneous and therefore confusing.

Other improved error reporting systems. Attempting to solve this problem, constraint-based type inference algorithms [22, 23, 24] separate the two following processes: the generation of type constraints for a given term and their unification. Many works are based on this idea to improve error reporting (a probably incomplete list includes [15, 10, 11, 9, 25, 26, 27]). Independently from this separation, there exist many different approaches toward improving error reporting [32]: error explanation systems [2, 31] and error reporting systems [28]. Another approach to type error reporting is the one of Lerner et al. [19] or Hage and Heeren [12] suggesting changes to perform in the untypable code to solve type errors.

Type error slicing. Haack and Wells [11] noted that “*Identifying only one node or subtree of the program as the error location makes it difficult for programmers to understand type errors. To choose the correct place to fix a type error, the programmer must find all of the other program points that participate in the error.*” They locate type errors at *program slices* which include all parts of an untypable piece of code where changes can be made to fix the error and exclude the parts where changes cannot fix the error.

Haack and Wells gave their method of *type error slicing* (TES) for a tiny subset of SML barely larger than the λ -calculus. The TES of Haack and Wells generates constraints for SML code, enumerates minimal unsatisfiable subsets of the constraint set, and then computes type error slices. Generation and solving of constraints are not interleaved. To identify program slices responsible for type errors, each constraint is labeled by the location responsible for its generation. Error slices are portions of a program where all subterms with no responsibility for the error are elided (e.g., replaced by dots). Slices can also be shown by highlighting the source code. These slices are intended to contain all and only the information needed to solve the type errors.

The method of Haack and Wells meets the following criteria [32] for good type error reports: it reports only errors for ill-typed code (*correct*), it reports no more than the conflicting portions of code (*precise*), it reports short messages (*succinct*), it does not report internal information such as internal types generated during type inference (*a-mechanical*), it reports only code written by the programmer which has not been transformed as happens with existing SML implementations (*source-based*), it does not privilege any location over the others (*unbiased*), and it reports all the conflicting portions of code (*comprehensive*).

Slicing for a full language. We aim toward a TES method that (1) covers the full SML language, (2) is practical on real programs, and (3) has a simple and general design. As would happen for any programming language, we encountered challenges.

One challenge was avoiding a combinatorial explosion in the number of constraints. Naive constraint generation could duplicate the environment of a polymorphic declaration such as in SML’s let-expressions. Our solution is related to a constraint system by Pottier and Rémy [24, 23] although it has evolved significantly beyond that, especially to handle the challenge of SML’s *open* declaration. The most interesting constraints in their constraint system are “let-constraints” (generated for let-bindings). They are to some extent inspired by constraint-based type systems such as the one by Oder-sky, Sulzmann and Wehr [22] (and mainly by the type schemes used in that system). As explained by Pottier such constraints “allow building a constraint of linear size” [23]. We have generalised the structure of these constraints to deal with sequences of diverse (both polymorphic and monomorphic) identifier declarations (for values, types, structures, and signatures), and we developed a compatible slicing machinery. Our new constraint system replaced the type duplicating approach of Haack and Wells and gained scalability at the cost of losing compositional analysis.

Another major challenge was SML’s *open* feature which splices the declarations of a structure into the current environment. This feature has been criticized in the literature [13, 3, 4, 1]. Harper writes [13]: “it is hard to control its behaviour, since it incorporates the entire body of a structure, and hence may inadvertently shadow identifiers that happen to be also used in the structure”. Blume writes [3]: “Programs are not only read by analysis tools; human read them as well. A language construct like *open* that serves to confuse the analysis tool is also likely to confuse the human reader”. Our TES provides useful type error reports when *open* is involved, clarifying otherwise obscure type errors, and enhancing the usability of *open*. To handle errors involving *open*, we designed a simple and general machinery of “constrained environments” (definition in Sec. 4.2, example in Sec. 2.2) which goes well beyond what is supported by Pottier and Rémy’s let-constraints.

Another challenge is SML’s value identifier statuses. In SML, a value identifier can be a value variable (the only status considered by Haack and Wells), a datatype constructor, or an exception constructor (omitted in this paper’s formalism). For example, if identifier *c* has value variable status in the context, $\text{fn } c \Rightarrow (c \ 1, \ c())$ has a unique minimal error which is that *c* has a monomorphic type but is applied to two expressions with different types: *int* and *unit*. However, this error would not exist if the code was preceded by, e.g., `datatype t = c` because the *fn*-binding would not bind *c*, but instead there would be a minimal error that *c* is declared as a nullary datatype constructor and is applied to an argument in *c* 1. To compute correct type error slices, we annotate constraints by context dependencies on identifier statuses. For the *fn*-binding presented above we generate during unification constraints relating the occurrences of *c* annotated by the dependency that *c* is a value variable and not a datatype constructor. These constraints are not generated if a context confirms that *c* must be a datatype constructor. The constraints but not the context dependency are generated if a context confirms that *c* cannot be a datatype constructor. When handling incomplete programs, we report conditional errors (warnings) that assume a sensible default truth status for the dependencies.

Later sections detail solving these and other challenges.

2. Key motivating examples

This section gives motivating examples of TES. Type error slices are highlighted with very light grey. Dark grey highlights error *end points* (e.g., the sources of conflicting types constrained to be equal). A color version has been made available.

2.1 Datatypes, pattern matching and type functions. Fig. 1 shows how TES is important for intricate errors. The code declares the datatype *t* and the function *trans* to deal with user defined

Figure 1 Datatypes, pattern matching and type functions

```
datatype ('a, 'b, 'c) t = Red    of 'a * 'b * 'c
                       | Blue  of 'a * 'b * 'c
                       | Pink  of 'a * 'b * 'c
                       | Green  of 'a * 'b * 'b①
                       | Yellow of 'a * 'b * 'c
                       | Orange of 'a * 'b * 'c
fun trans (Red    (x, y, z)) = Blue  (y, x, z)
  | trans (Blue  (x, y, z)) = Pink  (y, x, z)
  | trans (Pink  (x, y, z)) = Green  (y, x, z)
  | trans (Green  (x, y, z)) = Yellow (y, x, z)③
  | trans (Yellow (x, y, z)) = Orange (y, x, z)
  | trans (Orange (x, y, z)) = Red    (y, x, z)
type ('a, 'b) u = ('a, 'a, 'b) t * 'b
val x = (Red (2, 2, false), true)⑤
val y : (int, bool) u = (trans (#1 x), #2 x)④
```

Figure 2 Chained *opens* and nested structures

```
structure S = struct
  structure Y = struct
    structure A = struct val x = false end
    structure X = struct val x = false end
    structure M = struct val x = true end
  end
  open Y
  val m = M.x
  val x = if m then true else false
end
structure T = struct
  structure X = struct val x = 1 end
  open S
  open X
  val y = if m then 1 else x
end
```

colours. This function is then applied to an instance of a colour (the first element in the pair *x*). Assume that our programming error is that we wrote *'b* instead of *'c* in *Green*’s definition at location ①. SML/NJ (version 110.72) reports a type constructor clash at ④:

```
operator domain: (int,int,int) t
operand:         (int,int,bool) t
in expression:
  trans ((fn {1=<pat>,...} => 1) x)
```

The reported code does not resemble our code and is far away from the programming error location. SML/NJ gives the same error message if, instead of the error described above, we write *x* instead of *z* in the right-hand-side of any branch of *trans*. This means with SML/NJ one must check the entire program to find the error.

Fig. 1 shows one of the type error slices reported by our type error slice, highlighted in the code. This error is context-dependent: it assumes that *y* and *z* are value variables and not datatype constructors. The programming error location being in the slice, we track it down by considering only the highlighted portions of code, starting from the clashing types on the last line. The type $(\text{int}, \text{bool}) \ u$ constrains the type of *trans*’s application and the highlighted portion of *trans* is when applied to a *Green* object. At ①, *Green*’s second and third arguments are constrained to be of the same type. At ②, *y* is incidentally constrained to be of the same type as *z*. At ③, because *y* and *z* are respectively the first and third arguments of *Yellow* and using *Yellow*’s definition, we infer that the type of the application of *Yellow* to its three arguments (returned by *trans*) is *t* where its first and third parameters have to be equal. At ④ and ⑤ we can see that *trans* is constrained to return a *t* where its first (*int*) and third (*bool*) parameters differ.

2.2 Chained *opens* and nested structures. Fig. 2 presents a type error involving nested *opens* leading to intricate type errors. Let us describe what the code was meant to do. In the structure *T*, we declare a structure *X* declaring an integer *x*. We then open the structure *S* to access the Boolean *m*. We then open *X* to access the

integer x . Finally, if m is true then we return 1 otherwise we return x . Unfortunately, this piece of code is untypable and SML/NJ reports the following error message which blames y 's body:

```
Error: types of if branches do not agree [literal]
  then branch: int
  else branch: bool
  in expression:
    if m then 1 else x
```

The programming error here, as our type error slice explains clearly, is that opening S causes S 's declarations to shadow the current typing environment. Because Y is opened in S , the three structures A , X and M are part of S 's declarations. Hence, when opening S in T , the structure X which was in our current typing environment is shadowed by the one defined in Y . One can solve this programming error by replacing “open S open X ” by “open S X ”.

Our type error slice rules out x 's declarations in X and S and clearly shows why x does not have the expected type. SML/NJ's report leaves us to track down x 's binding by hand.

2.3 Merged minimal error slices. We have found cases needing the display of many minimal errors at once. One important case is in record field name clashes where, e.g., the highlighting `val {foo,bar} = {foo1=0,bar=1}` reports two minimal errors at once: that `foo1` is not in `{foo,bar}` and `foo` is not in `{foo1,bar}`. This merged error is preferable over the minimal errors because of the explosion in the number of minimal slices. Light grey highlights the fields that are common to different minimal slices. For merged slices minimality is understood as follows: retain a single dark grey field name in one of the two clashing records and all field names in the other.

3. Mathematical definitions and notations

Let i, j, n, m be metavariables ranging over \mathbb{N} , the set of natural numbers. If a metavariable v ranges over a class C , then the metavariables v_x (where x can be anything) and the metavariables v', v'' , etc., also range over C . Let s range over sets. If v ranges over s , then let \bar{v} range over $\mathbb{P}(s)$, the power set of s . Let $\text{dj}(s_1, \dots, s_n)$ (“disjoint”) hold iff for all $i, j \in \{1, \dots, n\}$, if $i \neq j$ then $s_i \cap s_j = \emptyset$. Let $s_1 \uplus s_2$ be $s_1 \cup s_2$ if $\text{dj}(s_1, s_2)$ and undefined otherwise. Let R range over binary relations (we write $\langle x, y \rangle$ for a pair). Given a relation R let $\text{dom}(R) = \{x \mid \langle x, y \rangle \in R\}$ and $\text{ran}(R) = \{y \mid \langle x, y \rangle \in R\}$. Let $s \triangleleft R = \{\langle x, y \rangle \in R \mid x \notin s\}$. Let f range over functions, let $s \rightarrow s' = \{f \mid \text{dom}(f) \subseteq s \wedge \text{ran}(f) \subseteq s'\}$, and let $x \mapsto y$ be an alternative notation for $\langle x, y \rangle$ used when writing some functions. A tuple t is a function such that $\text{dom}(t) \subseteq \mathbb{N}$ and if $1 \leq k \in \text{dom}(t)$ then $k-1 \in \text{dom}(t)$. Let t range over tuples. We write the tuple $\{0 \mapsto x_0, \dots, n \mapsto x_n\}$ as $\langle x_0, \dots, x_n \rangle$. We define the appending $\langle x_1, \dots, x_i \rangle @ \langle y_1, \dots, y_j \rangle$ of two tuples as the tuple $\langle x_1, \dots, x_i, y_1, \dots, y_j \rangle$. If v ranges over s , \vec{v} is defined to range over $\text{tuple}(s) = \{t \mid \text{ran}(t) \subseteq s\}$.

4. Technical design of our TES

The different modules of our TES are: constraint generation (Sec. 4.3), constraint solving (Sec. 4.4), minimisation (Sec. 4.5), enumeration (Sec. 4.5), and slicing (Sec. 4.6). Sec. 4.8 discusses minimality, Sec. 4.7 defines the overall algorithm of our type slicer, and Sec. 4.9 discusses the principles of our approach.

4.1 External syntax. Fig. 3 (upper half) defines our external syntax which is a subset of the SML syntax. Many syntactic forms are annotated with labels (l). These labels are generated by our TES to track locations responsible for inferences made during analysis. To provide a visually convenient place for labels, expression applications are surrounded by $\lceil \ \rceil$ which are not seen by programmers but are part of an internal representation used to avoid confusion with

$()$ as part of SML syntax. Value identifiers (vid) are subscripted to distinguish between occurrences in expressions (vid_e^l), datatype constructor definitions (vid_c^l), and patterns (vid_p^l).

4.2 Constraint syntax. Constraint terms. Fig. 3 (lower half) defines our constraint terms.

In addition to distinguishing identifier classes (VId for value identifiers, TyCon for type constructor names, etc.), SML assigns statuses within the value identifier class to distinguish value variables, datatype constructors, and exception constructors. Because SML has no lexical distinction between, e.g., a datatype constructor and a value variable, a value identifier's status cannot always be inferred from any context smaller than the entire program.

In our constraint system, an identifier status can either be a raw status (ris) or a raw status annotated with dependencies (is). The v status is for value variables (e.g., the recursive function f in `val rec f = fn x => x` is a value variable and not a datatype constructor). Statuses c and d are for unary and nullary datatype constructors respectively (e.g., c in `datatype 'a t = C of 'a` and d in `datatype 'a t = D`). Status u is for unconfirmed context-dependent statuses (e.g., in `fn x => x`, the identifier x could be a value variable or a nullary datatype constructor). Status p is for unresolvable statuses such as in `let open S in fn x => x end`, where x could be declared as a value variable as well as a datatype constructor in the free structure S . Finally, status a is similar to a variable as it can be any status (used by our constraint filtering function `filt` defined in Fig. 8 in Sec. 4.4 to generate dummy environments that cannot participate in type errors).

Some syntactic forms, called *dependent forms*, are annotated by dependencies: $\langle x, \bar{d} \rangle$. A dependency d can be a label l or a value identifier vid . During analysis, if a dependency d is a label l , the annotated syntactic form depends on the program node labelled by l . For example, if the dependent equality constraint

$\tau_1 \xrightarrow{\bar{d} \cup \{l\}} \tau_2$ is generated for the annotated code $\lceil \text{exp atexp} \rceil^l$, then the equality constraint $\tau_1 = \tau_2$ depends on the application root node of $\lceil \text{exp atexp} \rceil^l$. If d is a value identifier vid , then the syntactic form depends on vid 's status in the code being a v or u . Because identifiers' statuses are resolved during constraint solving, such dependencies (value identifiers) are only generated during constraint solving and not during initial constraint generation. For example, if constraint solving generates the dependent equality constraint $\tau_1 \xrightarrow{\bar{d} \cup \{vid\}} \tau_2$, then the equality constraint $\tau_1 = \tau_2$ needs to be satisfied only if vid cannot be a datatype constructor. Let `strip` be the function that strips off the outer dependencies (not nested under another constructor than the dependency constructor) of any syntactic form: `strip(x) = strip(y)` if $x = \langle y, \bar{d} \rangle$ and x otherwise.

An internal type of the form $\tau \mu$ is called a *type construction* and is built from an internal type constructor μ and its argument τ (such as the polymorphic list type `'a list`, where `'a` is an explicit type variable in SML). We only allow type constructors to take one parameter in this paper and so we only allow internal type constructors to take one parameter in our constraint system. The internal type constructor `ar` is used during constraint solving to represent the arrow type constructor so that we can generate type constraints between the arrow type constructor and any other unary type constructor. This is necessary to compute the necessary portions of code when generating type errors. A type scheme can either be a universal quantification or an internal type.

We use \downarrow to represent bindings (as in $\downarrow id = x$ that associates the semantics x to the binding occurrence id) and \uparrow to constrain the semantics of non-binding occurrences (also called accessors) of identifiers (as in $\uparrow id = x$ that constrains id 's semantics to be x). A binder of the form $\downarrow vid = \alpha$, is an unconfirmed binder that can either be confirmed to be a binder of a value variable at constraint solving

Figure 3 External syntax and constraint system

external syntax		(what the programmer sees, plus labels)	
$l \in \text{Label}$	(labels)	$dec \in \text{Dec}$	$::= \text{val rec pat} \stackrel{l}{=} \text{exp} \mid \text{datatype dn} \stackrel{l}{=} \text{cb} \mid \text{open}^l \text{sid}$
$vid \in \text{VId}$	(value identifiers)	$atexp \in \text{AtExp}$	$::= vid_e^l \mid \text{let}^l \text{dec in exp end}$
$sid \in \text{StrId}$	(structure identifiers)	$exp \in \text{Exp}$	$::= atexp \mid \text{fn pat} \stackrel{l}{\Rightarrow} \text{exp} \mid [\text{exp atexp}]^l$
$tv \in \text{TyVar}$	(type variables)	$atpat \in \text{AtPat}$	$::= vid_p^l$
$tc \in \text{TyCon}$	(type constructors)	$pat \in \text{Pat}$	$::= atpat \mid vid^l \text{atpat}$
$ty \in \text{Ty}$	$::= tv^l \mid ty_1 \stackrel{l}{\rightarrow} ty_2 \mid ty \text{ tc}^l$	$sdec \in \text{StrDec}$	$::= dec \mid \text{structure sid} \stackrel{l}{=} \text{sexp}$
$cb \in \text{ConBind}$	$::= vid_c^l \mid vid \text{ of}^l \text{ ty}$	$sexp \in \text{StrExp}$	$::= sid^l \mid \text{struct}^l \text{sdec}_1 \dots \text{sdec}_n \text{ end}$
$dn \in \text{DatName}$	$::= tv \text{ tc}^l$		
constraint terms (syntax of entities used internally by the type error slicer and which the programmer never sees)			
$ev \in \text{EnvVar}$	(environment variables)	$\mu \in \text{ITyCon}$	$::= \delta \mid \gamma \mid \text{ar} \mid \langle \mu, \bar{d} \rangle$
$\delta \in \text{TyConVar}$	(type constructor variables)	$\tau \in \text{ITy}$	$::= \alpha \mid \tau \mu \mid \tau_1 \rightarrow \tau_2 \mid \langle \tau, \bar{d} \rangle$
$\gamma \in \text{TyConName}$	(type constructor names)	$\sigma \in \text{Scheme}$	$::= \tau \mid \forall \bar{\alpha}. \tau$
$\alpha \in \text{ITyVar}$	(internal type variables)	$bind \in \text{Bind}$	$::= \downarrow tc = \mu \mid \downarrow sid = e \mid \downarrow tv = \alpha \mid \downarrow vid = \sigma \mid \downarrow vid = is \mid \uparrow vid = \alpha$
$d \in \text{Dependency}$	$::= l \mid vid$	$acc \in \text{Accessor}$	$::= \uparrow tc = \delta \mid \uparrow sid = ev \mid \uparrow tv = \alpha \mid \uparrow vid = \alpha \mid \uparrow vid = ris$
$ris \in \text{RawIdStatus}$	$::= v \mid c \mid d \mid u \mid p \mid a$	$c \in \text{Constraint}$	$::= \mu_1 = \mu_2 \mid e_1 = e_2 \mid \tau_1 = \tau_2 \mid is_1 = is_2$
$is \in \text{IdStatus}$	$::= ris \mid \langle is, \bar{d} \rangle$	$e \in \text{Env}$	$::= \square \mid ev \mid bind \mid acc \mid c \mid \text{poly}(e) \mid e_2; e_1 \mid \langle e, \bar{d} \rangle$
extra metavariables mostly used in side conditions			
$id \in \text{Id}$	$::= vid \mid sid \mid tv \mid tc$	$var \in \text{Var}$	$::= \alpha \mid \delta \mid ev$
		$dep \in \text{Dependent}$	$::= \langle \tau, \bar{d} \rangle \mid \langle \mu, \bar{d} \rangle \mid \langle e, \bar{d} \rangle \mid \langle is, \bar{d} \rangle$

time, and so be turned into a binder of the form $\downarrow vid = \alpha$ or be turned into an accessor $\uparrow vid = \alpha$ if it turns out that vid is a datatype constructor. This mechanism is further illustrated in Sec. 4.4.

The keystone of our constraint system is the constrained environment $e_1; e_2$ where e_1 constrains e_2 . The environment $e_1; e_2$ builds a new environment from its two components where references of the form $\uparrow id$ (accessors) in e_2 can depend on occurrences of $\downarrow id$ (binders) in e_1 . For example, in $\downarrow vid = \sigma; \uparrow vid = \alpha$, α is constrained to be σ through the binding of vid . The motivation for these environments is to have a general mechanism to build environments for sequential declarations.

In addition to a constrained environment of the form $e_1; e_2$, an environment can also be an empty environment \square , an environment variable ev , a binder $bind$ associating static semantics to identifiers, an accessor to look identifiers' static semantics up in environments, an equality constraint c , a special form $\text{poly}(e)$ (explained below) which grants e the possibility to be polymorphic, or a conditional environment $\langle e, \bar{d} \rangle$ depending on \bar{d} . Environments are special kinds of constraint (on internal types, internal type constructors, environments and statuses).

Let $e_1; \dots; e_n \in \square$ if $n = 0$ and $(e_1; \dots; e_{n-1}); e_n$ if $n > 0$.

“Atomic” syntactic forms. Let $\text{atoms}(x)$ be the set of syntactic forms belonging to $\text{Var} \cup \text{TyConName} \cup \text{Dependency}$ and occurring in x whatever x is. We define the following functions:

$$\begin{aligned} \text{vars}(x) &= \text{atoms}(x) \cap \text{Var} && \text{(set of variables)} \\ \text{labs}(x) &= \text{atoms}(x) \cap \text{Label} && \text{(set of labels)} \\ \text{deps}(x) &= \text{atoms}(x) \cap \text{Dependency} && \text{(set of dependencies)} \end{aligned}$$

Freshness of variables. We use distinguished dummy variables: $\text{DumVar} = \{\alpha_{\text{dum}}, ev_{\text{dum}}, \delta_{\text{dum}}\}$. Each use of a dummy variable acts like a fresh variable. These variables are used to generate dummy environments and constraints. For example, the equality constraint $ev_{\text{dum}} = e$ means that the environment e must be solved and does not constrain any other environment. The relation dja ensures the freshness of the generated variables and type constructor names: $\text{dja}(x_1, \dots, x_n) \Leftrightarrow \text{dj}(f(x_1), \dots, f(x_n), \text{DumVar})$, where $f(x) = \text{atoms}(x) \setminus \text{VId}$. This also ensures that each label occurs at most once in a labelled program.

Syntactic sugar. We write $\langle x, d \rangle$ for $\langle x, \{d\} \rangle$. Let y be a d or a \bar{d} . We write x^y for $\langle x, y \rangle$. We write $x_1 \stackrel{y}{=} x_2$ for $\langle x_1 = x_2, y \rangle$, and similarly for $bind$'s and acc 's. We write $\downarrow vid \stackrel{y}{=} \langle \sigma, is \rangle$ for $\downarrow vid \stackrel{y}{=} is; \downarrow vid \stackrel{y}{=} \sigma$, and $\uparrow vid \stackrel{y}{=} \langle \alpha, ris \rangle$ for $\uparrow vid \stackrel{y}{=} \alpha$.

$ris; \uparrow vid \stackrel{y}{=} \alpha$ We write $[e]$ for $(ev_{\text{dum}} = e)$. Such a constraint defines a local environment e which is not visible from outside the constraint. This is used for local bindings by rules (G2) and (G4) of our constraint generation algorithm defined in Fig. 4.

4.3 Constraint generation. Value bindings. At constraint generation (Fig. 4), in the pattern rule (G5), we generate monomorphic, unconfirmed binders of the form $\uparrow vid = \alpha$ where no type variable is yet quantified over. These binders are monomorphic because in SML, e.g., the type of a recursive function such as f in the let-expression $\text{let val rec } f = \text{fn } x \Rightarrow f \ x \text{ in } f \ \text{end}$, is monomorphic within its definition (f 's first and second occurrences' types are equal) and generalised into a polymorphic *for all* type scheme when typing the declaration (f 's third occurrence's type is an instance of the generalisation of f 's first occurrence's type). An environment e is then turned into a polymorphic one during constraint solving (using toPoly defined in Fig. 6 in Sec. 4.4) if marked as follows: $\text{poly}(e)$. Such forms are generated by the recursive value declaration rule (G12) and the datatype declaration rule (G13). In (G5) again, the binder is unconfirmed and no status constraint is generated (as opposed to, e.g., rule (G10) which forces the analysed identifier to be a nullary datatype constructor) because in SML, e.g., in $\text{fn } x \Rightarrow x$, without any more context, the identifier x could be a value variable or a datatype constructor. The status of x is then unknown. Because recursive functions are forced to be value variables (v) even when in the scope of a datatype constructor binding, toV (used by (G12)) generates a status constraint:

$$\begin{aligned} \text{toV}(e_1; e_2) &= \text{toV}(e_1); \text{toV}(e_2) \\ \text{toV}(\uparrow vid \stackrel{l}{=} \alpha) &= (\downarrow vid \stackrel{l}{=} \langle \alpha, v \rangle) \\ \text{toV}(e) &= e, \text{ if none of the above applies} \end{aligned}$$

As explained in Sec. 4.4, at constraint solving, an unconfirmed binder of the form $\uparrow vid = \alpha$ eventually turns into a binder of the form $\downarrow vid = \alpha$ or an accessor of the form $\uparrow vid = \alpha$. (In some cases, a status constraint is also generated from an unconfirmed binder.)

Algorithm. Fig. 4 defines our constraint generator. At initial constraint generation, the only labelled environments are equality constraints (c), binders ($bind$), accessors (acc), and environment variables (ev).

In rule (G12) for recursive function declarations, the environment $\text{toV}(e_1)$ generated for the pattern part of the declaration constrains the environment e_2 generated for the expression part. This order is necessary to handle the recursivity of such declarations.

Figure 4 Constraint generation rules

Expressions	(G1) $vid_e^l \triangleright \langle \alpha, \uparrow vid \stackrel{l}{=} \alpha \rangle$		
	(G2) $let^l dec \text{ in } exp \text{ end} \triangleright \langle \alpha, [e_1; e_2; (\alpha \stackrel{l}{=} \alpha_2)] \rangle$	$\Leftarrow dec \triangleright e_1 \wedge exp \triangleright \langle \alpha_2, e_2 \rangle \wedge dja(e_1, e_2, \alpha)$	
	(G3) $[exp \text{ atexp}]^l \triangleright \langle \alpha, e_1; e_2; (\alpha \stackrel{l}{=} \alpha_2 \rightarrow \alpha) \rangle$	$\Leftarrow exp \triangleright \langle \alpha_1, e_1 \rangle \wedge atexp \triangleright \langle \alpha_2, e_2 \rangle \wedge dja(e_1, e_2, \alpha)$	
	(G4) $fn \text{ pat} \stackrel{l}{\Rightarrow} exp \triangleright \langle \alpha, [(ev=e_1); ev^l; e_2; (\alpha \stackrel{l}{=} \alpha_1 \rightarrow \alpha_2)] \rangle$	$\Leftarrow pat \triangleright \langle \alpha_1, e_1 \rangle \wedge exp \triangleright \langle \alpha_2, e_2 \rangle \wedge dja(e_1, e_2, \alpha, ev)$	
Patterns	(G5) $vid_p^l \triangleright \langle \alpha, \downarrow vid \stackrel{l}{=} \alpha \rangle$	(G6) $vid^l \text{ atpat} \triangleright \langle \alpha_2, \alpha_1 \stackrel{l}{=} \alpha \rightarrow \alpha_2; \uparrow vid \stackrel{l}{=} \langle \alpha_1, c \rangle; e \rangle$	$\Leftarrow atpat \triangleright \langle \alpha, e \rangle \wedge dja(e, \alpha_1, \alpha_2)$
Types	(G7) $tv^l \triangleright \langle \alpha, \uparrow tv \stackrel{l}{=} \alpha \rangle$	(G8) $ty \text{ tc}^l \triangleright \langle \alpha', (\uparrow tc \stackrel{l}{=} \delta); (\alpha' \stackrel{l}{=} \alpha \delta); e \rangle$	$\Leftarrow ty \triangleright \langle \alpha, e \rangle \wedge dja(e, \alpha', \delta)$
	(G9) $ty_1 \stackrel{l}{\rightarrow} ty_2 \triangleright \langle \alpha, e_2; e_1; (\alpha \stackrel{l}{=} \alpha_1 \rightarrow \alpha_2) \rangle$	$\Leftarrow ty_1 \triangleright \langle \alpha_1, e_1 \rangle \wedge ty_2 \triangleright \langle \alpha_2, e_2 \rangle \wedge dja(e_1, e_2, \alpha)$	
Constructor bindings	(G10) $vid_c^l \triangleright \langle \alpha, \downarrow vid \stackrel{l}{=} \langle \alpha, d \rangle \rangle$		
	(G11) $vid \text{ of }^l ty \triangleright \langle \alpha_1, e; \alpha_2 \stackrel{l}{=} \alpha \rightarrow \alpha_1; \downarrow vid \stackrel{l}{=} \langle \alpha_2, c \rangle \rangle$	$\Leftarrow ty \triangleright \langle \alpha, e \rangle \wedge dja(e, \alpha_1, \alpha_2)$	
Declarations	(G12) $val \text{ rec } pat \stackrel{l}{=} exp \triangleright (ev = poly(\text{toV}(e_1); e_2; (\alpha_1 \stackrel{l}{=} \alpha_2))); ev^l$	$\Leftarrow pat \triangleright \langle \alpha_1, e_1 \rangle \wedge exp \triangleright \langle \alpha_2, e_2 \rangle \wedge dja(e_1, e_2, ev)$	
	(G13) $datatype \text{ dn} \stackrel{l}{=} cb \triangleright (ev = ((\alpha_1 \stackrel{l}{=} \alpha_2); e_1; poly(e_2))); ev^l$	$\Leftarrow dn \triangleright \langle \alpha_1, e_1 \rangle \wedge cb \triangleright \langle \alpha_2, e_2 \rangle \wedge dja(e_1, e_2, ev)$	
	(G14) $open^l \text{ sid} \triangleright (\uparrow sid \stackrel{l}{=} ev); ev^l$		
Datatype names	(G15) $tv \text{ tc}^l \triangleright \langle \alpha', (\alpha' \stackrel{l}{=} \alpha \gamma); (\downarrow tc \stackrel{l}{=} \gamma); (\downarrow tv \stackrel{l}{=} \alpha) \rangle$	$\Leftarrow \alpha \neq \alpha'$	
Structure declarations	(G16) $structure \text{ sid} \stackrel{l}{=} sexp \triangleright (ev' = (e; (\downarrow sid \stackrel{l}{=} ev))); ev'^l$	$\Leftarrow sexp \triangleright \langle ev, e \rangle \wedge dja(e, ev')$	
Structure expressions	(G17) $sid^l \triangleright \langle ev, \uparrow sid \stackrel{l}{=} ev \rangle$		
	(G18) $struct^l \text{ sdec}_1 \dots \text{ sdec}_n \text{ end} \triangleright \langle ev, (ev \stackrel{l}{=} ev'); (ev' = (e_1; \dots; e_n)) \rangle$	$\Leftarrow \text{ sdec}_1 \triangleright e_1 \wedge \dots \wedge \text{ sdec}_n \triangleright e_n \wedge dja(e_1, \dots, e_n, ev, ev')$	

In rule (G13) for datatype declarations, the environment e_1 generated for the declared type constructor constrains the environment $poly(e_2)$ generated for the datatype constructor of the declared type constructor. This order is necessary to handle the recursivity of such datatype declarations (in `datatype nat = z | s of nat`, the second occurrence of `nat` refers to its first occurrence).

Rule (G14) for structure opening (as for rules (G4), (G12), (G13) and (G16)) labels an environment variable, so that a sliced out declaration does not shadow its environment. Without this label, the environment variable would be a constraint that always have to be satisfied. With the label, the environment variable is a constraint that has to be satisfied only when the declaration is not sliced out. The link between the environment variable and the structure to open is made via the labelled accessor $\uparrow sid \stackrel{l}{=} ev$.

Rule (G18) for structure expressions (as for rules (G4), (G12), (G13) and (G16)) generates unlabelled equality constraints. An unlabelled equality constraint such as a constraint of the form $ev' = (e_1; \dots; e_n)$ generated by (G18) needs to be unlabelled because each of the e_i is not dependent on the analysed structure expression itself but is dependent on the corresponding declaration packed together with other declarations in the structure expression. The information related to the analysed structure expression, carried by the unlabelled constraint is the fact that a sequence of declarations (corresponding to the constrained environment $e_1; \dots; e_n$) is packed into a structure. This information depends on the analysed structure expression via the extra labelled equality constraint $ev \stackrel{l}{=} ev'$. In (G4), (G12), (G13) and (G16), we use labelled environment variables of the form ev^l .

4.4 Constraint solving. Syntax. Fig. 5 defines the syntactic forms used by our constraint solver (Fig. 7) where one unification step is defined by the relation \rightarrow , where \rightarrow^* is its reflexive and transitive closure. To each state of a unification computation (except the errors), a unification context $\Delta = \langle u, e \rangle$ is associated. In a state $\text{solve}(\langle u, e \rangle, \bar{d}, e')$, $\langle u, e \rangle$ is the context in which e' must be solvable for the algorithm to succeed. Let $\langle u, e \rangle(\text{var})$ be $u(\text{var})$, let $\langle u, e \rangle; e'$ be $\langle u, e; e' \rangle$, and let $u_1 \boxplus u_2$ be $u_1 \cup (\text{DumVar} \triangleleft u_2)$ if $\text{dj}(\text{dom}(u_1), \text{dom}(u_2))$, and undefined otherwise.

Given constraints (of the form e), our constraint solver either succeeds with $\text{succ}(\Delta)$ returning its current unification context Δ ,

or fails with $\text{err}(er)$ returning an error which can be (see ek in Fig. 5) a type constructor clash, a circularity error or a status clash (discussed below with the compatible relation). The application of a renaming ren to an internal type τ is defined as usual and, is denoted $\tau[ren]$. Renamings are used to instantiate type schemes.

Environment application. Constraint solving maintains a type environment (e in Δ) where some parts might be shadowed and so inaccessible. For example, in $bind_2; ev; bind_1^d$, the usable part is $bind_1$ and ev shadows $bind_2$ because an environment variable stands for any environment and could potentially bind any identifier. During unification, no c or acc occurs in the e stored in a Δ because they are transformed into unifiers u (rules (U3) and (U4) in Fig. 7). Similarly, the $poly(e)$ and $\downarrow vid = \alpha$ forms are eliminated. Concerning \square , we never add it to a unification context Δ while unifying constraints, but we always start a unification process with the initial unification context $\langle \emptyset, \square \rangle$. Thus, during unification if $\Delta = \langle u, e \rangle$ then e is of the form $\square; e_1 \dots; e_n$, where each e_i is either an ev or a $\downarrow id \stackrel{\bar{d}}{=} x$. Let the predicate `hiding` be defined as follows: `hiding(e)` be true iff $e \in \text{Var}$ or e is of the form $e_1; e_2$ and `hiding(e_i)` for $i = 1$ or $i = 2$. Let `hiding`($\langle u, e \rangle$) be true iff `hiding(e)`. We define the application $e[id]$ as follows:

$$\begin{aligned}
 \text{(EA1)} \quad & (e'; \downarrow id \stackrel{\bar{d}}{=} \forall \bar{\alpha}. \tau)(id) = \forall \bar{\alpha}. \tau^{\bar{d}} \\
 \text{(EA2)} \quad & (e'; \downarrow id \stackrel{\bar{d}}{=} x)(id) = x^{\bar{d}}, \quad \text{if } x \text{ of the form } \tau, \mu \text{ or } e \\
 \text{(EA3)} \quad & (e'; \downarrow id \stackrel{\bar{d}}{=} x)(id) = e'(id), \text{ if } id \neq id' \text{ or } x \in \text{IdStatus}
 \end{aligned}$$

We define $e[[id]]$ to access value identifiers' statuses:

$$\begin{aligned}
 \text{(EAS1)} \quad & (e'; \downarrow id \stackrel{\bar{d}}{=} is)[[id]] = is^{\bar{d}} \\
 \text{(EAS2)} \quad & (e'; \downarrow id' \stackrel{\bar{d}}{=} x)[[id]] = e'[[id]], \text{ if } id \neq id' \text{ or } x \notin \text{IdStatus}
 \end{aligned}$$

For example, $(\downarrow vid \stackrel{\bar{d}_3}{=} \sigma; \downarrow vid \stackrel{\bar{d}_2}{=} v; \downarrow sid \stackrel{\bar{d}_1}{=} e)(vid) = \sigma$ but $(\downarrow vid \stackrel{\bar{d}_3}{=} \sigma; \downarrow vid \stackrel{\bar{d}_2}{=} v; ev^{\bar{d}}; \downarrow sid \stackrel{\bar{d}_1}{=} e)(vid)$ and $(\downarrow vid \stackrel{\bar{d}_3}{=} \sigma; \downarrow vid \stackrel{\bar{d}_2}{=} v; \downarrow sid \stackrel{\bar{d}_1}{=} e)(tc)$ are undefined.

Let $\Delta(id) = e(id)$ and $\Delta[[id]] = e[[id]]$, where $\Delta = \langle u, e \rangle$.

Context dependencies solving. Context dependencies are solved during unification. An unconfirmed binder of the form $\downarrow vid = \alpha$ is then either turned into a binder of the form $\downarrow vid = \alpha$ or

Figure 5 Syntactic forms used by the constraint solver

$er \in \text{Error} ::= \langle ek, \bar{d} \rangle$	$u \in \text{Unifier} = \{f_1 \cup f_2 \cup f_3 \mid f_1 \in \text{ITyVar} \rightarrow \text{ITy} \wedge f_2 \in \text{TyConVar} \rightarrow \text{ITyCon} \wedge f_3 \in \text{EnvVar} \rightarrow \text{Env}\}$
$\Delta \in \text{UnifEnv} ::= \langle u, e \rangle$	$ek \in \text{ErrKind} ::= \text{tyConsClash}(\mu_1, \mu_2) \mid \text{statusClash}(is_1, is_2) \mid \text{circularity}$
	$state \in \text{State} ::= \text{solve}(\Delta, \bar{d}, e) \mid \text{succ}(\Delta) \mid \text{err}(er)$
	$ren \in \text{Ren} = \{ren \in \text{ITyVar} \rightarrow \text{ITyVar} \mid ren \text{ is injective} \wedge \text{dj}(\text{dom}(ren), \text{ran}(ren))\}$

an accessor of the form $\uparrow vid = \alpha$ by one of these rules: (B2)-(B5). These rules make use of the function `ifNotDum` which is defined as follows: `ifNotDum(α_{dum}, is) = a` and `ifNotDum(α, is) = is` if $\alpha \notin \text{DumVar}$. This function is required to ensure that a dummy binder cannot bind something else than a dummy status. Rule (B2) discards binders generated under unsatisfied context dependencies, e.g., in `let datatype t = x in fn x => x end`, x 's second occurrence does not bind x 's third occurrence because of x 's declaration as a datatype constructor. The unconfirmed binder is then turned into an accessor. In all three other rules, the unconfirmed binder is turned into a confirmed one. Rule (B3) validates context dependencies, e.g., in `val rec x = fn x => x`, x is confirmed to be a value variable, x 's second occurrence being in the scope of x 's first occurrence which is a recursive function and so in SML is forced to be a value variable and not a datatype constructor. Rule (B4) generates context dependencies, e.g., in `fn x => x`, because x can be a value variable as well as a datatype constructor then x 's second occurrence is bound to x 's first occurrence under the context dependency that x is not a datatype constructor. Rule (B5) generates dummy environments when there is not enough information to check whether a context dependency is satisfied or not, e.g., in `let open S in fn x => x end`, if S is free, it might declare x as a datatype constructor or as a recursive function. Thus, we do not allow x to be a monomorphic binding but we still generate a dummy binding to catch status clashes (e.g., if instead of the second x we had `fn (x y) => y` where x is a unary datatype constructor, we would then have x occurring in patterns both at a nullary position and a unary position).

Status compatibility. Two identifier statuses are incompatible iff a unary datatype constructor, occurring in a pattern, is bound to a (context-dependent or independent) value variable as in `let val rec f = fn x => x in fn (f x) => x end` where f 's first occurrence is a value variable and f 's second occurrence is a unary datatype constructor (taking an argument in a pattern); or if a value variable in a pattern (not applied) is bound to a unary datatype constructor as in `let datatype t = x of int in fn x => x end`.

$$\text{compatible}(is_1, is_2) \Leftrightarrow \{is_1, is_2\} \not\subseteq \{\{c, v\}, \{c, u\}, \{c, p\}\}$$

Status compatibility is checked by constraint solving rules (S6) and (S7) (in Fig. 7). Compatibility is only defined on raw statuses because constraint solving rule (S8) removes dependencies on (among other things) statuses.

Building of constraint terms. The constraint solver uses `build` to build, w.r.t. a given unifier, polymorphic types (Fig. 6), check circularity errors (in order not to generate a unifier where, e.g., $\alpha = \tau \rightarrow \alpha$), and build environment:

$$\begin{aligned} \text{build}(u, var) &= \begin{cases} \text{build}(u, x), & \text{if } u(var) = x \\ var, & \text{otherwise} \end{cases} \\ \text{build}(u, \tau \mu) &= \text{build}(u, \tau) \text{ build}(u, \mu) \\ \text{build}(u, \tau_1 \rightarrow \tau_2) &= \text{build}(u, \tau_1) \rightarrow \text{build}(u, \tau_2) \\ \text{build}(u, x^{\bar{d}}) &= \text{build}(u, x)^{\bar{d}} \\ \text{build}(u, x) &= x, \text{ if none of the above applies} \end{aligned}$$

As explained at the end of this Section, types have to be built up when generating polymorphic environments for efficiency issues. Because SML does not allow infinite types, we also use `build` to detect circularity issues. During unification, before augmenting any unification context, we check if it would allow allow generat-

ing infinite types (see rule (U1) of our unification algorithm defined in Fig. 7). For example, given the unifier $\{\alpha_1 \mapsto \alpha_2^{\bar{d}_1}, \alpha_2 \mapsto \langle \alpha_3^{\bar{d}_3} \rightarrow \alpha_4^{\bar{d}_4}, \bar{d}_2 \rangle\}$, we do not allow its augmentation with, e.g., $\{\alpha_3 \mapsto \langle \alpha_5^{\bar{d}_6} \rightarrow \alpha_1^{\bar{d}_7}, \bar{d}_5 \rangle\}$ because it would allow generating infinite types.

Environment extraction. The function `diff` is used by rules (U4), (P1) and (P2) of our constraint solver to extract environments generated during unification. It allows, when solving an environment, getting back its “solved version” once all of its constraints have been dealt with. By “solved version” of an environment e , we mean the sequence of environments that has been added to the unification context of the state in which the unification process was when it started to solve e . For example, if `solve($\langle u, e \rangle, \bar{d}, e_0$) \rightarrow^* succ($\langle u', e' \rangle$)` then $e' = e; e_1 \cdots ; e_n$ and `diff(e, e') = $\square; e_1 \cdots ; e_n$` which is the “solved version” of e_0 .

$$\begin{aligned} \text{diff}(e, e) &= \square \\ \text{diff}(e_1, e_2; e_3) &= \text{diff}(e_1, e_2); e_3, \text{ if } e_1 \neq (e_2; e_3) \end{aligned}$$

Polymorphic environments. Fig. 6 defines `toPoly` which is used by rule (P1) of our constraint solver to generate a polymorphic environment by quantifying the type variables not occurring in the types of the monomorphic bindings of the unification environment of the current state.

In Fig. 6, τ is the type from which we want to generate a type scheme. First, we build up the type, using the unifier of the unification context of the current state (u), to obtain the type τ' . The set $\bar{\tau}$ is the set of types of the monomorphic bindings for which, the binding currently being generalised, is in the scope. The set $\bar{\alpha}$ is the set of type variables that are allowed to be quantified over because they do not depend on the types of the monomorphic bindings. Finally, \bar{d} is the set of dependencies “explaining” why the type variables not in $\bar{\alpha}$ but occurring in τ' (the type variables occurring in τ' and also depending on the monomorphic bindings) are not allowed to be quantified over.

Let us illustrate how this mechanism works with the `fn-expression` `exp: fn x => let val rec f = fn z => x z in f end`. The constraint generation algorithm generates an environment of the form `poly(e_1)` for the recursive declaration `val rec f = fn z => x z`. When solving the constraints generated for `exp`, the constraint solver eventually applies `toPoly` to a unification context $\langle u, e \rangle$ and a binding of the form $\downarrow f \stackrel{\bar{d}}{=} \alpha_1$ (which is the “solved version” of e_1). Building up α_1 results in a type τ' of the form $\langle \alpha_2^{\bar{d}_2} \rightarrow \alpha_3^{\bar{d}_3}, \bar{d}_1 \rangle$. Because x 's type is monomorphic, a monomorphic binding of the form $\downarrow x \stackrel{\bar{d}_1}{=} \alpha_0$ occurs in e (the only monomorphic binding occurring in e) and so we build a $\bar{\tau}$ (see Fig. 6) of the form $\{\tau_0\}$ where τ_0 is obtain building up α_0 and is of the form $\langle \alpha_2^{\bar{d}_5} \rightarrow \alpha_3^{\bar{d}_6}, \bar{d}_4 \rangle$ (equivalent to τ' up to dependencies because $f \eta$ -reduces to x). We therefore build a $\bar{\alpha}$ (see Fig. 6) of the form \emptyset because α_2 and α_3 both occur in τ' . We also build a \bar{d}' of the form $\bar{d}_4 \cup \bar{d}_5 \cup \bar{d}_6$ which is the set of “reasons” for not allowing α_2 and α_3 to be in $\bar{\alpha}$ (set of type variables that are allowed to be generalised over when building the type scheme returned by `toPoly`). Finally, e is augmented with $\downarrow f \stackrel{\bar{d}}{=} \forall \emptyset. \langle \alpha_2^{\bar{d}_2} \rightarrow \alpha_3^{\bar{d}_3}, \bar{d}_1 \cup \bar{d}' \rangle$.

Figure 6 Monomorphic to polymorphic environment

$$\text{toPoly}(\langle u, e \rangle, \downarrow \text{vid} \stackrel{\bar{d}}{=} \langle \tau, is \rangle) = \langle u, e; (\downarrow \text{vid} \stackrel{\bar{d}}{=} \langle \forall \bar{\alpha}. \tau \bar{\alpha}', is \rangle) \rangle \quad \text{where} \quad \begin{cases} \tau' = \text{build}(u, \tau) \\ \bar{\tau} = \{\tau_0 \mid \exists \text{vid}. \tau_1 = e(\text{vid}) \wedge \tau_0 = \text{build}(u, \tau_1)\} \\ \bar{\alpha} = (\text{vars}(\tau') \cap \text{ITyVar}) \setminus (\text{vars}(\bar{\tau}) \cup \{\alpha_{\text{dum}}\}) \\ \bar{d}' = \{d \mid \tau_0 \in \bar{\tau} \wedge d \in \text{deps}(\tau_0) \wedge \neg \text{dj}(\text{vars}(\tau_0) \cap \text{ITyVar}, \text{vars}(\tau') \setminus (\bar{\alpha} \cup \{\alpha_{\text{dum}}\}))\} \end{cases}$$

When solving constraints generated by our constraint generator, `toPoly` is only applied to $\text{bind}^{\bar{d}}$'s resulting from the solving of an environment wrapped by `poly` which in turn is only used to wrap environments built from: dependencies, a unique monomorphic binding and equality constraints.

Extracting the monomorphic type variables of a binding's type is an expansive computation. We only perform it once per polymorphic binding by, provided a unification context, first building the type of a given binding and by then looking up in the environment (in the unification context) which type variables cannot be quantified over because they are monomorphic. When accessing the type of a polymorphic binding we then only have to generate an instance of its type scheme (see rule (A1) of our unification algorithm).

Algorithm. Fig. 7 defines our constraint solver.

The accessor rule (A4) can also be used to report free identifiers. If $\text{solve}(\Delta, \bar{d}, \uparrow \text{id} = x) \rightarrow \text{succ}(\Delta)$ and $\neg \text{hiding}(\Delta)$ then it means that there is no binder for id and so that it is a free identifier.

Free identifiers are in any case important to report, but it is especially vital for structure identifiers in *open* declarations. In our approach, a free opened structure is considered as potentially redefining its entire context. Hence, `val x = 1 open S val y = x 1` is typable because x 's first occurrence is hidden by the declaration `open S`. This might be confusing if S was not reported as being free. Let us explain how a free opened structure hides its context. Given a declaration `open S`, our constraint generation algorithm generates an environment of the form $(\uparrow S \stackrel{l}{=} ev); ev^l$ where l is the label labelling the declaration. Because S is free, rule (A4) applies when solving $\uparrow S = ev$. The environment variable ev is then not constrained to anything. Hence when solving ev , rule (V2) applies and $\Delta; ev$ (from the right-hand-side of rule (V2)) results in the hiding of Δ by ev : all the binders in Δ are hidden by ev .

Let the relations `isErr` and `solvable` be defined as follows:

$$\begin{aligned} e \stackrel{\text{isErr}}{\rightarrow} er &\Leftrightarrow \text{solve}(\langle \emptyset, \square \rangle, \emptyset, e) \rightarrow^* \text{err}(er) \\ \text{solvable}(e) &\Leftrightarrow \exists \Delta. \text{solve}(\langle \emptyset, \square \rangle, \emptyset, e) \rightarrow^* \text{succ}(\Delta) \\ \text{solvable}(sdec) &\Leftrightarrow \exists e. sdec \triangleright e \wedge \text{solvable}(e) \end{aligned}$$

4.5 Minimisation and enumeration. Extraction of environment labels. Given an environment e , `IBinds` extracts the labels labelling binders (*bind*) occurring in e . It is used during the first phase of our minimisation algorithm which consists in trying to remove entire sections of code (datatype declarations, functions, structures, ...) by “disconnecting” accessors from their binders:

$$\text{IBinds}(e) = \{l \mid \text{bind}^l \text{ occurs in } e\}$$

Constraint filtering. Fig. 8 defines our constraint filtering function `filt`, used to check the solvability of constraints in which some constraints are discarded. It is only applied to constraints generated by our constraint generator. This is why we only filter environment equality constraints of the form $ev = e$ and not of the general from $e_1 = e_2$. In $\text{filt}(e, \bar{l}_1, \bar{l}_2)$, \bar{l}_1 is the label set for which we want to keep the annotated environments (first case of the filtering rule for e^l), and \bar{l}_2 is the label set for which we do not want to keep the equality constraints and accessors but for which we want to turn the binders into dummy ones and keep the environment variables (second case of the filtering rule for e^l). The environments annotated by labels that are not in $\bar{l}_1 \cup \bar{l}_2$ are then discarded (third case of the filtering rule for e^l). In the context of constraint filtering, label sets are

sometimes called filters. Being able to distinguish between binders to discard (not labelled by a label in $\bar{l}_1 \cup \bar{l}_2$) and binders to turn into dummy ones (labelled by a label in \bar{l}_2) is necessary because during minimisation, throwing away any environment might result in different bindings in the filtered constraints (corresponding to a different SML code). For example, removing the binder labelled by l_2 in $(\downarrow x \stackrel{l_1}{=} \langle \tau_1, is_1 \rangle); (\downarrow x \stackrel{l_2}{=} \langle \tau_2, is_2 \rangle); (\text{tx} \stackrel{l}{=} \tau)$ would result in x 's accessor being bound to x 's first binder instead of its second. Similarly, removing the binding labelled by the label associated to f 's second occurrence in the environment generated for

```
let val rec f = fn x => x 1
in let val rec f = fn x => x + 1 in f true end
end
```

would result in f 's third occurrence to be bound to its first occurrence and so to the enumeration algorithm to find a type error that does not exist in the original piece of code. When a binding is labelled by a label from \bar{l}_2 , it becomes a dummy unlabelled one that cannot be involved in any error and it results that the same holds for its accessors.

Minimisation algorithm. Fig. 9 defines our minimisation algorithm: the relation `min` that uses the relation $\rightarrow_{\text{test}}$ to test if a label can be removed from a slice and where $\rightarrow_{\text{test}}^*$ is its reflexive and transitive closure. It consists of two main phases. The first one $(\langle e, \text{labs}(er) \setminus \bar{l}, \text{labs}(er) \cap \bar{l} \rangle \rightarrow_{\text{test}}^* \langle e, \bar{l}_1, \emptyset \rangle)$ tries to remove entire sections of code at once by turning bindings into dummy ones using `IBinds`. In a fine-grained second phase $(\langle e, \emptyset, \bar{l}_1 \rangle \rightarrow_{\text{test}}^* \langle e, \bar{l}_2, \emptyset \rangle)$ the algorithm tries to remove the remaining labels (\bar{l}_1) one at a time.

A step of our unification algorithm is as follow: $\langle e, \bar{l}_1, \{l\} \uplus \bar{l}_2 \rangle \rightarrow_{\text{test}} \langle e, \bar{l}_3, \bar{l}_4 \rangle$ where \bar{l}_3 and \bar{l}_4 depend on the solvability of $\text{filt}(e, \bar{l}_1 \cup \bar{l}_2, \{l\})$. The set $\bar{l}_1 \cup \bar{l}_2 \cup \{l\}$ is the label set of the error that the minimisation algorithm is minimising and $\{l\} \uplus \bar{l}_2$ is the label set yet to try to discard. The environment $\text{filt}(e, \bar{l}_1 \cup \bar{l}_2, \{l\})$ is obtained from e by filtering out the constraints that are not labelled by $\bar{l}_1 \cup \bar{l}_2 \cup \{l\}$ and by turning the binders labelled by l into dummy ones. If the obtained filtered environment is solvable it means that l is necessary and $\bar{l}_3 = \bar{l}_1 \cup \{l\}$ and $\bar{l}_4 = \bar{l}_2$. If it is unsolvable (solving the filtered environment failed and we obtained a new smaller error), it means that l is unnecessary for an error to occur and that any environment labelled by l can be completely filtered out in the next step. The label sets \bar{l}_3 and \bar{l}_4 are then restricted to the newly found error (see rule (M1)).

Environments (bindings, environment variables, ...) can be completely filtered out from one step to another because our constraint generator and solver, together ensure that if a binder is turned into a dummy one then none of its accessors will be part of any error. This invariant could explicitly be enforced during constraint solving by adding side conditions to rules (A1)-(A3) checking if the types of the accessed identifiers are not dummy variables (`DumVar`).

Enumeration algorithm. Fig. 9 also defines our enumeration algorithm: the relation \rightarrow_e where \rightarrow_e^* is its reflexive and transitive closure. Enumerating the minimal type errors in a piece of code consists of trying to solve diverse results of filtering the constraints generated for the piece of code. The tested filters (label sets) form the search space which is built while searching for errors. The enumeration algorithm starts with a unique filter: the empty set, to

Figure 7 Constraint solver

equality simplification		equality constraint reversing	
(S1) $\text{solve}(\Delta, \bar{d}, x=x)$	$\rightarrow \text{succ}(\Delta)$	(R) $\text{solve}(\Delta, \bar{d}, x=y)$	$\rightarrow \text{solve}(\Delta, \bar{d}, y=x)$, if $y \in s$ and $x \notin s$, where $s = \text{Var} \cup \text{Dependent}$
(S2) $\text{solve}(\Delta, \bar{d}, \tau \mu = \tau' \mu')$	$\rightarrow \text{solve}(\Delta, \bar{d}, (\mu = \mu'); (\tau = \tau'))$		
(S3) $\text{solve}(\Delta, \bar{d}, \tau_1 \rightarrow \tau_2 = \tau_3 \rightarrow \tau_4)$	$\rightarrow \text{solve}(\Delta, \bar{d}, (\tau_1 = \tau_3); (\tau_2 = \tau_4))$		
(S4) $\text{solve}(\Delta, \bar{d}, \tau_1 = \tau_2)$	$\rightarrow \text{solve}(\Delta, \bar{d}, \mu = \text{ar})$		if $\{\tau_1, \tau_2\} = \{\tau \mu, \tau_3 \rightarrow \tau_4\}$ for some μ, τ, τ_3 and τ_4
(S5) $\text{solve}(\Delta, \bar{d}, \mu_1 = \mu_2)$	$\rightarrow \text{err}(\langle \text{tyConsClash}(\mu_1, \mu_2), \bar{d} \rangle)$,		if $\mu_1 \neq \mu_2$ and $\{\mu_1, \mu_2\} \in \{\{\gamma, \gamma'\}, \{\gamma, \text{ar}\}\}$ for some γ and γ'
(S6) $\text{solve}(\Delta, \bar{d}, is_1 = is_2)$	$\rightarrow \text{err}(\langle \text{statusClash}(is_1, is_2), \bar{d} \rangle)$,		if $\neg \text{compatible}(is_1, is_2)$
(S7) $\text{solve}(\Delta, \bar{d}, is_1 = is_2)$	$\rightarrow \text{succ}(\Delta)$,		if $\text{compatible}(is_1, is_2)$ and $is_1, is_2 \in \text{RawldStatus}$
(S8) $\text{solve}(\Delta, \bar{d}, x^{\bar{d}'} = y)$	$\rightarrow \text{solve}(\Delta, \bar{d} \cup \bar{d}', x=y)$		
unifier access/updating Rules (U1) through (U6) have also these common side conditions: $var \neq x$ and $y = \text{build}(u, x)$			
(U1) $\text{solve}(\langle u, e \rangle, \bar{d}, var=x)$	$\rightarrow \text{err}(\langle \text{circularity}, \bar{d} \cup \text{deps}(y) \rangle)$,		if $var \in \text{vars}(y) \setminus (\text{dom}(u) \cup \text{Env})$ and $var \neq \text{strip}(y)$
(U2) $\text{solve}(\langle u, e \rangle, \bar{d}, var=x)$	$\rightarrow \text{succ}(\langle u, e \rangle)$,		if $var \in \text{vars}(y) \setminus (\text{dom}(u) \cup \text{Env})$ and $var = \text{strip}(y)$
(U3) $\text{solve}(\langle u, e \rangle, \bar{d}, var=x)$	$\rightarrow \text{succ}(\langle u \boxplus \{var \mapsto x^{\bar{d}'}\}, e \rangle)$,		if $var \notin \text{vars}(y) \cup \text{dom}(u) \cup \text{Env}$
(U4) $\text{solve}(\langle u, e \rangle, \bar{d}, var=x)$	$\rightarrow \text{succ}(\langle u' \boxplus \{var \mapsto \text{diff}(e, e')^{\bar{d}'}\}, e \rangle)$,		if $var \in \text{Env} \setminus \text{dom}(u)$ and $\text{solve}(\langle u, e \rangle, \bar{d}, x) \rightarrow^* \text{succ}(\langle u', e' \rangle)$
(U5) $\text{solve}(\langle u, e \rangle, \bar{d}, var=x)$	$\rightarrow \text{err}(er)$,		if $var \in \text{Env} \setminus \text{dom}(u)$ and $\text{solve}(\langle u, e \rangle, \bar{d}, x) \rightarrow^* \text{err}(er)$
(U6) $\text{solve}(\langle u, e \rangle, \bar{d}, var=x)$	$\rightarrow \text{solve}(\langle u, e \rangle, \bar{d}, z=x)$,		if $u(var) = z$
constrained environments		dependent/empty/variables	
(C1) $\text{solve}(\Delta, \bar{d}, e_1; e_2)$	$\rightarrow \text{solve}(\Delta', \bar{d}, e_2)$, if $\text{solve}(\Delta, \bar{d}, e_1) \rightarrow^* \text{succ}(\Delta')$	(D) $\text{solve}(\Delta, \bar{d}, e^{\bar{d}'})$	$\rightarrow \text{solve}(\Delta, \bar{d} \cup \bar{d}', e)$
(C2) $\text{solve}(\Delta, \bar{d}, e_1; e_2)$	$\rightarrow \text{err}(er)$, if $\text{solve}(\Delta, \bar{d}, e_1) \rightarrow^* \text{err}(er)$	(N) $\text{solve}(\Delta, \bar{d}, \square)$	$\rightarrow \text{succ}(\Delta)$
		(V) $\text{solve}(\langle u, e \rangle, \bar{d}, ev)$	$\rightarrow \text{succ}(\langle u, e; \text{build}(u, ev) \rangle)$
binders			
(B1) $\text{solve}(\Delta, \bar{d}, \downarrow id=x)$	$\rightarrow \text{succ}(\Delta; (\downarrow id \stackrel{\bar{d}}{=} x))$		
(B2) $\text{solve}(\Delta, \bar{d}, \uparrow vid=\alpha)$	$\rightarrow \text{solve}(\Delta, \bar{d}, \uparrow vid=\langle \alpha, \text{ifNotDum}(\alpha, u) \rangle)$,		if $\text{strip}(\Delta[\![vid]\!]) \in \{c, d\}$
(B3) $\text{solve}(\Delta, \bar{d}, \uparrow vid=\alpha)$	$\rightarrow \text{succ}(\Delta; (\downarrow id \stackrel{\bar{d} \cup \bar{d}'}{=} \alpha))$,		if $\Delta[\![vid]\!] = is$ and $\text{strip}(is) = v$ and $\text{deps}(is) = \bar{d}'$
(B4) $\text{solve}(\Delta, \bar{d}, \uparrow vid=\alpha)$	$\rightarrow \text{succ}(\Delta; (\downarrow id \stackrel{\bar{d} \cup \{vid\}}{=} \langle \alpha, \text{ifNotDum}(\alpha, u) \rangle))$,		if $\text{strip}(\Delta[\![vid]\!]) = u$ or $(\neg \text{hiding}(\Delta))$ and $\Delta[\![vid]\!]$ undefined
(B5) $\text{solve}(\Delta, \bar{d}, \uparrow vid=\alpha)$	$\rightarrow \text{succ}(\Delta; (\downarrow id \stackrel{\bar{d}}{=} \langle \alpha_{\text{dum}}, \text{ifNotDum}(\alpha, p) \rangle))$,		if $\text{strip}(\Delta[\![vid]\!]) \in \{a, p\}$ or $(\text{hiding}(\Delta))$ and $\Delta[\![vid]\!]$ undefined
accessors			
(A1) $\text{solve}(\Delta, \bar{d}, \uparrow vid=\alpha)$	$\rightarrow \text{solve}(\Delta, \bar{d}, \tau[\text{ren}]=\alpha)$, if $\Delta[\![vid]\!] = \forall \bar{\alpha}. \tau$ and $\text{dom}(\text{ren}) = \bar{\alpha}$ and $\text{dja}(\text{vars}(\langle \Delta, \alpha \rangle) \setminus \{\alpha_{\text{dum}}\}, \text{ran}(\text{ren}))$		
(A2) $\text{solve}(\Delta, \bar{d}, \uparrow vid=ris)$	$\rightarrow \text{solve}(\Delta, \bar{d}, is=ris)$,		if $\Delta[\![vid]\!] = is$
(A3) $\text{solve}(\Delta, \bar{d}, \uparrow id=var)$	$\rightarrow \text{solve}(\Delta, \bar{d}, x=var)$,		if $\Delta[\![id]\!] = x$ and x is not of the form $\forall \bar{\alpha}. \tau$
(A4) $\text{solve}(\Delta, \bar{d}, \uparrow id=x)$	$\rightarrow \text{succ}(\Delta)$,		if $(x \in \text{IdStatus}$ and $\Delta[\![id]\!]$ undefined) or $(x \notin \text{IdStatus}$ and $\Delta[\![id]\!]$ undefined)
polymorphic environments			
(P1) $\text{solve}(\langle u_1, e_1 \rangle, \bar{d}, \text{poly}(e))$	$\rightarrow \text{succ}(\text{toPoly}(\langle u_2, e_1 \rangle, e'))$, if $\text{solve}(\langle u_1, e_1 \rangle, \bar{d}, e) \rightarrow^* \text{succ}(\langle u_2, e_2 \rangle)$ and $\text{diff}(e_1, e_2) = \square; e'$		
(P2) $\text{solve}(\langle u_1, e_1 \rangle, \bar{d}, \text{poly}(e))$	$\rightarrow \text{succ}(\langle u_2, e_2 \rangle)$,		if $\text{solve}(\langle u_1, e_1 \rangle, \bar{d}, e) \rightarrow^* \text{succ}(\langle u_2, e_2 \rangle)$ and $\text{diff}(e_1, e_2) = \square$
(P3) $\text{solve}(\langle u_1, e_1 \rangle, \bar{d}, \text{poly}(e))$	$\rightarrow \text{err}(er)$,		if $\text{solve}(\langle u_1, e_1 \rangle, \bar{d}, e) \rightarrow^* \text{err}(er)$

Figure 8 Constraint filtering

$\text{filt}(e^l, \bar{l}_1, \bar{l}_2)$	$= \begin{cases} e^l, & \text{if } l \in \bar{l}_1 \setminus \bar{l}_2 \\ \text{dum}(e), & \text{if } l \in \bar{l}_2 \\ \square, & \text{otherwise} \end{cases}$	$\text{dum}(\downarrow id=x) = (\downarrow id = \text{toDumVar}(x))$	$\text{toDumVar}(\sigma) = \alpha_{\text{dum}}$
$\text{filt}(ev=e, \bar{l}_1, \bar{l}_2)$	$= (ev = \text{filt}(e, \bar{l}_1, \bar{l}_2))$	$\text{dum}(\uparrow id=x) = (\uparrow id = \text{toDumVar}(x))$	$\text{toDumVar}(\mu) = \delta_{\text{dum}}$
$\text{filt}(e_1; e_2, \bar{l}_1, \bar{l}_2)$	$= \text{filt}(e_1, \bar{l}_1, \bar{l}_2); \text{filt}(e_2, \bar{l}_1, \bar{l}_2)$	$\text{dum}(ev) = ev_{\text{dum}}$	$\text{toDumVar}(e) = ev_{\text{dum}}$
$\text{filt}(\text{poly}(e), \bar{l}_1, \bar{l}_2)$	$= \text{poly}(\text{filt}(e, \bar{l}_1, \bar{l}_2))$	$\text{dum}(c) = \square$	$\text{toDumVar}(\alpha) = \alpha_{\text{dum}}$
$\text{filt}(\square, \bar{l}_1, \bar{l}_2)$	$= \square$	$\text{dum}(acc) = \square$	$\text{toDumVar}(is) = a$

Figure 9 Minimisation and enumeration algorithms

minimisation			
(M1) $\langle e, \bar{l}_1, \{l\} \boxplus \bar{l}_2 \rangle$	$\rightarrow_{\text{test}} \langle e, \bar{l}_1 \cap \bar{d}, \bar{l}_2 \cap \bar{d} \rangle$, if $\text{filt}(e, \bar{l}_1 \cup \bar{l}_2, \{l\}) \stackrel{\text{isErr}}{\rightarrow} \langle ek, \bar{d} \rangle$		
(M2) $\langle e, \bar{l}_1, \{l\} \boxplus \bar{l}_2 \rangle$	$\rightarrow_{\text{test}} \langle e, \bar{l}_1 \cup \{l\}, \bar{l}_2 \rangle$, if $\text{solvable}(\text{filt}(e, \bar{l}_1 \cup \bar{l}_2, \{l\}))$		
(M3) $\langle e, er \rangle$	$\xrightarrow{\text{min}} er'$, if $\text{IBinds}(e) = \bar{l}$ and $\langle e, \text{labs}(er) \setminus \bar{l}, \text{labs}(er) \cap \bar{l} \rangle \rightarrow_{\text{test}}^* \langle e, \bar{l}_1, \emptyset \rangle$ and $\langle e, \emptyset, \bar{l}_1 \rangle \rightarrow_{\text{test}}^* \langle e, \bar{l}_2, \emptyset \rangle$ and $\text{filt}(e, \bar{l}_2, \emptyset) \stackrel{\text{isErr}}{\rightarrow} er'$		
enumeration EnumState ::= enum(e) enum(e, \overline{er} , \bar{l}) errors(\overline{er})			
(E1) enum(e)	$\rightarrow_e \text{enum}(e, \emptyset, \{\emptyset\})$	(E4) enum(e, \overline{er} , $\bar{l} \boxplus \{\bar{l}\}$)	$\rightarrow_e \text{enum}(e, \overline{er} \cup \{\langle ek, \bar{d} \rangle\}, \bar{l} \cup \bar{l})$,
(E2) enum(e, \overline{er} , \emptyset)	$\rightarrow_e \text{errors}(\overline{er})$		if $\text{filt}(e, \text{labs}(e), \bar{l}) \stackrel{\text{isErr}}{\rightarrow} er$ and $\langle e, er \rangle \xrightarrow{\text{min}} \langle ek, \bar{d} \rangle$
(E3) enum(e, \overline{er} , $\bar{l} \boxplus \{\bar{l}\}$)	$\rightarrow_e \text{enum}(e, \overline{er}, \bar{l})$, if $\text{solvable}(\text{filt}(e, \text{labs}(e), \bar{l}))$		and $\bar{l}' = \{\bar{l} \cup \{l\} \mid l \in \bar{d} \wedge \forall \bar{l}_0 \in \bar{l}. \bar{l}_0 \not\subseteq \bar{l} \cup \{l\}\}$

solve all the generated constraints. Then, when an error is found and minimised, the labels of the error are used to build new filters (see \bar{l}' in rule (E4)). Once all the filters are exhausted the enu-

meration algorithm stops. The found errors are then all the minimal type errors of the analysed piece of code (see rule (E2)). For example, assume that $\text{sdec} \triangleright e$ for an untypable given piece of

code $sdec$. Then, the first enumeration state is (see rule (E1)): $\text{enum}(e, \emptyset, \{\emptyset\})$ where the first empty set is the set of found errors (empty at the beginning) and where the second empty set is the first filter. Because $sdec$ is untypable, the constraint solver fails and returns a type error er . The minimisation algorithm minimises er and returns a minimal error er' such that $er'(1) \subseteq er(1)$. The error er' can be er if it was already in a minimal form when found by the enumerator. New filters are then computed based on the filter used to find this new error (\emptyset in our example) and the new error itself (er'): $\{\{l\} \mid l \in er'(1)\}$. The enumeration keeps searching for errors using this updated search space: the new state is $\text{enum}(e, \{er'\}, \{\{l\} \mid l \in er'(1)\})$. For the next step, one of the $\{l\}$ where $l \in er'(1)$ will be picked as the filter to try to find another error. When a filter leads to a solvable filtered environment, the filter is discarded (rule (E3)) otherwise the filter is used to update the search space as explained above (rule (E4))

4.6 Slicing. The last phase of our TES consists in the computation of a minimal type error slice from an untypable piece of code and a minimal error found by the enumeration algorithm. The nodes labelled by the labels not involved in the error are discarded and replaced by “dot” terms. For example, if we remove the node associated to the label l_2 (the unit expression) in $[1^{l_1} ()^{l_2}]^{l_3}$ then we obtain $[1^{l_1} \text{dot-}e(\emptyset)]^{l_3}$, displayed as $1 \langle . \rangle$ in our implementation. Dots are used as a visually convenient way to show that information has been discarded. Fig. 10 extends our syntax and constraint generator to “dot” terms. Our constraint generator is extended to dot terms so that every piece of (our extended) syntax can be type checked (by generating constraints and by then solving the constraints), which is needed to state our minimality criteria in Sec. 4.8. We call *slice*, any syntactic form that can be produced using the grammar rules defined in Fig. 3 and Fig 10 combined. We call *type error slice*, any slice for which our constraint generation algorithm (defined in Fig 4 and Fig. 10 combined) only generates unsolvable constraints. Let us restrict our slice definitions to structure declarations. Formally, a slice is a $sdec$ and a type error slice is a $sdec$ such that $\neg \text{solvable}(sdec)$.

Flattening. Turning nodes not participating in errors into dot nodes is not enough. Our slicing algorithm uses two tidying functions *flat* and *tidy*. The flattening function *flat* flattens sequences of parts (pt). For example, flattening $\langle . . 1 . . \langle . \rangle . . \rangle$ results in $\langle . . 1 . . \rangle$. Not all nested dot terms are flattened: in order not to mix up bindings in a slice, we do not allow a declaration to be an ept (expression term as opposed to a declaration) and only allow ept 's to be flattened, so that declarations cannot escape the scope defined by a dot term. For example, we do not flatten $\langle . . \text{val } x = () . . \langle . . \text{val } x = 1 . . \rangle . . x + 1 . . \rangle$ to $\langle . . \text{val } x = () . . \text{val } x = 1 . . x + 1 . . \rangle$ because they have different semantics: the first slice is not typable but the second is. Let *flat* be defined as follows (where x can be any of e, p, s, d):

$$\begin{aligned} \text{flat}(\langle \rangle) &= \langle \rangle \\ \text{flat}(\langle pt \rangle @ \overrightarrow{pt}) &= \begin{cases} \overrightarrow{ept} @ \text{flat}(\overrightarrow{pt}), & \text{if } pt = \text{dot-}x(\overrightarrow{ept}) \\ \langle pt \rangle @ \text{flat}(\overrightarrow{pt}), & \text{otherwise} \end{cases} \end{aligned}$$

The function *tidy* tidies sequences of structure declarations ($sdec$) when slicing structure expressions:

$$\begin{aligned} \text{tidy}(\langle \rangle) &= \langle \rangle \\ \text{tidy}(\langle \text{dot-d}(\overrightarrow{ept}), \text{dot-d}(\overrightarrow{pt}) \rangle @ \overrightarrow{sdec}) &= \text{tidy}(\langle \text{dot-d}(\overrightarrow{ept} @ \overrightarrow{pt}) \rangle @ \overrightarrow{sdec}) \\ \text{tidy}(\langle \text{dot-d}(\emptyset) \rangle @ \overrightarrow{sdec}) &= \text{tidy}(\overrightarrow{sdec}), \text{ if } \overrightarrow{sdec} \text{ not of the form } \text{dot-d}(\overrightarrow{ept}) @ \overrightarrow{sdec}' \\ \text{tidy}(\langle \overrightarrow{sdec} \rangle @ \overrightarrow{sdec}) &= \langle \overrightarrow{sdec} \rangle @ \text{tidy}(\overrightarrow{sdec}), \text{ if none of the above applies} \end{aligned}$$

Algorithm. Our slicing algorithm can be presented in a simple fashion if our syntax forms defined in Fig. 3 and Fig. 10 are regarded as abstract syntax trees. In such a tree $tree$, leaves are identifiers id and otherwise a node is labelled by a node kind $node$ and a label l (denoted $node^l \langle tree_1, \dots, tree_n \rangle$). Using this notation, we define our slicing function *sl* in Fig. 11.

4.7 Overall algorithm for Type Error Slicing. First, given a SML structure declaration $sdec$, our constraint generation algorithm defined in Fig. 4 generates constraints structured in an environment e . Then, type errors of e are enumerated using the enumeration algorithm defined in Fig. 9. Once an error is found by the enumeration algorithm, it is minimised using the minimisation algorithm also defined in Fig. 9. Then a slice is computed from the minimised error and the original piece of code using the slicing algorithm defined in Fig. 11. Both enumeration and minimisation rely on the constraint solver defined in Fig. 7. The computed type error slices are finally reported to the user. In addition to a type error slice, a type error report also includes a highlighting of the slice in the SML user code, a message explaining the kind of the error (see Fig. 5), and a set of identifier status context dependencies. Formally, our overall algorithm *tes* is defined as follows:

$$\begin{aligned} \text{tes}(sdec) &= \{ \langle sdec', ek, \overrightarrow{vid} \rangle \mid sdec \rightarrow e \\ &\quad \wedge \text{enum}(e) \rightarrow^* \text{errors}(\overrightarrow{er}) \\ &\quad \wedge \langle ek, \overrightarrow{l} \cup \overrightarrow{vid} \rangle \in \overrightarrow{er} \\ &\quad \wedge \text{sl}(sdec, \overrightarrow{l}) = sdec' \} \end{aligned}$$

4.8 Minimality. Let us informally define the function bindings on environments. This function uses a modified version of our constraint solver that keeps track of the bindings generated by the accessor rules ((A1)-(A3)). Given a piece of code, bindings, using the constraint generation algorithm, generates an environment, filters out all the labelled equality constraint in the generated environment, runs the modified constraint solver on the filtered environment, and finally returns the recorded bindings. For example, if exp is `let val x = true in let val x = 1 in x end end`, and the label l_i is associated to the i th occurrence of x then $\text{bindings}(exp) = \{\langle l_2, l_3 \rangle\}$.

We define the sub-slice relation as follows: $sdec_1 \sqsubseteq_{\overrightarrow{l}} sdec_2$ iff $\text{sl}(sdec_2, \overrightarrow{l}) = sdec_1$ and $\text{bindings}(sdec_1) \subseteq \text{bindings}(sdec_2)$.

We say $sdec_2$ is a minimal type error slice of $sdec_1$ iff $sdec_2 \sqsubseteq_{\overrightarrow{l}} sdec_1$, $\neg \text{solvable}(sdec_2)$ and for all $sdec'$ if $sdec' \sqsubseteq_{\overrightarrow{l}} sdec_2$ and $sdec' \neq sdec_2$ for some $\overrightarrow{l'}$ then $\text{solvable}(sdec')$.

We consider minimality as a design principle for our TES even though minimal slices do not always seem to be the correct answer to type error reporting. For example for record field name clashes, we do not want to provide minimal slices, as presented in Sec. 2.3.

For the subset of our TES presented in the present paper, we believe the following holds: a slice $sdec'$ is a minimal slice of $sdec$ iff $\langle sdec', ek, \overrightarrow{vid} \rangle \in \text{tes}(sdec)$. We do not formally prove this statement for diverse reasons. First, our TES is constantly being updated and proving the minimality of one of its versions would not guaranty the minimality of the others. Then, as mentioned above, minimality is only a design principle. Let us finally stress that we feel improving the range and quality of our slices is more important than ensuring their minimality in particular.

4.9 Design principles. While developing our type error slicer we discovered, developed, and followed the following principles. *Constraint terms* are those pieces of syntax that can occur anywhere inside a constraint. In our system, this is any is, μ, τ, σ , or e .

1. Each syntactic sort of constraint terms should have a case ranging over an infinite variable set. This allows incomplete information in every possible place in constraints, which allows considering every possible way of slicing away parts of the program

Figure 10 Extension of our syntax and constraint generator to “dot” terms**extension of the syntax**

$pt \in \text{Part} ::= ept \mid sdec$	$\text{ConBind} ::= \dots \mid \text{dot-e}(\vec{pt})$	$\text{AtExp} ::= \dots \mid \text{dot-e}(\vec{pt})$	$\text{Pat} ::= \dots \mid \text{dot-p}(\vec{pat})$
$ept \in \text{ExpPart} ::= exp \mid ty \mid sexp \mid pat$	$\text{DatName} ::= \dots \mid \text{dot-e}(\vec{pt})$	$\text{Exp} ::= \dots \mid \text{dot-e}(\vec{pt})$	$\text{StrDec} ::= \dots \mid \text{dot-d}(\vec{pt})$
$\text{Ty} ::= \dots \mid \text{dot-e}(\vec{pt})$	$\text{Dec} ::= \dots \mid \text{dot-d}(\vec{pt})$	$\text{AtPat} ::= \dots \mid \text{dot-p}(\vec{pat})$	$\text{StrExp} ::= \dots \mid \text{dot-s}(\vec{pt})$

extension of the constraint generator

Parts	$ept \triangleright e \Leftarrow ept \triangleright \langle var, e \rangle$
Declarations	$\text{dot-d}(\langle pt_1, \dots, pt_n \rangle) \triangleright \langle e_1; \dots; e_n \rangle \Leftarrow pt_1 \triangleright e_1 \wedge \dots \wedge pt_n \triangleright e_n \wedge \text{dja}(e_1, \dots, e_n)$
Patterns	$\text{dot-p}(\langle pat_1, \dots, pat_n \rangle) \triangleright \langle \alpha, e_1; \dots; e_n \rangle \Leftarrow pat_1 \triangleright e_1 \wedge \dots \wedge pat_n \triangleright e_n \wedge \text{dja}(e_1, \dots, e_n, \alpha)$
Structure expressions	$\text{dot-s}(\langle pt_1, \dots, pt_n \rangle) \triangleright \langle ev, [e_1; \dots; e_n] \rangle \Leftarrow pt_1 \triangleright e_1 \wedge \dots \wedge pt_n \triangleright e_n \wedge \text{dja}(e_1, \dots, e_n, ev)$
Expressions/Types/Constructor bindings/Datatype names	$\text{dot-e}(\langle pt_1, \dots, pt_n \rangle) \triangleright \langle \alpha, [e_1; \dots; e_n] \rangle \Leftarrow pt_1 \triangleright e_1 \wedge \dots \wedge pt_n \triangleright e_n \wedge \text{dja}(e_1, \dots, e_n, \alpha)$

Figure 11 Slicing algorithm

$node \in \text{Node} ::= \text{tyV} \mid \text{tyC} \mid \text{tyA} \mid \text{cbN} \mid \text{cbU} \mid \text{dn} \mid \text{decV} \mid \text{decD} \mid \text{decO} \mid \text{expI} \mid \text{expL} \mid \text{expF} \mid \text{expA} \mid \text{patI} \mid \text{patA} \mid \text{sdec} \mid \text{sexpI} \mid \text{sexpS}$

The node kinds correspond to the cases of the external syntax grammar. In particular, sexpS is the kind for a structure expression of the form $\text{struct}^l \text{sdec}_1 \dots \text{sdec}_n \text{end}$. In (SL3) and (SL7), x is chosen depending on the argument provided to the slicing function. For example, in (SL3), if $node = \text{sexpS}$ then x is s .

(SL1)	$\text{sl}(node^l \langle tree_1, \dots, tree_n \rangle, \vec{l}) = node^l \langle \text{sl}_1(tree_1, \vec{l}), \dots, \text{sl}_1(tree_n, \vec{l}) \rangle,$	if $l \in \vec{l}$ and $node \neq \text{sexpS}$ or $\text{sl}_1(tree_1, \vec{l}) = node'^l \vec{tree}$ where $node' \in \{\text{patI}, \text{patA}\}$
(SL2)	$\text{sl}(node^l \langle tree_1, \dots, tree_n \rangle, \vec{l}) = node^l \text{tidy}(\langle \text{sl}_1(tree_1, \vec{l}), \dots, \text{sl}_1(tree_n, \vec{l}) \rangle),$	if $l \in \vec{l}$ and $node = \text{sexpS}$
(SL3)	$\text{sl}(node^l \langle tree_1, \dots, tree_n \rangle, \vec{l}) = \text{dot-x}(\text{flat}(\langle \text{sl}_2(tree_1, \vec{l}), \dots, \text{sl}_2(tree_n, \vec{l}) \rangle)),$	if none of the above applies
(SL4)	$\text{sl}_1(node^l \vec{tree}, \vec{l}) = \text{sl}(node^l \vec{tree}, \vec{l})$	(SL6) $\text{sl}_1(id, \vec{l}) = id$
(SL5)	$\text{sl}_2(node^l \vec{tree}, \vec{l}) = \text{sl}(node^l \vec{tree}, \vec{l})$	(SL7) $\text{sl}_2(id, \vec{l}) = \text{dot-x}(\langle \rangle).$

syntax tree. This is essential to get precise type error slices that include all relevant details and exclude the irrelevant.

For us, this means the sorts μ , τ , and e have the variable cases δ , α , and ev . Our implementation has a variable case of raw identifier statuses (ris) which is omitted from this paper to save space.

2. Each syntactic sort of constraint terms should have a dependency annotation case. This allows precise blame tracking, which in turn enables precise slicing, which we already motivate above.

For us, this means the sorts is , μ , τ , and e have the dependency cases $\langle is, \vec{d} \rangle$, $\langle \mu, \vec{d} \rangle$, $\langle \tau, \vec{d} \rangle$, and $\langle e, \vec{d} \rangle$. We omit type scheme dependencies (σ) because handling schemes is already complex and only dependencies on plain types (τ) are needed in this paper.

3. In our system, when processing a program syntax tree node, a constraint generation rule will return a main result (either a type or an environment) and in some cases also an environment result (used for constraints and bindings when the main result is not an environment). The rule may connect information from the results for the node’s subtrees to the other subtrees or to the node’s results.

The principle is that these connections should generally be via constraints that carry the syntax tree node’s label and that are “shallow”, i.e., contain only connection details and not constraints from program subtrees. Fresh variables should be used as needed. This allows a program syntax node to be “disconnected” for type errors that depend on the node’s details, while still keeping type errors that arise solely due to connections between environment accessors and bindings that pass through the node.

A good example in our system is rule (G18) that handles structures (SML’s modules). The environment for the structure is built by the unlabelled constraint $ev' = (e_1; \dots; e_n)$. This “deep” constraint holds a complex structure in order to pack together a sequence of environments from the declarations making up a structure body. The structure environment is connected to the main result by the labelled shallow constraint $ev \stackrel{l}{=} ev'$.

4. Duplicating constraints should be unnecessary.

This seems obvious, but some previous constraint formalisms seem to be too weak to allow the needed sharing. Again, we use the example of rule (G18), where our system allows the environment

for the structure can be written as the sequential composition of the environments for the component declarations: $e_1; \dots; e_n$. Here the environment e_1 from the first declaration is available both in the subsequent declarations and also in the result (provided its bindings are not shadowed). A previous version of our system had a weaker constraint system with let-constraints similar to those of Pottier and Rémy [24], and the best solution we could devise required duplicating the environments for the pattern portion of each declaration, which resulted in an exponential slowdown.

5. Dependencies must be propagated during constraint solving exactly when needed. If dependencies are not propagated to places they should be, minimization will over-minimize producing non-errors. This can be detected. More insidiously, propagating dependencies where they are unneeded will keep alive unneeded parts of error slices much longer during minimization, resulting in serious slowdowns. Because correct results are eventually produced, detecting such bugs is harder so this issue requires great care.

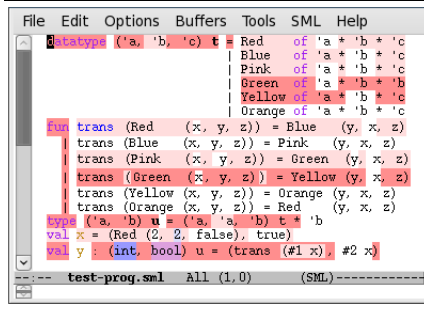
For example, an earlier version of our constraint solver copied dependencies from declarations in a structure to the structure’s result, forcing the minimizer to remove each declaration one at a time. Debugging was hard because only speed suffered.

More generally, the constraint solver should be designed to produce error slices (before minimization) that are as close to minimal as can be reasonably achieved. If constraint solving yields a non-minimal error slice, then solving steps must have annotated a constraint with a location on which it does not uniquely depend.

5. Implementation details

5.1 Supported features. Our TES handles most SML features, e.g., structures, signatures, datatype replications (handled like *open*), operator and constant overloading, many uses of functors, etc. We also handle imperative features such as exceptions and the value polymorphism restriction. We also slice context-sensitive syntax errors, which comes naturally from handling identifier statuses and doing context-independent type checking, e.g., x occurring twice in the pattern in $\text{fn } (x, x) \Rightarrow x$ is an error only if x has value variable status. We do not yet handle fixity changes. Type

Figure 12 Highlighting of an SML type error in Emacs



and structure sharing is incomplete. Some errors involving equality types and flexible record patterns are not reported.

5.2 Performance. Our implementation is currently usable for small projects (a few thousand lines) and is steadily improving. Our latest constraint system and solver is 10 to 100 times faster in many cases than before we switched to using constrained environments ($e_2; e_1$). Our previous TES version was already enormously faster than the original by Haack and Wells due to avoiding duplication of polymorphic types.

5.3 User interface. An Emacs interface (and a preliminary one for Vim) highlights slices in the edited source code. There is also a terminal command-line interface. Figure 12 presents a screenshot of the type error presented in Sec. 2.1 highlighted in Emacs. The light pink corresponds to slices other than the focused one.

5.4 The Standard ML basis library. Our examples have used operators like `::` and `+`. For now, we allow defining the Standard ML basis in a file, and we provide a file declaring a portion of the basis. For the future, we have begun implementing a way to use library types extracted from a running instance of SML/NJ, but there are still technical challenges to overcome.

6. Related work

6.1 Methods making use of slices. After the first version of TES presented by Haack and Wells, many researchers began to present type errors as program slices obtained from unsolvable sets of constraints.

Neubauer and Thiemann [21] use flow analysis to compute type dependencies for a small ML-like language to report type errors. Their system uses discriminative sum types and can analyze any term. Their first step (“collecting phase”) labels the studied term and infers type information. This analysis generates a set of program point sets. These program points are directly stored in the discriminative sum types. A conflicting type (“multivocal”) is then paired with the locations responsible for its generation. Their second step (“reporting phase”) consists in generating error reports from the conflicts generated during the first phase. Slices are built from which highlighting are produced. An interesting detail is that a type derivation can be viewed as the description of all type errors in an untypable piece of code, from which another step then computes error reports.

Similar to ours is work by Stuckey, Sulzmann and Wazny [27, 29] (based on earlier work without slices [25, 26]). They do type inference, type checking and report type errors for the Chameleon language (a modified Haskell subset). Chameleon includes algebraic data types, type-class overloading, and functional dependencies. They code the typing problem into a constraint problem and attach labels to constraints to track program locations and highlight parts of untypable pieces of code. First they compute a minimal unsatisfiable set of generated constraints from which they select one

of the type error locations to provide their type explanation. They finally provide a highlighting and an error message depending on the selected location. They provide slice highlighting but using a different strategy from ours. They focus on explaining conflicts in the inferred types at one program point inside the error location set. It is not completely clear, but they do not seem to worry much about whether the program text they are highlighting is exactly (no more and no less) a complete explanation of the type error. For example, they do not highlight applications because “they have no explicit tokens in the source code”. It is then stated: “We leave it to the user to understand when we highlight a function position we may also refer to its application”. This differs from our strategy because we think it is preferable to highlight all the program locations responsible for an error even if these are white spaces. Moreover, they do not appear to highlight the parts of datatype declarations relevant to type errors.

When running on a translation of the code presented in Sec. 2.1 into Haskell, ChameleonGecko outputs the error report partially displayed below (the rest of the output seems to be internal information computed during unification). This highlighting identifies the same error location as SML/NJ and would not help solve the error.

```
ERROR: Type error; conflicting sites:
y = (trans x1, x2)
```

Significantly, because they handle a Haskell-like language, they face challenges for accurate type error location that are different from the ones for SML.

Gast [9] generates “detailed explanations of ML type errors in terms of data flows”. His method is in three steps: generation of subtyping constraints annotated by reasons for their generation; gathering of reasons during constraint unification; transformation of the gathered reasons into explanations by data flows. He provides a visually convenient display of the data flows with arrows in XEmacs. Gast’s method (which seems to be designed only for a small portion of OCaml) can be considered as a slicing method with data flow explanations.

Braßel [7] presents a generic approach (implemented for the language Curry) for type error reporting that consists in two different procedures. The first one tries to replace portions of code by dummy terms that can be assigned any type. If an untypable piece of code becomes typable when one of its subtrees has been replaced by a dummy term then the process goes on to apply the same strategy inside the subtree. The second procedure consists in the use of a heuristic to guide the search of type errors. The heuristic is based on two principles: it will always “prefer an inner correction point to an outer one” and will always “prefer the point which is located in a function farther away in the call graph from the function which was reported by the type checker as the error location”. Braßel’s method does not seem to compute proper slices but instead singles out different locations that might be the cause of a type error inside a piece of code.

6.2 Significant non-slicing type explanation methods. Heeren et al. designed a method used in the Helium project [15, 14, 16, 12] to provide error messages for the Haskell language relying on a constraint-based type inference. First, a constraint graph is generated from a piece of code. For an ill-typed piece of code, a conflicting path called an inconsistency is extracted from the constraint graph. Such a conflicting path is a structured unsolvable set of type constraints. Heuristics are used to remove inconsistencies. To each type constraint is associated a trust value and depending on these values and the defined heuristics, some constraints are discarded until the inconsistency is removed. They also propose some “program correcting heuristics” used to search for a typable piece of code from an untypable one. Such a heuristic is for example the permutation of parameters which is a common mistake in program-

ming. Their approach has been used with students learning functional programming. Using pieces of code written by students and their expertise of the language they refined their heuristics. This approach differs from ours by privileging locations over others by the use of some heuristics. They do not compute minimal slices and highlightings.

We present below the most interesting part of the error report obtained using Helium on a translation of the code presented in Sec. 2.1 into Haskell. It is reported that `x1` and `trans` don't have the expected types. The application, which is at the end of the code, is then blamed when our programming error is at the very beginning of the code.

```
(16,6): Type error in application
expression      : trans x1
term            : trans
type            : T a a a -> T a a a
does not match : T Int Int Bool -> T Int Int Bool
```

Compilation failed with 1 error

They have also recently tackled the task to report type errors for Java [5, 6]. Error reports provided by usual compilers can be of little help, especially in the presence of generics. El Boustani and Haage try to do a better job by keeping track of more information during type checking. Having more information at hands when analysing an untypable piece of code allows a more global view of its type errors and leads to more informative error reports. The main difference between type error reporting for SML and for Java is that in Java “types are instantiated based on local information only and not through a long and complicated sequence of unifications” [5].

Lerner, Flower, Grossman and Chambers [19] present type error messages by constructing well-typed programs from ill-typed ones using different techniques (like Heeren et al. [12]), e.g., switching two parameters. Automatically conceived modifications to the ill-typed piece of code are checked for typability. They target Caml, and also developed a prototype for C++. The new typable generated code is presented as possible code that the programmer might have intended. It could be interesting to study the combination of this with TES.

7. Conclusion

7.1 Summary of contributions.

1. We solve a previous efficiency problem of TES (combinatorial explosion of the number of generated constraints) and also support features such as declaration sequences, structures, and *open* with the techniques of constrained environments and environment variables.

2. We solve SML's identifier status ambiguity (value variable vs. datatype constructor) while also computing minimal type error slices by using type constraints with context dependencies on identifier statuses.

3. We solve many other problems to provide clear and helpful type error slices for many different kinds of SML errors, which have been carefully designed to provide just the information the programmer needs.

4. This paper reports for a stripped-down core of SML the essential technical details of the TES machinery that solve the trickiest of the above-mentioned problems, and discusses some implementation issues.

5. We have an implementation that covers most of the SML language. A web demo is available, and there are downloadable packages for Ubuntu (covering also some other Debian-based systems) and CentOS (covering also some other Red-Hat-based systems). The web site is: <http://www.macs.hw.ac.uk/ultra/compositional-analysis/type-error-slicing>. We also have another implementation that faithfully implements just this paper.

7.2 Future work. We have already implemented some merging of minimal slices and are extending this idea to other kind of errors than record clashes, such as for unmatched signature specifications.

In the near future, we plan to finish extending our TES to the full SML language. This includes finishing handling the key feature of functors and less vital features such as flexible records or equality types, which can cause errors we currently do not detect.

We also plan to extend our ideas to other languages such as the F# programming language or the C++ template language.

Finally, we have done user evaluations and begun designing proper scientific experiments to compare the effectiveness in improving the productivity of real users of TES vs. more traditional type error messages.

References

- [1] A. W. Appel. A critique of Standard ML. *J. Funct. Program.*, 3(4), 1993.
- [2] M. Beaven, R. Stansifer. Explaining type errors in polymorphic languages. *ACM Letters on Programming Languages and Systems*, 2(1-4), 1993.
- [3] M. Blume. *Hierarchical Modularity and Intermodule Optimization*. PhD thesis, Princeton University, 1997.
- [4] M. Blume. Dependency analysis for Standard ML. *ACM Trans. Program. Lang. Syst.*, 21(4), 1999.
- [5] N. E. Boustani, J. Hage. Improving type error messages for generic java. In *PEPM*. ACM, 2009.
- [6] N. E. Boustani, J. Hage. Corrective hints for type incorrect generic java programs. In *PEPM*. ACM, 2010.
- [7] B. Braßel. TypeHope - there is hope for your type errors. In IFL04 [17].
- [8] L. Damas, R. Milner. Principal type-schemes for functional programs. In *POPL82*, New York, NY, USA, 1982. ACM.
- [9] H. Gast. Explaining ML type errors by data flows. In IFL04 [17].
- [10] C. Haack, J. B. Wells. Type error slicing in implicitly typed higher-order languages. In *ESOP*, vol. 2618 of *LNCS*. Springer, 2003.
- [11] C. Haack, J. B. Wells. Type error slicing in implicitly typed higher-order languages. *Science of Computer Programming*, 50(1-3), 2004.
- [12] J. Hage, B. Heeren. Heuristics for type error discovery and recovery. In *18th Int'l Symp., IFL 2006*, vol. 4449 of *LNCS*. Springer, 2007.
- [13] R. Harper. Programming in Standard ML, 2009. Working draft of August 20, 2009.
- [14] B. Heeren, J. Hage. Type class directives. In *7th Int'l Symp., PADL 2005*, vol. 3350 of *LNCS*. Springer, 2005.
- [15] B. Heeren, J. Jeuring, D. Swierstra, P. A. Alcocer. Improving type-error messages in functional languages. Technical report, Utrecht University, 2002.
- [16] B. J. Heeren. *Top Quality Type Error Messages*. PhD thesis, Universiteit Utrecht, The Netherlands, 2005.
- [17] *16th Int'l Workshop, IFL 2004*, vol. 3474 of *LNCS*. Springer, 2005.
- [18] O. Lee, K. Yi. Proofs about a folklore let-polymorphic type inference algorithm. *ACM Transactions on Programming Languages and Systems*, 20(4), 1998.
- [19] B. S. Lerner, M. Flower, D. Grossman, C. Chambers. Searching for type-error messages. In *ACM SIGPLAN 2007 Conference PLDI*. ACM, 2007.
- [20] B. J. McAdam. On the unification of substitutions in type inference. In *10th Int'l Workshop, IFL'98*, vol. 1595 of *LNCS*. Springer, 1999.
- [21] M. Neubauer, P. Thiemann. Discriminative sum types locate the source of type errors. In *8th ACM SIGPLAN Int'l Conference, ICFP 2003*. ACM, 2003.
- [22] M. Odersky, M. Sulzmann, M. Wehr. Type inference with constrained types. *Theor. Pract. Object Syst.*, 5(1), 1999.
- [23] F. Pottier. A modern eye on ML type inference: old techniques and recent developments. Lecture notes for the APPSEM Summer School, 2005.
- [24] F. Pottier, D. Rémy. The essence of ML type inference. In B. C. Pierce, ed., *Advanced Topics in Types and Programming Languages*, chapter 10. MIT Press, 2005.
- [25] P. J. Stuckey, M. Sulzmann, J. Wazny. Interactive type debugging in Haskell. In *Haskell '03: Proceedings of the 2003 ACM SIGPLAN workshop on Haskell*, New York, NY, USA, 2003. ACM.
- [26] P. J. Stuckey, M. Sulzmann, J. Wazny. Improving type error diagnosis. In *Haskell '04: Proceedings of the 2004 ACM SIGPLAN workshop on Haskell*, New York, NY, USA, 2004. ACM.
- [27] P. J. Stuckey, M. Sulzmann, J. Wazny. Type processing by constraint reasoning. In *4th Asian Symp., APLAS 2006*, vol. 4279 of *LNCS*. Springer, 2006.
- [28] M. Wand. Finding the source of type errors. In *13th ACM SIGACT-SIGPLAN Symp., POPL'86*, New York, NY, USA, 1986. ACM.
- [29] J. Wazny. *Type inference and type error diagnosis for Hindley/Milner with extensions*. PhD thesis, University of Melbourne, Australia, 2006.
- [30] J. Yang. Explaining type errors by finding the source of a type conflict. In *1st Workshop, SFP'99*, Exeter, UK, 2000. Intellect Books.
- [31] J. Yang, G. Michaelson, P. Trinder. Explaining polymorphic types. *The Computer Journal*, 45(4), 2002.
- [32] J. Yang, J. Wells, P. Trinder, G. Michaelson. Improved type error reporting. In *12th Int'l Workshop, IFL 2000*, vol. 2011 of *LNCS*. Springer, 2001.

A. Additional details for Section 4

A.1 Details for Section 4.1 (External syntax)

Remark about syntactic restrictions on our external syntax. Note that we do not enforce all the syntactic restrictions of the SML syntax. For example, in SML, in a recursive declaration such as $\text{val rec pat} \stackrel{l}{=} \text{exp}$, the expression exp must be a fn-expression.

A.2 Details for Section 4.2 (Constraint syntax)

Comparison with Pottier and Rémy's let-constraints. Our constraint system has evolved through many versions. One earlier version of our constraint system had a kind of constraint that was very close to the let-constraints¹ of systems of Pottier and Rémy [24, 23]. Pottier and Rémy define one system [24] that is an instance of HM(X) [22], and Pottier defines a very similar system [23] as a variation of the Damas/Milner type system. In our discussion, we will collectively refer to these two systems as the PR (Pottier/Rémy) system and ignore their technical differences, although our presentation will follow more closely the presentation of Pottier and Rémy [24].

In PR, a constraint can, among other things, be a let-constraint, a subtyping constraint, a type scheme instantiation constraint, a conjunction of constraints, or the constant (and satisfied) `true` constraint. A PR let-constraint looks like $\text{let } id:\dot{\sigma} \text{ in } C$ where $\dot{\sigma}$ ranges over type schemes, and C ranges over constraints. In PR, type schemes are of the form $\forall \bar{X}[C].T$ where \bar{X} is a type variable set, C is a constraint, and T is a type. We borrow for our discussion two abbreviations that Pottier and Rémy define: (1) the form $\forall \bar{X}.T$ stands for the type scheme $\forall \bar{X}[\text{true}].T$, and (2) the form $\text{let } id:T \text{ in } C$ stands for $\text{let } id:\forall \emptyset.T \text{ in } C$.

The idea of let-constraints is that a constraint of the form

$$\text{let } id:\forall \bar{X}[C].T \text{ in } (id = T_1 \wedge id = T_2)$$

is (roughly) equivalent to a constraint of this form:

$$(\exists \bar{X}.(C \wedge T = T_1)) \wedge (\exists \bar{X}.(C \wedge T = T_2)) \wedge (\exists \bar{X}.C)$$

The key point is that one can get the effect of making the appropriate number of copies of C and T while keeping the size of the constraint proportional to the program size. The constraints will need to be copied and each copy solved independently, but each copy can be solved immediately before the next copy is made, avoiding an exponential increase in the amount of memory used during constraint solving. To get the full benefit of this, an implementation should be eager in simplifying C and calculating T as much as possible before making any copies. (In our application, it could be good to also be lazy in simplifying and calculating only those portions of C and T that are actually needed by the uses of id , because our type error slicer needs to spend most of its time finding minimal portions of unsatisfiable constraints. We leave investigating this idea for future work.)

In our latest system, the equivalent of let-constraints can be represented as a special case of what our system supports. Informally, a let-constraint of the form $\text{let } id:\forall \bar{X}[C_1].T \text{ in } C_2$ generated for a SML recursive `let`-binding would be represented in our system by (using a combination of rules (G2) and (G12) in Fig. 4)

$$[\text{poly}(\downarrow id = \langle \tau, \nu \rangle; e_1); e_2]$$

where C_i is represented by e_i and T is represented by τ . (Let-constraints generated for other SML forms would not necessarily get the same representation.) There is no explicit representation of \bar{X} in the representation in our system; instead the correct set of

¹Technically, the let-constraints of Pottier and Rémy are based on their more primitive `def`-constraints.

type variables that can be quantified is calculated by `toPoly` which generates type schemes when it handles environments of the form $\text{poly}(e)$ (see Fig. 6).

We now give an example comparing the constraints that would be generated for SML recursive value declarations in the PR system and our system. Consider the SML expression

$$\text{let val rec } f = \text{fn } z \Rightarrow \text{exp}_1 \text{ in } \text{exp}_2$$

where exp_1 and exp_2 are two sub-expressions. The constraint generated in PR for this let-expression would be

$$\text{let } f:\forall XY[\text{let } f:X \rightarrow Y \text{ in let } z:X \text{ in } C_1].X \rightarrow Y \text{ in } C_2$$

where X and Y are internal type variables, where XY is PR notation for the set $\{X, Y\}$, where C_i for $i \in \{1, 2\}$ is the constraint generated for exp_i , and where Y is the result type of exp_1 . Due to the way let-constraints declare a local environment, the PR system needs two binders for f . The outer one polymorphically binds the occurrences of f in exp_2 and the inner one monomorphically binds the occurrences of f in exp_1 .

Some of the differences between PR and our system can be seen when comparing how this example is handled. Our constraint generator builds roughly² the following constraint (technically, an environment) for the example let-expression:

$$[\text{poly}(\downarrow f = \langle \alpha_1 \rightarrow \alpha_2, \nu \rangle; [(\downarrow z = \alpha_1); e_1]); e_2]$$

In contrast to how PR handles this example, only one binder for f is needed in our system. Two features of our system interact to allow this. First, in a constrained environment $(e_1; e_2)$, the bindings from e_1 are available in e_2 , but also form part of the result (except where bindings in e_2 shadow them). Second, in an environment of the form $\text{poly}(e)$, the `poly` operator changes the status of binders in the result from the status they had internally. In the example constraint (environment) above, f 's binder is monomorphic within the scope of the `poly` operator (in e_1) and polymorphic outside (in e_2).

There is a sense in which what the PR system does is similar to what would happen in our system if the `poly` operator worked on just single types or single bindings rather than entire environments. It is significant that we can form environments of the form $\text{poly}(\downarrow vid = \langle \tau, is \rangle; e_1); e_2$, in which the type for vid is available monomorphically in e_1 and polymorphically in e_2 .

The differences between the PR system and our system gain greater significance when we consider how to handle the SML module system. The most basic construct of the module system is what forms the body of a structure, namely a sequence of declarations $dec_1 \cdots dec_n$. For this discussion, assume each dec_i declares exactly one identifier x_i . Consider how declaration sequences can be handled by the PR system and our system. PR can handle such a sequence with nested let-constraints as follows:

$$\text{let } x_1:\dot{\sigma}_1 \text{ in } (\cdots \text{let } x_n:\dot{\sigma}_n \text{ in } C_0 \cdots)$$

The constraints must be nested as indicated because each x_i is only visible in the “in” part of the corresponding let-constraint, where an identifier binding occurrence is visible when constraints can refer to it. In contrast, our system handles the same declaration sequence with the environment

$$e_1; \cdots; e_n$$

²We have omitted labels and simplified a bit. The actual constraint that is generated (still omitting labels though) is

$$[(ev_2 = \text{poly}(\downarrow f = \langle \alpha_1, \nu \rangle; [(ev_1 = \langle \downarrow z = \alpha_2 \rangle); ev_1; e_1; c_1]; c_2)); ev_2; e_2; c_3]$$

where $c_1 = \langle \alpha_3 = \alpha_2 \rightarrow \alpha_4 \rangle$, $c_2 = \langle \alpha_1 = \alpha_3 \rangle$, $c_3 = \langle \alpha_5 = \alpha_6 \rangle$, $\langle \alpha_4, e_1 \rangle$ is generated for exp_1 , $\langle \alpha_6, e_2 \rangle$ is generated for exp_2 , and α_5 is the type of the entire let-expression.

where e_i is the environment generated for the declaration dec_i for each $i \in \{1, \dots, n\}$.

The importance of the difference becomes clearer when we consider how to represent full structures and structure bindings. Take the above example declaration sequence and wrap it up in a structure definition:

```
structure sid = struct dec1 ··· decn end
```

A structure expression packs into a unit a sequence of declarations. The normal scope of the declarations ends at the end of the structure, and subsequent access to the declarations must go through the structure itself, which must first be bound to a name via either a structure declaration like above or a functor application. When performing type inference for SML structure expressions, it is most natural and straightforward that the type inferred for a structure will be a sequence of individual mappings from declared names to their types³. Such sequences are often called *environments*. It seems clear that any type inference method will need to handle environments.

The PR system has never been extended to handle ML-style structures⁴, but let us imagine how it might be extended to do this. First, let us point out that Pottier and Rémy allow abbreviating the above example of nested let-constraints as follows:

```
let  $\Gamma_d$  in  $C_0$ , where  $\Gamma_d = x_1:\sigma_1; \dots; x_n:\sigma_n$ 
```

Let us call this constraint C_d where the “d” means “declarations”. Given an SML structure definition, this kind of constraint can represent the constraints required for typability of the sequence of declarations in the structure body, and it is the only easy way to do so in the context of the PR system.

Now, how do we represent the connection of the structure’s body to the structure’s name? The immediately (and naively) obvious idea is to extend PR with let-constraints of a form similar to **let** $sid:\Gamma_s$ **in** C , where sid is a structure identifier, and Γ_s is an environment (the type of a structure). Let us call this new constraint C_s . This is not enough, because there needs to be some way to connect the constraint C_d to the environment Γ_s . In fact, the environment Γ_d inside C_d is just what we need, but there is no easy way to get at it, because there is no mechanism in PR for generating an environment from a constraint. The easiest thing to do is to nest the entire constraint C_s inside the constraint C_0 inside of C_d , because the types of the x_i ’s are not accessible from outside C_d , but this seems like turning the program inside out, because the entire rest of the program must be nested inside the scope of the constraints for just the structure’s body.

So one might then want to extend the PR constraint system with an exporting mechanism and generate a constrained environment of the form $[C_d].\Gamma_s$ for the structure expression where C_d would export the type schemes of the x_i ’s and where Γ_s would refer to these exported type schemes. But, all this technicality really shouldn’t be needed because Γ_d is already the environment that we would want to generate for the structure expression.

The way our constraint system achieves that is by instead of having only one mechanism (the let-constraints) to bind identifiers and to restrict their scope (let-constraints define a local scope), it has two separate mechanisms: one for bindings that does not restrict the scope of the binders (we obtain this behaviour by having binding constraints of form $\downarrow id=x$ and by having our general con-

strained environments forms $e_1;e_2$ where the accessors occurring in e_2 can depend on the binders occurring in e_1), and another one for constraining the scope of a type environment (obtained thanks to our environments of the form $[e]$). The environment we generate for the structure expression presented above is then similar to the environment Γ_d .

A.3 Details for Section 4.3 (Constraint generation)

An additional view of the constraints generated initially.

Our constraint generator (Fig. 4) only generates restricted forms of environments (eg where “g” stands for “generation”). Let us present these restricted forms, where t is a restriction of τ , and the other forms are restrictions of e (where “p” stands for “poly” and “l” for “labelled”):

$$\begin{aligned} t &\in \text{ShallowTypes} ::= \alpha \mid \alpha \delta \mid \alpha \gamma \mid \alpha_1 \rightarrow \alpha_2 \\ lbind \in \text{LabBind} & ::= \downarrow tc \stackrel{l}{=} \gamma \mid \downarrow sid \stackrel{l}{=} ev \mid \downarrow tv \stackrel{l}{=} \alpha \\ & \quad \mid \downarrow vid \stackrel{l}{=} \alpha \mid \downarrow vid \stackrel{l}{=} ris \mid \downarrow vid \stackrel{l}{=} \alpha \\ lc &\in \text{LabCs} ::= ev_1 \stackrel{l}{=} ev_2 \mid \alpha \stackrel{l}{=} t \\ lacc &\in \text{LabAcc} ::= acc^l \\ lev &\in \text{LabEnvVar} ::= ev^l \\ eip &\in \text{InPolyEnv} ::= lacc \mid lc \mid eip_1; eip_2 \\ ep &\in \text{PolyEnv} ::= \downarrow vid \stackrel{l}{=} \langle \alpha, ris \rangle \mid ep; eip \mid eip; ep \\ eg &\in \text{GenEnv} ::= \square \mid lev \mid lbind \mid lacc \mid lc \mid ev = e \\ & \quad \mid poly(ep) \mid eg_1; eg_2 \end{aligned}$$

The rules of our constraint generator either return environments e (rules (G12)-(G14),(G16)) or constrained variables of the form $\langle var, e \rangle$ where e constrains var . In such a constrained variable, var is in some cases an internal type variable α (rules (G1)-(G11),(G15)) and in some other cases an environment variable ev (rules (G17)-(G18)). We chose not to have a constructor of constrained types that would build an internal type from an environment and an internal type (as $e_1;e_2$ builds a constrained environment from two environments), as it simplifies the presentation of our system by not having deep type structures. Such a system with constrained types could be investigated. Having chosen to return pairs of the form $\langle \alpha, e \rangle$ for expressions, we then decided to follow the same pattern for structure expressions and return pairs of the form $\langle ev, e \rangle$ instead of returning constrained environments of the form $e;ev$.

A.4 Details for Section 4.4 (Constraint solving)

An additional view of the environments generated at constraint solving.

During constraint solving (see Fig. 7), a unification context of the form $\langle u, e \rangle$ is maintained. Such an environment e has a restricted form as follows (it is of the form es):

$$\begin{aligned} sbind \in \text{SolvBind} & ::= \downarrow tc = \mu \mid \downarrow sid = es \mid \downarrow tv = \alpha \\ & \quad \mid \downarrow vid = \sigma \mid \downarrow vid = is \\ esrhs \in \text{SolvEnvRHS} & ::= ev \mid sbind^{\bar{d}} \mid es_1; es_2 \\ es &\in \text{SolvEnv} ::= \square \mid \square; esrhs \end{aligned}$$

Moreover, if $\langle u, e \rangle$ occurs in a unification state $state$ and ev occurs in e then $ev \notin \text{dom}(u)$. It is also the case that, for any environment variable ev , if $ev \in \text{dom}(u)$ then $u(ev) \in \text{SolvEnv}$.

We sometimes call an environment of the form es , a “solved” environment.

Improvement of the generation of polymorphic environments.

Fig. 6 defines `toPoly` which is used by rule (P1) of our constraint solver to generate a polymorphic environment from a monomorphic one by quantifying the type variables not occurring in the types of the monomorphic bindings of the current unification context. In this figure $\bar{\tau}$ is the set of types of the monomorphic bindings from the current unification context. The set $\bar{\alpha}$ is the set

³The order of the sequence is important because a type scheme for one value identifier in a structure can refer to a type constructor name defined by the structure, while at the same time a type scheme for a different value identifier can use the same type constructor name to refer to a definition outside the structure.

⁴François Pottier told us this on 2010-08-09.

of type variables occurring in τ' (the type that we want to generalise to a *for all* type scheme) that can be generalised and quantified over. The dependencies in the dependency set \bar{d}' are the reasons for not generalising the type variables occurring in τ' that are not in $\bar{\alpha}$ (these dependencies are the reasons why some type variables are not allowed to be quantified over).

The computation of \bar{d}' and our constraining of τ' with \bar{d}' , even though a correct approximation (that cannot generate false errors and that will eventually allow obtaining minimal type errors), could be refined, thereby speeding up minimisation. We will now present how this can be done.

Let us first define two functions `getDepsVar` and `putDepsVar`. The application `getDepsVar(α, τ, \emptyset)` will result in the dependency set occurring in τ on the paths from its root node to any occurrence of α . The application `putDepsVar(τ, α, \bar{d})` will result in the constraining of the occurrences of the type variable α in τ with the dependency set \bar{d} . The function `getDepsVar` is defined as follows:

$$\begin{aligned} \text{getDepsVar}(\alpha, \alpha', \bar{d}) &= \begin{cases} \bar{d}, & \text{if } \alpha = \alpha' \\ \emptyset, & \text{otherwise} \end{cases} \\ \text{getDepsVar}(\tau \mu, \alpha, \bar{d}) &= \text{getDepsVar}(\tau, \alpha, \bar{d}) \\ \text{getDepsVar}(\tau_1 \rightarrow \tau_2, \alpha, \bar{d}) &= \cup_{i=1}^2 \text{getDepsVar}(\tau_i, \alpha, \bar{d}) \\ \text{getDepsVar}(\tau^{\bar{d}}, \alpha, \bar{d}') &= \text{getDepsVar}(\tau, \alpha, \bar{d} \cup \bar{d}') \end{aligned}$$

The function `putDepsVar` is defined as follows:

$$\begin{aligned} \text{putDepsVar}(\alpha, \alpha', \bar{d}) &= \begin{cases} \alpha^{\bar{d}}, & \text{if } \alpha = \alpha' \\ \alpha, & \text{otherwise} \end{cases} \\ \text{putDepsVar}(\tau \mu, \alpha, \bar{d}) &= \text{putDepsVar}(\tau, \alpha, \bar{d}) \mu \\ \text{putDepsVar}(\tau_1 \rightarrow \tau_2, \alpha, \bar{d}) &= \tau_1' \rightarrow \tau_2' \\ &\text{where for } i \in \{1, 2\}, \tau_i' = \text{putDepsVar}(\tau_i, \alpha, \bar{d}) \\ \text{putDepsVar}(\tau^{\bar{d}}, \alpha, \bar{d}') &= \text{putDepsVar}(\tau, \alpha, \bar{d}' \cup \bar{d}) \end{aligned}$$

Let us now present another way of constraining τ' in Fig. 6 (different from constraining it with \bar{d}'). In the following, $\tau', \bar{\tau}$ and $\bar{\alpha}$ are the same as in Fig. 6. First,

$$\{\alpha_1, \dots, \alpha_n\} = (\text{vars}(\tau') \cap \text{ITyVar}) \setminus (\bar{\alpha} \cup \{\alpha_{\text{dum}}\})$$

is the set of type variables that are not allowed to be quantified in the generated type scheme. Then,

$$\forall i \in \{1, \dots, n\}. \bar{d}_i = \{d \mid \tau_0 \in \bar{\tau} \wedge d \in \text{getDepsVar}(\tau_0, \alpha_i, \emptyset)\}$$

is the set of reasons for α_i for not being quantified over. Finally,

$$\forall i \in \{1, \dots, n\}. \tau_i' = \text{putDepsVar}(\tau_{i-1}', \alpha_i, \bar{d}_i)$$

where $\tau_0' = \tau'$. The function `toPoly` would then generate the following type scheme: $\forall \bar{\alpha}. \tau_n'$.

A.5 Details for Section 4.5 (Minimisation and enumeration)

Clarification on the domain of the constraint filtering function. Note that our filtering function (Fig. 8) is not defined on all environments. These forms on which the function is defined correspond to the ones generated by our constraint generator (defined in Sec. A.3). When applied to unlabelled equality constraints on environments, our filtering function is only applied to unlabelled equality constraints of the form $ev=e$ because our constraint generator only generates variables as the left-hand-side of an equality constraint on environments. Similarly, we only apply our filtering function on constrained environments of the form e^l (constrained by a unique label).

The intended meaning of a labelled constraint is that it only must hold if the condition represented by the label is true. The machinery of this paper is designed to implement this intended semantics.

Given that, we then allow our filtering function to entirely discard labelled equality constraints, bindings, accessors and environ-

ment variables because when generated, these forms are always shallow. As a matter of fact, by definition, the right-hand-side of an accessor can only be a variable (*var*) or a raw status (*ris*). When generated, the right-hand-side of a binding is either a variable (*var*), a type constructor name (γ), or a raw status (*ris*). Concerning the generated equality constraints, by shallow we mean a *lc* constraint as defined in Sec. A.3. The non-shallow generated equality constraints are the non-labelled ones generated by rules (G4), (G12), (G13), (G14), (G16) and (G18). Because these constraints are not labelled, they are then never filtered out but the filtering function is recursively called on the right-hand-sides of these constraints as they can be non shallow.

Further explanations on minimisation and binding discarding. A step of the first phase of our minimisation algorithm is as follows: we test if we can remove a label l associated to a binder *bind* from the slice we want to minimise (and still obtain a type error slice) by first filtering the constraints of the original piece of code as follows: $\text{filt}(e, \bar{l}, \{l\})$, to obtain e' and where e is the environment generated for the original piece of code and \bar{l} is the label set labelling the current slice. In order not to mix up the bindings, the binder *bind* associated to l is then replaced by a non labelled dummy binder that cannot participate to any error but that still acts as a binder. If we then solve e' and obtain an error then no label labelling (in e') an accessor to (the dummy version of) our binder *bind* will occur in the found error (we give below an informal argument as why none of these accessors will be part of the new error). The bindings in this new error are then not mixed up. The found error is then the new slice to try to minimise further and next time the constraints will be filtered w.r.t. this slice, the binding *bind* will be completely thrown away (as well as the other constraints not participating in the new error).

Note that filtering itself does not prevent bindings to get mixed up because, for example, filtering allows throwing away the binder generated for the second occurrence of x in $\text{fn } x \Rightarrow \text{fn } x \Rightarrow x$ while not throwing away the binder generated for the first occurrence of x and not throwing away its accessor. However, we give below an informal argument as why we never filter a binder without filtering its accessor.

Let us now explain why when our unification algorithm returns an error, the error does not involve accessors to dummy binders or accessors without their corresponding binders.

According to rules (A1)-(A4), during unification the label labelling an accessor only gets recorded in a unification context if the accessed identifier is in the type environment stored in the unification context in the current state. In the environment (1) either the accessed identifier has a non labelled dummy static semantics (resulting from filtering) and then, according to rules (U3), (U4) and (S7), the label of the accessor does not get recorded into the unification environment. Given an accessor $\uparrow id=x$, according to rules (A1)-(A3), a constraint of the form $var=x$ where $var \in \text{DumVar}$ comes from the corresponding binder and $x \in \text{Var}$ or of the form $a=x$ where $x \in \text{RawldStatus}$, is generated. Then (U3), (U4) or (S7) applies and the newly generated constraint is discarded without generating anything. (2) Or the accessed identifier has a labelled non dummy static semantics and both the label associated to the binder and the one associated the bound occurrence will always occur together in the unification context.

This is why we believe that an identifier occurring at a non-binding position in a piece of code (an accessor) only occurs in a slice if it is bound and its binder occurs in the slice as well.

This argument would be enough if only the rules (G1), (G7), and (G17) were generating accessors because each of them generates a unique labelled accessor. This is unfortunately not the case for the rules (G6) and (G8). These rules generate labelled accessors as well as labelled equality constraints. We might then think that

these labelled equality constraints can participate in an error without having the accessor and its binder participating. We could then potentially have the label of an identifier occurring at a bound position participate in an error without having its binding occurrence participate. We do not believe so. First, let us point out that this issue could be fixed by enforcing in our labelled syntax that each identifier must be labelled by a unique label that does not label any larger piece of syntax. This enforcement can be considered as a design principle concerning the labelled syntax, that was not followed for clarity purposes. Let us now explain why it works even without the enforcement described above. Let us first consider (G6). The type returned by the rule (the type α_2 in (G6)) does not directly occur on the left or right-hand-side of the generated constraint. For α_2 (which is the link with the context of the pattern) to be constrained to be equal to a type, first α_1 has to be constrained to be equal to a type. The only way for α_1 to be constrained is to solve the generated accessor which would result in the case of an error in having the binder of the identifier participate in the error. The same reasoning applies to (G8). The type returned by the rule (α') is constrained to be equal to $\alpha \delta$ which at this stage can be equal to any other type. This is true because we have the internal type constructor `ar` which corresponds to the arrow type constructor \rightarrow . Without this type constructor `ar`, during unification we could infer that $\alpha \delta$ cannot be an arrow type (of the form $\tau_1 \rightarrow \tau_2$) and generate an error. Because of that, we would obtain:

```

(..datatype 'a t = U of 'a
 ..datatype (..) = T of (..) t
 ..val rec g = fn (..) => (..)
 ..(..val rec h = fn U x => T x
 ..h (U g)..)

```

as a minimal slice for:

```

structure S = struct
  datatype 'a t = U of 'a
  datatype 'a t = T of (..) t
  val rec g = fn v => v
  val rec f = fn v => let val rec h = fn U x => T x
                      in h (U g)
                    end
end
end

```

In the first slice `t`'s occurrence in `T`'s declaration is not bound to the same occurrence of `t` as in the original code.

But because we allow $\alpha \delta$ to be equal to any type as long as δ is not constrained to be equal to a type constructor name or to `ar` we then need to resolve the accessor generated in rule (G8) to obtain a type error if any. For the piece of code presented above we then obtain the following minimal type error slice instead of the one displayed above:

```

(..datatype 'a t = U of 'a
 ..datatype 'a t = T of (..) t
 ..val rec g = fn (..) => (..)
 ..(..val rec h = fn U x => T x
 ..h (U g)..)

```

This type error slice differs from the previous one by the presence of the second binding occurrence of `t` in the slice.

Because of the invariant that if a binder is filtered out then its bound occurrences are also filtered out, we can then easily compute the free identifiers thanks to rule (A4) which is the rule for an accessor for which no binder exists in the current unification environment (free identifier) or for which the binder is hidden.

We could also enforce this invariant by (1) ensuring that identifiers are labelled independently from any other piece of syntax (as explained above), (2) ensuring that at constraint generation, if a label l labels an accessor then it does not label any other constraint, and (3) discarding accessors when the corresponding binders are dummy binders (binding a dummy variable or the status `a`).

Alternatively, we could enforce this invariant by adding an extra component to unification contexts as follows: $\langle u, e, \bar{l} \rangle$, where \bar{l} indicates the labels that are not allowed to participate in an error. If an error is found involving a label in \bar{l} then this error is not reported.

A.6 Details on Section 4.6 (Slicing)

An alternative formal presentation of the slicing algorithm.

We will now provide an alternative generic definition of the external syntax presented in Fig. 3. We also extend our algorithm to “dot” terms. (This alternative definition is not the one presented in the main body of this paper for readability issues.) First we define abstract syntax trees as follows:

```

class ∈ Class ::= ty | conbind | datname | dec | atexp
                | exp | atpat | pat | strdec | strexp
prod ∈ Prod ::= tyVar | tyArr | tyCon
                | conbindOf | datnameCon
                | decRec | decDat | decOpn
                | atexpLet | expFn
                | strdecDec | strdecStr
                | strexpId | strexpSt
                | vid | app
dot ∈ Dot ::= dotE | dotP | dotD | dotS
node ∈ Node ::= ⟨class, prod⟩
tree ∈ Tree ::= ⟨node, l, tree⟩ | ⟨dot, tree⟩ | id

```

A node in a tree $tree$ can either be a labelled node of the form $\langle node, l, tree \rangle$, an unlabelled “dot” node of the form $\langle dot, tree \rangle$, or a leaf of the form id .

Let a *term* be any term that can be derived from Fig. 3:

```

term ∈ Term ::= ty | cb | dn | dec | atexp | exp
                | atpat | pat | sdec | seap

```

Fig. 13 defines the function toTree that associates a *tree* to each *term*.

We are now going to redefine our flattening and tidying functions, and our slicing algorithm. First, we need a mechanism to distinguish between declarations and non-declarations, similar to the distinction between `ExpPart` and `Part \ ExpPart`. Let $isClass(tree, \overrightarrow{class})$ be true iff $tree = \langle \langle class, prod \rangle, l, \overrightarrow{tree} \rangle$ and $class \in \overrightarrow{class}$, used to check the class of the root node of a tree. Let $declares(tree)$ be true iff $isClass(tree, \{\text{dec}, \text{strdec}\})$.

Let us redefine our flattening function `flat`:

$$\text{flat}(\langle \rangle) = \langle \rangle$$

$$\text{flat}(\langle (tree) @ \overrightarrow{tree} \rangle) = \begin{cases} \langle (tree_1, \dots, tree_n) @ \text{flat}(\overrightarrow{tree}), \\ \text{if } tree = \langle dot, \langle tree_1, \dots, tree_n \rangle \rangle \\ \text{and } (\forall i \in \{1, \dots, n\}. \neg \text{declares}(tree_i)) \\ \text{or } tree = \langle \rangle \rangle \\ \langle (tree) @ \text{flat}(\overrightarrow{tree}), \text{otherwise} \end{cases}$$

We slightly altered our flattening function from the one defined in Sec. 4.6, by adding the condition “or $\overrightarrow{tree} = \langle \rangle$ ”. As a matter of fact, the condition “ $\forall i \in \{1, \dots, n\}. \neg \text{declares}(tree_i)$ ” is there to ensure that bindings are not mixed up as explained in Sec. 4.6. However, flattening the last dot term (if it actually is a dot term) cannot mix up the bindings because there is no identifier left to bind. Therefore, flattening $\langle ..val x = 1..val x = true.. \rangle$ would lead to $\langle ..val x = 1..val x = true.. \rangle$. We do not believe that this is an improvement of the function because we have not found a concrete example where this situation occurs.

We also redefine our function that tidies sequences of declarations in structure expressions as follows:

Figure 13 From terms to trees

Types	$\text{toTree}(tv^l)$	$= \langle \langle \text{ty}, \text{tyVar} \rangle, l, \langle tv \rangle \rangle$
	$\text{toTree}(ty_1 \xrightarrow{l} ty_2)$	$= \langle \langle \text{ty}, \text{tyArr} \rangle, l, \langle \text{toTree}(ty_1), \text{toTree}(ty_2) \rangle \rangle$
	$\text{toTree}(ty \ tc^l)$	$= \langle \langle \text{ty}, \text{tyCon} \rangle, l, \langle \text{toTree}(ty), tc \rangle \rangle$
Constructor bindings	$\text{toTree}(vid_c^l)$	$= \langle \langle \text{conbind}, \text{vid} \rangle, l, \langle vid \rangle \rangle$
	$\text{toTree}(vid \text{ of }^l ty)$	$= \langle \langle \text{conbind}, \text{conbindOf} \rangle, l, \langle vid, \text{toTree}(ty) \rangle \rangle$
Datatype names	$\text{toTree}(tv \ tc^l)$	$= \langle \langle \text{datname}, \text{datnameCon} \rangle, l, \langle tv, tc \rangle \rangle$
Declarations	$\text{toTree}(\text{val } \text{rec } pat \stackrel{l}{=} exp)$	$= \langle \langle \text{dec}, \text{decRec} \rangle, l, \langle \text{toTree}(pat), \text{toTree}(exp) \rangle \rangle$
	$\text{toTree}(\text{datatype } dn \stackrel{l}{=} cb)$	$= \langle \langle \text{dec}, \text{decDat} \rangle, l, \langle \text{toTree}(dn), \text{toTree}(cb) \rangle \rangle$
	$\text{toTree}(\text{open}^l sid)$	$= \langle \langle \text{dec}, \text{decOpn} \rangle, l, \langle sid \rangle \rangle$
Expressions	$\text{toTree}(vid_c^l)$	$= \langle \langle \text{atexp}, \text{vid} \rangle, l, \langle vid \rangle \rangle$
	$\text{toTree}(\text{let}^l dec \text{ in } exp \text{ end})$	$= \langle \langle \text{atexp}, \text{atexpLet} \rangle, l, \langle \text{toTree}(dec), \text{toTree}(exp) \rangle \rangle$
	$\text{toTree}(\text{fn } pat \xrightarrow{l} exp)$	$= \langle \langle \text{exp}, \text{expFn} \rangle, l, \langle \text{toTree}(pat), \text{toTree}(exp) \rangle \rangle$
	$\text{toTree}([\text{exp } \text{atexp}]^l)$	$= \langle \langle \text{exp}, \text{app} \rangle, l, \langle \text{toTree}(exp), \text{toTree}(\text{atexp}) \rangle \rangle$
Patterns	$\text{toTree}(vid_p^l)$	$= \langle \langle \text{atpat}, \text{atpatVid} \rangle, l, \langle vid \rangle \rangle$
	$\text{toTree}(vid^l \text{ atpat})$	$= \langle \langle \text{pat}, \text{app} \rangle, l, \langle vid, \text{toTree}(\text{atpat}) \rangle \rangle$
Structure declarations	$\text{toTree}(\text{structure } sid \stackrel{l}{=} sexp)$	$= \langle \langle \text{strdec}, \text{strdecStr} \rangle, l, \langle sid, \text{toTree}(sexp) \rangle \rangle$
Structure expressions	$\text{toTree}(sid^l)$	$= \langle \langle \text{strexpr}, \text{strexprId} \rangle, l, \langle sid \rangle \rangle$
	$\text{toTree}(\text{struct}^l sdec_1 \dots sdec_n \text{ end})$	$= \langle \langle \text{strexpr}, \text{strexprSt} \rangle, l, \langle \text{toTree}(sdec_1), \dots, \text{toTree}(sdec_n) \rangle \rangle$
Dot terms	$\text{toTree}(\text{dot-e}(\langle pt_1, \dots, pt_n \rangle))$	$= \langle \text{dotE}, \langle \text{toTree}(pt_1), \dots, \text{toTree}(pt_n) \rangle \rangle$
	$\text{toTree}(\text{dot-d}(\langle pt_1, \dots, pt_n \rangle))$	$= \langle \text{dotD}, \langle \text{toTree}(pt_1), \dots, \text{toTree}(pt_n) \rangle \rangle$
	$\text{toTree}(\text{dot-p}(\langle pat_1, \dots, pat_n \rangle))$	$= \langle \text{dotP}, \langle \text{toTree}(pat_1), \dots, \text{toTree}(pat_n) \rangle \rangle$
	$\text{toTree}(\text{dot-s}(\langle pt_1, \dots, pt_n \rangle))$	$= \langle \text{dotS}, \langle \text{toTree}(pt_1), \dots, \text{toTree}(pt_n) \rangle \rangle$

$\text{tidy}(\langle \rangle) = \langle \rangle$
 $\text{tidy}(\langle \langle \text{dotD}, \text{tree}_1 \rangle, \langle \text{dotD}, \text{tree}_2 \rangle \rangle @ \text{tree})$
 $= \text{tidy}(\langle \langle \text{dotD}, \text{tree}_1 @ \text{tree}_2 \rangle \rangle @ \text{tree}),$
 if $\forall \text{tree} \in \text{ran}(\text{tree}_1). \neg \text{declares}(\text{tree})$
 $\text{tidy}(\langle \langle \text{dotD}, \emptyset \rangle \rangle @ \text{tree}) = \text{tidy}(\text{tree}),$ if none of the above applies
 $\text{tidy}(\langle \text{tree} \rangle @ \text{tree}) = \langle \text{tree} \rangle @ \text{tidy}(\text{tree}),$ if none of the above applies

We also need the function `getDot` that generates a dot marker (a term in Dot) from a node kind *node*:

```

getDot((ty, prod)) = dotE
getDot((conbind, prod)) = dotE
getDot((datname, prod)) = dotE
getDot((dec, prod)) = dotD
getDot((atexp, prod)) = dotE
getDot((exp, prod)) = dotE
getDot((atpat, prod)) = dotP
getDot((pat, prod)) = dotP
getDot((strdec, prod)) = dotD
getDot((strexpr, prod)) = dotS
  
```

This function is, among other things, used by rule (SL1) when discarding a labelled node and so generating a new dot node.

Fig. 14 redefines our slicing algorithm. Note that rule (SL9) generates the dot marker `dotE`, but we could have used any of the terms in Dot because the flattening function `flat` discards such terms. The functions `sl1` and `sl2` are defined on trees but also on sequences of trees in rules (SL6) and (SL7). Patterns are treated specially because in our system we do not add the label associated to the fn-expression to the following type error slice (the error being that *x* is declared as a unary datatype constructor and occurs at a nullary position in a pattern):

```

{..datatype(..) = x of(..)
 ..fn x => (<..>..)}
  
```

This is because the unconfirmed binder generated for *x*'s occurrence in the fn-expression turns into an accessor at constraint solving (*x* being declared as a datatype constructor) and this ac-

cessor can directly refer to *x*'s binder without using any constraint labelled by the label associated to the fn-expression.

B. Case study: modify user data types using TES.

Our TES is generally of great help when coding in SML. It is particularly helpful when one wants to modify a user data type in a well-typed program. Let us consider the very simple program provided in Fig. 15a where we define a structure `Id` to deal with labelled identifiers (see the type `idlab`). In this structure we define some functions to handle labelled identifiers such as a function to compare two labelled identifiers (`compare`), or a function to build a labelled identifier from a label and an identifier (`cons`).

Now, let us change the type `idlab`, for a more convenient type: `type idlab = {id : id, lab : lab}` which is a record type containing two fields, one named `id` of type `id` and a second one named `lab` of type `lab`. Records are usually preferred over tuples because they are more flexible and meaningful thanks to the field names.

For example, one can access the field named `id` in an expression *x* of type `idlab` (the new type `idlab`) as follows: `#id(x:idlab)`. Records are more flexible than tuples because the order of the fields does not matter in a record. For example, `{id = 0, lab = 0}` is equivalent to `{lab = 0, id = 0}`. Note that a tuple `(id, lab)` is equivalent to a record `{1 = id, 2 = lab}`.

First of all, let us mention that when compiling the updated code with SML/NJ v.110.72, one obtains a type error report for each function defined in the structure `Id`. The report concerning the `compare` function is as follows (where the first line has been split into two lines to fit in the column):

```

test-prog.sml:14.1-31.4
Error: value type in structure doesn't match signature spec
  name: compare
  spec: ?.Id.idlab * ?.Id.idlab -> order
  actual: (int * int) * (int * int) -> order
  
```

Note that the reported region is the entire structure `Id`.

Figure 14 Slicing algorithm

$$(SL1) \text{sl}(\langle \text{node}, l, \overrightarrow{\text{tree}} \rangle, \bar{l}) = \begin{cases} \langle \text{node}, l, \text{sl}_1(\overrightarrow{\text{tree}}, \bar{l}) \rangle, & \text{if } (l \in \bar{l} \text{ and } \text{getDot}(\text{node}) \neq \text{dotS}) \text{ or isClass}(\text{sl}_1(\overrightarrow{\text{tree}}(0), \bar{l}), \{\text{pat}, \text{atpat}\}) \\ \langle \text{node}, l, \text{tidy}(\text{sl}_1(\overrightarrow{\text{tree}}, \bar{l})) \rangle, & \text{if } l \in \bar{l} \text{ and } \text{getDot}(\text{node}) = \text{dotS} \\ \langle \text{getDot}(\text{node}), \text{flat}(\text{sl}_2(\overrightarrow{\text{tree}}, \bar{l})) \rangle, & \text{otherwise} \end{cases}$$

$$(SL2) \text{sl}_1(\langle \text{dot}, \langle \text{tree}_1, \dots, \text{tree}_n \rangle \rangle, \bar{l}) = \langle \text{dot}, \text{flat}(\langle \text{sl}_2(\text{tree}_1, \bar{l}), \dots, \text{sl}_2(\text{tree}_n, \bar{l}) \rangle) \rangle$$

$$(SL3) \text{sl}_2(\langle \text{dot}, \langle \text{tree}_1, \dots, \text{tree}_n \rangle \rangle, \bar{l}) = \langle \text{dot}, \text{flat}(\langle \text{sl}_2(\text{tree}_1, \bar{l}), \dots, \text{sl}_2(\text{tree}_n, \bar{l}) \rangle) \rangle$$

$$(SL4) \text{sl}_1(\langle \text{node}, l, \overrightarrow{\text{tree}} \rangle, \bar{l}) = \text{sl}(\langle \text{node}, l, \overrightarrow{\text{tree}} \rangle, \bar{l})$$

$$(SL5) \text{sl}_2(\langle \text{node}, l, \overrightarrow{\text{tree}} \rangle, \bar{l}) = \text{sl}(\langle \text{node}, l, \overrightarrow{\text{tree}} \rangle, \bar{l})$$

$$(SL6) \text{sl}_1(\langle \text{tree}_1, \dots, \text{tree}_n \rangle, \bar{l}) = \langle \text{sl}_1(\text{tree}_1, \bar{l}), \dots, \text{sl}_1(\text{tree}_n, \bar{l}) \rangle$$

$$(SL7) \text{sl}_2(\langle \text{tree}_1, \dots, \text{tree}_n \rangle, \bar{l}) = \langle \text{sl}_2(\text{tree}_1, \bar{l}), \dots, \text{sl}_2(\text{tree}_n, \bar{l}) \rangle$$

$$(SL8) \text{sl}_1(\text{id}, \bar{l}) = \text{id}$$

$$(SL9) \text{sl}_2(\text{id}, \bar{l}) = \langle \text{dotE}, \langle \rangle \rangle$$

Figure 15 Using TES to modify user data types

(a) Structure defining labelled identifiers

```
signature ID = sig
  type id
  type lab
  type idlab

  val compare : idlab * idlab -> order
  val cons : id -> lab -> idlab
  val getId : idlab -> id
  val getLab : idlab -> lab
  val updId : idlab -> id -> idlab
  val updLab : idlab -> lab -> idlab
end

structure Id : ID = struct
  type id = int
  type lab = int
  type idlab = id * lab

  fun compare ((id1, lab1), (id2, lab2)) =
    case Int.compare (id1, id2) of
      EQUAL => Int.compare (lab1, lab2)
    | x => x

  fun cons id lab = (id, lab)

  fun getId (id, lab) = id
  fun getLab (id, lab) = lab

  fun updId (_, lab) id = cons id lab
  fun updLab (id, _) lab = cons id lab
end
```

(b) Type error obtained after a type change

```
File Edit Options Buffers Tools Errors SML
signature ID = sig
  type id
  type lab
  type idlab

  val compare : idlab * idlab -> order
  val cons : id -> lab -> idlab
  val getId : idlab -> id
  val getLab : idlab -> lab
  val updId : idlab -> id -> idlab
  val updLab : idlab -> lab -> idlab
end

structure Id : ID = struct
  type id = int
  type lab = int
  type idlab = {id : id, lab : lab}

  fun compare ((id1, lab1), (id2, lab2)) =
    case Int.compare (id1, id2) of
      EQUAL => Int.compare (lab1, lab2)
    | x => x

  fun cons id lab = (id, lab)

  fun getId (id, lab) = id
  fun getLab (id, lab) = lab

  fun updId (_, lab) id = cons id lab
  fun updLab (id, _) lab = cons id lab
end
```

(c) Program obtained after solving all the type errors

```
File Edit Options Buffers Tools Errors SML Help
signature ID = sig
  type id
  type lab
  type idlab

  val compare : idlab * idlab -> order
  val cons : id -> lab -> idlab
  val getId : idlab -> id
  val getLab : idlab -> lab
  val updId : idlab -> id -> idlab
  val updLab : idlab -> lab -> idlab
end

structure Id : ID = struct
  type id = int
  type lab = int
  type idlab = {id : id, lab : lab}

  fun compare ({id, lab}, {id' = id', lab' = lab'}) =
    case Int.compare (id, id') of
      EQUAL => Int.compare (lab, lab')
    | x => x

  fun cons id lab = {id = id, lab = lab}

  fun getId (idlab : idlab) = #id idlab
  fun getLab (idlab : idlab) = #lab idlab

  fun updId idlab id = cons id (getLab idlab)
  fun updLab idlab lab = cons (getId idlab) lab
end
```

test-prog.sml All (31,3) (SML)-----
(SML-TES) SLICING FINISHED WITH STATUS: slicer worked OK, program is typable

In contrast, Fig. 15b presents the highlighting that one obtains when running our TES on the updated piece of code. The error in focus (highlighted with a darker red) shows that the parameter of `compare` is a pair of pairs. The second pair (equivalent to a record with two fields named 1 and 2) clashes with the type of `compare`'s second parameter given in the signature `ID`, which is `idlab`, declared as a record with field names `id` and `lab` in the structure `Id`. In the parameter of `compare`, the second pair has its elements surrounded by grey boxes. We do so, because tuples do not have explicitly written field names. The first grey box surrounds the first element of a pair that corresponds to a record where the element would be in field with field name 1 (and similarly for the second box). Note that the number of boxes indicates the arity of the tuple. In addition to the highlighting, we also report a type error slice and the following message for this type error:

Record clash, the fields {id,lab} conflict with {1,2}

This error is not context dependent, so no context dependency is reported.

The light pink corresponds to slices other than the focused one. One can then start solving the errors one at a time by just editing the highlighted portions of code, to get from a well-typed program to another well-typed program (see Fig. 15c).

C. Extensions to handle more of SML

Let us now present some extensions of our TES in order to handle features such as local declarations, type functions or signatures. Some of these features were already used in the examples provided so far. We will now formally present how to handle them.

Some syntactic forms will sometimes need to be redefined. In this section, we will sometimes write $x \xrightarrow{s} y$ to mean that in the set s , syntactic forms of the form x are replaced by syntactic forms of the form y .

C.1 Local declarations

External syntax. First, let us extend our external syntax with local declarations as follows:

$$\text{dec} ::= \dots \mid \text{local}^l \text{dec}_1 \text{ in } \text{dec}_2 \text{ end}$$

For example,

```
val x = true
local val x = 1 in val y = x end
val z = x + 1
```

is untypable because `x`'s last occurrence is bound to its first occurrence and not to its second (assuming that `+` is the one from the Standard ML basis library).

Let us present another example:

```
val x = true
local val x = 1 in val y = x end
val z = fn w => (w y, w x)
```

Only the declaration z differs from the previous example. This piece of code is also untypable because w has a monomorphic type and is applied to y which is an integer and x which is a Boolean. This example will be reused later in this section.

Constraint syntax. We extend environments with local environments as follows:

$$e ::= \dots \mid \text{loc } e_1 \text{ in } e_2$$

The meaning of such an environment is that it builds an environment e_2 which depends on e_1 and only exports the binders of e_2 (only e_2 's binders can be accessed from outside the local environment). Such environments differ from environments of the form $e_1; e_2$ because an environment of the form $e_1; e_2$ builds a new environment from both e_1 and e_2 and exports both the binders of e_1 (not shadowed by e_2) and e_2 .

Environments of the form $[e]$ are not enough to handle local declarations because they do not allow partially exporting an environment. The requirement imposed by local declarations is that only e_1 and e_2 should be able to access e_1 's binders. Unfortunately, $[e_1; e_2]$ does not export e_2 's binders, and $[e_1]; e_2$ does not allow e_2 's accessors to refer to e_1 's binders. The solution was to introduce environments of the form $\text{loc } e_1 \text{ in } e_2$.

Note that these environments are not only used to generate constraints for local declarations, they are also used to, e.g., handle bindings of external type variables (see Sec. C.2). In Sec. 4 we allow binding occurrences of explicit type variables to have a larger scope than they should, which is harmless in the tiny language of Sec. 4, but needs to be (and is) fixed to work for full SML in Sec. C.2.

Constraint generation. We extend our constraint generator with the following rule:

$$(G19) \text{local}^l \text{dec}_1 \text{ in } \text{dec}_2 \text{ end} \triangleright (ev = e_1); \text{loc } ev^l \text{ in } e_2 \Leftarrow \text{dec}_1 \triangleright e_1 \wedge \text{dec}_2 \triangleright e_2 \wedge \text{dja}(e_1, e_2, ev)$$

Because our initial constraint generation algorithm generates new forms of constraints, we extend the eg forms as follows (see Sec. A.3):

$$eg ::= \dots \mid \text{loc } eg_1 \text{ in } eg_2$$

The forms generated by our initial constraint generator are in fact more restricted than that, but we already anticipate the forms generated by further extensions later, e.g., for type functions.

Constraint solving. We extend our constraint solver as follows:

- (L1) $\text{solve}(\langle u, e \rangle, \bar{d}, \text{loc } e_1 \text{ in } e_2) \rightarrow \text{succ}(\langle u'', e_0 \rangle)$,
if $\text{solve}(\langle u, e \rangle, \bar{d}, e_1) \rightarrow^* \text{succ}(\langle u', e' \rangle)$
and $\text{solve}(\langle u', e' \rangle, \bar{d}, e_2) \rightarrow^* \text{succ}(\langle u'', e'' \rangle)$
and $\text{diff}(e', e'') = \square; e'_1; \dots; e'_n$
and $e_0 = e; e'_1; \dots; e'_n$
- (L2) $\text{solve}(\langle u, e \rangle, \bar{d}, \text{loc } e_1 \text{ in } e_2) \rightarrow \text{err}(er)$,
if $\text{solve}(\langle u, e \rangle, \bar{d}, e_1) \rightarrow^* \text{succ}(\langle u', e' \rangle)$
and $\text{solve}(\langle u', e' \rangle, \bar{d}, e_2) \rightarrow^* \text{err}(er)$
- (L3) $\text{solve}(\langle u, e \rangle, \bar{d}, \text{loc } e_1 \text{ in } e_2) \rightarrow \text{err}(er)$,
if $\text{solve}(\langle u, e \rangle, \bar{d}, e_1) \rightarrow^* \text{err}(er)$

Constraint filtering (Minimisation and enumeration). We extend our filtering function as follows:

$$\text{filt}(\text{loc } e_1 \text{ in } e_2, \bar{l}_1, \bar{l}_2) = \text{loc } \text{filt}(e_1, \bar{l}_1, \bar{l}_2) \text{ in } \text{filt}(e_2, \bar{l}_1, \bar{l}_2)$$

Slicing. Finally, our slicing algorithm does not need to be extended but we need to update the tree syntax for programs as follows:

$$\text{Prod} ::= \dots \mid \text{decLoc}$$

We also need to extend the toTree function that associates trees of the form $tree$ to terms of the form $term$ as follows:

$$\text{toTree}(\text{local}^l \text{dec}_1 \text{ in } \text{dec}_2 \text{ end}) = \langle \langle \text{dec}, \text{decLoc} \rangle, l, \langle \text{toTree}(\text{dec}_1), \text{toTree}(\text{dec}_2) \rangle \rangle$$

Minimality. Let us illustrate what would happen if we were not generating an extra labelled environment variable in rule (G19). Consider the last example presented above. With our current system, we would obtain a type error slice involving the local declaration itself in addition to the nested declarations of x and y . If we were not to label the environment variable in rule (G19) or if we were to use e_1 instead of ev^l in the local constraint (and omit $ev = e_1$ which becomes useless), then we would obtain a type error slice that would look like:

```
<..val x = true
..val x = 1
..val y = x
..val z = fn w => <..w y..w x..>>
```

which is typable and therefore not a minimal type error slice of the piece of code presented above: both bound occurrences of x are bound to x 's second declaration.

C.2 Type declarations

External syntax. First, let us extend our external syntax with type functions as follows:

$$\text{Dec} ::= \dots \mid \text{type } dn \stackrel{l}{=} ty$$

For example,

```
type 'a t = 'a -> 'a -> 'a
datatype 'a u = U of 'a t
val x = U (fn x => x)
```

is untypable because U is applied to the identity function which cannot have the type $'a \rightarrow 'a \rightarrow 'a$.

Constraint syntax. We extend our constraint system with type functions:

$$\begin{aligned} \phi &\in \text{TypFunVar} && \text{(type function variables)} \\ \theta &\in \text{TypFun} && ::= \phi \mid \Lambda \alpha. \tau \mid \langle \theta, \bar{d} \rangle \\ \tau &\in \text{ITy} && ::= \dots \mid \theta \cdot \tau \\ c &\in \text{Constraint} && ::= \dots \mid \theta_1 = \theta_2 \\ var &\in \text{Var} && ::= \dots \mid \phi \end{aligned}$$

Let ϕ_{dum} be a distinguished dummy type function variable. We redefine DumVar to be $\text{DumVar} = \{\alpha_{\text{dum}}, ev_{\text{dum}}, \delta_{\text{dum}}, \phi_{\text{dum}}\}$.

We then have to change the binders and accessors for type constructors:

$$\downarrow tc = \mu \text{ Bind} \downarrow tc = \theta \quad \uparrow tc = \delta \text{ Accessor} \uparrow tc = \phi$$

Constraint generation. Fig. 16 modifies our rules for datatype names (G15), datatype declarations (G13) and type constructions (G8), and defines a new rule (G20) for type function declarations. Note the use of local environments (of the form $\text{loc } e_1 \text{ in } e_2$) in rules (G13) and (G20), used to handle binding occurrences of explicit type variables.

Because our initial constraint generation algorithm generates new forms of constraints, we extend the t forms as follows (see Sec. A.3):

$$t ::= \dots \mid \phi \cdot \alpha$$

Figure 16 Constraint generation rules for type functions

(G8) $ty\ tc^l \triangleright \langle \alpha', (\uparrow tc \stackrel{l}{=} \phi); (\alpha' \stackrel{l}{=} \phi \cdot \alpha); e \rangle$	$\Leftarrow ty \triangleright \langle \alpha, e \rangle \wedge \text{dja}(e, \alpha', \phi)$
(G15) $tw\ tc^l \triangleright \langle \alpha, \alpha', \downarrow tc \stackrel{l}{=} \Lambda \alpha'. \alpha, \downarrow tw \stackrel{l}{=} \alpha' \rangle$	$\Leftarrow \alpha \neq \alpha'$
(G13) $\text{datatype}\ dn \stackrel{l}{=} cb \triangleright ev = ((\alpha_1 \stackrel{l}{=} \alpha'_1 \ \gamma); (\alpha_1 \stackrel{l}{=} \alpha_2); e_1; \text{loc}\ e'_1 \text{ in poly}(e_2)); ev^l$	$\Leftarrow dn \triangleright \langle \alpha_1, \alpha'_1, e_1, e'_1 \rangle \wedge cb \triangleright \langle \alpha_2, e_2 \rangle \wedge \text{dja}(e_1, e_2, \gamma, ev)$
(G20) $\text{type}\ dn \stackrel{l}{=} ty \triangleright ev = ((\alpha_1 \stackrel{l}{=} \alpha_2); \text{loc}\ e'_1 \text{ in } (e_2; e_1)); ev^l$	$\Leftarrow dn \triangleright \langle \alpha_1, \alpha'_1, e_1, e'_1 \rangle \wedge ty \triangleright \langle \alpha_2, e_2 \rangle \wedge \text{dja}(e_1, e_2, ev)$

We also replace the initially generated type constructor binders as follows:

$$\downarrow tc \stackrel{l}{=} \gamma \xrightarrow{\text{LabBind}} \downarrow tc \stackrel{l}{=} \Lambda \alpha. \alpha'$$

Constraint solving. Because we replaced our binders for type constructors, we need to modify our environment application as follows ((EA1) and (EA3) are the same as before but repeated here, and (EA2) differs by the replacement of μ by θ):

$$\begin{aligned} \text{(EA1)} \quad (e'; \downarrow vid \stackrel{\bar{d}}{=} \forall \bar{\alpha}. \tau)(vid) &= \forall \bar{\alpha}. \tau^{\bar{d}} \\ \text{(EA2)} \quad (e'; \downarrow id \stackrel{\bar{d}}{=} x)(id) &= x^{\bar{d}}, \quad \text{if } x \text{ of the form } \tau, e, \text{ or } \theta \\ \text{(EA3)} \quad (e'; \downarrow id' \stackrel{\bar{d}}{=} x)(id) &= e'(id), \text{ if } id \neq id' \text{ or } x \in \text{IdStatus} \end{aligned}$$

We also extend unifiers as follows:

$$u \in \text{Unifier} ::= \{f_1 \cup f_2 \cup f_3 \cup f_4 \mid \begin{aligned} f_1 &\in \text{ITyVar} \rightarrow \text{ITy} \\ f_2 &\in \text{TyConVar} \rightarrow \text{ITyCon} \\ f_3 &\in \text{EnvVar} \rightarrow \text{Env} \\ f_4 &\in \text{TypFunVar} \rightarrow \text{TypFun} \end{aligned}\}$$

Because we added types of the form $\theta \cdot \tau$, we need to update our type building function with the following cases:

$$\text{build}(u, \theta \cdot \tau) = \begin{cases} \tau^{\bar{d}}[\{\alpha \mapsto \text{build}(u, \tau)\}], \\ \text{if } \text{build}(u, \theta) = (\Lambda \alpha. \tau')^{\bar{d}} \\ \alpha_{\text{dum}}, \text{ otherwise} \end{cases}$$

$$\text{build}(u, \Lambda \alpha. \tau) = \Lambda \alpha. \text{build}(\{\alpha\} \triangleleft u, \tau)$$

Our building function builds internal types of the form $\theta \cdot \tau$ by first building the type function. If building the type function leads to a type function variable then our building function gives up building the application and returns the dummy type variable α_{dum} . This behaviour is correct in our system because the constraint generation rule (G8) is the only rule generating type function applications and it constrains ϕ before generating the application $\phi \cdot \alpha$. Thus, at constraint solving, when dealing with $\phi \cdot \alpha$, the constraints on ϕ will already have been dealt with.

Fig. 17 extends our constraint solver with two new rules to handle our new internal types of the form $\theta \cdot \tau$.

Constraint filtering (Minimisation and enumeration). We update our filtering function to handle the semantics of type constructors, by replacing $\text{toDumVar}(\mu) = \delta_{\text{dum}}$ by:

$$\text{toDumVar}(\theta) = \phi_{\text{dum}}$$

Slicing. Because we have changed our constraint generation rule for dn 's, we need to replace the dot terms in DatName as follows:

$$\text{dot-e}(\vec{pt}) \xrightarrow{\text{DatName}} \text{dot-n}(\vec{pt})$$

We define the new constraint generation rule for terms of the form $\text{dot-n}(\vec{pt})$ as follows:

$$\begin{aligned} \text{dot-n}(\langle pt_1, \dots, pt_n \rangle) &\triangleright \langle \alpha, \alpha', [e_1; \dots; e_n], \square \rangle \Leftarrow \\ pt_1 \triangleright e_1 \wedge \dots \wedge pt_n \triangleright e_n \wedge \text{dja}(e_1, \dots, e_n, \alpha) \end{aligned}$$

Finally, our slicing algorithm does not need to be extended but we need to update the tree syntax for programs as follows:

$$\begin{aligned} \text{Prod} &::= \dots \mid \text{decTyp} \\ \text{Dot} &::= \dots \mid \text{dotN} \end{aligned}$$

We also need to modify the getDot function that associates dot markers to node kinds as follows (the function now returns a dotN marker and not a dotE marker anymore when applied to a datname node):

$$\text{getDot}(\langle \text{datname}, \text{prod} \rangle) = \text{dotN}$$

We also need to extend the toTree function that associates trees of the form $tree$ to terms of the form $term$ as follows:

$$\begin{aligned} \text{toTree}(\text{type}\ dn \stackrel{l}{=} ty) &= \langle \langle \text{dec}, \text{decTyp} \rangle, l, \langle \text{toTree}(dn), \text{toTree}(ty) \rangle \rangle \\ \text{toTree}(\text{dot-n}(\langle pt_1, \dots, pt_n \rangle)) &= \langle \text{dotN}, \langle \text{toTree}(pt_1), \dots, \text{toTree}(pt_n) \rangle \rangle \end{aligned}$$

C.3 Signatures

This section shows how to design a type error slicer that handles some signature related features. This section deals with value, type, datatype and structure specifications. It does not deal with include or sharing specifications, and does not deal with type realisations (where clauses) either. Type realisations and include specifications are “almost fully” supported by our implementation. We also partially support sharing specifications in our implementation.

Some kinds of errors are not handled by the system presented in this section. For example we do not handle unmatched errors: when an identifier is specified in a signature but not declared in a structure constrained by the signature. These errors will be dealt with in Sec. C.4. Another kind of error which is not dealt with in this section is when a type constructor is defined as a type function in a structure and as a datatype in the structure's signature. Even though this kind of error is handled by our implementation but we have not yet written the details.

External syntax. First, let us extend our external syntax with signatures as follows:

$sigid \in \text{SigId}$	(signature identifiers)
$sigdec \in \text{SigDec}$	$::= \text{signature } sigid \stackrel{l}{=} sigexp$ $\mid \text{dot-d}(\vec{pt})$
$sigexp \in \text{SigExp}$	$::= sigid^l \mid sig^l\ spec_1 \dots spec_n \text{ end}$ $\mid \text{dot-s}(\vec{pt})$
$spec \in \text{Spec}$	$::= \text{val } vid :^l ty$ $\mid \text{type } dn^l$ $\mid \text{datatype } dn \stackrel{l}{=} cd$ $\mid \text{structure } sid :^l sigexp$ $\mid \text{dot-d}(\vec{pt})$
$cd \in \text{ConDesc}$	$::= vid_c^l \mid vid \text{ of }^l ty$ $\mid \text{dot-e}(\vec{pt})$
$id \in \text{Id}$	$::= \dots \mid sigid$
$sexp \in \text{StrExp}$	$::= \dots \mid sexp :^l sigexp \mid sexp :>^l sigexp$
$topdec \in \text{TopDec}$	$::= sdec \mid sigdec$
$prog \in \text{Program}$	$::= \text{topdec}_1; \dots; \text{topdec}_n, \text{ where } n \geq 0$

The symbol $:>$ is used for opaque constraints and $:$ for translucent constraints. The structure $sexp :>^l sigexp$ is the structure $sexp$ constrained by the signature $sigexp$ where each of $sigexp$'s specifications has to be matched by one of $sexp$'s declarations (and similarly for $sexp :^l sigexp$). The structure $sexp$ can declare more identifiers than are specified in $sigexp$. In the structure $sexp :>^l sigexp$, only the identifiers specified in $sigexp$ can be

Figure 17 Constraint solving rules for type functions

$$\begin{array}{l}
\text{(S9)} \quad \text{solve}(\langle u, e \rangle, \vec{d}, \theta_1 \cdot \tau_1 = \theta_2 \cdot \tau_2) \rightarrow \text{solve}(\langle u, e \rangle, \vec{d}, \text{build}(u, \theta_1 \cdot \tau_1) = \text{build}(u, \theta_2 \cdot \tau_2)) \\
\text{(S10)} \quad \text{solve}(\langle u, e \rangle, \vec{d}, \tau_1 = \tau_2) \rightarrow \text{solve}(\langle u, e \rangle, \vec{d}, \tau = \text{build}(u, \theta \cdot \tau')) \\
\quad \text{if } \{\tau_1, \tau_2\} = \{\theta \cdot \tau', \tau\} \text{ and } \tau \text{ is of the form } \tau_3 \rightarrow \tau_4 \text{ or } \tau_3 \mu
\end{array}$$

accessed from sexp (only the sigexp part from sexp is visible to the outside world). The difference between $\text{sexp} :>^l \text{sigexp}$ and $\text{sexp} :^l \text{sigexp}$ is that in the first one if sigexp specifies a type constructor tc then in $\text{sexp} :>^l \text{sigexp}$ it is not constrained by its declaration in sexp , whereas in $\text{sexp} :^l \text{sigexp}$ the type constructor would be constrained by its declaration in sexp . Opaque signatures are used to abstract types from structures and are usually preferred over translucent ones for this reason.

Let us now present an example involving an opaque signature:

```
(EX1) signature s = sig val x : 'a end
      structure S = struct val x = 1 end
      structure T = S :> s
```

This piece of code is untypable because the type variable $'a$ is more general than the type int . Types of declarations in structures have to be at least as general as the corresponding specifications in signatures. This kind of error will be referred as a *too general* error henceforth.

Constraint syntax. We now extend our constraint system to handle signatures:

$$\begin{array}{l}
\text{bind} \in \text{Bind} \quad ::= \dots \mid \downarrow \text{sigid} = e \\
\text{acc} \in \text{Accessor} ::= \dots \mid \uparrow \text{sigid} = ev \\
\tau \in \text{ITy} \quad ::= \dots \mid \widehat{tv} \\
\mu \in \text{ITyCon} ::= \dots \mid \widetilde{tv} \\
e \in \text{Env} \quad ::= \dots \mid e_1 : e_2 \mid e_1 :> e_2
\end{array}$$

We also extend the form of the explicit type variable binders as follows:

$$\downarrow tv = \alpha \xrightarrow{\text{Bind}} \downarrow tv = \tau$$

We add the explicit type variables to the internal type set and extend the explicit type variable binders to internal types because we want to allow explicit type variables to bind explicit type variables and not only internal type variables. This helps catching *too general* errors as presented above.

We also add the explicit type variables to the internal type constructor set because, in order to help catch *too general* errors, we do not generalise the external type variables occurring in a value specification in a signature until we match the signature against a structure. Inside a signature, an explicit type variable is then considered as a constant type. To such a constant type we associate an internal type constructor which is the explicit type variable itself. Let us explain our reason for doing so using the following piece of code (the same as (EX1) where we replaced $'a$ by bool in x 's specification):

```
(EX2) signature s = sig val x : bool end
      structure S = struct val x = 1 end
      structure T = S :> s
```

Given this piece of code, our enumeration algorithm would find the type error that x is specified as a Boolean in s , which is the signature constraining S , and that x is declared as an integer in S . The issue is that our minimisation algorithm would eventually try to slice out the type bool in x 's specification. This would result in x having a type scheme of the form $\forall \{\alpha\}. \alpha$ in its specification. This type scheme is more general than int which is x 's type in its declaration. If we were to generalise the explicit type variables occurring in value specifications, we would also generate the type scheme $\forall \{\alpha\}. \alpha$ for x 's specification in (EX1). We then would not be able to distinguish between a type scheme which is genuinely too general (in (EX1)) and a type scheme which is too general

because some information has been discarded (in (EX2) where bool has been filtered out). In order to avoid that, explicit type variables occurring in a signature are not generalised until we match the signature against a structure.

Because we extended our internal types, we also need to extend our building function as follows:

$$\text{build}(u, \widehat{tv}) = \widehat{tv}$$

Constraint generation. Fig. 18 presents the new constraint generation rules for the syntactic forms introduced above. Rule (G24) uses the function tvBind which is defined as follows: $\text{tvBind}(ty, l) = (\downarrow tv_1 \stackrel{l}{=} \widehat{tv}_1; \dots; \downarrow tv_n \stackrel{l}{=} \widehat{tv}_n)$ such that $\{tv_1, \dots, tv_n\}$ is the set of external type variables occurring in ty and where if $i < j$ then tv_j does not occur on the left of tv_i in ty . This function is used to generate explicit type variable binders for explicit type variables occurring in value specifications such as in the specification $\text{val } f : 'a \rightarrow 'a$, for which we would generate a binder of the form $\downarrow 'a = \widehat{a}$.

Note that rules (G21), (G22) and (G23) for signature declarations and expressions are similar to rules (G16), (G17) and (G18), defined in Fig. 7, for structure declarations and expressions. Rule (G24) is a simplified version of rule (G12) for recursive value declarations (defined in Fig. 7), where the expression is replaced by an external type and where the pattern is reduced to a single value identifier. Note that even though explicit type variables occurring in a signature are not generalised until the signature is matched against a structure, rule (G24) generates a poly environment to generalise internal type variables that are unconstrained due to constraint filtering. The novelty in this rule, as described above is the necessity to bind the explicit type variables occurring in the external type. Rule (G25) is similar to rule (G20) defined in Fig. 16, but instead of binding the specified type constructor to an internal type computed from an external type, it generates a new type constructor name. Such a name might then be renamed during constraint solving when matching a signature against a structure. Rule (G27) is similar to rule (G13) defined in Fig. 16 and rule (G26) is similar to rule (G16) defined in Fig. 7. The constraint generation rules for constructor descriptions are the same as the ones for constructor declarations. Finally, rules (G28) and (G29) are the most interesting rules. They are the ones generating our new environments of the forms $e_1 : e_2$ (generated by rule (G28) for transparent signature constraints) and $e_1 :> e_2$ (generated by rule (G29) for opaque signature constraints).

Because our initial constraint generation algorithm generates new forms of constraints, we extend the lbind and eg forms as follows (see Sec. A.3):

$$\begin{array}{l}
\text{ietv} \in \text{IETyVar} ::= \alpha \mid \widehat{tv} \\
\text{lbind} \in \text{LabBind} ::= \dots \mid \downarrow \text{sigid} \stackrel{l}{=} ev \\
\text{eg} \in \text{GenEnv} ::= \dots \mid ev_1 : ev_2 \mid ev_1 :> ev_2
\end{array}$$

An ietv (used below) can either be an internal (“I”) or an external (“E”) type variable.

We also replace the initially generated external type variable binders as follows:

$$\downarrow tv \stackrel{l}{=} \alpha \xrightarrow{\text{LabBind}} \downarrow tv \stackrel{l}{=} \text{ietv}$$

Constraint solving. Let us extend unification states and error kinds as follows:

Figure 18 Constraint generation rules for signatures

Signature declarations	(G21) signature $sigid \stackrel{l}{=} sigexp \rightarrow ev' = (e; \downarrow sigid \stackrel{l}{=} ev); ev'^l$	$\Leftarrow sigexp \rightarrow \langle ev, e \rangle \wedge dja(e, ev')$
Signature expressions	(G22) $sigid^l \rightarrow \langle ev, \uparrow sigid \stackrel{l}{=} ev \rangle$	
	(G23) $sig^l spec_1 \dots spec_n \text{ end} \rightarrow \langle ev, (ev \stackrel{l}{=} ev'); (ev' = (e_1; \dots; e_n)) \rangle$ $\Leftarrow spec_1 \rightarrow e_1 \wedge \dots \wedge spec_n \rightarrow e_n \wedge dja(e_1, \dots, e_n, ev, ev')$	
Specifications	(G24) $val \text{ vid} :^l ty \rightarrow (ev = \text{poly}(\text{loc } tv \text{ Bind}(ty, l) \text{ in } (e; \downarrow vid \stackrel{l}{=} \langle \alpha, v \rangle))); ev'^l$	$\Leftarrow ty \rightarrow \langle \alpha, e \rangle \wedge dja(e, ev)$
	(G25) $type \text{ dn}^l \rightarrow (ev = ((\alpha \stackrel{l}{=} \alpha' \gamma); e)); ev'^l$	$\Leftarrow dn \rightarrow \langle \alpha, \alpha', e, e' \rangle \wedge dja(e, e', ev, \gamma)$
	(G26) $structure \text{ sid} :^l sigexp \rightarrow (ev' = (e; \downarrow sid \stackrel{l}{=} ev)); ev'^l$	$\Leftarrow sigexp \rightarrow \langle ev, e \rangle \wedge dja(e, ev')$
	(G27) $datatype \text{ dn} \stackrel{l}{=} cd \rightarrow (ev = ((\alpha_1 \stackrel{l}{=} \alpha'_1 \gamma); (\alpha_1 \stackrel{l}{=} \alpha_2); e_1; \text{loc } e'_1 \text{ in } \text{poly}(e_2))); ev'^l$ $\Leftarrow dn \rightarrow \langle \alpha_1, \alpha'_1, e_1, e'_1 \rangle \wedge cd \rightarrow \langle \alpha_2, e_2 \rangle \wedge dja(e_1, e_2, \gamma, ev)$	
Structure expressions	(G28) $sexp :^l sigexp \rightarrow \langle ev, e_2; e_1; (ev \stackrel{l}{=} ev_1; ev_2) \rangle$	$\Leftarrow sexp \rightarrow \langle ev_1, e_1 \rangle \wedge sigexp \rightarrow \langle ev_2, e_2 \rangle \wedge dja(e_1, e_2, ev)$
	(G29) $sexp :>^l sigexp \rightarrow \langle ev, e_2; e_1; (ev \stackrel{l}{=} ev_1; >ev_2) \rangle$	$\Leftarrow sexp \rightarrow \langle ev_1, e_1 \rangle \wedge sigexp \rightarrow \langle ev_2, e_2 \rangle \wedge dja(e_1, e_2, ev)$
Programs	(G30) $topdec_1; \dots; topdec_n \rightarrow e_1; \dots; e_n$	$\Leftarrow topdec_1 \rightarrow e_1 \wedge \dots \wedge topdec_n \rightarrow e_n \wedge dja(e_1, \dots, e_n, ev)$

```

tfm ∈ TypFunMap = TyConName → TypFun
state ∈ State ::= ...
| prematch(Δ, d̄, e₁, e₂)
| match(Δ, d̄, tfm, e₁, e₂)
| succ(Δ, tfm)
ek ∈ ErrKind ::= ...
| TyVarClash(tv₁, tv₂)
| TooGeneral(μ₁, μ₂)

```

Type functions of the form tfm are used to gather the type functions defined in a structure, to then apply them to the types generated for a signature constraint. Roughly speaking, when solving an environment of the form $e_1 : e_2$ or of the form $e_1 : > e_2$, the type functions defined in e_1 (related to a structure) are gathered and applied to e_2 (related to a signature) using the appTFM function defined below. This step is required because our initial constraint generation algorithm might generate different type constructor names for two type constructors that might turn out to be the same type constructor. For example, in the following (typable) piece of code

```

signature s = sig datatype 'a t = T end
structure S = struct datatype 'a t = T end : s

```

Our initial constraint generation algorithm will generate two distinct type constructor names for the two occurrence of t . But, when checking that the signature s , matches the structure S , these two type constructor names have to be equated. This is done by extracting the one defined in the structure and by then renaming the one from the signature into the one from the structure.

Error kinds of the form $\text{TooGeneral}(\mu_1, \mu_2)$ are for type errors as the one described above (*too general* errors), where a signature constrains a structure and is more general than the structure. Error kinds of the form $\text{TyVarClash}(tv_1, tv_2)$ are for type errors such that the one in the following piece of code:

```

signature s = sig val f : 'a -> 'b end
structure S = struct val rec f = fn x => x end
structure T = S :> s

```

In this piece of code, f is specified in the signature s as a function where its argument's type can differ from its body's type. In the structure S , the function f is declared as the identity function and so its argument's type has to be the same as its body's type. Finally S is constrained by s . Therefore, we report an explicit type variable clash between $'a$ and $'b$. This is a special kind of *too general* errors.

Rules (SM4) and (SM5) of the extension of our constraint solver defined below in Fig. 20, make use of the function appTFM that applies type functions (extracted from a structure) to a type or an environment (related to a signature) and which is defined as follows:

$$\text{appTFM}(\tau \mu, tfm) = \begin{cases} \tau'[\{\alpha \mapsto \text{appTFM}(\tau, tfm)\}], & \text{if } tfm(\mu) = \Lambda \alpha. \tau' \\ \alpha_{\text{dum}}, \text{ if } tfm(\mu) \in \text{Var} \\ \text{undefined, if } \mu \text{ is of the form } \mu^{\bar{d} \cup d} \\ \text{appTFM}(\tau, tfm) \mu, \text{ otherwise} \end{cases}$$

$$\begin{aligned} \text{appTFM}(\tau_1 \rightarrow \tau_2, tfm) &= \text{appTFM}(\tau_1, tfm) \rightarrow \text{appTFM}(\tau_2, tfm) \\ \text{appTFM}(\forall \alpha. \tau, tfm) &= \forall \alpha. \text{appTFM}(\tau, tfm) \\ \text{appTFM}(\Lambda \alpha. \tau, tfm) &= \Lambda \alpha. \text{appTFM}(\tau, tfm) \\ \text{appTFM}(e_1; e_2, tfm) &= \text{appTFM}(e_1, tfm); \text{appTFM}(e_2, tfm) \\ \text{appTFM}(\downarrow id = x, tfm) &= \downarrow id = \text{appTFM}(x, tfm) \\ \text{appTFM}(x^{\bar{d}}, tfm) &= \text{appTFM}(x, tfm)^{\bar{d}} \\ \text{appTFM}(x, tfm) &= x, \text{ if none of the above applies} \end{aligned}$$

Let us define tyvars and nonDumVars that are used by some functions and predicates defined below. The function tyvars is defined as follows: $\text{tyvars}(x)$ is the set of syntactic forms belonging to TyVar and occurring in x whatever x is. The function nonDumVars is defined as follows: $\text{nonDumVars}(x) = \text{vars}(x) \setminus \text{DumVar}$.

Rule (SC2) of the extension of our constraint solver defined below in Fig. 20 uses the predicate abstract which is used to rename the type constructor names declared in an environment and defined as follows:

$$\langle \Delta, e, \{\gamma_1\} \uplus \dots \uplus \{\gamma_n\} \rangle \xrightarrow{\text{abstract}} \text{appTFM}(e, tfm)$$

if $dja(\text{nonDumVars}(\Delta), \gamma'_1, \dots, \gamma'_n, \alpha_1, \dots, \alpha_n)$
and $tfm = \cup_{i \in \{1, \dots, n\}} \{\gamma_i \mapsto \Lambda \alpha_i. \alpha_i \gamma'_i\}$

Fig. 19 defines our algorithm genExTyVar that generalises explicit type variables in environments. It also uses renamings of the form renTv defined as follows:

$$\text{renTv} \in \text{RenTv} ::= \{ \text{renTv} \in \text{TyVar} \rightarrow \text{ITyVar} \mid \text{renTv is injective} \wedge dja(\text{dom}(\text{renTv}), \text{ran}(\text{renTv})) \}$$

These renamings are applied using the renTv function which is defined as follows:

$$\begin{aligned} \text{renTv}(\alpha, \text{renTv}) &= \alpha \\ \text{renTv}(\widehat{tv}, \text{renTv}) &= \text{renTv}(tv) \\ \text{renTv}(\tau \mu, \text{renTv}) &= \text{renTv}(\tau, \text{renTv}) \mu \\ \text{renTv}(\tau_1 \rightarrow \tau_2, \text{renTv}) &= \text{renTv}(\tau_1, \text{renTv}) \rightarrow \text{renTv}(\tau_2, \text{renTv}) \\ \text{renTv}(\tau^{\bar{d}}, \text{renTv}) &= \text{renTv}(\tau, \text{renTv})^{\bar{d}} \end{aligned}$$

This function is partially defined. It is not defined on types of the form $\theta \cdot \tau$ because these forms cannot occur in environments in “solved” forms. Moreover when applying renTv to an internal type and a renaming of explicit type variables, then the explicit type variables occurring in the internal type have to be the domain of the renaming. This is always the case when calling renTv in Fig. 19.

Fig. 20 extends our constraint solver to deal with our new constraint terms.

Figure 19 Generalisation of explicit type variables

$\langle \Delta, e_1; e_2 \rangle \xrightarrow{\text{genExTyVar}} e'_2; e'_2$	$\Leftrightarrow \langle \Delta, e_1 \rangle \xrightarrow{\text{genExTyVar}} e'_1$ and $\langle \Delta; e'_1, e_2 \rangle \xrightarrow{\text{genExTyVar}} e'_2$
$\langle \Delta, \downarrow id = e \rangle \xrightarrow{\text{genExTyVar}} \downarrow id = e'$	$\Leftrightarrow \langle \Delta, e \rangle \xrightarrow{\text{genExTyVar}} e'$
$\langle \Delta, \downarrow vid = \tau \rangle \xrightarrow{\text{genExTyVar}} \downarrow vid = \forall \bar{\alpha}_0. \text{renTv}(\tau, \text{rentv})$	$\Leftrightarrow \text{tyvars}(\tau) = \text{dom}(\text{rentv})$ and $\bar{\alpha}_0 = \text{ran}(\text{rentv})$ and $\text{dja}(\text{nonDumVars}(\Delta), \bar{\alpha}_0)$
$\langle \Delta, \downarrow vid = \forall \bar{\alpha}. \tau \rangle \xrightarrow{\text{genExTyVar}} \downarrow vid = \forall (\bar{\alpha} \cup \bar{\alpha}_0). \text{renTv}(\tau, \text{rentv})$	$\Leftrightarrow \text{tyvars}(\tau) = \text{dom}(\text{rentv})$ and $\bar{\alpha}_0 = \text{ran}(\text{rentv})$ and $\text{dja}(\text{nonDumVars}(\Delta), \bar{\alpha}_0)$
$\langle \Delta, e \bar{d} \rangle \xrightarrow{\text{genExTyVar}} e' \bar{d}$	$\Leftrightarrow \langle \Delta, e \rangle \xrightarrow{\text{genExTyVar}} e'$
$\langle \Delta, e \rangle \xrightarrow{\text{genExTyVar}} e,$	if none of the above applies

Rule (S10), originally defined in Sec. C.2, is updated to handle explicit type variables as internal types.

Rules (S11)-(S13) are to handle our new cases of internal types and internal type constructors.

Rule (B1) is redefined so that it builds up the semantics of binders. The reason is that when checking if a signature matches a structure, we want the corresponding environments fully built up. Building the binders when solving them allows having a shallow building function, as the one we currently have, that does not need to go down structure or signature binders.

Rules (SM1)-(SM13) check whether a signature matches a structure. These rules are used for both translucent and opaque constraints. If $\text{match}(\Delta, \bar{d}, \text{tfm}, e_1, e_2) \rightarrow^* \text{match}(\Delta', \bar{d}', \text{tfm}', e'_1, e'_2)$ using rules (SM1)-(SM13) then $e_1 = e'_1$.

Rule (SM4) checks that a type scheme from the signature does not allow generating types that the corresponding type scheme in the structure cannot generate. Before checking that, the gathered type functions are applied to the type extracted from the binding coming from the signature (τ_1). This is where the instantiation of a signature is performed in our system. Finally, explicit type variables occurring in the generated binder are generalised. The generated type scheme is built from τ_1 and not from τ_2 in case the binding from the signature is a dummy binding. If the binding from the signature was a dummy binding and the corresponding binding from the structure was not a dummy binding, then we do not want to generate an unlabelled non-dummy binding that could therefore lead to a type error because this type error might then not be dependent on the label associated to the specification for which the binder has been generated (the specification would not be part of the reported error).

Rules (SM6) and (SM11) gather type functions defined by the structure. These type functions are applied to the signature during the process of checking the matching (in rules (SM4) and (SM5)). Once again, extra care has to be taken when the binder from the signature is a dummy binder, so that the algorithm does not generate a non-dummy binder.

The other (SM i) rules are fairly straightforward.

Rule (SC5) is just so that the same mechanism can be used for opaque and translucent signatures.

Rule (SC1) for translucent signature just checks that the signature matches the structure. It does not need further computation because the resulting structure is computed while checking the matching.

Rule (SC3) for opaque signatures checks that the signature matches the structure and generates a new structure based on the signature. The generated structure is the signature where the internal type constructor names are renamed and where the explicit type variables are generalised.

Constraint filtering (Minimisation and enumeration). We extend our filtering function as follows:

$$\begin{aligned} \text{filt}(e_1; e_2, \bar{l}_1, \bar{l}_2) &= \text{filt}(e_1, \bar{l}_1, \bar{l}_2) : \text{filt}(e_1, \bar{l}_1, \bar{l}_2) \\ \text{filt}(e_1 :> e_2, \bar{l}_1, \bar{l}_2) &= \text{filt}(e_1, \bar{l}_1, \bar{l}_2) :> \text{filt}(e_1, \bar{l}_1, \bar{l}_2) \\ \text{filt}(ev, \bar{l}_1, \bar{l}_2) &= ev \end{aligned}$$

We now need the filtering of unlabelled environment variables because we now allow unlabelled environment variables to occur within environments of the form $e_1; e_2$ or $e_1 :> e_2$.

Because explicit type variables can now bind internal types and not only internal type variables, we also need to update our filtering function by replacing $\text{toDumVar}(\alpha) = \alpha_{\text{dum}}$ by:

$$\text{toDumVar}(\tau) = \alpha_{\text{dum}}$$

Slicing. We extend our tree syntax for programs as follows:

```

Class ::= ... | sigdec | sigexp | spec
Prod ::= ...
        | sigdecDec
        | sigexpSig
        | specVal | specTyp | specDat | specStr
        | strexpTr | strexpOp

```

We also extend our function `getDot` that associates dot markers to node kinds as follows:

```

getDot((sigdec, prod)) = dotD
getDot((sigexp, prod)) = dotS
getDot((spec, prod))   = dotD

```

Finally, Fig. 21 extends our function `toTree` that transforms a term *term* into a tree *tree*.

C.4 Reporting unmatched errors

There is a kind of error involving signatures that is not handled by the constraint solver as defined above: the “unmatched” errors.

For example, in

```

signature s = sig val fool : int end
structure S = struct val foo = 1 val bar = 2 end
structure T = S :> s

```

the specification `fool` from the signature `s` is not matched in the structure `S`, but `s` constrains `S` in `T`. This error could be solved in many ways, such as: (1) one could replace `fool` by `foo` in `s`, (2) one could replace `foo` by `foo1` in `S`, (3) one could constrain `S` using a different signature, (4) one could bind `s` or `S` to other expressions.

For this error we would like to report that `foo1` specified in `s` is not any of `foo` or `bar` declared in `S`, but `s` constrains `S`.

For that we need to be able to check that indeed `foo1` is not any of the declarations of `S`.

With the system as described above, we cannot report such errors because we do not have any way of knowing whether an environment is constituted by the binders corresponding to *all* the declarations of a structure. As a matter of fact, this is not possible with the current system because of the way constraint filtering can replace environment variables and binders by \square .

We will now show how to extend our system to report such errors.

Constraint syntax. We extend our environment with a new empty environment as follows:

$$\text{Env} ::= \dots \mid \odot$$

The meaning of the environment \odot lies in between the meaning of \square and the meaning of environment variables.

Figure 20 Constraint solving for signature related constraints

Some kinds of errors are not handled by the system presented in this section, although our implementation handles them. For more information please refer to the introductory paragraph of this section (Sec. C.3).

equality simplification

- (S10) $\text{solve}(\langle u, e \rangle, \bar{d}, \tau_1 = \tau_2) \rightarrow \text{solve}(\langle u, e \rangle, \bar{d}, \tau = \text{build}(u, \theta \cdot \tau'))$ if $\{\tau_1, \tau_2\} = \{\theta \cdot \tau', \tau\}$ and τ is of the form $\tau_3 \rightarrow \tau_4, \tau_3 \mu$, or \widehat{tv}
(S11) $\text{solve}(\Delta, \bar{d}, \widehat{tv}_1 = \widehat{tv}_2) \rightarrow \text{err}(\langle \text{TyVarClash}(tv_1, tv_2), \bar{d} \rangle)$, if $tv_1 \neq tv_2$
(S12) $\text{solve}(\Delta, \bar{d}, \tau_1 = \tau_2) \rightarrow \text{solve}(\Delta, \bar{d}, \widehat{tv} = \mu)$, if $(\{\tau_1, \tau_2\} = \{\widehat{tv}, \tau \rightarrow \tau'\} \text{ and } \mu = \text{ar})$
or $\{\tau_1, \tau_2\} = \{\widehat{tv}, \tau \mu\}$, for some tv, τ, τ', μ
(S13) $\text{solve}(\Delta, \bar{d}, \mu_1 = \mu_2) \rightarrow \text{err}(\langle \text{TooGeneral}(\mu_1, \mu_2), \bar{d} \rangle)$, if $\{\mu_1, \mu_2\} \in \{\{tv, \text{ar}\}, \{\widehat{tv}, \gamma\}\}$, for some tv and γ

binders

- (B1) $\text{solve}(\langle u, e \rangle, \bar{d}, \downarrow id = x) \rightarrow \text{succ}(\langle u, e \rangle; (\downarrow id \stackrel{\bar{d}}{=} \text{build}(u, x)))$

signature constraints

- (SC1) $\text{solve}(\langle u, e \rangle, \bar{d}, e_1 : e_2) \rightarrow \text{succ}(\Delta)$, if $\text{prematch}(\langle u, e \rangle, \bar{d}, e_1, e_2) \rightarrow^* \text{succ}(\Delta, tfm)$,
(SC2) $\text{solve}(\langle u, e \rangle, \bar{d}, e_1 : e_2) \rightarrow \text{err}(er)$, if $\text{prematch}(\langle u, e \rangle, \bar{d}, e_1, e_2) \rightarrow^* \text{err}(er)$,
(SC3) $\text{solve}(\langle u, e \rangle, \bar{d}, e_1 : > e_2) \rightarrow \text{succ}(\langle u', e; e_2' \rangle)$, if $\text{prematch}(\langle u, e \rangle, \bar{d}, e_1, e_2) \rightarrow^* \text{succ}(\langle u', e' \rangle, tfm)$
and $\langle \langle u', e' \rangle, \text{build}(u, e_2), \text{dom}(tfm) \stackrel{\text{abstract}}{\rightarrow} e_2' \text{ and } \langle \langle u', e' \rangle, e_2' \rangle \stackrel{\text{genExTyVar}}{\rightarrow} e_2''$
(SC4) $\text{solve}(\langle u, e \rangle, \bar{d}, e_1 : > e_2) \rightarrow \text{err}(er)$, if $\text{prematch}(\langle u, e \rangle, \bar{d}, e_1, e_2) \rightarrow^* \text{err}(er)$
(SC5) $\text{prematch}(\langle u, e \rangle, \bar{d}, e_1, e_2) \rightarrow \text{state}$, if $\text{match}(\langle u, e \rangle, \bar{d}, \emptyset, \text{build}(u, e_1), \text{build}(u, e_2)) \rightarrow^* \text{state}$

structure/signature matching

- (SM1) $\text{match}(\Delta, \bar{d}, tfm, e, \square) \rightarrow \text{succ}(\Delta, tfm)$
(SM2) $\text{match}(\Delta, \bar{d}, tfm, e, e_1; e_2) \rightarrow \text{match}(\Delta', \bar{d}, tfm', e, e_2)$,
if $\text{match}(\Delta, \bar{d}, tfm, e, e_1) \rightarrow^* \text{succ}(\Delta', tfm')$
(SM3) $\text{match}(\Delta, \bar{d}, tfm, e, e_1; e_2) \rightarrow \text{err}(er)$,
if $\text{match}(\Delta, \bar{d}, tfm, e, e_1) \rightarrow^* \text{err}(er)$
(SM4) $\text{match}(\Delta, \bar{d}, tfm, e, \downarrow vid = \sigma_1) \rightarrow \text{succ}(\Delta; e_0, tfm)$,
if $e(\text{vid}) = \sigma_2$ and $\forall i \in \{1, 2\}. (\sigma_i = \forall \bar{\alpha}_i. \tau_i \text{ or } (\sigma_i = \tau_i \text{ and } \bar{\alpha}_i = \emptyset))$
and $\tau'_1 = \text{appTFM}(\tau_1, tfm)$ and $\text{solve}(\Delta, \bar{d}, \tau'_1 = \tau_2) \rightarrow^* \text{succ}(\langle u', e' \rangle)$
and $\tau = \text{build}(u', \tau'_1)$ and $\langle \langle u', e' \rangle, \downarrow vid \stackrel{\bar{d}}{=} \forall (\bar{\alpha}_1 \cup \bar{\alpha}_2) \cap \text{vars}(\tau). \tau \rangle \stackrel{\text{genExTyVar}}{\rightarrow} e_0$
(SM5) $\text{match}(\Delta, \bar{d}, tfm, e, \downarrow vid = \sigma_1) \rightarrow \text{err}(er)$,
if $e(\text{vid}) = \sigma_2$ and $\forall i \in \{1, 2\}. (\sigma_i = \forall \bar{\alpha}_i. \tau_i \text{ or } (\sigma_i = \tau_i \text{ and } \bar{\alpha}_i = \emptyset))$
and $\text{solve}(\Delta, \bar{d}, \text{appTFM}(\tau_1, tfm) = \tau_2) \rightarrow^* \text{err}(er)$
(SM6) $\text{match}(\Delta, \bar{d}, tfm, e, \downarrow tc = \theta_1) \rightarrow \text{succ}(\Delta; (\downarrow tc \stackrel{\bar{d}}{=} \theta'_2), tfm')$,
if $e(tc) = \theta_2$ and (if $\theta_1 \in \text{DumVar}$ then $\theta'_2 = \theta_1$ else $\theta'_2 = \theta_2^{\text{deps}(\theta_1)}$)
and (if $\theta_1 = \Lambda \alpha. (\alpha \gamma)^{\bar{d}}$ then $tfm' = tfm \boxplus \{\gamma \mapsto \theta'_2\}$ else $tfm' = tfm$)
(SM7) $\text{match}(\langle u_1, e_1 \rangle, \bar{d}, tfm, e, \downarrow sid = e_0) \rightarrow \text{succ}(\langle u_2, e_1; (\downarrow sid \stackrel{\bar{d}}{=} \text{diff}(e_1, e_2)) \rangle, tfm')$,
if $e(\text{sid}) = e'_0$ and $\text{match}(\langle u_1, e_1 \rangle, \bar{d}, tfm, e'_0, e_0) \rightarrow^* \text{succ}(\langle u_2, e_2 \rangle, tfm')$
(SM8) $\text{match}(\Delta, \bar{d}, tfm, e, \downarrow sid = e_0) \rightarrow \text{err}(er)$,
if $e(\text{sid}) = e'_0$ and $\text{match}(\Delta, \bar{d}, tfm, e'_0, e_0) \rightarrow^* \text{err}(er)$
(SM9) $\text{match}(\Delta, \bar{d}, tfm, e, \downarrow vid = is_1) \rightarrow \text{succ}(\Delta; (\downarrow vid \stackrel{\bar{d}}{=} vid), tfm)$,
if $e[\text{vid}] = is_2$ and $\text{deps}(is_2) = \bar{d}$ and $(\text{solve}(\Delta, \bar{d}, is_1 = is_2) \rightarrow^* \text{succ}(\Delta'))$ or $\text{strip}(is_1) = v$
(SM10) $\text{match}(\Delta, \bar{d}, tfm, e, \downarrow vid = is_1) \rightarrow \text{err}(er)$,
if $e[\text{vid}] = is_2$ and $\text{strip}(is_1) \neq v$ and $\text{solve}(\Delta, \bar{d}, is_1 = is_2) \rightarrow^* \text{err}(er)$
(SM11) $\text{match}(\Delta, \bar{d}, tfm, e, \downarrow id = x) \rightarrow \text{succ}(\Delta; (\downarrow id = \text{toDumVar}(x)), tfm')$,
if $e(id)$ is undefined and (if $x = \Lambda \alpha. (\alpha \gamma)^{\bar{d}}$ then $tfm' = tfm \boxplus \{\gamma \mapsto \alpha_{\text{dum}}\}$ else $tfm' = tfm$)
(SM12) $\text{match}(\Delta, \bar{d}, tfm, e, ev) \rightarrow \text{succ}(\Delta; ev, tfm)$
(SM13) $\text{match}(\Delta, \bar{d}, tfm, e, e'^{\bar{d}'}) \rightarrow \text{match}(\Delta, \bar{d} \cup \bar{d}', tfm, e, e')$

Figure 21 Extension of our slicing algorithm with signatures

Signature declarations	$\text{toTree}(\text{signature } \text{sigid} \stackrel{l}{=} \text{sigexp}) = \langle \langle \text{sigdec}, \text{sigdecDec} \rangle, l, \langle \text{sigid}, \text{toTree}(\text{sigexp}) \rangle \rangle$
Signature expressions	$\text{toTree}(\text{sigid}^l) = \langle \langle \text{sigexp}, \text{sigexpId} \rangle, l, \langle \text{sigid} \rangle \rangle$ $\text{toTree}(\text{sig}^l \text{ spec}_1 \cdots \text{spec}_n \text{ end}) = \langle \langle \text{sigexp}, \text{sigexpSig} \rangle, l, \langle \text{toTree}(\text{spec}_1), \dots, \text{toTree}(\text{spec}_n) \rangle \rangle$
Specifications	$\text{toTree}(\text{val } \text{vid} :^l \text{ty}) = \langle \langle \text{spec}, \text{specVal} \rangle, l, \langle \text{vid}, \text{toTree}(\text{ty}) \rangle \rangle$ $\text{toTree}(\text{type } \text{dn}^l) = \langle \langle \text{spec}, \text{specTyp} \rangle, l, \langle \text{toTree}(\text{dn}) \rangle \rangle$ $\text{toTree}(\text{datatype } \text{dn} \stackrel{l}{=} \text{cd}) = \langle \langle \text{spec}, \text{specDat} \rangle, l, \langle \text{toTree}(\text{dn}), \text{toTree}(\text{cd}) \rangle \rangle$ $\text{toTree}(\text{structure } \text{sid} :^l \text{sigexp}) = \langle \langle \text{spec}, \text{specStr} \rangle, l, \langle \text{sid}, \text{toTree}(\text{sigexp}) \rangle \rangle$
Structure expressions	$\text{toTree}(\text{sexp} :^l \text{sigexp}) = \langle \langle \text{strexpr}, \text{strexprTr} \rangle, l, \langle \text{toTree}(\text{sexp}), \text{toTree}(\text{sigexp}) \rangle \rangle$ $\text{toTree}(\text{sexp} :>^l \text{sigexp}) = \langle \langle \text{strexpr}, \text{strexprOp} \rangle, l, \langle \text{toTree}(\text{sexp}), \text{toTree}(\text{sigexp}) \rangle \rangle$
Programs	$\text{toTree}(\text{topdec}_1 ; \cdots ; \text{topdec}_n) = \langle \text{dotD}, \langle \text{toTree}(\text{topdec}_1), \dots, \text{toTree}(\text{topdec}_n) \rangle \rangle$

The difference between \boxplus and \odot is that the second one will be used to indicate that we filtered out an environment which has the potential to bind (either an environment variable or a binder) and not just, say, an equality constraint.

The difference between \odot and an environment variable is that in an environment of the form $\odot; e$, \odot does not shadow e .

Constraint solving. The extra environment \odot will be allowed to exist within unification contexts. Given a unification state, if Δ occurs in it, then Δ is of the form $\langle u, e \rangle$ where $e = \boxplus; e_1; \cdots; e_n$ where each e_i can either be an environment variable, a labelled binder or \odot .

Because \odot can occur in unification contexts, we extend our environment application functions as follows:

$$\begin{aligned} (e;\odot)(id) &= e(id) \\ (e;\odot)[id] &= e[id] \end{aligned}$$

Let us extend error kinds as follows:

$$ek \in \text{ErrKind} ::= \dots \mid \text{unmatched}(id, \overline{id})$$

Fig. 22 extends our constraint solver with rules to handle unmatched errors: rule (SM11) replaces the previous rule (SM11) from Fig. 20 and rules (SM14) and (N2) are new.

Rules (SM11) and (SM14) make use of the predicate *complete* (similar to *hiding*) which is defined as follows:

$$\text{complete}(e) \Leftrightarrow \begin{cases} (e \text{ of the form } \downarrow id \stackrel{\overline{a}}{=} x \\ \text{and } x \notin \text{DumVar} \cup \{a\}) \\ \text{or } (e \text{ of the form } e_1; e_2 \\ \text{and } \forall i \in \{1, 2\}. \text{complete}(e_i)) \\ \text{or } e = \boxplus \end{cases}$$

A “solved” environment (occurring in a unification context) is said to be *complete* if it is not composed by an environment variable, a filtered binder or a dummy binder.

Rule (SM14) makes use of the function *getBinders* which gathers the identifiers bound in its argument:

$$\begin{aligned} \text{getBinders}(e_1; e_2) &= \text{getBinders}(e_1) \cup \text{getBinders}(e_2) \\ \text{getBinders}(\boxplus) &= \emptyset \\ \text{getBinders}(\downarrow id \stackrel{\overline{a}}{=} x) &= \{id\} \end{aligned}$$

Constraint filtering (Minimisation and enumeration). We add a new rule to filter \odot and update the filtering of labelled environment as follows:

$$\text{filt}(e^l, \overline{l}_1, \overline{l}_2) = \begin{cases} e^l, & \text{if } l \in \overline{l}_1 \setminus \overline{l}_2 \\ \text{dum}(e), & \text{if } l \in \overline{l}_2 \\ \odot, & \text{if } l \notin \overline{l}_1 \cup \overline{l}_2 \text{ and } e \in \text{Var} \cup \text{Bind} \\ \boxplus, & \text{otherwise} \end{cases}$$

$$\text{filt}(\odot, \overline{l}_1, \overline{l}_2) = \odot$$

Slicing. We also need to modify our slicing algorithm. Consider the following piece of code:

```
signature s = sig val x : int val y : bool end
structure S : s = struct val x = 1 val y = true end
structure T :> s = struct val x = 1 val y = true end
val u = let open T val z = y open S
        in fn w => (w z, w x)
end
```

where in the *fn*-expression, z is the y from T and x comes from S via the structure opening. The structures S and T have the same structure body constrained by the same signature s , but S has a translucent signature while T 's signature is opaque.

This piece of code is untypable because w has a monomorphic type and is applied to z which is the Boolean y defined in T , and it is also applied to x which is an integer defined in S .

With our current slicing algorithm, one of the type error slice we obtain would be as follows:

```
<..signature s = sig val x : <..> val y : bool end
..structure S : s = struct val x = 1 end
..structure T :> s = <..>
..<..open T..val z = y..open S..fn w => <..w z..w x..>..>
```

which is not minimal: s does not match S because y is not declared in S .

The problem comes from our tidying of declarations in structure expressions. We therefore need to update our tidying function so that it does not discard empty dot declarations:

$$\begin{aligned} \text{tidy}(\langle \rangle) &= \langle \rangle \\ \text{tidy}(\langle \langle \text{dotD}, \overrightarrow{tree}_1 \rangle, \langle \text{dotD}, \overrightarrow{tree}_2 \rangle \rangle @ \overrightarrow{tree}) &= \text{tidy}(\langle \langle \text{dotD}, \overrightarrow{tree}_1 @ \overrightarrow{tree}_2 \rangle \rangle @ \overrightarrow{tree}), \\ &\quad \text{if } \forall \text{tree} \in \text{ran}(\overrightarrow{tree}_1). \neg \text{declares}(\text{tree}) \\ \text{tidy}(\langle \text{tree} \rangle @ \overrightarrow{tree}) &= \langle \text{tree} \rangle @ \text{tidy}(\overrightarrow{tree}), \text{ if none of the above applies} \end{aligned}$$

With this new tidy function, we would then obtain a slice as follows:

```
<..signature s = sig val x : <..> val y : bool end
..structure S : s = struct <..> val x = 1 end
..structure T :> s = <..>
..<..open T..val z = y..open S..fn w => <..w z..w x..>..>
```

We also have to replace our constraint generation rule for dot declarations, in order to generate markers of discarded binders:

$$\text{dot-d}(\langle pt_1, \dots, pt_n \rangle) \triangleright [e_1; \dots; e_n]; \odot \Leftarrow pt_1 \triangleright e_1 \wedge \cdots \wedge pt_n \triangleright e_n \wedge \text{dja}(e_1, \dots, e_n)$$

However, this modification is not enough because binders are generated for *cb*'s, *pat*'s, and *dn*'s.

For example, we would like to generate a marker of discarded binder for the following declaration: `datatype 'a t = <..>`.

Figure 22 Constraint solving rules handling unmatched errors

Some kinds of errors are not handled by the system presented in this section, although our implementation handles them. For more information please refer to the introductory paragraph of Sec. C.3.

structure/signature matching

$$\begin{aligned}
\text{(SM11)} \quad \text{match}(\Delta, \bar{d}, \text{tfm}, e, \downarrow \text{id}=x) &\rightarrow \text{succ}(\Delta; (\downarrow \text{id} \stackrel{\bar{d}}{=} y), \text{tfm}'), && \text{if } e(\text{id}) \text{ is undefined and } \neg \text{complete}(e) \text{ and } y = \text{toDumVar}(x) \\
&&& \text{and (if } x = \Lambda\alpha. \alpha \gamma \text{ then } \text{tfm}' = \text{tfm} \boxplus \{\gamma \mapsto \alpha_{\text{dum}}\} \text{ else } \text{tfm}' = \text{tfm}) \\
\text{(SM14)} \quad \text{match}(\Delta, \bar{d}, \text{tfm}, e, \downarrow \text{id}=x) &\rightarrow \text{err}(\langle \text{unmatched}(\text{id}, \bar{\text{id}}, \bar{d}) \rangle), && \text{if } e(\text{id}) \text{ is undefined and } \text{complete}(e) \text{ and where } \bar{\text{id}} = \text{getBinders}(e) \\
\text{(SM15)} \quad \text{match}(\Delta, \bar{d}, \text{tfm}, e, \odot) &\rightarrow \text{succ}(\Delta; \odot, \text{tfm}) \\
\text{empty} \\
\text{(N2)} \quad \text{solve}(\Delta, \bar{d}, \odot) &\rightarrow \text{succ}(\Delta; \odot)
\end{aligned}$$

First, let us replace the dot terms for *cb*'s. We need to do so because we want to generate markers of discarded binders only for *cb* dot terms, but not for expressions and types. We replace these dot terms as follows:

$$\text{dot-}e(\bar{pt}) \xrightarrow{\text{ConBind}} \text{dot-}c(\bar{pt})$$

We redefine the constraint generation rules for the forms $\text{dot-n}(\bar{pt})$ and $\text{dot-p}(\bar{pt})$, and we introduce a new constraint generation rule for the forms $\text{dot-c}(\bar{pt})$ as follows:

$$\begin{aligned}
\text{dot-n}(\langle pt_1, \dots, pt_n \rangle) &\triangleright \langle \alpha, \alpha', [e_1; \dots; e_n]; \odot, \square \rangle \\
&\Leftarrow pt_1 \triangleright e_1 \wedge \dots \wedge pt_n \triangleright e_n \wedge \text{dja}(e_1, \dots, e_n, \alpha, \alpha') \\
\text{dot-p}(\langle pat_1, \dots, pat_n \rangle) &\triangleright \langle \alpha, e_1; \dots; e_n; \odot \rangle \\
&\Leftarrow pat_1 \triangleright e_1 \wedge \dots \wedge pat_n \triangleright e_n \wedge \text{dja}(e_1, \dots, e_n, \alpha) \\
\text{dot-c}(\langle pt_1, \dots, pt_n \rangle) &\triangleright \langle \alpha, [e_1; \dots; e_n]; \odot \rangle \\
&\Leftarrow pt_1 \triangleright e_1 \wedge \dots \wedge pt_n \triangleright e_n \wedge \text{dja}(e_1, \dots, e_n, \alpha)
\end{aligned}$$

We add a new dot marker to the set *Dot* as follows:

$$\text{Dot} ::= \dots \mid \text{dotC}$$

Finally, we extend the *toTree* function as follows:

$$\begin{aligned}
\text{toTree}(\text{dot-n}(\langle pt_1, \dots, pt_n \rangle)) &= \\
\langle \text{dotC}, \langle \text{toTree}(pt_1), \dots, \text{toTree}(pt_n) \rangle \rangle
\end{aligned}$$

C.5 Further extensions

We are currently extending the formal presentation of our TES to handle features such as functors, non-recursive value declarations, type annotations, or long identifiers. These features are already handled by our implementation and we invite the reader to try it and read its source code for more details on how the features are handled.

D. Extensions for a better error handling

D.1 Merged minimal slices

With the constraint solver as defined above, our TES would report two minimal *unmatched* type error slices for the following piece of code:

```

structure S = struct val (foo1, barr, x, y) = (1, 2, 3, 4) end
signature s = sig val foo : int val bar : int val x : int end
structure T = S :> s

```

One of the type error is that the specification *foo* in *s* is not matched in the structure *S* (that declares *foo1*, *barr*, *x* and *y*), but *s* constrains *S* in *T*. The other error is similar but concerns the specification *bar*.

This is another typical example where finding and reporting merged minimal error slices would be useful (see Sec. 2.3, that presents another example of merged minimal error slices). For the example above, instead of the two reports described above, we would prefer a highlighting that would look like:

```

structure S = struct val (foo1, barr, x, y) = (1, 2, 3) end
signature s = sig val foo : int val bar : int val x : int end
structure T = S :> s

```

This highlighting shows that *foo* and *barr* are not matched in the structure *S*, but also suppose that *x* might not be the matching for *foo* or *bar* as *x* is specified in the signature *s*. Note that *x* is still reported because we can't know if *x* in the structure *S* is definitely not the matching of, e.g., *foo* in the signature *s*.

Note that we do not want to find the two minimal error reports and then merge them into a single report, but we directly want to generate the merged error.

We could obtain this slice by altering the part of our constraint solver defined in Fig. 20 and Fig. 22.

First, we want unmatched error kinds to be as follows instead (we replace the previous form by this new one):

$$ek \in \text{ErrKind} ::= \dots \mid \text{unmatched}(\bar{\text{id}}_1, \bar{\text{id}}_2, \bar{\text{id}}_3)$$

For the highlighting presented above, the error kind would then be $\text{unmatched}(\bar{\text{id}}_1, \bar{\text{id}}_2, \bar{\text{id}}_3)$, where $\bar{\text{id}}_1$ is the set of identifiers highlighted in dark grey (the identifiers specified in *s* that are not declared in *S*), $\bar{\text{id}}_2$ is the set of identifiers highlighted with the darkest grey (the identifiers declared in *S* that are not specified in *s*) and $\bar{\text{id}}_3$ is the set of identifiers highlighted in light grey (the identifiers both specified in *s* and declared in *S*).

Then, when checking if a signature matches a structure, in order to gather (1) the identifiers that are specified in the signature but not declared in the structure, (2) the identifiers that are declared in the structure but not specified in the signature, and (3) the identifiers that are both specified in the signature and declared in the structure, we extend our “match” states as follows:

$$\begin{aligned}
\text{unm} \in \text{Unmatched} &::= \langle \bar{\text{id}}_1, \bar{\text{id}}_2 \rangle \\
\text{state} \in \text{State} &::= \dots \\
&\quad \left| \begin{array}{l} \text{match}(\Delta, \bar{d}, \text{tfm}, \text{unm}, e_1, e_2) \\ \text{succ}(\Delta, \text{tfm}, \text{unm}) \end{array} \right.
\end{aligned}$$

Finally, Fig. 23 updates the rules defined in Fig. 20 and Fig. 22 to handle the reporting of merged unmatched errors. Rules (SC1), (SC2), (SC3) and (SC4) are as before and are not repeated here. Rules (SC6) and (SM16) are new and replace rule (SM14).

The main difference between this new algorithm and the one presented in Fig. 20 and Fig. 22, is that our new algorithm gathers the identifiers that are both specified in the signature and declared in the structure (rules (SM4), (SM6), and (SM7)) and also gathers the identifier that are not matched in the structure (rule (SM11)). If there exists such an identifier, it means that there is an unmatched error. We then wait to check the matching of the entire signature against the structure to finally report all such unmatched identifiers in a single error report (rules (SC6) and (SM16)).

Note that such type error reports (for unmatched errors) are still imperfect. For example, the highlighting above does not show that $\{\text{foo1}, \text{barr}, \text{x}, \text{y}\}$ is precisely the set of identifiers declared in the structure *S*. Similarly the highlighting does not show that $\{\text{foo}, \text{bar}, \text{x}\}$ is precisely the set of identifiers specified in the signature *s*. We could then consider the following convention when highlighting a type error: if all the identifiers declared in a structure or *s* specified in

Figure 23 Constraint solving to handle merged unmatched errors

Some kinds of errors are not handled by the system presented in this section, although our implementation handles them. For more information please refer to the introductory paragraph of Sec. C.3.

signature constraints

- (SC5) $\text{prematch}(\langle u, e \rangle, \bar{d}, e_1, e_2) \rightarrow \text{succ}(\Delta', \text{tfm})$, if $\text{build}(u, e_1) = e_1''$ and $\text{build}(u, e_2) = e_2''$
and $\text{match}(\langle u, e \rangle, \bar{d}, \emptyset, \langle \emptyset, \emptyset \rangle, e_1'', e_2'') \rightarrow^* \text{succ}(\Delta', \text{tfm}, \text{unm})$
and $(\text{unm} = \langle \emptyset, \bar{id}_2 \rangle$ or $\neg \text{complete}(e_1'', e_2'')$)
- (SC6) $\text{prematch}(\langle u, e \rangle, \bar{d}, e_1, e_2) \rightarrow \text{err}(\langle ek, \bar{d} \rangle)$, if $\text{build}(u, e_1) = e_1''$ and $\text{build}(u, e_2) = e_2''$
and $\text{match}(\langle u, e \rangle, \bar{d}, \emptyset, \langle \emptyset, \emptyset \rangle, e_1'', e_2'') \rightarrow^* \text{succ}(\Delta', \text{tfm}, \text{unm})$
and $\text{unm} = \langle \bar{id}_1, \bar{id}_2 \rangle$ and $\bar{id}_1 \neq \emptyset$ and $\text{complete}(e_1'', e_2'')$
and $ek = \text{unmatched}(\bar{id}_1, \text{getBinders}(e_1'') \setminus \bar{id}_2, \bar{id}_2)$
- (SC7) $\text{prematch}(\langle u, e \rangle, \bar{d}, e_1, e_2) \rightarrow \text{err}(er)$, if $\text{match}(\langle u, e \rangle, \bar{d}, \emptyset, \langle \emptyset, \emptyset \rangle, \text{build}(u, e_1), \text{build}(u, e_2)) \rightarrow^* \text{err}(er)$

structure/signature matching

- (SM1) $\text{match}(\Delta, \bar{d}, \text{tfm}, \text{unm}, e, \square) \rightarrow \text{succ}(\Delta, \text{tfm}, \text{unm})$
- (SM2) $\text{match}(\Delta, \bar{d}, \text{tfm}, \text{unm}, e, e_1; e_2) \rightarrow \text{match}(\Delta', \bar{d}, \text{tfm}', \text{unm}', e, e_2)$,
if $\text{match}(\Delta, \bar{d}, \text{tfm}, \text{unm}, e, e_1) \rightarrow^* \text{succ}(\Delta', \text{tfm}', \text{unm}')$
- (SM3) $\text{match}(\Delta, \bar{d}, \text{tfm}, \text{unm}, e, e_1; e_2) \rightarrow \text{err}(er)$,
if $\text{match}(\Delta, \bar{d}, \text{tfm}, \text{unm}, e, e_1) \rightarrow^* \text{err}(er)$
- (SM4) $\text{match}(\Delta, \bar{d}, \text{tfm}, \langle \bar{id}_1, \bar{id}_2 \rangle, e, \downarrow \text{vid} = \sigma_1) \rightarrow \text{succ}(\Delta; e_0, \text{tfm}, \langle \bar{id}_1, \bar{id}_2 \cup \{\text{vid}\} \rangle)$,
if $e(\text{vid}) = \sigma_2$ and $\forall i \in \{1, 2\}. (\sigma_i = \forall \bar{\alpha}_i. \tau_i$ or $(\sigma_i = \tau_i$ and $\bar{\alpha}_i = \emptyset))$
and $\tau_1' = \text{appTFM}(\tau_1, \text{tfm})$ and $\text{solve}(\Delta, \bar{d}, \tau_1' = \tau_2) \rightarrow^* \text{succ}(\langle u', e' \rangle)$
and $\tau = \text{build}(u', \tau_1')$ and $\langle \langle u', e' \rangle, \downarrow \text{vid} \stackrel{\bar{d}}{=} \forall (\bar{\alpha}_1 \cup \bar{\alpha}_2) \cap \text{vars}(\tau). \tau \rangle \xrightarrow{\text{genExTyVar}} e_0$
- (SM5) $\text{match}(\Delta, \bar{d}, \text{tfm}, \text{unm}, e, \downarrow \text{vid} = \sigma_1) \rightarrow \text{err}(er)$,
if $e(\text{vid}) = \sigma_2$ and $\forall i \in \{1, 2\}. (\sigma_i = \forall \bar{\alpha}_i. \tau_i$ or $(\sigma_i = \tau_i$ and $\bar{\alpha}_i = \emptyset))$
and $\text{solve}(\Delta, \bar{d}, \text{appTFM}(\tau_1, \text{tfm}) = \tau_2) \rightarrow^* \text{err}(er)$
- (SM6) $\text{match}(\Delta, \bar{d}, \text{tfm}, \langle \bar{id}_1, \bar{id}_2 \rangle, e, \downarrow tc = \theta_1) \rightarrow \text{succ}(\Delta; (\downarrow tc \stackrel{\bar{d}}{=} \theta_2'), \text{tfm}', \langle \bar{id}_1, \bar{id}_2 \cup \{tc\} \rangle)$,
if $e(tc) = \theta_2$ and (if $\theta_1 \in \text{DumVar}$ then $\theta_2' = \theta_1$ else $\theta_2' = \theta_2^{\text{deps}(\theta_1)}$)
and (if $\theta_1 = \Lambda\alpha. (\alpha \gamma) \bar{d}'$ then $\text{tfm}' = \text{tfm} \boxplus \{\gamma \mapsto \theta_2'\}$ else $\text{tfm}' = \text{tfm}$)
- (SM7) $\text{match}(\langle u_1, e_1 \rangle, \bar{d}, \text{tfm}, \langle \bar{id}_1, \bar{id}_2 \rangle, e, \downarrow \text{sid} = e_0) \rightarrow \text{succ}(\langle u_2, e_1; (\downarrow \text{sid} \stackrel{\bar{d}}{=} \text{diff}(e_1, e_2)) \rangle, \text{tfm}', \langle \bar{id}_1, \bar{id}_2 \cup \{\text{sid}\} \rangle)$,
if $e(\text{sid}) = e_0'$
and $\text{match}(\langle u_1, e_1 \rangle, \bar{d}, \text{tfm}, \langle \emptyset, \emptyset \rangle, e_0', e_0) \rightarrow^* \text{succ}(\langle u_2, e_2 \rangle, \text{tfm}', \text{unm}')$
and $(\text{unm}' = \langle \emptyset, \bar{id}_2 \rangle$ or $\neg \text{complete}(e_0', e_0))$
- (SM16) $\text{match}(\Delta, \bar{d}, \text{tfm}, \text{unm}, e, \downarrow \text{sid} = e_0) \rightarrow \text{err}(\langle \text{unmatched}(\bar{id}_1, \text{getBinders}(e_0') \setminus \bar{id}_2, \bar{id}_2), \bar{d} \rangle)$,
if $e(\text{sid}) = e_0'$
and $\text{match}(\Delta, \bar{d}, \text{tfm}, \langle \emptyset, \emptyset \rangle, e_0', e_0) \rightarrow^* \text{succ}(\Delta', \text{tfm}', \langle \bar{id}_1, \bar{id}_2 \rangle)$ and $\bar{id}_1 \neq \emptyset$ and $\text{complete}(e_0', e_0)$
- (SM8) $\text{match}(\Delta, \bar{d}, \text{tfm}, \text{unm}, e, \downarrow \text{sid} = e_0) \rightarrow \text{err}(er)$,
if $e(\text{sid}) = e_0'$ and $\text{match}(\Delta, \bar{d}, \text{tfm}, e_0', e_0) \rightarrow^* \text{err}(er)$
- (SM9) $\text{match}(\Delta, \bar{d}, \text{tfm}, \text{unm}, e, \downarrow \text{vid} = is_1) \rightarrow \text{succ}(\Delta; (\downarrow \text{vid} \stackrel{\bar{d}'}{=} \text{vid}), \text{tfm}, \text{unm})$,
if $e[\text{vid}] = is_2$ and $\text{deps}(is_2) = \bar{d}'$ and $(\text{solve}(\Delta, \bar{d}, is_1 = is_2) \rightarrow^* \text{succ}(\Delta')$ or $\text{strip}(is_1) = v)$
- (SM10) $\text{match}(\Delta, \bar{d}, \text{tfm}, \text{unm}, e, \downarrow \text{vid} = is_1) \rightarrow \text{err}(er)$,
if $e[\text{vid}] = is_2$ and $\text{strip}(is_1) \neq v$ and $\text{solve}(\Delta, \bar{d}, is_1 = is_2) \rightarrow^* \text{err}(er)$
- (SM11) $\text{match}(\Delta, \bar{d}, \text{tfm}, \langle \bar{id}_1, \bar{id}_2 \rangle, e, \downarrow id = x) \rightarrow \text{succ}(\Delta; (\downarrow id = \text{toDumVar}(x)), \text{tfm}', \langle \bar{id}_1 \cup \{id\}, \bar{id}_2 \rangle)$,
if $e(id)$ is undefined and (if $x = \Lambda\alpha. (\alpha \gamma) \bar{d}'$ then $\text{tfm}' = \text{tfm} \boxplus \{\gamma \mapsto \alpha_{\text{dum}}\}$ else $\text{tfm}' = \text{tfm}$)
- (SM12) $\text{match}(\Delta, \bar{d}, \text{tfm}, \text{unm}, e, ev) \rightarrow \text{succ}(\Delta; ev, \text{tfm}, \text{unm})$
- (SM13) $\text{match}(\Delta, \bar{d}, \text{tfm}, \text{unm}, e, e' \bar{d}') \rightarrow \text{match}(\Delta, \bar{d} \cup \bar{d}', \text{tfm}, \text{unm}, e, e')$

a signature are involved in the reported error and this information is necessary for the error to occur then we highlight the blank spaces (if any) preceding the corresponding `val`, `type`, `datatype` and `structure` keywords.

We would then obtain the following highlighting which is a bit more informative than the one presented above:

```
structure S = struct val (foo, bar, x, y) = (1, 2, 3) end
signature s = sig val foo : int val bar : int val x : int end
structure T = S :> s
```