Smartest Recompilation *

Zhong Shao zsh@princeton.edu

Andrew W. Appel appel@princeton.edu

CS-TR-395-92

Princeton University

October 1992

Abstract

To separately compile a program module in traditional statically-typed languages, one has to manually write down an import interface which explicitly specifies all the external symbols referenced in the module. Whenever the definitions of these external symbols are changed, the module has to be recompiled. In this paper, we present an algorithm which can automatically infer the "minimum" import interface for any module in languages based on the Damas-Milner type discipline (e.g., ML). By "minimum", we mean that the interface specifies a set of assumptions (for external symbols) that are just enough to make the module type-check and compile. By compiling each module using its "minimum" import interface, we get a separate compilation method that can achieve the following optimal property: *A compilation unit never needs to be recompiled unless its own implementation changes.*

*A short version of this paper will appear in the Twentieth Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, January, 1993.

1 Introduction

Most traditional separate compilation methods rely on manually created contexts (e.g., Modula-3 interfaces, "include-files" in C, and Ada package specifications) to enforce type correctness across module boundaries. Using the proper contexts, the compiler can check that each module uses its imported interfaces properly, and implements its exported interface as expected. The disadvantage of using these manually created contexts is that to guarantee consistency, all modules using a changed context must be recompiled, no matter how small the change is. The conventional recompilation rule (as described in Tichy [30]) is stated as follows: *"A compilation unit must be recompiled whenever* (1) *its own implementation changes, or* (2) *a context changes upon which the compilation unit depends."* This is obviously not satisfactory because adding a comment or adding a new declaration to a pervasive context may cause the unnecessary recompilation of the entire system. Tichy [30] presents an effective technique called "smart recompilation" that eliminates most of the redundant recompilations triggered by (2). In Tichy's scheme, a compilation unit is recompiled only if its implementation changes, or if it references a symbol defined elsewhere whose definition has changed. Schwanke and Kaiser [29] define "smarter recompilation" which can eliminate even more (but not all) redundant recompilations caused by (2). So a natural question to ask is: Can we eliminate all redundant recompilations? that is, can we achieve the following **smartest recompilation rule**: *A compilation unit never needs to be recompiled unless its own implementation (source code) changes*?

Standard ML (SML) [20] has a rather elaborate module system, but SML compilers have not supported separate compilation very well. The problem is that in SML, modules such as structures and functors can liberally reference externally defined identifiers without even mentioning what are their specifications. For example, by using "qualified" (dotted) identifiers, a structure `FOO` can use `BAR.QUX.f` to reference the function `f` defined in the substructure `QUX` of the structure `BAR`, without even knowing what the type of `f` is. Because of the lack of explicit import interfaces, structures and functors with free variables (let's call them open-formed modules) are not considered as separately compilable units. So how can we separately compile open-formed modules in SML?

This paper presents a new separate compilation method which actually answers both of the above two questions. Surprisingly, not only can we separately compile arbitrary structures and functors in SML, but we can also accomplish the "smartest recompilation rule." Our idea is simple: in order to separately compile a module with references to external identifiers, we have to know the specifications (e.g., types) of these external identifiers; since they are not explicitly specified, we infer them by looking at how these external identifiers are used inside the module; then we compile the module by using this inferred import interface as its context; finally, when all the modules are linked together, cross-module type errors are reported by checking whether the surroundings match (or satisfy) the specifications in each inferred import interface. The catch here is that in order to achieve the "smartest recompilation rule", we have to infer the "minimum" import interface. Informally speaking, this "minimum" import interface specifies a set of assumptions (on those external identifiers) that are just enough to make the module type-check and compile; at link time, if the module's surroundings satisfy this set of assumptions, the compiled code can be reused, otherwise there must be cross-module type errors. The inference algorithm is discussed in detail in section 2 and 3.

Now let's see an example of how our method works. From the following SML structure declaration,

```
structure FOO = struct val x = BAR.f
                       val y = (BAR.g 4, BAR.g true)
                end
```

we know that in order to compile `FOO`, the context should contain a structure named `BAR`. Inside `BAR`, there

should be at least two `val` declarations: one is `f`, which can have any type, say α; the other is `g`, which should be a function that can be applied to both integers and booleans, that is, `g` should have a type more general than $int \rightarrow \beta$ and $bool \rightarrow \gamma$. Here, α, β and γ are just type variables we used to denote unknown types. Compiling `FOO` in this inferred interface will result in a structure with two components: the variable `x` has type α and the variable `y` has type $\beta * \gamma$. Now suppose that the real structure `BAR` is defined as follows:

```
structure BAR = struct val f = 3
                       val h = true
                       val g = fn z => z
                end
```

then at link time, when the real `BAR` is matched against the import interface of `FOO`, we find that the type variables α and β should be int and γ should be $bool$, thus `FOO.x` will have type int and `FOO.y` will have type $int * bool$. This is exactly what we will get if we compile `FOO` in the environment that would result from compiling `BAR`.

To achieve the "smartest recompilation rule", the back end of the compiler must use only the type information specified in the module's inferred import interface. This limitation is not a big problem. The back end of the current SML/NJ compiler [4] uses almost no type information from the front end but it still produces quite efficient code. Many optimization techniques that do use the type information, such as Leroy's representation analysis [17], can still be partially incorporated into our separate compilation system. The details will be described in section 4 and 5 of this paper.

Our separate compilation method immediately has the following advantages over traditional methods:

- Because of the smartest recompilation rule, each module never needs to be recompiled unless its own implementation changes; so maximum reusability is achieved.
- Because all modules are compiled independently of each other, they can be compiled in any order. This also means that programmers need no longer maintain dependency files (e.g., `Makefile` [9]).
- Open-formed modules can also be separately compiled.
- Cross-module type errors are now symmetric. In traditional methods, if module A references identifiers defined in module B, type inconsistencies between module A and module B will show up when A is compiled. If the programmer fixes the error by editing B, both A and B must be recompiled. But in our method, because cross-module type errors are reported at link time, only B will be recompiled.
- The compiler based on our method will automatically be a standalone compiler. As far as we know, no one has yet built a standalone compiler for the complete SML module system.

1.1 Closed vs. Open-formed Modules

Standard ML allows programming in open-formed modules. The essential difference between closed and open-formed modules can be seen from rewriting the above open-formed structure `FOO` in closed form, the SML functor `FOO'`.

```
functor FOO' (BAR : sig val f : int
                        val g : 'a -> 'a
                    end)
  = struct val x = BAR.f
           val y = (BAR.g 4, BAR.g true)
    end
```

This shows that programmers have to give assumptions about the types of `BAR.f` and `BAR.g` based on a *pro forma* implementation of the structure declaration `BAR`. If later, `BAR.g` is changed to have a type scheme $\forall\alpha.\alpha \rightarrow int$, the above declaration will no longer be useful and will have to be modified and recompiled. An open-formed module such as structure `FOO` does not have this problem because it does not require that the specifications of imported identifiers be explicitly given. People may want to write `FOO'` using its minimum import interface as its argument signature, but this "minimum" is not expressible in the SML type system. Moreover, inferring the minimum import interface cannot be easily done by hand.

We do not advocate writing large programs all in open-formed modules. SML strongly encourages that every structure declaration should be written with a result signature constraint (as its export interface). Programmers can write their programs all in the form of closed functors such as `FOO'`. On the other hand, in practice, we find it extremely convenient and flexible to write parts of our programs as open-formed modules.

For languages based on the Damas-Milner type discipline [7] such as SML and Haskell [11], there is another reason in favor of writing certain modules in opened forms. One of the most important features of the Damas-Milner type discipline is that the most general type for arbitrary expressions can be automatically inferred by compilers. It is, however, nontrivial to infer the most general type simply by hand, especially with the presence of polymorphic references in SML or type classes in Haskell. This makes it also nontrivial to write explicit import interfaces for many modules. The SML commentary [19] also suggests that programmers will probably need to write down many sharing equations if they want to close every module and that it will be too restrictive to write everything in closed functors.

2 Assumption Inference in Core ML

ML has a sophisticated type inference system. Given an ML expression $e = \lambda x.f(x+1)$, even though the type of newly introduced variable x is not specified, we can still find the most general type of e if we know the type of f and $+$. For example, if f has type $\forall\alpha.\alpha \rightarrow \alpha$ and $+$ has type $int * int \rightarrow int$, e will have type $int \rightarrow int$. Milner's type inference (or type reconstruction) algorithm W (as in Tofte [31]) takes two arguments, a type environment TE and an ML expression e; all the free variables in e (such as f and $+$) must be specified with a type in TE, and $W(TE, e)$ will return the most general type for e.

To support "smartest recompilation," we face the challenge of doing type inference without even knowing the type of external identifiers. For example, can we find out the most general type of the above expression e if we do not know the type of f and $+$? This seems impossible. But in the case of separate compilation, we assume that the types of external identifiers will be known at link time. We can divide the type inference into two phases: first (at compile time), we infer a type τ for e and a set of assumptions A for the free variables in e, which essentially means that e will have type τ if the free variables in e satisfy the constraints in A; then (at link time) when the types of those free variables (i.e., TE) are known, we match them against those in A and "magically" recover the most general type of e in TE. To distinguish it from usual ML type inference, we call the inference done in the first phase "assumption inference".

In this section and section 3 we discuss the details of our assumption inference algorithm and show how the matching done at link time can successfully recover the correct type for each expression. To simplify the presentations, we divide our algorithms into two parts: this section for Core ML and the next section for the ML module language. We only give the details of our algorithm for the mini-ML language *Exp* and the skeletal module language *ModL* used by Tofte [31]. However it is easy to extend our algorithm to the rest of SML.

The expressions in the mini-ML language *Exp* are defined by the following grammar:

$$e ::= x \mid \lambda x.e_1 \mid e_1\ e_2 \mid \texttt{let}\ x\ \texttt{=}\ e_1\ \texttt{in}\ e_2$$

Here is a brief review of the notation. Suppose TyVar is an infinite set of type variables and TyCon is a set of nullary type constructors, the set of *types*, Type, ranged over by τ and the set of *type schemes*, TypeScheme, ranged over by σ are defined by $\tau ::= \pi \mid \alpha \mid \tau_1 \rightarrow \tau_2$ and $\sigma ::= \tau \mid \forall\alpha.\sigma_1$. A *type environment* is a finite map from program variables to type schemes. $tyvars(\tau)$, $tyvars(\sigma)$ and $tyvars(TE)$ are the set of type variables that occur *free* in τ, σ and TE respectively. A type τ' is a *generic instance* of a type scheme $\sigma = \forall\alpha_1, ..., \alpha_n.\tau$, written as $\tau' \prec \sigma$, if there exists a substitution S with its domain being a subset of $\{\alpha_1, ..., \alpha_n\}$ and $\tau' = S(\tau)$. A type scheme σ_1 is more general than σ_2, denoted as $\sigma_2 \prec \sigma_1$, if all generic instances of σ_2 are also generic instances of σ_1. The generalization of a type τ in a type environment TE is denoted by $gen(TE, \tau)$, it is the type scheme $\forall\alpha_1, ..., \alpha_n.\tau$ where $\{\alpha_1, ..., \alpha_n\} = tyvars(\tau) \backslash tyvars(TE)$. The core ML type system, in the form of type deduction rules as $TE \vdash e : \tau$, is listed in the appendix at the end of this paper. It is directly copied from Tofte's thesis [31].

2.1 The assumption inference algorithm W*

We define a *type assumption* to be a pair (x,τ) where x is a program variable and τ is a type. An *assumption environment*, ranged over by A, is a set of type assumptions; it is usually represented by a finite mapping from program variables to lists of types. In the following, we use $A\backslash\{x\}$ to denote the set of type assumptions in A except those for variable x, and $A(x)$ to denote the set of types associated with variable x in A. We also use *Unify* to denote Robinson's original unification algorithm on classical term algebras [26]. *Unify* takes a set of pairs of types and returns a substitution (the most general unifier).

Figure 1 gives the assumption inference algorithm W^* and the matching algorithm *Match*. W^* takes an ML expression, and returns a type and an assumption environment. *Match* takes an ML type environment and an assumption environment, and returns a substitution. The other two procedures in figure 1 are *MonoUnify* and *PolyUnify*. *MonoUnify* takes a type and a set of types, and returns a substitution. *PolyUnify* takes two arguments: a triple of a TyVar set and a type and an assumption environment, and a set of types; it returns a substitution and an assumption environment.

Given an ML expression e, $W^*(e)$ delays the type-checking of all free variables in e by recording their monomorphic type instances in an assumption environment A. In the case of lambda abstraction $\lambda x.e_1$, the type α of x is treated as monomorphic; the procedure *MonoUnify* checks whether the set of assumptions collected for x from e_1 satisfies this constraint. One the other hand, in the `let` expression, the type of x is treated as polymorphic; for each use of x in e_2, the type and the assumption environment from e_1 is renamed with new type variables; the procedure *PolyUnify* then checks the typing of x in e_2 and merges the assumption environments collected from e_1 and e_2. When the real type environment TE for the free variables is known (at link time), the matching algorithm $Match(TE, A)$ precisely recovers everything, including the result type of elaborating e in TE.

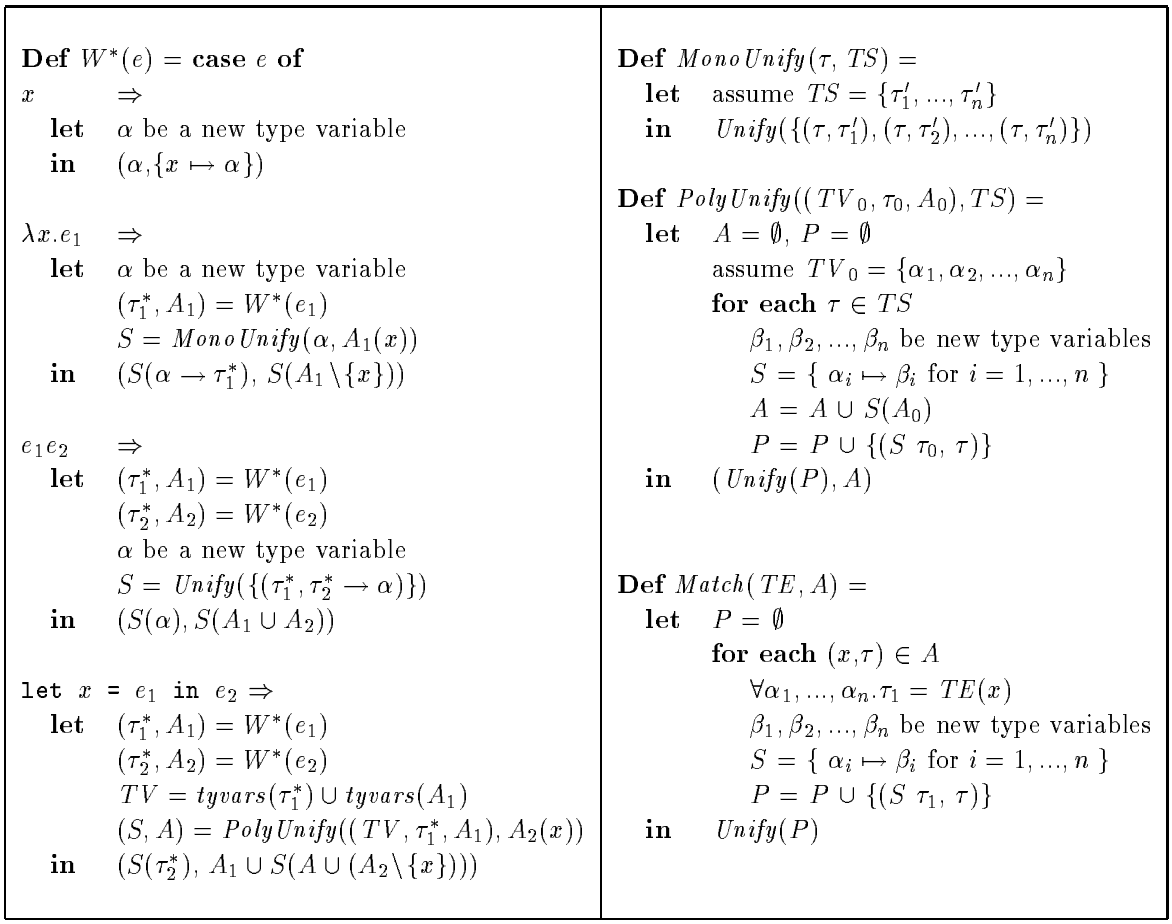

Figure 1: The Inference Algorithm W^* and *Match*

For example, given an expression e = "`let` $g = \lambda x.fx$ `in` $g\ g$", the free variable of e is f; $W^*(e)$ will return an assumption environment $A = \{f \mapsto (\alpha_6 \to \alpha_7) \to \alpha_8, f \mapsto (\alpha_6 \to \alpha_7), f \mapsto (\alpha_1 \to \alpha_3)\}$ and a type α_8 for e. If the real type environment TE is $\{f \mapsto \forall\alpha.\alpha \to \alpha\}$, $Match(TE, A)$ will result in the substitution $S^* = \{\alpha_6 \mapsto \beta_1,\ \alpha_7 \mapsto \beta_1, \alpha_1 \mapsto \beta_1,\ \alpha_3 \mapsto \beta_1,\ \alpha_8 \mapsto (\beta_1 \to \beta_1)\}$. Thus the expression e will have the type $S^*(\alpha_8) = (\beta_1 \to \beta_1)$. This is exactly what we will get if we apply Tofte's algorithm W to TE and e.

In fact we can show that the algorithm W^* is equivalent to Milner's W [31] in the following sense:

Theorem 2.1 *Given a type environment TE and an expression e, then $(S, \tau) = W(TE, e)$ succeeds if and only if both $(\tau^*, A) = W^*(e)$ and $S^* = Match(TE, A)$ succeed; Moreover, there exists two substitutions R_1 and R_2, such that the following are true: (1) $R_1 \circ R_2 = R_2 \circ R_1 = ID$; (2) $R_1(S(TE), \tau) = (S^*(TE), S^*\tau^*)$; (3) $(S(TE), \tau) = R_2(S^*(TE), S^*\tau^*)$.*

Proof By structural induction on the expression e. For details, see appendix. **QED**.

Notice that theorem 2.1 is not trying to show the soundness and completeness of W^* directly. It is just proving that the result of W^* and *Match* is equivalent to the result of W. Proving this kind of equivalence is

relatively easier. From the soundness and completeness of the algorithm W(which is proved in Damas's Ph.D thesis [6]), and the above theorem 2.1, we can easily get the following soundness and completeness results for our algorithm W^*.

Corollary 2.2 (Soundness of W^*) *Given a type environment TE and an ML expression e, if both $(\tau^*, A) = W^*(e)$ and $S^* = Match(TE, A)$ succeed, then $S^*(TE') \vdash e : S^*\tau^*$.*

Corollary 2.3 (Completeness of W^*) *Given a type environment TE and an ML expression e, suppose that $TE_1 = S_1(TE)$ and $TE_1 \vdash e : \tau_1$, then both $(\tau^*, A) = W^*(e)$ and $S^* = Match(TE, A)$ will succeed; Moreover there exists a substitution S' such that $TE_1 = S'(S^*(TE))$ and $\tau_1 \prec S'(\mathrm{gen}(S^*(TE), S^*\tau^*))$.*

The algorithm W^* itself is interesting. Recursive calls to W^* in the algorithm will not interfere with each other so they can be called in any order. If concurrency is used, W^* can be efficiently implemented. The case for the `let` x `=` e_1 `in` e_2 expression implies that we can link two pieces of programs, i.e., e_1 and e_2, even though both of them contain free variables; this is done by the algorithm *PolyUnify* in figure 1.

The assumption environment A returned from W^* may be big. A possible optimization is to insert a simplifying procedure at each recursive call to W^* in the algorithm. This simplifying procedure will identify all "isolated" type assumptions in A. Given $(\tau, A) = W^*(e)$, we define an equivalence relation $\sim$ on type variables: "$\alpha \sim \beta$ if there exists a type τ' such that either $\tau' = \tau$ or $(x, \tau') \in A$ for some x is true, and both α and β are type variables of τ'." Let TV be the transitive closure of $tyvars(\tau)$ under $\sim$, then all pairs (x, t) in A where $tyvars(t) \cap TV = \emptyset$ are denoted as "isolated" assumptions. All type variables occurred in "isolated" assumptions need not to be renamed in *PolyUnify* and most redundant "isolated" assumptions can be eliminated.

2.2 The assumption inference algorithm D

One disadvantage of W^* is that its sequential implementation may be not very efficient in practice. In most compilers, there is a pervasive basis (or initial library) which tends to be referenced very frequently by user programs, thus the resulting assumption environment from W^* may be quite big (even if it uses certain optimizations mentioned above). It turns out that this problem can be elegantly solved by extending ML types with type predicates and assumptions. This extension, which is called "constrained type" in Kaes [13] and "qualified type" in Jones [12], is normally used to reason about the ML type system in the presence of overloading and subtyping. The algorithm D presented in this section is the type reconstruction algorithm for Kaes's constrained type system. By using a special set of type predicates, the algorithm D can efficiently solve the assumption inference problem even when there is a pervasive basis.

2.2.1 An extension of ML with constrained types

The extension of ML type system with constrained types (denoted as ML^+) is discussed in detail by Kaes [13] and Jones [12] to solve the type inference problem in languages that support overloading and subtyping. It turns out that it can also be used to solve our assumption inference problem. The language syntax they use is essentially same as the mini-ML language *Exp*. In the following, we give a quick review of Kaes's framework of extending ML with constrained types. To ease the notation, $\overline{x_n}$ is used to denote a sequence $x_1, ..., x_n$.

Definition 2.1 *Let P be a finite n-indexed family of predicate symbols. The set of* predicate constraints *over Type is defined as $\{p(\overline{\tau_n}) \mid p \in P_n, \tau_i \in Type \text{ where } i = 1, ..., n\}$. An interpretation of P is a family of total computable functions $(\hat{p})_{p \in P}$, such that for $p \in P_n, \hat{p} : (\text{Type})^n \rightarrow \mathbf{2}$.*

Definition 2.2 *A set C of constraints is* satisfiable *if there exists a substitution S such that if $p(\overline{\tau_n}) \in C$ then $\hat{p}(S(\overline{\tau_n}))$ is true. Satisfiablity will be denoted as $S \models C$.*

Definition 2.3 *A substitution S is a* solution *of C, if $S' \circ S \models C$ for all substitutions S'. A solution S is called the* most general solution *of a constraint set C, if for any solution R of C, there exists a substitution S', such that $R = S' \circ S$. In most cases, the most general solution does not exist.*

The *entailment* relation on constraint sets, written $C_1 \Vdash C_2$, may be defined once a particular predicate system is given. In the following, we only consider those predicate systems which satisfy the following properties: if $C_1 \Vdash C_2$ then $\forall S : S \models C_1 \Rightarrow S \models C_2$.

Definition 2.4 *A* constrained type *is a pair $\tau \mid C$, consisting of a type τ and a set of constraints C. A* constrained type scheme *is of the form $\forall\overline{\alpha_n}.\tau|C$. A* constrained type environment *is now just a finite map from variables to constrained type schemes.*

Definition 2.5 *A constrained type $\tau'|C'$ is a* generic instance *of a constrained type scheme $\sigma = \forall\overline{\alpha_n}.\tau|C$, written as $\tau'|C' \prec \sigma$, if there exists a substitution S with domain being a subset of $\overline{\alpha_n}$ such that $\tau' = S(\tau)$ and $C' \Vdash S(C)$. We also denote $\sigma' \prec \sigma$ if all generic instances of σ' are generic instances of σ; and $\sigma' \equiv \sigma$ if $\sigma' \prec \sigma$ and $\sigma \prec \sigma'$.*

Typability of an expression e in ML^+ is expressed as a judgement $C, TE \vdash e : \tau$, which can be read as "e has type τ under (constrained) type environment TE, provided C is satisfiable." The generalization of a constrained type $\tau|C$ in the context of a type environment TE is denoted by $gen(TE, \tau|C)$, it is the constrained type scheme $\forall\overline{\alpha_n}.\tau|C'$ where $\{\alpha_1, ..., \alpha_n\} = tyvars(\tau|C) \backslash tyvars(TE)$ and $C' = \{p(\overline{\tau}) \in C \text{ where } tyvars(p(\overline{\tau})) \cap \overline{\alpha_n} \neq \emptyset\}$.

Definition 2.6 *A typing $C, TE \vdash e : \tau$ is* more general *than $C', TE' \vdash e : \tau'$, if there exists a substitution S, such that (1) $x \in dom(TE) \Rightarrow TE'(x) \prec S(TE(x))$ (2) $gen(TE', \tau'|C') \prec S(gen(TE, \tau|C))$.*

Kaes [13] presented the type deduction rules (also listed in the appendix at the end of this paper for reference) and the type inference algorithm D (as in figure 2) for the above extension. It can be proved that his type inference algorithm D is sound and (syntactically) complete in the following sense:

Theorem 2.4 *Let an instance of a constraint based inference system be given, e be an expression, TE and TE' be type environments. Suppose for certain substitution S_1, for each $x \in dom(TE)$, $TE'(x) \prec S_1(TE(x))$; and $C', TE' \vdash e : \tau'$ is a valid typing, then $(S, \tau|C) = D(TE, e)$ succeeds. Moreover, $C, S(TE) \vdash e : \tau$ is valid and more general than $C', TE' \vdash e : \tau'$.*

2.2.2 Application to assumption inference

We can use the above extension to solve our assumption inference problem. The set of predicates we use, denoted by P_m, is $\{p_x(\tau)$ where x is any program variable$\}$. The interpretation of p_x is "$\hat{p_x}(\tau)$ = true if

Def $D(TE, e)$ = **case** e **of**

x $\Rightarrow$ **let** $TE(x) = \forall\overline{\alpha_n}.\tau|C$ and $\overline{\beta_n}$ be new type variables and $S = \{\alpha_i \mapsto \beta_i$ for $i = 1, ..., n\}$
in $(ID\ ,\ S(\tau)|S(C))$

$\lambda x.e_1$ $\Rightarrow$ **let** α be a new type variable and $(S_1, \tau_1|C_1) = D(TE \pm \{x \mapsto \alpha|\emptyset\}, e_1)$
in $(S_1\ ,\ (S_1(\alpha) \rightarrow \tau_1)|C_1)$

$e_1 e_2$ $\Rightarrow$ **let** $(S_1, \tau_1|C_1) = D(TE, e_1)$ and $(S_2, \tau_2|C_2) = D(S_1(TE), e_2)$
α be a new type variable and $S_3 = Unify(S_2\tau_1, \tau_2 \rightarrow \alpha)$
in $(S_3 \circ S_2 \circ S_1\ ,\ (S_3\alpha)|(S_3(S_2(C_1) \cup C_2)))$

`let` x `=` e_1 `in` e_2 $\Rightarrow$
let $(S_1, \tau_1|C_1) = D(TE, e_1)$ and $(S_2, \tau_2|C_2) = D(S_1(TE) \pm \{x \mapsto gen(S_1(TE), \tau_1|C_1)\}, e_2)$
in $(S_2 \circ S_1\ ,\ \tau_2|(S_2(C_1) \cup C_2))$

Figure 2: The Type Inference Algorithm D

and only if $\tau \prec \sigma$, assuming that the type of x is a closed ML type scheme σ." The entailment relation on constraint sets is defined as: $C_1 \Vdash C_2$ if and only if C_1 is satisfiable and $\forall S : S \models C_1 \Rightarrow S \models C_2$. This relation is decidable for our particular predicate system P_m because of the following lemma:

Lemma 2.5 *For any constraint set C formed in the predicate system P_m, either there is no solution or there exists a most general solution S.*

Proof If we consider each $p_x(\tau)$ as an assumption (x, τ), also assume that the type of x is known, the *Match* algorithm in figure 1 can be used to find the most general solution of C. The lemma then follows from Robinson's unification theorem. **QED**.

Given a closed ML type scheme $\sigma = \forall\overline{\alpha_n}.\tau$, it can be written as ML$^+$ constrained type schemes $\sigma_1 = \forall\overline{\alpha_n}.\tau|\emptyset$ or $\sigma_2 = \forall\alpha.\alpha|\{p_x(\alpha)\}$, but obviously $\sigma_1 \prec \sigma_2$ in ML$^+$.

The great thing about the algorithm D is that when it is running, it does not need any knowledge about the interpretation of the predicate system. This leads to the following theorem:

Theorem 2.6 *Given a* ML *type environment $TE = TE_1 \pm TE_2$, where* $\mathrm{Dom}(TE_1) \cap \mathrm{Dom}(TE_2) = \emptyset$ *and $tyvars(TE_2) = \emptyset$, we construct a* ML$^+$ *constrained type environment $TE' = TE'_1 \pm TE'_2$ where $TE'_1 = \{x \mapsto \forall\overline{\beta_n}.(\tau|\emptyset)$ where $x \in \mathrm{Dom}(TE_1)$ and $TE_1(x) = \forall\overline{\beta_n}.\tau\}$ and $TE'_2 = \{x \mapsto \forall\alpha.(\alpha|\{p_x(\alpha)\})$ where $x \in \mathrm{Dom}(TE_2)\}$ and the interpretation of p_x is "$\hat{p_x}(\tau)$ = true if and only if $\tau \prec TE_2(x)$". Then $(S, \tau) = W(TE, e)$ succeeds if and only if $(S', \tau'|C') = D(TE', e)$ succeeds and there exists a most general solution S^* for C'. Moreover, there exists two substitutions R_1 and R_2, such that the following are true: (1) $R_1 \circ R_2 = R_2 \circ R_1 = ID$; (2) $R_1(S(TE), \tau) = (S^*(S'(TE)), S^*\tau')$; (3) $(S(TE), \tau) = R_2(S^*(S'(TE)), S^*\tau')$.*

Proof Follows from lemma 2.5 and theorem 2.4. For details, see the appendix. **QED**.

In practice algorithm D will be more useful than W^* because it resembles the algorithm W and it also works more efficiently when the types of some free variables are known. Moreover, D can be easily extended to

dec	::=		m	$\in$ StrName
	|	*strdec*	N	$\in$ NameSet = Fin(StrName)
	|	*strdec dec1*	GE	$\in$ SigEnv = SigId $\xrightarrow{\text{fin}}$ Sig
			FE	$\in$ FunEnv = FunId $\xrightarrow{\text{fin}}$ FunSig
strexp	::=	*strid*	SE	$\in$ StrEnv = StrId $\xrightarrow{\text{fin}}$ Str
	|	`struct` *dec* `end`	S or (m, E)	$\in$ Str = StrName $\times$ Env
	|	*strexp.strid*	E	$\in$ Env = StrEnv
	|	*fctid(strexp)*	Σ or $(N)S$	$\in$ Sig = NameSet $\times$ Str
			$N(S, N'(S'))$	$\in$ FunSig = NameSet $\times$ (Str $\times$ Sig)
strdec	::=	`structure` *strid* = *strexp*	B	$\in$ Basis = Nameset $\times$ SigEnv $\times$ FunEnv $\times$ Env

Figure 3: *left:* Grammar; *right:* Semantic objects

work on various extensions of the ML type system with overloading and subtyping such as those in Kaes [13].

3 Assumption Inference in the SML Module Language

In this section, we present an assumption inference algorithm for the SML module language. To simplify the presentation, we only consider the skeletal language *ModL* (as in Tofte [31]) in figure 3. Notice that signature expressions and declarations are intentionally left out because their elaborations can be delayed to link time, thus are irrelevant to our assumption inference. Functor declarations are not considered in our language either because only their body, which is a structure expression, is elaborated at compile time. However, functor applications are considered in our language because they are structure expressions which must be elaborated at compile time.

The static semantics of *ModL* is discussed in detail in the definition [20] and Tofte [31]. Its deduction rule is in the form of "$B \vdash$ *phrase* $\Rightarrow A$" meaning that *phrase* is elaborated into a semantic object A in the basis B. The semantic objects are also defined in figure 3. The appendix at the end of this paper lists the set of static semantic rules for *ModL*. Here we give a quick review on the main concepts used in the static semantics.

Definition 3.1 *A* structure S *is a pair* (m, E)*, where* m *is the name of the structure and* E *is an environment, which gives the static information about the components of the structure. To make the presentation clear, from now on, we shall use* `str`(m,E) *to denote a structure* (m,E)*. A* signature *is an object of the form* $(N)S$*, where* S *is a structure and* N *is a finite set of names. A* functor signature Φ *is an object of the form* $N(S, N'(S'))$ *where* $N(S)$ *is the principal signature for the parameter signature expression of the functor and* S' *is the body structure of the functor, the names bound in* S' *are the names in* S' *which have to be generated afresh upon each functor application.*

Definition 3.2 *A* structure environment SE *is a finite map from structure identifiers to structures, similarly for* signature environment GE *and* functor environment FE*.*

Definition 3.3 (Names) StrName *is an infinite set of names: names that are specified in a signature expression and are not shared with already declared structures are called* flexible names, *denoted as* FlexStrName*; names of declared structures are called* rigid names, *denoted as* RigStrName.

Definition 3.4 (Realization) *A* realization *is a finite mapping from* FlexStrName *to* StrName, *and a* renaming realization $\varphi=\{m_i \mapsto m_i' \text{ where } i = 1,...,k\}$ *is a realization where* $m_i's$ *are distinct rigid names.*

Definition 3.5 (Enrichment) *A structure* $S_1 = \mathtt{str}(m_1, SE_1)$ enriches *a structure* $S_2 = \mathtt{str}(m_2, SE_2)$ *if* $m_1 = m_2$ *and the structure environment* SE_1 *enriches the* SE_2. *A structure environment* SE_1 *enriches a structure environment* SE_2 *if* $\mathrm{Dom}(SE_2) \subseteq \mathrm{Dom}(SE_1)$ *and for each* $s \in \mathrm{Dom}(SE_2)$, $SE_1(s)$ *enriches* $SE_2(s)$.

Definition 3.6 (Signature Matching) *A structure* S' *matches a signature* $\Sigma = N(S)$ *if there exists a realization* φ *such that* S' *enriches* $\varphi(S)$.

Because the SML module language is explicitly typed, the elaboration of a module expression simply involves type-checking. The static semantics in the Definition [20] can be viewed as a type checking algorithm. Given a structure expression with free identifiers, we want to infer the minimum constraints on these free identifiers with which the expression will just type-check. Again the minimum constraints are not expressible if we only use semantic objects in figure 3. We introduce a new kind of *structure variable* which is similar to row variables used in typing record calculi. Let StrVar be an infinite set of structure variables; structures and structure environments are now extended as $\mathrm{Str}' = (\mathrm{StrName} \times \mathrm{StrEnv}') \cup \mathrm{StrVar}$ and $\mathrm{StrEnv}' = \mathrm{StrId} \rightarrow \mathrm{Str}'$. Each structure variable t must have a kind. Kinds are defined as (StrName $\times$ KindEnv) where KindEnv is just a finite mapping StrId $\rightarrow$ (Str $\cup$ StrVar). To distinguish it from structures, a kind (m, KE) is represented as $\mathtt{STR}(m, KE)$. A kind assignment is a finite mapping from structure variables to kinds. A structure $S = \mathtt{str}(m, SE)$ has the kind k under the kind assignment K, written as $K \vdash S :: k$, if it is derivable from the following set of kinding rules:

(1) $K \vdash t :: \mathtt{STR}(m,KE)$, if $K(t) = \mathtt{STR}(m, KE)$
(2) $K \vdash \mathtt{str}(m, SE) :: \mathtt{STR}(m, KE)$ if $\mathrm{Dom}(SE) \supseteq \mathrm{Dom}(KE)$ and $\forall \mathtt{s} \in \mathrm{Dom}(KE), SE(\mathtt{s}) = KE(\mathtt{s})$.

A substitution now consists of two parts: one from StrVar to Str$'$, another from FlexStrName to StrName (i.e., realization). A kinded substitution is a pair consisting of a kind assignment and a substitution. A kinded substitution (K_1,R) respects a kind assignment K_2 if, for all t in dom(K_2), $K_1 \vdash R(t) :: R(K_2(t))$ is a derivable kinding. A kinded substitution (K_1,R_1) is more general than (K_2,R_2) if $R_2 = R_3 \circ R_1$ for some R_3 such that (K_2,R_3) respects K_1. A kinded substitution (K_1,R) is an unifier of a kinded set of equations (K_2,P) if it respects K_2 and $R(t_1) = R(t_2)$ for all (t_1, t_2) in P.

Figure 4 gives our inference algorithms W_{strexp} on structure expressions, W_{dec} on declarations, W_{fctid} on functor identifiers, W_{strdec} on structure declarations and the matching algorithm *ModlMatch*. The argument V and M records those already-used structure variables and flexible names. All functor applications are done by the matching algorithm at link time. The "thinning effect" in functor applications (on the argument signature) is achieved by the set of constraints generated by the *GenRec* algorithm. The inferred assumption environment A automatically records the "minimum" sharing constraints required to make the structure expression elaborate.

Def $W_{strexp}(se, V, M)$ = **case** se **of**

`x` $\Rightarrow$
  let $t \notin V,\ m \notin M,$
    $V' = V \cup \{t\},\ M' = M \cup \{m\}$
  in $(\{t{::}\texttt{STR}(m,\emptyset)\},\ t,\ \{\texttt{x} \mapsto t\},\ V',\ M')$

`s.a` $\Rightarrow$
  let $(K_1, u_1, A_1, V_1, M_1) = W_{strexp}(\texttt{s}, V, M)$
    $t_1, t_2 \notin V_1$, and $m_1, m_2 \notin M_1$
    $K_2 = K_1 \cup \{t_1{::}\texttt{STR}(m_1, \{\texttt{a} \mapsto t_2\})\}$
    $K_3 = K_2 \cup \{t_2{::}\texttt{STR}(m_2, \emptyset)\}$
    $V_2 = V_1 \cup \{t_1, t_2\},\ M_2 = M_1 \cup \{m_1, m_2\}$
    $(K_4, R) = KindUnify(K_3, \{(u_1, t_1)\})$
  in $(K_4,\ R(t_2),\ R(A_1),\ V_2,\ M_2)$

`f(s)` $\Rightarrow$
  let $(K_1, u_1, A_1, V_1, M_1) = W_{fctid}(\texttt{f}, V, M)$
    $(K_2, u_2, A_2, V_2, M_2) = W_{strexp}(\texttt{s}, V_1, M_1)$
    $t \notin V_2, m \notin M_2$
    $V_3 = V_2 \cup \{t\},\ M_3 = M_2 \cup \{m\}$
    $K_3 = K_1 \cup K_2 \cup \{t{::}\texttt{STR}(m, \emptyset)\}$
    $(K_4, R) = KindUnify(K_3, \{(u_1, u_2 \rightarrow t)\})$
  in $(K_4,\ R(t),\ R(A_1 \cup A_2),\ V_3,\ M_3)$

`struct d end` $\Rightarrow$
  let $(K_1, Env_1, A_1, V_1, M_1) = W_{dec}(\texttt{d})$
    $m \notin M_1, M_2 = M_1 \cup \{m\}$
  in $(K_1,\ \texttt{str}(m, Env_1),\ A_1,\ V_1,\ M_2)$

Def $W_{fctid}(f, V, M)$ = **case** f **of**

`x` $\Rightarrow$
  let $t_1, t_2 \notin V; m_1, m_2 \notin M$
    $V_1 = V \cup \{t_1, t_2\}; M_1 = M \cup \{m_1, m_2\}$
    $K = \{t_1{::}\texttt{STR}(m_1, \emptyset),\ t_2{::}\texttt{STR}(m_2, \emptyset)\}$
  in $(K, t_1 \rightarrow t_2, \{x \mapsto t_1 \rightarrow t_2\}, V_1, M_1)$

Def $W_{strdec}(sd, V, M)$ = **case** sd **of**

`structure s = se` $\Rightarrow$
  let $(K_1, u_1, A_1, V_1, M_1) = W_{strexp}(\texttt{se}, V, M)$
  in $(K_1,\ \{\texttt{s} \mapsto u_1\},\ A_1,\ V_1,\ M_1)$

Def $W_{dec}(d, V, M)$ = **case** d **of**

`sd` $\Rightarrow W_{strdec}(\texttt{sd}, V, M)$

`sd d` $\Rightarrow$
  let $(K_1, Env_1, A_1, V_1, M_1) = W_{strdec}(\texttt{sd}, V, M)$
    assume $Env_1 = \{\texttt{s} \mapsto u_1\}$
    $(K_2, Env_2, A_2, V_2, M_2) = W_{dec}(\texttt{d}, V_1, M_1)$
    assume $A_2(\texttt{s}) = \{t_1, ..., t_k\}$
    $A_3 = (A_1 \cup (A_2 \backslash \{\texttt{s}\}))$
    $P = \{(u_1, t_1), ..., (u_1, t_k)\}$
    $(K_3, R) = KindUnify(K_1 \cup K_2, P)$
  in $(K_3,\ R(Env_1 \pm Env_2),\ R(A_3),\ V_2,\ M_2)$

Def $GenRec(\alpha, \texttt{str}(m, Env), K, V)$
  let $KE = \emptyset$
    for each $\texttt{s} \in \mathrm{Dom}(Env)$
      $\beta \notin V$ and $V = V \cup \{\beta\}$
      $(K, V) = GenRec(\beta, Env(\texttt{s}), K, V)$
      $KE = KE \pm \{\texttt{s} \mapsto \beta\}$

    $K = K \cup \{\alpha :: \texttt{STR}(m, KE)\}$
  in (K, V)

Def $ModlMatch(B,T)$ =
  let $(N, GE, FE, SE) = B$
    $(K, u, A, V, M) = T$ and $P = \emptyset$
    for each $(x, t_x) \in A$ and $x \in$ StrId
      $P = P \cup \{(t_x, SE(x))\}$

    for each $(x, t_x) \in A$ and $x \in$ FunId
      assume $N_1(S_1, N_1'(S_1')) = FE(x)$
      $\alpha \notin V$ and $V = V \cup \{\alpha\}$
      assume $\{m_1, ..., m_k\} = N_1 \cup N_1'$
      $m_1', ..., m_k' \notin M$
      $M = M \cup \{m_1', ..., m_k'\}$
      $\varphi = \{m_i \mapsto m_i'$ where $i = 1, ..., k\}$
      $P = P \cup \{(\alpha \rightarrow \varphi(S_1'), t_x)\}$
      $(K, V) = GenRec(\alpha, \varphi(S_1), K, V)$

    $(K', R) = KindUnify(K, P)$
  in $R(u)$

Figure 4: Assumption Inference in *ModL*

Def $KindUnify(K, P) =$ **case** (K, P) **of**

$(K_1, \emptyset) \Rightarrow (K_1, ID)$
$(K_1, P_1 \cup \{(t, t)\}) \Rightarrow KindUnify(K_1, P_1)$
$(K_1, P_1 \cup \{(t_1 \rightarrow t_2, t_1' \rightarrow t_2')\}) \Rightarrow KindUnify(K_1, P_1 \cup \{(t_1, t_1'), (t_2, t_2')\})$

$(K_1 \cup \{t_1 :: \mathtt{STR}(m_1, Env_1)\}, P_1 \cup \{(t_1, \mathtt{str}(m_2, Env_2))\}) \Rightarrow$
  let check $\mathrm{Dom}(Env_1) \subseteq \mathrm{Dom}(Env_2)$, otherwise fail
    $(R_m, m) = NameUnify(m_1, m_2)$ and $Env_1' = R_m(Env_1)$ and $Env_2' = R_m(Env_2)$
    $R = (\{t_1 \mapsto \mathtt{str}(m, Env_2')\})(R_m)$ and $K_2 = R(K_1)$
    $P_2 = R(P_1) \cup \{(Env_1'(s), Env_2'(s)) \mid \forall s \in \mathrm{Dom}(Env_1)\}$
    $(K_3, R') = KindUnify(K_2, P_2)$
  in $(K_3, R' \circ R)$

$(K_1 \cup \{t_1 :: \mathtt{STR}(m_1, Env_1), t_2 :: \mathtt{STR}(m_2, Env_2)\}, P_1 \cup \{(t_1, t_2)\}) \Rightarrow$
  let $(R_m, m) = NameUnify(m_1, m_2)$ and $R = (\{t_1 \mapsto t_2\})(R_m)$
    $Env_1' = R(Env_1)$ and $Env_2' = R(Env_2)$ and $Env' = Env_2' \cup (Env_1' \setminus \mathrm{Dom}(Env_2'))$
    $K_2 = R(K_1) \cup \{t_2 :: \mathtt{STR}(m, Env')\}$
    $P_2 = R(P_1) \cup \{(Env_1'(s), Env_2'(s)) \mid \forall s \in \mathrm{Dom}(Env_1') \cap \mathrm{Dom}(Env_2')\}$
    $(K_3, R') = KindUnify(K_2, P_2)$
  in $(K_3, R' \circ R)$

$(K_1, P_1 \cup \{(\mathtt{str}(m_1, Env_1), \mathtt{str}(m_2, Env_2))\}) \Rightarrow$
  let check $\mathrm{Dom}(Env_1) = \mathrm{Dom}(Env_2)$, otherwise fail
    $(R_m, m) = NameUnify(m_1, m_2)$ and $Env_1' = R_m(Env_1)$ and $Env_2' = R_m(Env_2)$
    $K_2 = R_m(K_1)$ and $P_2 = R_m(P_1) \cup \{(Env_1'(s), Env_2'(s)) \mid \forall s \in Dom(Env_1')\}$
    $(K_3, R') = KindUnify(K_2, P_2)$
  in $(K_3, R' \circ R_m)$

Def $NameUnify(m_1, m_2) =$
    if $m_1 = m_2$ **then** (ID, m_1)
    else if $m_1, m_2 \in$ RigStrName **then** *fail*
    else if $m_1 \in$ RigStrName **then** $(\{m_2 \mapsto m_1\}, m_1)$
    else $(\{m_1 \mapsto m_2\}, m_2)$

Figure 5: Kinded unification algorithm

Figure 5 gives the unification algorithms *KindUnify* and *NameUnify*. The kinded unification algorithm *KindUnify* presented there extends the one in Ohori [24] with considerations on ML structure names. The following theorem can be proved in the same way as Ohori [24].

Theorem 3.1 *Given any kinded set of equations, the algorithm KindUnify computes a most general unifier if one exists and reports failure otherwise.*

The following lemma shows how the "thinning" effect is achieved in our algorithm.

Lemma 3.2 *Given a signature* $\Sigma = N(S)$ *and a structure* S'*, suppose that* S' *does not contain any flexible names; let* α *be a structure variable and* V_1 *be any set of structure variables; suppose that* $(K, V) = GenRec(\alpha, S, \emptyset, V_1)$*, then* $KindUnify(K, \{(\alpha, S')\})$ *succeeds if and only if the structure* S' *matches the signature* Σ*. Moreover, if* $R = KindUnify(K, \{(\alpha, S')\})$*, then* S' *enriches* $R(S)$*.*

The following theorem can be proved by structural induction on structure expressions.

Theorem 3.3 *Given an* ML *basis* B *and a structure expression strexp, then* $B \vdash strexp : S$ *succeeds if and only if both* $(K, u, A, V, M) = W_{strexp}(strexp, \emptyset, \text{RigStrName})$ *and* $S' = ModlMatch(B, (K, u, A, V, M))$ *succeed. Moreover, there exists a renaming realization* φ *such that* $S = \varphi(S')$*.*

The algorithm W_{strexp} possesses most properties that W^* has. It can also be modified to take a basis B as its argument (just as algorithm D) so that it can work more efficiently when we compile a structure expression in the pervasive basis.

4 Code Generation Issues

A compiling process usually contains two parts: elaboration (i.e., type inference or type-checking) and code generation (also code optimization). The assumption inference algorithms presented in the last two sections successfully solve the problem in the elaboration phase. In order to achieve the "smartest recompilation rule", our compiler should generate code that will be reusable as long as the surroundings satisfy the "minimum" import interface (i.e., match the assumption environment). This requires that our code generator should use no more type information than is specified in the "minimum" import interface.

Fortunately there are very few dependencies between the static semantics and the dynamic semantics in SML. Moreover, although Leroy's representation analysis [17] shows that the compiler can benefit a lot by using type information in the front end, the SML/NJ compiler [4] uses almost no type information in its back end but it still produces quite efficient code. In SML/NJ, the only things that the back end needs to know from the front end are the corresponding dynamic interface for each signature and the identifier status for each identifier. By delaying these dependencies to be resolved at link time, a program can be translated into machine code even before it is elaborated.

In the following, we only informally discuss the solutions to these issues.

Functor application In SML, a functor F with argument signature *SIG* can be applied to any structure S that *matches SIG*. A structure does not have to agree exactly with a signature in order for it to match the signature, instead it can contain more components than required. In such cases, signature matching

will coerce the structure against the signature, producing a "thinned" structure that exactly agrees with the signature in terms of number of components and their types. Suppose the corresponding dynamic code for the functor F and the structure S is fd and sd, the code generated for a function application $F(S)$ will be $fd(th(sd))$ where th is a thinning function from S to SIG. In our separate compilation scheme, it is possible that we still do not know the argument signature of F or the exact specification of structure S when we have to generate code for the functor application $F(S)$. This is simply solved by adding an abstraction on th, and the code becomes $\lambda th.(...fd(th(sd))...)$. The correct thinning function is filled in at link time when the real specifications of SIG and S are known.

Pattern matching In the SML/NJ compiler, the representation of a user defined datatype is determined by its definition. For example,

```
structure A = struct datatype color = RED | GREEN | BLUE | MIX of real * real * real
              end

structure B = struct fun redp(A.RED) = 1.0
                      | redp(A.MIX(x,_,_)) = x
                      | redp(_) = 0.0
              end
```

the datatype `color` in `A` may be represented with integer tags 0,1,2,3 for the data constructors `BLUE`, `GREEN`, `MIX`, and `RED`. However this imposes some problems if we want to separately compile structure `B`. What representations are we going to use for `A.RED` and `A.MIX` in the `redp` function? Again this is solved by making the representation of data constructors abstract (as in Aitken and Reppy's recent work [2]). A "constant" data constructor (such as `A.RED`) is compiled as a variable. A value carrying constructor (such as `A.MIX`) is compiled as a pair of injection and projection functions. These details are filled in at link time when the definition of the datatype is known.

Polymorphic equality functions Nothing needs to be done to support our separate compilation scheme if the equality function is implemented as it currently is in the SML/NJ compiler. In SML/NJ (as in all ML compilers to our knowledge), the polymorphic equality function is implemented as a runtime "equality interpreter" which checks equality of two objects based on their runtime tags. Another way to implement polymorphic equality, which is used in Haskell [11], is to pass an equality function for each formal parameter that is a polymorphic equality type variable. The code produced by this scheme closely depends on the derivation tree of the elaboration phase. In our separate compilation scheme, because the types of some external identifiers are not known at compile time, the derivation tree we get at compile time is not accurate. For example,

```
fun f x = S.g (3,x)
```

from assumption inference, we know `S.g`'s type must be in the form of $int * \alpha \rightarrow \beta$ and `f`'s in the form of $\alpha \rightarrow \beta$. Because `S.g` may want to test the equality on its 2nd argument, the function `f` here has to be implemented with an equality function for type α as its extra argument. This will have some runtime overhead in the common case that `S.g` actually never does equality test on its 2nd argument.

Representation analysis Leroy [17] presented a program transformation that allows polymorphic languages to be implemented with unboxed, multi-word data representation. The main idea is to introduce coercions between various representations based on the typing derivation tree. In our separate compilation system, accurate type information for external identifiers are not available at compile time, so the typing derivation tree is not very specific. However the representation analysis can still be carried out since all type instances of external identifiers are recorded in the assumption environment. At link time, the

matching algorithm *Match* will find out the accurate type information of all external identifiers and coerce them into different type instances in the assumption environment. For example, the above function `f` will be implemented as a polymorphic function $\alpha \rightarrow \beta$. When we find that `S.g` has type $int * int \rightarrow int$, `S.g` has to be coerced to type $int * \alpha \rightarrow \beta$ and `f` has to be coerced from $\alpha \rightarrow \beta$ to $int \rightarrow int$. The code produced in this way will be less efficient, but it should be acceptable if in practice there are not too many external identifiers in a module (especially when we use algorithm D).

Open declaration The `open` declaration in SML causes several nasty problems for our separate compilation scheme. For example, in the following structure declaration, there is no way to figure out the identifier status of `RED` without looking at the definition of structure `A`.

```
structure S = struct open A
                      fun redt(RED) = RED
                        | redt _ = BLUE
              end
```

Here `RED` can either be a data constructor if it is defined in a data type declaration in structure `A`, or `RED` could simply be a pattern variable if it is not defined elsewhere (this is one of the ugliest features in ML). A solution to this problem is to introduce a boolean status variable for each "ambiguous" identifier. This status variable specifies whether the identifier is a data constructor or a pattern variable. Assume that `status-RED` is the status variable of `RED`, then the above `redt` function can be rewritten as follows:

```
val redt = if status-RED then fn x => x   (* RED is not a constructor *)
           else (fn RED => RED            (* RED is a constructor *)
                 | _ => BLUE)
```

However this solution may cause code explosion if not careful. Another solution is to introduce a predicate for each ambiguous identifier. Assume that `pred-RED` is the corresponding predicate for `RED`, then at link time, `pred-RED` will be set to `fn x => true` if `RED` is a pattern variable, and `fn x => (x = RED)` otherwise. Then the above `redt` function can be rewritten as follows:

```
fun redt(x) = if pred-RED(x) then x
              else BLUE
```

This solution may incur some runtime overhead for the case that `RED` is a pattern variable. But in that case, the above code is written in such a bad style that it well deserves such a penalty. Other nasty problems related with `open` can be solved in the same way by resolving certain information at link time. These solutions may increase the complexity of compiling and linking but most of them do not incur any runtime overhead (however, they may stop some inline-expansion optimizations).

5 Implementation

We are currently prototyping a separate compilation system based on our algorithms into the SML/NJ compiler. In our system, a large ML program is composed of a set of top-level structure declarations, signature declarations and functor declarations. No two top-level structures (or signatures, functors) can have the same identifier name so that we can uniquely determine which definition each external identifier refers to. Every top-level declaration is considered as a compilation unit. Because signatures are usually small and compiling signature declarations does not take much time, their elaborations are delayed to be done at link time. To compile a structure or functor declaration, we apply the assumption inference algorithm to its body (which is always a structure expression), generate the machine code for the body, and then write both the inferred

interface (i.e., the assumption environment) and the machine code into its binary file. The final linking phase is done in certain order according to the dependency relation among different modules. This dependency relation has been already recorded in the inferred interface in each binary file. For each module, the linker simply reads the binary file, elaborates every signature expression, applies the matching algorithm (i.e., *Match* and *ModlMatch*) to recover the correct static environment and detect cross-module type errors if there are any, and then concatenates the machine code with correct thinning functions.

6 Related Work

Most dynamically-typed languages such as Lisp also allow independent compilations and can achieve the same kind of "smartest recompilation" in the sense that a module never needs to be recompiled unless its implementation changes. However, this is based on a big sacrifice: cross-module type errors will be detected only at runtime. Our method, however, will detect all cross-module type errors at link time.

Levy [18] presents a separate compilation method very similar to ours for PASCAL-like languages. Its compiler also automatically infers the import interface for each compilation unit. Cross-module type errors are reported at link time. However he does not mention whether he achieves the smartest recompilation rule, and the type systems of PASCAL-like languages are much simpler than that of SML.

Traditional separate compilation systems adopted in most statically typed languages all use manually created interface files. Each compilation unit contains an implementation plus several interface files. It has to be fully closed up to the pervasive basis so that the specifications of all external symbols will be found at compile time. The `make` system [9] is the simplest one along this line. It will trigger recompilations if the interface file a module depends on changes. Tichy [30] and Schwanke [29] eliminated most recompilations by examining finer-levels of dependency relations between interfaces and implementations. In their methods, if the interface file a module depends on changes, but the set of symbols the module imports does not change, then the module does not need to be recompiled. SRC Modula-3 [22, 14] implements exactly the same idea: a version stamp which encodes the specification of a symbol is produced for each exported symbol in an interface; modules import the version stamps of the symbols that they import; a module only needs to be recompiled if any of its imported version stamps are no longer exported. Languages with very powerful module systems such as Mesa [21], the System Modeller in Cedar [15], and FX-87 [8] also adopt similar separate compilation methods which are only applicable to closed modules. The compiler for Russell [5] does partially support separate compilations on "open-formed" expressions, however its "module system" is very restrictive and all "modules" must be loaded and compiled in an order determined by their dependencies. In summary, these previous methods cannot achieve the smartest recompilation rule, neither can they be applied to compile open-formed modules in SML.

In SML, two kinds of separate compilation methods have been proposed: Rothwell and Tofte's `import` scheme [28] and Rollins's `SourceGroup` scheme [27]; both methods apply only to closed functors. Recently, Emden Gansner [10] is implementing a `make`-like separate compilation system for open-formed modules in the interactive SML/NJ compiler. In his method, all modules are loaded and compiled in a top level environment in an order determined by their dependencies. Whenever a module is compiled, a new time stamp is generated; both the binary and the time stamp are then written out to the binary file. A module has to be recompiled whenever its source changes or any of its predecessors (in the dependency graph) have been recompiled. Gansner is also planning to export the static semantics of each module into the binary file so that redundant recompilations can be detected and avoided if the static semantics of a module has not been changed.

Aditya and Nikhil [1] have been working on similar kinds of assumption inference algorithms for their incremental compiler for Id [23]. However as far as we know, their algorithm does not infer the minimum constraints, thus fails to achieve our theorem 2.1. Because their system allows mutually recursive top-level declarations, it cannot fully recover the correct type information by simply using our assumption inference and matching algorithm. In the SML module language, however, top-level declarations cannot be mutually recursive.

Damas [6] gave an inference algorithm called T which is very similar to our W^* in section 2. His type system permits that a variable can be bound to several distinct types in the type environment (just like our assumption environment). However since he mainly used the system to handle overloading, he did not try to prove our theorem 2.1. The soundness and syntactic completeness results he proved for T are only for his particular type system, not for the usual ML type system [31], so they are not in the same sense as our corollary 2.2 and 2.3. The algorithm V in Leivant [16] is just Damas's T restricted to the type system without ML-polymorphism. Its extension V_2 is for the polymorphic discipline of rank 2 and the relation between W and V_2 is not clear. On the side of the SML module language, Aponte [3] presented a type checking algorithm for *ModL* based on Remy's approach to record typing [25]. Her approach is very elegant; however, in practice it is probably very difficult to implement efficiently. It is also not clear whether her algorithm can be modified to do our assumption inference.

7 Concluding Remarks

We have presented a separate compilation method that achieves the "smartest recompilation rule" for open-formed modules in Standard ML. In our method, each module is compiled independently without knowing the specifications of its external identifiers; its import interface, instead, is inferred by looking at how each external identifier is used inside the module. Cross-module type inconsistencies are detected at link time by simply matching the real specifications against the inferred import interface (this process should be very fast because it only involves an unification of a set of types). The independent compilation of each module may disable some inter-module optimizations, but we believe that the code generated by our recompilation method will be comparable to the quite efficient code generated by the current SML/NJ compiler [4]. We plan to implement and measure our algorithm in SML/NJ in the future.

The smartest recompilation technique in this paper is presented in the framework of SML; however, it can be easily applied to other polymorphic languages based on the Damas-Milner type discipline. It should be straightforward to extend the algorithm D in section 2.2 to work on the extension of ML type system with parametric overloadings [13]. The assumption inference algorithms presented in section 2 and 3 can also be used as a basis to build incremental compilers for similar languages. On the other hand, we still do not know how to extend the algorithm for *ModL* to work on the extension of ML module system with high-order functors [32].

The smartest recompilation technique should also be applicable to languages in the Algol family. The type system in those languages are much simpler than that in ML, so it is not hard to infer the "minimum" import interface for each module. However, the code produced by smartest recompilation will be less efficient because the code generators of these languages usually rely much more on the inter-procedure type and data flow information than those of polymorphic languages. For applications where reusability and reconfiguration are more important than efficiency, the smartest recompilation property is still very desirable.

Acknowledgements

We would like to thank William Aitken, Carl Gunter, and QingMing Ma for many valuable comments on an early version of this paper. We are also grateful to David MacQueen and Pierre Cregut for interesting discussions on related subjects. This research is supported by the National Science Foundation Grant CCR-9002786 and CCR-9200790, and by the first author's summer research internship at AT&T Bell Laboratories.

References

[1] Shail Aditya and Rishiyur S. Nikhil. Incremental polymorphism. In *The Fifth International Conference on Functional Programming Languages and Computer Architecture*, pages 378–405, New York, August 1991. Springer-Verlag.

[2] William E. Aitken and John H. Reppy. Abstract value constructors. In *ACM SIGPLAN Workshop on ML and its Applications*, June 1992.

[3] Maria Virginia Aponte. *Typage d'un systeme de modules parametriques avec partage: une application de l'unification dans les theories equationnelles.* PhD thesis, Universite de Paris, February 1992.

[4] Andrew W. Appel and David B. MacQueen. Standard ML of New Jersey. In Martin Wirsing, editor, *Third Int'l Symp. on Prog. Lang. Implementation and Logic Programming*, pages 1–13, New York, August 1991. Springer-Verlag.

[5] Hans Boehm and Alan J. Demers. Implementing Russell. In *Symposium on Compiler Construction*, pages 186–195. ACM Sigplan, June 1986.

[6] Luis Damas. *Type Assignment in Programming Languages.* PhD thesis, University of Edinburgh, Department of Computer Science, Edinburgh, UK, 1985.

[7] Luis Damas and Robin Milner. Principal type-schemes for functional programs. In *Ninth Annual ACM Symp. on Principles of Prog. Languages*, New York, Jan 1982. ACM Press.

[8] David K. Gifford et al. FX-87 reference manual. Technical Report MIT/LCS/TR-407, M.I.T. Laboratory for Computer Science, September 1987.

[9] Stuart I. Feldman. Make – a program for maintaining computer programs. *Software – Practice and Experience*, 9(4):255–265, Apirl 1979.

[10] Emden R. Gansner. AT&T Bell Labs, personal communication, 1992.

[11] Paul Hudak, Simon Peyton Jones, and Philip Wadler *et al.* Report on the programming language Haskell a non-strict, purely functional language version 1.2. *SIGPLAN Notices*, 21(5), May 1992.

[12] Mark P. Jones. A theory of qualified types. In *The 4th European Symposium on Programming*, pages 287–306, Berlin, February 1992. Spinger-Verlag.

[13] Stefan Kaes. Type inference in the presence of overloading, subtyping and recursive types. In *1992 ACM Conference on Lisp and Fucntional Programming*, New York, June 1992. ACM Press.

[14] Bill Kalsow and Eric Muller. SRC Modula-3 version 1.6 manual, February 1991.

[15] Butler W. Lampson and Eric S. Schmidt. Pratical use of a polymorphic applicative language. In *Tenth Annual ACM Symp. on Principles of Prog. Languages*, New York, Jan 1983. ACM Press.

[16] Daniel Leivant. Polymorphic type inference. In *Tenth Annual ACM Symp. on Principles of Prog. Languages*, New York, Jan 1983. ACM Press.

[17] Xavier Leroy. Unboxed objects and polymorphic typing. In *Nineteenth Annual ACM Symp. on Principles of Prog. Languages*, New York, Jan 1992. ACM Press.

[18] Michael R. Levy. Type checking, separate compilation and reusability. *SIGPLAN Notices (Proc. Sigplan '84 Symp. on Compiler Construction)*, 19(6):285–289, June 1984.

[19] Robin Milner and Mads Tofte. *Commentary on Standard ML.* MIT Press, Cambridge, Massachusetts, 1991.

[20] Robin Milner, Mads Tofte, and Robert Harper. *The Definition of Standard ML.* MIT Press, Cambridge, Massachusetts, 1990.

[21] J. Mitchell, W. Maybury, and R. Sweet. Mesa language manual. Technical Report CSL-79-3, Xerox Palo Alto Research Center, Palo Alto, CA, 1979.

[22] Greg Nelson, editor. *Systems programming with Modula-3.* Prentice Hall, Englewood Cliffs, NJ, 1991.

[23] Rishiyur S. Nikhil. Id version 90.0 reference manual. Technical Report TR-CSG-Memo 284-1, MIT Laboratory for Computer Science, 1990.

[24] Atsushi Ohori. A compilation method for ML-style polymorphic record calculi. In *Nineteenth Annual ACM Symp. on Principles of Prog. Languages*, New York, Jan 1992. ACM Press.

[25] Didier Remy. Typechecking records and variants in a natural extension of ML. In *Sixteenth Annual ACM Symp. on Principles of Prog. Languages*, pages 77–87, New York, Jan 1989. ACM Press.

[26] J. Robinson. A machine-oriented logic based on the resolution principle. *Journal of the ACM*, 12(1):23–41, 1965.

[27] Eugene J. Rollins. SourceGroup: A selective recompilation system for SML. In *Third International Workshop on Standard ML*, Pittsburgh, September 1991. Carnegie Mellon University.

[28] Nick Rothwell and Mads Tofte. `Import` command source code. with Standard ML of New Jersey releases 0.65.

[29] Robert W. Schwanke and Gail E. Kaiser. Smarter recompilation. *ACM Transactions on Programming Languages and Systems*, 10(4):627–632, October 1988.

[30] Walter Tichy. Smart recompilation. *ACM Transactions on Programming Languages and Systems*, 8(3):273–291, July 1986.

[31] Mads Tofte. *Operational Semantics and Polymorphic Type Inference.* PhD thesis, University of Edinburgh, Edinburgh, UK, November 1987.

[32] Mads Tofte. Principal signatures for high-order ML functors. In *Nineteenth Annual ACM Symp. on Principles of Prog. Languages*, New York, Jan 1992. ACM Press.

8 Appendix: Proof of Theorem 2.1

The proof of theorem 2.1 is organized in five parts: the first part introduces some notations and lemmas used later in the proofs; the rest four parts give the detailed proofs.

8.1 Notations and Lemmas

Before we proceed, let's introduce some new notations and lemmas:

Definition 8.1 *A* substitution *is a finite map from type variables to types. The* domain *of a substitution S, denoted as* Dom*(S), is the set of type variables α such that $S(\alpha) \neq \alpha$. The* region *of a substitution S, denoted as* Reg*(S), is the union of all tyvars$(S(\alpha))$ for all $\alpha \in$* Dom*(S). A substitution is called a* variable-pure *substitution if for all $\alpha \in$* TyVar*, $S(\alpha)$ is a type variable. The* addition *of two substitutions S_1 and S_2, denoted as $S_1 + S_2$, is "$S(\alpha) =$ **if** $\alpha \in Dom(S_1)$ **then** $S_1(\alpha)$ **else** $S_2(\alpha)$." Assume that D is a set of type variables, the substitution $S{\downarrow}D$ denotes the substitution S with its domain restricted to $D \cap$* Dom*(S).*

We define an *object* to be either a type, a type scheme, a type environment, or a tuple of other objects. The set of type variables occurred free in an object A is denoted as *tyvars*(A).

Definition 8.2 *Given an object A, a* renaming substitution *for this object is a variable-pure substitution $\{\alpha_i \mapsto \beta_i \mid i = 1, ..., n\}$ such that $\{\alpha_i \mid i = 1, ..., n\} \subseteq tyvars(A)$, the β_is are distinct and $(tyvars(A) \setminus \{\alpha_i \mid i = 1, ..., n\}) \cap \{\beta_i \mid i = 1, ..., n\} = \emptyset$.*

Definition 8.3 *Two types τ_1 and τ_2 are called* variants *if there exist substitutions S_1 and S_2 such that $\tau_1 = S_1(\tau_2)$ and $\tau_2 = S_2(\tau_1)$, we can also define variants relations similarly on type schemes, type environments, or objects with same shapes. The variant relation is denoted as $\cong$.*

The variants relation is obviously reflexive, symmetric and transitive, thus it is an equivalence relation. "Variants" is very important in comparing two objects because of the following lemma:

Lemma 8.1 *If two objects A and B are variants, then there exist a renaming substitution S for A and a renaming substitution R for B such that $B = S(A)$ and $A = R(B)$ and $S \circ R = R \circ S = ID$.*

From lemma 8.1, to prove the second part of the theorem 2.1, we only need show that $(S(TE), \tau)$ and $(S^*(TE), S^*\tau^*)$ are variants. Here are some other lemmas we are going to use in the proof:

Lemma 8.2 *Given a finite set of pairs of types P, suppose both S_1 and S_2 are the most general unifiers of P, then for any objects A, $S_1(A)$ and $S_2(A)$ are variants.*

Lemma 8.3 *Given a finite set of pairs of types P, then the following statements are true:*

- *If $Unify(P)$ succeeds, it will return a mgu S with $\mathrm{Reg}(S) \subseteq tyvars(P)$.*
- *If $Unify(P)$ succeeds and $P_1 \subseteq P$, then $Unify(P_1)$ also succeeds; moreover, if $P = P_1 \cup P_2$ and S_1 is a mgu of P_1, then $Unify(S_1(P_2))$ also succeeds.*
- *Suppose $P = P_1 \cup P_2$, $Unify(P_1)$ succeeds and S_1 is a mgu of P_1, and $Unify(S_1(P_2))$ succeeds and S_2 is a mgu of $S_1(P_2)$, then $Unify(P)$ succeeds and $(S_2 \circ S_1)$ is a mgu of P.*
- *Given a renaming substitution R for P, suppose R' is an inverse substitution of R in the sense that $R \circ R' = R' \circ R = ID$, then $Unify(P)$ succeeds if and only if $Unify(R(P))$ succeeds. Moreover, if S is a mgu of P then $(S \circ R')$ is a mgu of $R(P)$.*

In the figure 6 we review Tofte [31]'s version of the classical ML type inference algorithm W. The following lemma can easily be proved by structural induction on the expression e.

Lemma 8.4 *Given TE and e, if $(S, \tau) = W(TE, e)$ succeeds, and (S', τ') is simply the result of another run of the algorithm W on TE and e by choosing different new type variables, then $(S(TE), \tau)$ and $(S'(TE), \tau')$ are variants.*

We use TyVar' to denote the set of new type variables used in running $W(TE, e)$ plus those type variables occurred free in TE. Then we let W^* and *Match* only use new type variables from $\mathrm{TyVar} \setminus \mathrm{TyVar}'$. From lemma 8.4, whichever set of new type variables W is using, the resulting $(S(TE), \tau)$ is always a variant of each other. Thus it suffices to prove that theorem 2.1 is true when W and W^* are using different set of new type variables. In the next for sections, we show it by structural induction on the expression e in *Exp*.

Def $W(TE, e)$ = **case** e **of**

x $\Rightarrow$ **let** $TE(x) = \forall\overline{\alpha_n}.\tau$ and $\overline{\beta_n}$ be new type variables and $S = \{\alpha_i \mapsto \beta_i$ for $i = 1, ..., n\}$
in $(ID\ ,\ S(\tau))$

$\lambda x.e_1$ $\Rightarrow$ **let** α be a new type variable and $(S_1, \tau_1) = W(TE \pm \{x \mapsto \alpha\}, e_1)$
in $(S_1\ ,\ (S_1(\alpha) \rightarrow \tau_1))$

$e_1 e_2$ $\Rightarrow$ **let** $(S_1, \tau_1) = W(TE, e_1)$ and $(S_2, \tau_2) = W(S_1(TE), e_2)$
α be a new type variable and $S_3 = Unify(S_2\tau_1, \tau_2 \rightarrow \alpha)$
in $(S_3 \circ S_2 \circ S_1\ ,\ (S_3\alpha))$

`let` x `=` e_1 `in` e_2 $\Rightarrow$
let $(S_1, \tau_1) = W(TE, e_1)$ and $(S_2, \tau_2) = W(S_1(TE) \pm \{x \mapsto gen(S_1(TE), \tau_1)\}, e_2)$
in $(S_2 \circ S_1\ ,\ \tau_2)$

Figure 6: The classical ML Type Inference Algorithm W

8.2 Case 1: variable $\boxed{e = x}$

(Part 1) $W(TE, x)$ will succeed only when $x \in \text{Dom}(TE)$, in which case both $(\tau^*, A) = W^*(x)$ and $Match(TE, A^*)$ will succeed, *vice versa*.

(Part 2) Now suppose $(\tau^*, A) = (\gamma, \{x \mapsto \gamma\})$ and $TE(x) = \forall\overline{\alpha_n}.\tau$, then $W(TE, x) = (ID, S_1'\tau)$ where $S_1' = \{\alpha_i \mapsto \beta_i \mid i = 1, ..., n\}$ and all β_i are new type variables. On the other hand, $S^* = Match(TE, A) = \{\beta \mapsto S_2'\tau\}$ where $S_2' = \{\alpha_i \mapsto \gamma_i \mid i = 1, ..., n\}$ and all γ_i are new type variables. Thus let's choose $R_1 = \{\beta_i \mapsto \gamma_i\}$ and $R_2 = \{\gamma_i \mapsto \beta_i\}$, then obviously $(S^*\ (TE),\ S^*\beta)$ and $(ID\ (TE),\ S_1'\tau)$ are variants through R_1 and R_2.

8.3 Case 2: lambda abstraction $\boxed{e = \lambda x.e_1}$

(Part 1) Suppose $W(TE, e)$ succeeds, we want to prove that both $(\tau^*, A) = W^*(\lambda x.e_1)$ and $Match(TE, A)$ will also succeed.

Let α be a new type variable used in W, then $(S_1, \tau_1) = W(TE \pm \{x \mapsto \alpha\}, e_1)$ should also succeed. Therefore by induction, both $(\tau_1^*, A_1) = W^*(e_1)$ and $S_1^* = Match(TE \pm \{x \mapsto \alpha\}, A_1)$ will succeed too. Let the set of type assumptions for x in A_1 be $\{x \mapsto t_i$ where $i = 1, ..., k\}$ and let β be a new type variable used in W^* and let $S_0^* = MonoUnify(\beta, A_1(x))$, then obviously S_0^* is simply $Unify(\{(\beta, t_i) \mid i = 1, ..., n\})$. Let $P = \{\alpha \mapsto \beta\}$ and $P' = \{\beta \mapsto \alpha\}$, by lemma 8.3, because both $P(TE) = TE$ and $P'(A_1) = A_1$ are true, $Match(TE \pm \{x \mapsto \beta\}, A_1)$ will also succeed and $S_1^* \circ P'$ is one of its mgu. On the other hand, because $Match(TE \pm \{x \mapsto \alpha\}, A_1)$ succeeds, by lemma 8.3, $Unify(\{(\alpha, t_i) \mid i = 1, ..., n\})$ should also succeed. Thus $S_0^* = Unify(\{(\beta, t_i) \mid i = 1, ..., n\})$ will succeed too. Thus because $S_0^*(TE) = TE$ is true, by lemma 8.3, $S^* = Match(TE, S_0^*(A_1 \setminus \{x\}))$ will also succeed. The reverse direction can be proved in the same way using lemma 8.3.

(Part 2) Let's still use the notations in part 1, by induction, we know that there exists two substitutions

R_{11} and R_{12} such that the following are true:

- $R_{11} \circ R_{12} = R_{12} \circ R_{11} = ID$,
- $S_1(TE \pm \{x \mapsto \alpha\}) \xrightarrow{R_{11}} S_1^*(TE \pm \{x \mapsto \alpha\})$ and $\tau_1 \xrightarrow{R_{11}} S_1^*(\tau_1^*)$,
- $S_1^*(TE \pm \{x \mapsto \alpha\}) \xrightarrow{R_{12}} S_1(TE \pm \{x \mapsto \alpha\})$ and $S_1^*(\tau_1^*) \xrightarrow{R_{12}} \tau_1$.

Then we can deduce the following by using lemma 8.3 and lemma 8.2,

$$\begin{aligned}
S^*(TE, S_0^*(\beta \to \tau_1^*)) &\cong (Match(TE, S_0^*(A_1 \setminus \{x\})))(TE, S_0^*(\beta \to \tau_1^*)) \\
&\cong (Match(S_0^*(TE), S_0^*(A_1 \setminus \{x\})) \circ S_0^*)\ (TE, \beta \to \tau_1^*) \\
&\cong (Match(TE \pm \{x \mapsto \beta\}, A_1))\ (TE, \beta \to \tau_1^*) \\
&\cong (Match(TE \pm \{x \mapsto P(\alpha)\}, A_1))\ (TE, \beta \to \tau_1^*) \\
&\cong (Match(TE \pm \{x \mapsto \alpha\}, A_1) \circ P')\ (TE, \beta \to \tau_1^*) \\
&\cong (Match(TE \pm \{x \mapsto \alpha\}, A_1))\ (TE, \alpha \to \tau_1^*) \\
&\cong (S_1^*(TE),\ S_1^*(\alpha \to \tau_1^*)) \\
&\cong (S_1(TE), (S_1\ \alpha \to \tau_1))
\end{aligned}$$

Thus we prove that $S^*(TE, S_0^*(\beta \to \tau_1^*))$ and $(S_1\ TE, (S_1\ \alpha \to \tau_1))$ are variants.

8.4 Case 3: application $\boxed{e = e_1\ e_2}$

(Part 1) Suppose $W(TE, e)$ succeeds, we want to prove that both (τ^*, A)=$W^*(e_1\ e_2)$ and $Match(TE, A)$ will also succeed.

Because $W(TE, e)$ succeeds, both $(S_1, \tau_1) = W(TE, e_1)$ and $(S_2, \tau_2) = W(S_1(TE), e_2)$ will also succeed. Let α be a new type variable used in W, then $S_3 = Unify(S_2(\tau_1), \tau_2 \to \alpha)$ will succeed. By induction we know that both (τ_1^*, A_1)=$W^*(e_1)$ and (τ_2^*, A_2)=$W^*(e_2)$ will succeed. The corresponding unifications $S_1^* = Match(TE, A_1)$ and S_2^*=$Match(S_1(TE), A_2)$ will succeed too. Moreover there exists four (renaming) substitutions R_{11}, R_{12}, R_{21} and R_{22} such that the following are true:

- $R_{11} \circ R_{12} = R_{12} \circ R_{11} = ID$, and $R_{21} \circ R_{22} = R_{22} \circ R_{21} = ID$.
- $S_1(TE) \xrightarrow{R_{11}} S_1^*(TE)$ and $\tau_1 \xrightarrow{R_{11}} S_1^*(\tau_1^*)$,
- $S_1^*(TE) \xrightarrow{R_{12}} S_1(TE)$ and $S_1^*(\tau_1^*) \xrightarrow{R_{12}} \tau_1$,
- $S_2(S_1(TE)) \xrightarrow{R_{21}} S_2^*(S_1(TE))$ and $\tau_2 \xrightarrow{R_{21}} S_2^*(\tau_2^*)$,
- $S_2^*(S_1(TE)) \xrightarrow{R_{22}} S_2(S_1(TE))$ and $S_2^*(\tau_2^*) \xrightarrow{R_{22}} \tau_2$,

First we prove $R_{22}(S_2^*(\tau_1)) = S_2(\tau_1)$. Suppose TT is the set of pairs of types to be unified in running $Match(S_1(TE), A_2)$, TT will contain type variables from $tyvars(S_1(TE)) \cup tyvars(A_2)$ and some new type variables. From the definition of a mgu, S_2^*=$Unify(TT)$ should be the identity on all type variables except those in $tyvars(TT)$. Because we are running Robinson's original unification algorithm, Reg(S_2^*) is a subset of $tyvars(TT)$. Similarly from the definition of renaming substitutions, R_{22} should be the identity on all type

variables except those in $\mathrm{Reg}(S_2^*)$. We also know S_2 will be the identity on $tyvars(\tau_1) \setminus tyvars(S_1(TE))$ because $\mathrm{Dom}(S_2)$ will be $tyvars(S_1(TE))$ plus some new type variables. Thus we have $R_{22}(S_2^*(\tau_1)) = S_2(\tau_1)$.

Now because $S_3 = Unify(S_2(\tau_1), \tau_2 \to \alpha)$ succeeds, $Unify(R_{22}(S_2^*(\tau_1)), R_{22}(S_2^*((\tau_2^* \to \alpha))))$ should also succeed. By lemma 8.3, $Unify(S_2^*\tau_1, S_2^*(\tau_2^* \to \alpha))$ will succeed too and $(S_3 \circ R_{22})$ is one of its mgu.

Because $S_2^* = Match(S_1(TE), A_2)$ succeeds, by lemma 8.3, $Match(S_1(TE) \pm \{x \mapsto \tau_1\}, A_2 \cup \{x \mapsto (\tau_2^* \to \alpha)\})$ will also succeed (here x is used to ease the notation; it is an unused program variable). Moreover also by lemma 8.3, $Q_1 = (S_3 \circ R_{22}) \circ S_2^*$ will be one of its mgu.

Because we are using different new type variables in W and W^*, and the call of $W^*(e_1)$ and $W^*(e_2)$ are using totally different type variables, the following equations are obviously true: $R_{12}S_1^*(TE) = S_1(TE)$ and $R_{12}S_1^*(A_2) = A_2$ and $R_{12}S_1^*\tau_2^* = \tau_2^*$ and $R_{12}S_1^*\tau_1^* = \tau_1$ and $R_{12}S_1^*\alpha = \alpha$.

Because $Match(TE \pm \{x \mapsto R_{12}S_1^*\tau_1^*\}, A_2 \cup \{x \mapsto (\tau_2^* \to \alpha)\})$ succeeds and Q_1 is one of its mgu, by lemma 8.3, $Match(TE \pm \{x \mapsto \tau_1^*\}, A_1 \cup A_2 \cup \{x \mapsto (\tau_2^* \to \alpha)\})$ will also succeed. Moreover $Q_2 = (Q_1 \circ R_{12}) \circ S_1^*$ is one of its mgu.

Let β be the new type variable used in W^*, let's define $P = \{\alpha \mapsto \beta\}$ and $P' = \{\beta \mapsto \alpha\}$. Again by lemma 8.3, $Match(TE \pm \{x \mapsto \tau_1^*\}, A_1 \cup A_2 \cup \{x \mapsto (\tau_2^* \to \beta)\})$ will also succeed. Moreover, $Q_2 \circ P'$ will be one of its mgu. Thus again by lemma 8.3, $S_3^* = Unify(\{(\tau_1^*, \tau_2^* \to \beta)\})$ also succeeds. Let's assume S_3^* be the mgu produced when running W^*. Thus because $S_3^*(TE) = TE$ is true, $S^* = Match(TE, S_3^*(A_1 \cup A_2)$ also succeeds.

The reverse direction can be proved in the same way as above by keeping applying lemma 8.3.

(Part 2) Let's still use the notations introduced in part 1, we can get the following by using lemma 8.3 and lemma 8.2:

$$
\begin{aligned}
S^*(TE, S_3^*(\beta)) &\cong (Match(TE, S_3^*(A_1 \cup A_2)))(TE, S_3^*(\beta)) \\
&\cong (Match(S_3^*(TE), S_3^*(A_1 \cup A_2)))(S_3^*(TE), S_3^*(\beta)) \\
&\cong (Match(TE \pm \{x \mapsto \tau_1^*\}, A_1 \cup A_2 \cup \{x \mapsto (\tau_2^* \to \beta)\}))\ (TE, \beta) \\
&\cong (Q_2 \circ P')\ (TE, \beta) \\
&\cong Q_2(TE, \alpha) \\
&\cong ((Q_1 \circ R_{12}) \circ S_1^*)(TE, \alpha) \\
&\cong Q_1(R_{12}(S_1^*(TE)), R_{12}(S_1^*(\alpha))) \\
&\cong Q_1(S_1(TE), \alpha) \\
&\cong ((S_3 \circ R_{22}) \circ S_2^*)(S_1(TE), \alpha) \\
&\cong S_3(R_{22}(S_2^*(S_1(TE))), R_{22}(S_2^*(\alpha))) \\
&\cong S_3(S_2S_1(TE), \alpha) \\
&\cong (S_3S_2S_1(TE), S_3\alpha)
\end{aligned}
$$

Thus we prove that $S^*(TE, S_3^*(\beta))$ and $(S_3S_2S_1(TE), S_3\alpha)$ are variants.

8.5 Case 4: let expression $\boxed{e = \textbf{let } x = e_1 \textbf{ in } e_2}$

(Part 1) We want to prove that if $W(TE, e)$ succeeds, then both $(\tau^*, A) = W^*(e)$ and $Match(TE, A)$ will also succeed. Suppose $W(TE, e)$ succeeds, then both $(S_1, \tau_1) = W(TE, e_1)$ and $(S_2, \tau_2) = W(S_1(TE) \pm \{x \mapsto gen(S_1(TE), \tau_1\}, e_2)$ will also succeed. By induction, both $(\tau_1^*, A_1) = W^*(e_1)$, $(\tau_2^*, A_2) = W^*(e_2)$ and $S_1^* =$

$Match(TE, A_1)$ and $S_2^* = Match(S_1(TE) \pm \{x \mapsto gen(S_1(TE), \tau_1)\}, A_2)$ will succeed too. Moreover, there exists four renaming substitutions R_{11}, R_{12}, R_{21} and R_{22} such that the following is true:

(1) $R_{11} \circ R_{12} = R_{12} \circ R_{11} = ID$, and $R_{21} \circ R_{22} = R_{22} \circ R_{21} = ID$;
(2) $S_1(TE) \xrightarrow{R_{11}} S_1^*(TE)$ and $\tau_1 \xrightarrow{R_{11}} S_1^*(\tau_1^*)$; $S_1^*(TE) \xrightarrow{R_{12}} S_1(TE)$ and $S_1^*(\tau_1^*) \xrightarrow{R_{12}} \tau_1$;
(3) $S_2(S_1(TE) \pm \{x \mapsto gen(S_1(TE), \tau_1)\}) \xrightarrow{R_{21}} S_2^*(S_1(TE) \pm \{x \mapsto gen(S_1(TE), \tau_1)\})$ and $\tau_2 \xrightarrow{R_{21}} S_2^*(\tau_2^*)$;
(4) $S_2^*(S_1(TE) \pm \{x \mapsto gen(S_1(TE), \tau_1)\}) \xrightarrow{R_{22}} S_2(S_1(TE) \pm \{x \mapsto gen(S_1(TE), \tau_1)\})$ and $S_2^*(\tau_2^*) \xrightarrow{R_{22}} \tau_2$.

First we define the following notations: let the type assumptions for x in A_2 be $A_2(x) = \{t_i \mid i = 1, ..., k\}$; let $tyvars(A_1)$ be $TV_{A_1} = \{\alpha_i \mid i = 1, ..., m\}$ and $tyvars(\tau_1^*)$ be $TV_{\tau_1^*} = \{\alpha_{i_j} \mid 1 \leq i_j \leq m,\ j = 1, ..., u$ and $u \leq m\} \cup \{\eta_i \mid i = 1, ..., v\}$; we assume that τ_1^* contains some type variables α_{i_j} from TV_{A_1} and some new type variables η_i. Assume that α_{ij}, where $j = 1, ..., k$ and $i = 1, ..., m$, are $m * k$ new type variables, for each $j = 1, ..., k$, we define a renaming substitution C_j for A_1 where $C_j = \{\alpha_i \mapsto \alpha_{ij} \mid i = 1, ..., m\}$, and a substitution $C_j' = \{\alpha_{ij} \mapsto \alpha_i \mid i = 1, ..., m\}$. Let $C' = C_1' + ... + C_k'$. Assume that η_{ij}, where $j = 1, ..., k$ and $i = 1, ..., v$, are $v * k$ new type variables, for each $j = 1, ..., k$, we define a renaming substitution H_j for τ_1^* where $L_j = \{\eta_i \mapsto \eta_{ij} \mid i = 1, ..., v\}$. Also we denote $H_j = C_j + L_j$ for each $j = 1, ..., k$.

Using the above notations, $(S_3^*, A_3) = PolyUnify((TV_{\tau_1^*} \cup TV_{A_1}, \tau_1^*, A_1), A_2(x))$ is equivalent to that $S_3^* = Unify(\{(H_j\tau_1^*, t_j) \mid j = 1, ..., k\})$ and $A_3 = C_1A_1 \cup ... \cup C_kA_1$. Therefore, to prove that both $S_3^* = Unify(\{(H_j\tau_1^*, t_j) \mid j = 1, ..., k\})$ and $S^* = Match(TE, A_1 \cup S_3^*(A_3 \cup A_2 \setminus \{x\}))$ succeed, by lemma 8.3, we simply need to prove that $Match(TE \pm \{x^{(j)} \mapsto H_j\tau_1^* \mid j = 1, ..., k\}, A_1 \cup C_1A_1 \cup ... \cup C_kA_1 \cup A_2)$ succeeds. Here we informally use $x^{(j)}$ to mean the k different instances of x, corresponding to those k assumptions t_j in A_2 for x.

We are going to prove that $Match(TE \pm \{x^{(i)} \mapsto H_i\tau_1^* \mid i = 1, ..., k\}, A_1 \cup C_1A_1 \cup ... \cup C_kA_1 \cup A_2)$ succeeds in the following two steps:

Step 1.1 First, we prove that the unification $S^+ = Match(TE, A_1 \cup C_1A_1 \cup ... \cup C_kA_1)$ succeeds. We shall construct such mgu S^+ from $S_1^* = Match(TE, A_1)$.

Assume that when doing $Match(TE, A_1)$, the set of new type variables introduced by instantiating bound variables in TE is $\{\beta_i \mid i = 1, ..., n\}$. Let β_{ij}, where $i = 1, ..., n$ and $j = 1, ..., k$, be $n * k$ new type variables. All these type variables virtually corresponds to different instantiations of bound variables in TE when doing $Match(TE, A_1 \cup C_1A_1 \cup ... \cup C_kA_1)$. Let's define a substitution B_j for each $j = 1, ..., k$ such that $B_j = \{\beta_i \mapsto \beta_{ij} \mid i = 1, ..., n\}$. We also define substitutions $B_j' = \{\beta_{ij} \mapsto \beta_i \mid i = 1, ..., n\}$ for each $j = 1, ..., k$ and $B' = B_1' + ... + B_k'$. Let's denote the set of free type variables $tyvars(TE)$ in TE by $\{\gamma_1, ..., \gamma_p\}$. Suppose $S_1^* = Match(TE, A_1) = Unify(\{(x_i, y_i) \mid i = 1, ..., q\})$. The domain of S_1^* will be $TV_{S_1^*} = \{\alpha_1, ..., \alpha_m\} \cup \{\beta_1, ..., \beta_n\} \cup \{\gamma_1, ..., \gamma_p\}$. The region of S_1^* will be a subset of $TV_{S_1^*}$. Therefore we know for all η_i, $S_1^*(\eta_i) = \eta_i$. We also define $TV_{S^+} = TV_{S_1^*} \cup \{\beta_{ij} \mid i = 1, ..., n; j = 1, ..., k\} \cup \{\alpha_{ij} \mid i = 1, ..., m; j = 1, ..., k\}$. Because $\mathrm{Dom}(B_j) \cap \mathrm{Dom}(C_j) = \emptyset$, we define a substitution D_j to be $B_j + C_j$ and D_j' to be $B_j' + C_j'$ for each $j = 1, ..., k$. Now $Match(TE, A_1 \cup C_1A_1 \cup ... \cup C_kA_1)$ is essentially equivalent to $Unify(TT)$ where

$$TT = (\{(x_i, y_i) \mid i = 1, ..., q\} \cup (\bigcup_{j=1}^{k} \{(B_jx_i, C_jy_i) \mid i = 1, ..., q\}))$$

The substitution S^+ is constructed as follows, note that the domain of S^+ is TV_{S^+} and the region of S^+ is a subset of TV_{S^+}.

- First we define a substitution F_j for each $j = 1, ..., k$ such that

- $F_j(\alpha_i) =$ **if** $\alpha_i \in \Phi$ **then** α_i **else** α_{ij} where $i = 1, ..., m$;
- $F_j(\beta_i) =$ **if** $\beta_i \in \Phi$ **then** β_i **else** β_{ij} where $i = 1, ..., n$;
- $F_j(\gamma_i) = \gamma_i$ where $i = 1, ..., p$.

- We define a substitution F'_j on $D_j(TV_{S_1^*})$ such that
  - $F'_j(\alpha_{ij}) =$ **if** $\alpha_i \in \Phi$ **then** α_{ij} **else** α_i where $i = 1, ..., m$;
  - $F'_j(\beta_{ij}) =$ **if** $\beta_i \in \Phi$ **then** β_{ij} **else** β_i where $i = 1, ..., n$;
  - $F'_j(\gamma_i) = \gamma_i$ where $i = 1, ..., p$.

  We also define $F' = F'_1 + ... + F'_k$. Obviously we have $F_j \circ F'_j = F'_j \circ F_j = ID$.

- Now we can define the substitution S^+ on TV_{S^+}:
  - $S^+(\alpha_i) = S_1^*(\alpha_i)$ where $i = 1, ..., m$;
  - $S^+(\alpha_{ij}) = (F_j \circ S_1^*)(\alpha_i)$ where $i = 1, ..., m$ and $j = 1, ..., k$;
  - $S^+(\beta_i) = S_1^*(\beta_i)$ where $i = 1, ..., n$;
  - $S^+(\beta_{ij}) = (F_j \circ S_1^*)(\beta_i)$ where $i = 1, ..., n$ and $j = 1, ..., k$;
  - $S^+(\gamma_i) = S_1^*(\gamma_i)$ where $i = 1, ..., p$. Notice here we also have $F_j(S_1^*(\gamma_i)) = S_1^*(\gamma_i)$ for each $j = 1, ..., k$ because $tyvars(S_1^*(\gamma_i))$ is a subset of Φ.

We can easily prove the following several statements:

- S^+ is an unifier of TT;
  Proof
  - $S^+(x_i) = S_1^*(x_i) = S_1^*(y_i) = S^+(y_i)$ where $i = 1, ..., q$;
  - From the definition of F_j, we know for each type variable $\zeta \in TV_{S_1^*}$, $S^+(D_j(\zeta)) = F_j(S_1^*(\zeta))$. Thus for each $j = 1, ..., k$, we have $S^+(B_j x_i) = F_j(S_1^*(x_i)) = F_j(S_1^*(y_i)) = S^+(C_j y_i)$ where $i = 1, ..., q$.

- The substitution $S_{1j}^* = S^+ \downarrow (D_j(TV_{S_1^*}))$ is a most general unifier of $\{(B_j x_i, C_j y_i) \mid i = 1, ..., q\}$;
  Proof We know for each type variable $\zeta \in TV_{S_1^*}$, $S^+(D_j(\zeta)) = F_j(S_1^*(\zeta))$. Thus $S_{1j}^* \circ D_j = F_j \circ S_1^*$ and therefore $S_1^* = F'_j \circ S_{1j}^* \circ D_j$.

  Given an unifier G of $\{(B_j x_i, C_j y_i) \mid i = 1, ..., q\}$, since $G(B_j x_i) = G(C_j y_i)$ for each $i = 1, ..., q$, $G \circ D_j$ is also an unifier of $\{(x_i, y_i) \mid i = 1, ..., q\}$. Thus there exists a substitution G' such that $G \circ D_j = G' \circ S_1^*$. Thus we can get $G = G' \circ S_1^* \circ D'_j = (G' \circ F'_j) \circ S_{1j}^*$ This exactly means that S_{1j}^* is a most general unifier of $\{(B_j x_i, C_j y_i) \mid i = 1, ..., q\}$;

- S^+ is a most general unifier of TT.
  Proof Given an unifier G for TT, $G_0 = G \downarrow TV_{S_1^*}$ will be an unifier for $\{(x_i, y_i) \mid i = 1, ..., q\}$ and for each $j = 1, ..., k$, $G_j = G \downarrow D_j(TV_{S_1^*})$ is an unifier of $\{(B_j x_i, C_j y_i) \mid i = 1, ..., q\}$. Thus from the definition of mgu, there exists $k + 1$ substitutions G'_j where $j = 0, ..., k$ such that $G_0 = G'_0 \circ S_1^*$ and $G_j = G'_j \circ S_{1j}^*$ for each $j = 1, ..., k$. Since the region of S_1^* is a subset of $TV_{S_1^*}$, and for each $j = 1, ..., k$, the region of S_{1j}^* is a subset of $D_j(TV_{S_1^*})$, we can construct a substitution G' from G'_j for $j = 0, ..., k$ as follows:
  - For each $\zeta \in \mathrm{Reg}(S^+)$, because $S^+(\zeta) = \zeta$, we define $G'(\zeta) = G(\zeta)$.

– For all $\zeta \notin \mathrm{Reg}(S^+)$, we simply define $G'(\zeta) = \zeta$.

Now it's easy to see for all $\zeta \in TV_{S^+}$, $G(\zeta) = G'(S^+(\zeta))$, that is, $G = G' \circ S^+$, thus S^+ is a most general unifier of TT.

From the definition of S^+, we know that the following are true:

- $S^+(TE) = S_1^*(TE) = R_{11}(S_1(TE))$;
- $S^+(\tau_1^*) = S_1^*(\tau_1^*) = R_{11}(\tau_1)$;
- $S^+(H_j\tau_1^*) = L_j(S_{1j}^*(D_j\tau_1^*)) = L_j(F_j(S_1^*(\tau_1^*))) = (L_j \circ (F_j \circ R_{11}))\tau_1$ where $j = 1, ..., k$;

Step 1.2 We prove that the unification $Match(S^+(TE) \pm \{x^{(j)} \mapsto S^+H_j\tau_1^* \mid j = 1, ..., k\}, S^+A_2)$ also succeeds.

We denote $TV_j = \{\alpha_{ij} \mid i = 1, ..., m\} \cup \{\beta_{ij} \mid i = 1, ..., n\} \cup \{\eta_{ij} \mid i = 1, ..., v\}$ for each $j = 0, ..., k$. Let

$$\Phi = \bigcup_{j=1}^{p} tyvars(S_1^*(\gamma_j)) = \{\phi_i \mid i = 1, ..., r\};$$

Notice that if $\gamma_j \in \mathrm{Reg}(S_1^*)$, then $S_1^*(\gamma_j) = \gamma_j$ thus $\gamma_j \in \Phi$. Because $S_1^*(TE) \xrightarrow{R_{12}} S_1(TE)$, let the image of ϕ_i in $S_1^*(TE)$ after R_{12} be denoted by ψ_i. Obviously we have $tyvars(S_1(TE)) = \Psi = \{\psi_i \mid i = 1, ..., r\}$ and $tyvars(S_1^*(TE)) = \Phi$. In doing the unification $S_2^* = Match(S_1(TE) \pm \{x \mapsto gen(S_1(TE), \tau_1)\}, A_2)$, there are k instantiations of the type $gen(S_1(TE), \tau_1)$. Suppose $\forall\mu_1...\mu_s.\tau_1 = gen(S_1(TE), \tau_1)$, for each $j = 1, .., k$, we define a substitution $P_j = \{\mu_i \mapsto \nu_{ij} \mid i = 1, ..., s\}$ where ν_{ij} are just new type variables used in W. And now the k instantiations of the type $gen(S_1(TE), \tau_1)$ will be simply $P_j\tau_1$ where $i = j, ..., k$. Note $\mathrm{Dom}(P_j) \cap tyvars(S_1(TE)) = \emptyset$ and $\mathrm{Reg}(P_j) \cap tyvars(S_1(TE)) = \emptyset$. We also define substitutions $P_j' = \{\nu_{ij} \mapsto \mu_i \mid i = 1, ..., s\}$ for each $j = 1, ..., k$ and a substitution $P' = P_1' + ... + P_k'$. Because Ψ is the set of free type variables in $S_1(TE)$, for each $j = 1, ..., k$, the set of type variables in $P_j\tau_1$ is essentially a subset of $PV_j = \Psi \cup \{\nu_{ij}\}$. we can construct the following renaming substitution for $P_j\tau_1$:

$$Q_j = (L_j \circ (F_j \circ (R_{11} \circ P_j'))) \downarrow PV_j.$$

We make the following two observations on Q_j:

- Notice that $\mathrm{Reg}(R_{11}) = tyvars(S_1^*(TE)) \cup tyvars(S_1^*\tau_1^*)$ is a subset of $\Phi \cup TV_0$. For each $\phi_i \in \Phi$, ϕ_i will be mapped to some free type variable in $S_1(TE)$ by R_{12}. Thus different bound type variables μ_i in $gen(S_1(TE), \tau_1)$ will be mapped by R_{11} to different type variables in $TV_0 \setminus \Phi$. Also for each $j = 1, ..., n$, according to the definition of L_j and F_j, different type variables in $TV_0 \setminus \Phi$ will be mapped to different type variables in TV_j by $(L_j \circ F_j)$. Thus Q_j is a renaming substitution for $P_j\tau_1$ because it virtually maps different type variables ν_{ij} to different type variables in TV_j.
- For all free type variables $\psi_i \in \Psi$ in $S_1(TE)$, $Q_j(\psi_i) = R_{11}(\psi_i)$. This is because for L_j and F_j are identities on all type variables in Φ and P_j' is an identity on all type variables in Ψ.

Because all Q_js have consistent definition on all type variables in Ψ, we can also define $Q = Q_1 + ...Q_k$. Similarly we let Q' to denote the inverse substitution of Q such that $Q' \circ Q = Q \circ Q' = ID$.

Now that we have already known that $S_2^* = Match(S_1(TE) \pm \{x^{(j)} \mapsto P_j\tau_1 \mid j = 1, ..., k\}, A_2)$ succeeds by inductions, by lemma 8.3, $Match(Q(S_1(TE)) \pm \{x^{(j)} \mapsto Q_j(P_j\tau_1) \mid j = 1, ..., k\}, Q(A_2))$ will also succeed,

moreover $S_2^* \circ Q'$ will be one of its most general unifier. Because $Q(S_1(TE)) = R_{11}(S_1(TE)) = S^+(TE)$ and $Q_j(P_j\tau_1) = L_j(F_j(R_{11}\tau 1)) = S^+(H_j(\tau_1^*))$ and $Q(A_2) = A_2 = S^+A_2$, we also get that $Match(S^+(TE) \pm \{x^{(j)} \mapsto S^+H_j\tau_1^* \mid j = 1, ..., k\}, S^+A_2)$ succeeds and $S_2^* \circ Q'$ is one of its most general unifier.

Therefore from the above two steps, we prove that $Match(TE + \{x^{(j)} \mapsto H_j\tau_1^* \mid j = 1, ..., k\}, A_1 \cup C_1A_1 \cup ... \cup C_kA_1 \cup A_2)$ succeeds, moreover by lemma 8.3, $(S_2^* \circ Q') \circ S^+$ is one of its most general unifier. Thus both S_3^* $= Unify(\{(H_i\tau_1^*, t_i) \mid i = 1, ..., k\})$ and $S^* = Match(TE, A_1 \cup S_3^*(A_3 \cup A_2 \setminus \{x\}))$ will succeed.

The reverse direction can be proved in the similar way.

(Part 2) Let's still use the notations in part 1, we can deduce the following by using lemma 8.3 and the results in part 1:

$$\begin{aligned}
S^*(TE, S_3^*(\tau_2^*)) &\cong Match(TE, A_1 \cup S_3^*(A_3 \cup A_2 \setminus \{x\}))\ (TE, S_3^*(\tau_2^*)) \\
&\cong Match(S_3^*(TE), S_3^*(A_1 \cup A_3 \cup A_2 \setminus \{x\}))\ (S_3^*(TE), S_3^*(\tau_2^*)) \\
&\cong ((S_2^* \circ Q') \circ S^+)(TE, \tau_2^*) \\
&\cong (S_2^* \circ Q')(R_{11}(S_1(TE)), \tau_2^*) \\
&\cong (S_2^*(S_1(TE), \tau_2^*)) \\
&\cong (S_2(S_1(TE)), \tau_2)
\end{aligned}$$

Thus we prove that $S^*(TE, S_3^*(\tau_2^*))$ and $(S_2(S_1(TE)), \tau_2)$ are variants. **QED**

9 Appendix: Proof of Theorem 2.6

In this appendix, we are going to prove theorem 2.6, i.e., the equivalence between the algorithm D and W. As mentioned in section 2, the set of predicates we use, denoted by P_m, is $\{p_x(\tau$ where x is any program variable$\}$. The interpretation of p_x is "$\hat{p_x}(\tau)$ = true if and only if $\tau \prec \sigma$, assuming that the type of x is a closed ML type scheme σ." The entailment relation on constraint sets is defined as: $C_1 \Vdash C_2$ if and only if C_1 is satisfiable and $\forall S : S \models C_1 \Rightarrow S \models C_2$. First, we prove the following several lemmas. They will be later used in the proof of theorem 2.6.

Lemma 9.1 *Given a substitution S and two constrained type schemes σ_1 and σ_2. If $\tau' | C' \prec \sigma_1$ and $\sigma_2 = S(\sigma_1)$, then $S(\tau' | C') \prec \sigma_2$.*

Proof Let $\sigma_1 = \forall\overline{\alpha_n}.(\tau | C)$ and let I be the instantiation substitution $I = \{\alpha_i \mapsto \tau_i$ where $i = 1, ..., n\}$ with $I(\tau | C) = \tau' | C'$. Assume that $S_0 = S \downarrow tyvars(\sigma_1)$ and $\beta_1, ..., \beta_n$ are n new type variables which are not in $\mathrm{Reg}(S_0)$. Let $R = \{\alpha_i \mapsto \beta_i\}$, then σ_2 must be of the form $\forall\overline{\beta_n}.S_0(R(\tau | C))$. Therefore defining $J = \{\beta_i \mapsto S(\tau_i)\}$ we get $J(S_0(R(\tau|C))) = S(I(\tau|C)) = S(\tau'|C')$ showing $S(\tau'|C') \prec \sigma_2$. **QED**.

Lemma 9.2 *If $C, TE \vdash e : \tau$ is a valid* ML^+ *typing, and S is a solution of C, then $S(C), S(TE) \vdash e : S(\tau)$ is also valid.*

Proof By structural induction on e. The case where e is a variable follows from Lemma 9.1. Of the remaining cases, only the case for let expressions is interesting.

Obviously $C, TE \vdash \texttt{let}\ x\ \texttt{=}\ e_1\ \texttt{in}\ e_2 : \tau$ must be inferred from the assumptions that both $C_1, TE \vdash e_1 : \tau_1$ and $C, TE \pm \{x \mapsto gen(TE, \tau_1|C_1)\} \vdash e_2 : \tau$ are valid, and $C \Vdash C_1$. Because S is a solution of C, S must be also a solution of C_1. By inductions, both $S(C_1), S(TE) \vdash e_1 : \tau_1$ and $S(C), S(TE) \pm \{x \mapsto S(gen(TE, \tau_1|C_1))\} \vdash e_2 : S(\tau)$ are also valid ML$^+$ typings. Let $\{\alpha_1, ..., \alpha_n\}$ be $tyvars(\tau_1|C_1) \setminus tyvars(TE)$. Assume that $S_0 = S{\downarrow}tyvars(TE)$ and $\beta_1, ..., \beta_n$ are n new type variables which are not in $\mathrm{Reg}(S_0)$. Let $R = \{\alpha_i \mapsto \beta_i\}$, then $S(gen(TE, \tau_1|C_1)) = S(\forall\overline{\alpha_n}.\tau_1|C_1) = \forall\overline{\beta_n}.(S_0(R(\tau_1|C_1)))$. No β_i is free in $S(TE)$. Moreover, any type variable that occurs free in $(S_0(R(\tau_1|C_1)))$ and is not a β_i must be free in $S(TE)$. Therefore we have $S(gen(TE, \tau_1|C_1)) = gen(S(TE), S(\tau_1)|S(C_1))$. Because $C \Vdash C_1$ and S is a solution of C, so $S(C) \Vdash S(C_1)$. Thus we prove that $S(C), S(TE) \vdash \texttt{let}\ x\ \texttt{=}\ e_1\ \texttt{in}\ e_2 : \mathrm{S}(\tau)$ is valid. **QED**.

Definition 9.1 *Given an* ML$^+$ *constrained type scheme* $\sigma = \forall\overline{\alpha_n}.(\tau|C)$, *its* image ML *type scheme, denoted by* $forget(\sigma)$, *is* $\forall\overline{\beta_m}.\tau$ *where* $\overline{\beta_m}$ *are all those* α_i*s that are free in* τ. *Similarly we use* $forget(TE)$ *to denote the* ML *type environment* $TE' = \{x \mapsto forget(TE(x))\}$.

Lemma 9.3 *If* $C, TE \vdash e : \tau$ *is a valid* ML$^+$ *typing, and* $\emptyset \Vdash C$, *let* $TE' = forget(TE)$, *then* $TE' \vdash e : \tau$ *is a valid* ML *typing.*

Proof By structural induction on the expression e. Only the case for let expressions is interesting. Obviously $C, TE \vdash \texttt{let}\ x\ \texttt{=}\ e_1\ \texttt{in}\ e_2 : \tau$ must be inferred from the assumptions that both $C_1, TE \vdash e_1 : \tau_1$ and $C, TE \pm \{x \mapsto gen(TE, \tau_1|C_1)\} \vdash e_2 : \tau$ are valid, and $C \Vdash C_1$. Because $\emptyset \Vdash C$, thus $\emptyset \Vdash C_1$. By inductions, both $TE' \vdash e_1 : \tau_1$ and $TE' \pm \{x \mapsto gen(TE, \tau_1)\} \vdash e_2 : \tau_2$ are valid. From the definition of *forget*, $tyvars(TE') \subseteq tyvars(TE)$ holds, thus $gen(TE, \tau_1) \prec gen(TE', \tau_1)$. Therefore $TE' \pm \{x \mapsto gen(TE', S(\tau_1))\} \vdash e_2 : S(\tau)$ is also a valid ML typing. So $TE' \vdash \texttt{let}\ x\ \texttt{=}\ e_1\ \texttt{in}\ e_2 : S(\tau)$ is a valid ML typing. **QED**.

Proof of theorem 2.6

(Part 1) Suppose that $(S, \tau) = W(TE, e)$ succeeds, we want to prove that $(S', \tau'|C') = D(TE', e)$ also succeeds. We first construct an ML$^+$ type environment $TE'' = \{x \mapsto \forall\overline{\beta_n}.(\tau|\emptyset)$ where $x \in \mathrm{Dom}(TE)$ and $TE(x) = \forall\overline{\beta_n}.\tau\}$. Obviously $tyvars(TE'') = tyvars(TE') = tyvars(TE)$. Because $(S, \tau) = W(TE, e)$ succeeds, from the soundness theorem of the algorithm W, we know that $S(TE) \vdash e : \tau$ is a valid ML typing. Then the type deduction tree of $S(TE) \vdash e : \tau$ can be easily transformed into a ML$^+$ type deduction tree for $\emptyset, S(TE'') \vdash e : \tau$, therefore we prove that $\emptyset, S(TE'') \vdash e : \tau$ is also valid in ML$^+$. Because for each $x \in \mathrm{Dom}(TE')$, $TE''(x) \prec TE'(x)$, and by theorem 2.4, $(S', \tau'|C') = D(TE', e)$ also succeeds. Moreover the typing $C', S'(TE') \vdash e : \tau'$ is more general than $\emptyset, S(TE'') \vdash e : \tau$. From the definition 2.6, there exists some substitution S_1 such that $\tau|\emptyset \prec S_1(gen(S'(TE'), \tau'|C'))$. This also means that there must exist a substitution S_2 such that $\emptyset \Vdash S_2(S_1(C'))$. Thus $S_2 \circ S_1$ is a solution of the constraint set C'. From lemma 2.5, there exists a most general solution for C'.

(Part 2) For the other direction, suppose that $(S', \tau'|C') = D(TE', e)$ succeeds and S^* is a most general solution for C', we want to prove that $(S, \tau) = W(TE, e)$ also succeeds. From theorem 2.4, because $(S', \tau'|C') = D(TE', e)$ succeeds, $C', S'(TE') \vdash e : \tau'$ is a valid ML$^+$ typing. By lemma 9.2, we know that $S^*(C'), S^*(S'(TE')) \vdash e : S^*(\tau')$ is also valid. We use the same notation TE'' as in part 1. Obviously $tyvars(TE'') = tyvars(TE') = tyvars(TE)$. We can easily prove that $S^*(C'), S^*(S'(TE'')) \vdash e : S^*(\tau')$ is valid by constructing an ML$^+$ type deduction tree from the typing $S^*(C'), S^*(S'(TE')) \vdash e : S^*(\tau')$. Then by lemma 9.3, because $forget(TE') = TE$, thus $S^*(S'(TE)) \vdash e : S^*(\tau')$ is a valid ML typing. From the completeness theorem of the algorithm W, $(S, \tau) = W(TE, e)$ should also succeed.

(Part 3) Suppose that both $(S', \tau'|C') = D(TE', e)$ and $(S, \tau) = W(TE, e)$ succeed, and S^* is a most general solution for C', we now prove that $(S(TE), \tau)$ and $(S^*(S'(TE)), S^*\tau')$ are variants.

(Part 3.1) From the discussion in part 1, we know that the typing $C', S'(TE') \vdash e : \tau'$ is more general than the typing $\emptyset, S(TE'') \vdash e : \tau$. Thus there exists a substitution S_1 such that the following is true:

- "$\tau|\emptyset \prec S_1(gen(S'(TE'), \tau'|C'))$." From the proof of lemma 9.2, we know that $S_1(gen(S'(TE'), \tau'|C'))$ $= gen(S_1(S'(TE')), S_1(\tau')|S_1(C'))$. Thus there exists a substitution S_2 with its domain being a subset of $tyvars(S_1(\tau')|S_1(C')) \setminus tyvars(S_1(S'(TE')))$ such that $\tau = S_2(S_1(\tau'))$ and $\emptyset \Vdash S_2(S_1(C'))$. This also means that $S_2 \circ S_1$ is a solution of C'. Assume that S^* is the most general solution of C', then there exists a substitution R_1 such that $S_2 \circ S_1 = R_1 \circ S^*$. Thus $\tau = R_1(S^*(\tau'))$.
- "$S(TE'') \prec S_1(S'(TE'))$." Because $tyvars(TE_2) = \emptyset$, thus $S(TE) \prec S_1(S'(TE))$ should also be true. But this can only be true if $S(TE) = S_1(S'(TE))$ Because S_2 is always an identity on $tyvars(S_1(S'(TE)))$, we also have $S(TE) = S_2(S_1(S'(TE)))$ and thus $S(TE) = R_1(S^*(S'(TE)))$.
- In summary we have proved that there exist a substitution R_2 such that $R_1(S^*(S'(TE)), S^*(\tau')) = (S(TE), \tau)$.

(Part 3.2) From the discussion in part 2, by the completeness theorem of the algorithm W, we know that the typing $S(TE) \vdash e : \tau$ is more general than $S^*(S'(TE)) \vdash e : S^*(\tau')$. This essentially means that there exists a substitution S_1 such that the following is true:

- "$S^*(\tau') \prec S_1(gen(S(TE), \tau))$." Because $S_1(gen(S(TE), \tau)) = gen(S_1(S(TE)), S_1(\tau))$, thus $S^*(\tau') \prec gen(S_1(S(TE)), S_1(\tau))$. This also means that there exists a substitution S_2 with its domain being a subset of $tyvars(S_1(\tau)) \setminus tyvars(S_1(S(TE)))$ such that $S^*(\tau') = S_2(S_1(\tau))$. Now let $R_2 = S_2 \circ S_1$, then $S^*(\tau') = R_2(\tau)$
- "$S_1(S(TE)) = S^*(S'(TE))$." Because S_2 is always an identity on $tyvars(S_1(S(TE)))$, we also have $S^*(S'(TE)) = R_2(S(TE))$
- In summary we have proved that there exist a substitution R_2 such that $(S^*(S'(TE)), S^*(\tau')) = R_2(S(TE), \tau)$.

From the definition 8.3, we know that $(S(TE), \tau)$ and $(S^*(S'(TE)), S^*\tau')$ are variants. **QED**.

10 Appendix: Type Deduction Rules

In this appendix, we list the type deduction rules for the mini-ML language *Exp*, the type deduction rules for the language ML$^+$ (i.e., ML with constrained types), and the static semantics of the skeletal module language *ModL*.

10.1 Type Deduction in *Exp*

$$\text{(VAR)} \qquad \frac{\tau \prec TE(x)}{TE \vdash x : \tau}$$

$$\text{(ABS)} \qquad \frac{TE \pm \{x \mapsto \tau'\} \vdash e : \tau}{TE \vdash \lambda x.e : \tau' \rightarrow \tau}$$

$$\text{(APP)} \qquad \frac{TE \vdash e_1 : \tau' \rightarrow \tau \quad TE \vdash e_2 : \tau'}{TE \vdash e_1 e_2 : \tau}$$

$$\text{(LET)} \qquad \frac{TE \vdash e_1 : \tau_1 \quad TE \pm \{x \mapsto gen(TE, \tau_1)\} \vdash e_2 : \tau}{TE \vdash \texttt{let } x \texttt{ = } e_1 \texttt{ in } e_2 : \tau}$$

10.2 Type Deduction in ML$^+$

$$\text{(VAR)} \qquad \frac{\tau | C \prec TE(x)}{C, TE \vdash x : \tau}$$

$$\text{(ABS)} \qquad \frac{C, TE \pm \{x \mapsto \tau' | C\} \vdash e : \tau}{C, TE \vdash \lambda x.e : \tau' \rightarrow \tau}$$

$$\text{(APP)} \qquad \frac{C, TE \vdash e_1 : \tau' \rightarrow \tau \quad C, TE \vdash e_2 : \tau'}{C, TE \vdash e_1 e_2 : \tau}$$

$$\text{(LET)} \qquad \frac{C_1, TE \vdash e_1 : \tau_1 \quad C, TE \pm \{x \mapsto gen(TE, \tau_1 | C_1)\} \vdash e_2 : \tau \quad C \Vdash C_1}{C, TE \vdash \texttt{let } x \texttt{ = } e_1 \texttt{ in } e_2 : \tau}$$

10.3 Static Semantics of *ModL*

The grammar of *ModL* and the semantic objects used here are those given in figure 3. Given a functor $\Phi = (N_1)(S_1, (N_1')(S_1'))$, a functor instance $(S_2, (N_2')(S_2'))$ is an *instance* of Φ, written $\Phi \geq (S_2, (N_2')(S_2'))$, if there exists a realization φ such that $\varphi(S_1, (N_1')(S_1')) = (S_2, (N_2')(S_2'))$ and $\mathrm{Dom}(\varphi) \subseteq N_1$. Also we use $Names(E)$ to denote the set of names occurred free in E and $B \oplus E$ to denote $B \pm (Names(E), \{\}, \{\}, E)$.

(DEC)
$$\frac{}{B \vdash \ \Rightarrow \{\}}$$

$$\frac{B \vdash strdec \Rightarrow E_1}{B \vdash strdec \Rightarrow E_1}$$

$$\frac{B \vdash strdec \Rightarrow E_1 \quad B \oplus E_1 \vdash dec1 \Rightarrow E_2}{B \vdash strdec\ dec1 \Rightarrow E_1 \pm E_2}$$

(STREXP)
$$\frac{B(strid) = S}{B \vdash strid \Rightarrow S}$$

$$\frac{B \vdash dec \Rightarrow E \quad m \notin ((\text{Nameset of } B) \cup Names(E))}{B \vdash \texttt{struct}\ dec\ \texttt{end} \Rightarrow (m, E)}$$

$$\frac{B \vdash strexp \Rightarrow (m, E) \quad E(strid) = S}{B \vdash strexp.strid \Rightarrow S}$$

$$\frac{B \vdash strexp \Rightarrow S \quad B(fctid) \geq (S'', (N')S') \quad S\ enriches\ S'' \quad (\text{Nameset of } B) \cap N' = \emptyset}{B \vdash fctid(strexp) \Rightarrow S'}$$

(STRDEC)
$$\frac{B \vdash strexp \Rightarrow S}{B \vdash \texttt{structure}\ strid = strexp \Rightarrow \{strid \mapsto S\}}$$