A Typed Intermediate Language for Flow-Directed Compilation[*]

J. B. Wells[**,1], Allyn Dimock[2], Robert Muller[3], and Franklyn Turbak[4]

[1] Boston University, Boston MA 02215, USA
[2] Harvard University, Cambridge MA 02138, USA
[3] Boston College, Chestnut Hill MA 02167, USA
[4] Wellesley College, Wellesley MA 02181, USA

Abstract. We present a typed intermediate language λ^{CIL} for optimizing compilers for function-oriented and polymorphically typed programming languages (e.g., ML). The language λ^{CIL} is a typed lambda calculus with product, sum, intersection, and union types as well as function types annotated with flow labels. A novel formulation of intersection and union types supports encoding flow information in the typed program representation. This flow information can direct optimization.

1 Introduction

Recently there has been much interest in the view of compilation as a composition of well typed program transformations. In this setting, the compiler maintains the invariant that at each step of the compilation process the intermediate representation of the source program is well typed. This invariant can be observed if the input program is well typed and each compiler transformation changes the intermediate representation and its typing in a consistent way. This approach requires using one or more typed intermediate languages.

Explicitly typed intermediate languages offer several benefits to the compiler writer [20, 19, 25, 21]. First, type information can guide program analyses and transformations. Second, some applications need accurate type information at run-time thereby requiring the compiler to preserve it. Finally, typed intermediate languages are useful as a debugging aid in the compiler development process.

This paper introduces a typed intermediate language for optimizing compilers for higher-order polymorphic programming languages. Our intermediate language[1] λ^{CIL} is an explicitly typed λ-calculus with product, sum, intersection, union as well as function types annotated with flow labels in the style of Heintze and Banerjee [13, 6].

The flow annotations on function types are sets of term labels that can encode control and data flow information as it would be computed by one of several

---

[1] In λ^{CIL}, "C" is for the Church Project (http://www.cs.bu.edu/groups/church/) and "IL" is for "intermediate language." The Church Project is investigating the use of intersection and union types in compiling ML–like languages.

typed flow analyses in the literature [13, 6]. If a flow analysis determines that subterm occurrence M has type $\sigma \xrightarrow[\psi]{\phi} \tau$, then the λ-abstractions flowing from M are those with labels in ϕ and they flow only to application sites with labels in ψ. The sets ϕ and ψ are sets of potential flow *sources* and *sinks*.

The formulation of λ^{CIL} allows flow information to be separated in a well typed manner to expose precise correspondences between sources and sinks of flow [9]. A λ-abstraction flowing to m application sites can be assigned an intersection type with m conjuncts. This is represented in λ^{CIL} by m virtual copies of the term. In a dual manner, an application to which n abstractions might flow can be assigned a union type with n disjuncts. This is represented in λ^{CIL} by a virtual case expression that dispatches to one of n clauses.

The program representation supported by λ^{CIL} can be exploited in generating efficient object code. One approach to compiling polymorphism is to generate *specialized* instances of a polymorphic definition based on its uses. Specialization not only avoids the overhead of boxing but more importantly enables subsequent optimizations such as inlining and common subexpression elimination. Empirical evidence suggests that the optimizations enabled by specialization can actually lead to smaller object programs than alternative approaches [16].

Typically the specialization approach is limited to non-escaping polymorphic functions where the required specializations of the definition are determined by its uses within the confines of a binding construct such as **let** [25, 8]. In λ^{CIL}, the required specializations can be determined by the flow analysis. Escaping polymorphic functions can be specialized for their uses in textually remote parts of the program. It is also easier in λ^{CIL} to provide multiple representations of a function for different types and for particular inputs. Inlining of functions can be performed even when multiple functions can flow to a call site. It can also be performed on open functions. This is further discussed in [9].

2 Flow-Directed Program Transformation

We informally illustrate the features of λ^{CIL} in the context of *closure conversion*, a key program transformation in optimizing compilers for function-oriented and object-oriented languages [30, 18, 11]. Closure conversion transforms programs that may contain open functions into equivalent programs that contain only closed functions. An important technical challenge in closure conversion is to generate efficient function representations without violating the invariant that all function representations flowing to a particular application site are consistent with that site's application protocol.

The simplest way to maintain this invariant is to give every function the same representation and use the same application protocol at every call site. In a naive strategy, closure conversion maps every source function to a *closure*, a pair of (1) the values of the function's free variables (the *environment*) and (2) a closed form of the function (the *code*) that takes the environment as an additional argument. However, the overhead of creating and applying a closure can often be avoided by choosing more efficient function representations.

We illustrate closure conversion with the following example: [2]

$$\textbf{let } f^{\text{int}\to\text{int}} = \lambda x^{\text{int}}.x * 2$$
$$g^{\text{int}\to\text{int}} = \lambda y^{\text{int}}.y + a^{\text{int}} \qquad (1)$$
$$\textbf{in } \times\big(f \ @ \ 5, (\textbf{if } b^{\text{bool}} \textbf{ then } f \textbf{ else } g) \ @ \ 7\big)$$

The closed function $(\lambda x^{\text{int}}.x * 2)$ flows to two call sites, the second of which is also a sink for the open function $(\lambda y^{\text{int}}.y + a^{\text{int}})$. The flow of this simple program is merely an example of more complex flow patterns arising in real programs.

A typed flow analysis of the example in λ^{CIL} might yield the flow graph: [3]

Each abstraction occurrence $(\lambda_\psi^l x^\tau.M)$ is identified by a label l and a set of labels ψ approximating the set of application occurrences that can consume it. Each application occurrence $(M \ @_k^\phi \ N)$ is identified by a label k and a set of abstraction occurrence labels ϕ approximating the set of abstraction occurrences that it can consume. Function types "$\text{int} \xrightarrow[\psi]{\phi} \text{int}$" are also annotated with sets of source and sink labels. [4]

Consider closure converting our example. The function $\lambda_{\{3,4\}}^1$ is already closed, so it is desirable to represent it as a function (not a closure) and to keep $@_3^{\{1\}}$ as a regular function application (not a closure application). This optimization is called *selective* closure conversion [30]. However, since $\lambda_{\{3,4\}}^1$ also flows to $@_4^{\{1,2\}}$ along with the open function $\lambda_{\{4\}}^2$, something must be done to ensure that the protocols at the call sites are consistent with the function representations that flow to them.
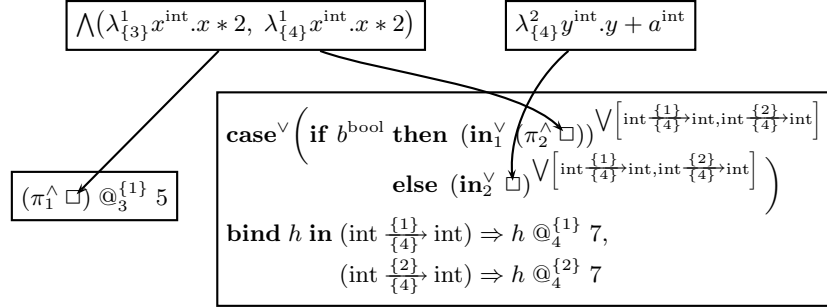
The flow-based features of λ^{CIL} are helpful in dealing with multiple representations that can be desirable in closure conversion. In λ^{CIL}, a term can be transformed to expose correspondences between sources and sinks via intersec-

---

[2] Remarks on notation: Variables are annotated with types, applications are marked by "@", and tuples are marked by "×". For readability, types on bound variable occurrences are omitted when the binding is present. We use base types (like int and bool), constants of these types, and familiar operators on these types, even though these are not formally defined.

[3] To emphasize that our approach addresses complex flow patterns, we present the example's flow graph diagramatically, detaching the abstractions and applications from their surrounding context.
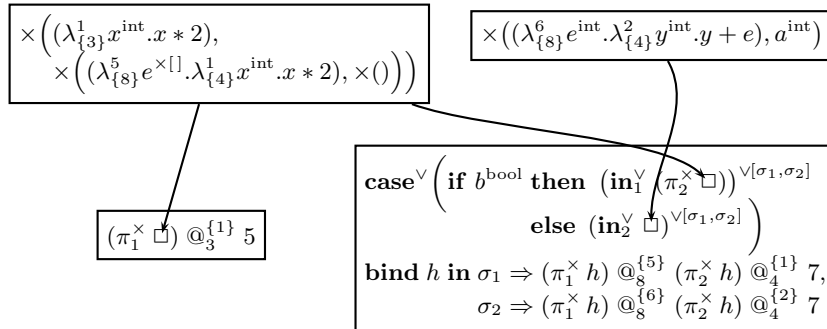
[4] For well-typedness, flow-label subtyping coercions may be needed. For readability, we omit these from examples.

tion ($\wedge$) and union ($\vee$) types:[5]

$$\bigwedge(\lambda^1_{\{3\}}x^{\text{int}}.x * 2,\ \lambda^1_{\{4\}}x^{\text{int}}.x * 2)$$

$$\lambda^2_{\{4\}}y^{\text{int}}.y + a^{\text{int}}$$

$$(\pi^\wedge_1\ \square)\ @^{\{1\}}_3\ 5$$

$$\mathbf{case}^\vee\left(\mathbf{if}\ b^{\text{bool}}\ \mathbf{then}\ (\mathbf{in}^\vee_1\ (\pi^\wedge_2\ \square))^{\vee\left[\text{int}\ \frac{\{1\}}{\{4\}}\text{int},\text{int}\ \frac{\{2\}}{\{4\}}\text{int}\right]}\right.$$
$$\left.\mathbf{else}\ (\mathbf{in}^\vee_2\ \square)^{\vee\left[\text{int}\ \frac{\{1\}}{\{4\}}\text{int},\text{int}\ \frac{\{2\}}{\{4\}}\text{int}\right]}\right)$$
$$\mathbf{bind}\ h\ \mathbf{in}\ (\text{int}\ \tfrac{\{1\}}{\{4\}}\ \text{int}) \Rightarrow h\ @^{\{1\}}_4\ 7,$$
$$(\text{int}\ \tfrac{\{2\}}{\{4\}}\ \text{int}) \Rightarrow h\ @^{\{2\}}_4\ 7$$

The abstraction occurrence $\lambda^1_{\{3,4\}}$ has been transformed into a *virtual tuple* (term of intersection type) containing two abstraction occurrences $\lambda^1_{\{3\}}$ and $\lambda^1_{\{4\}}$. Intuitively, a virtual tuple is a compile-time tuple containing copies of a term that differ only in their types. Since all of the components of a virtual tuple behave identically, no code will be generated to build or access its slots at run-time. Similarly, the application occurrence $@^{\{1,2\}}_4$ has been transformed into a *virtual case* expression that dispatches on the tag of a virtual variant to one of two application occurrences $@^{\{1\}}_4$ or $@^{\{2\}}_4$. All of the clauses of a virtual case expression will share the same code at run-time. The purpose of virtual tuples and variants is to make the term well typed and to provide a place to put type and flow annotations. However, a compiler can transform some virtual tuples ($\wedge$) to real tuples ($\times$) and some virtual variants ($\vee$) to real variants ($+$).
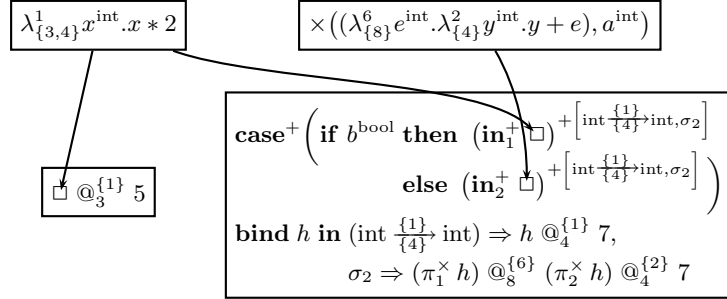
For example, one approach to closure converting our example is to *split* the virtual tuple for the closed function into two distinct functions representations, one which flows to $@^{\{1\}}_3$ and one which flows $@^{\{1\}}_4$. In this case, the virtual product becomes a real product, but the virtual variant stays virtual:

$$\times\left((\lambda^1_{\{3\}}x^{\text{int}}.x * 2),\right.$$
$$\left.\times\left((\lambda^5_{\{8\}}e^{\times[]}.\lambda^1_{\{4\}}x^{\text{int}}.x * 2), \times())\right)\right)$$

$$\times\left((\lambda^6_{\{8\}}e^{\text{int}}.\lambda^2_{\{4\}}y^{\text{int}}.y + e), a^{\text{int}}\right)$$

$$(\pi^\times_1\ \square)\ @^{\{1\}}_3\ 5$$

$$\mathbf{case}^\vee\left(\mathbf{if}\ b^{\text{bool}}\ \mathbf{then}\ (\mathbf{in}^\vee_1\ (\pi^\times_2\ \square))^{\vee[\sigma_1,\sigma_2]}\right.$$
$$\left.\mathbf{else}\ (\mathbf{in}^\vee_2\ \square)^{\vee[\sigma_1,\sigma_2]}\right)$$
$$\mathbf{bind}\ h\ \mathbf{in}\ \sigma_1 \Rightarrow (\pi^\times_1\ h)\ @^{\{5\}}_8\ (\pi^\times_2\ h)\ @^{\{1\}}_4\ 7,$$
$$\sigma_2 \Rightarrow (\pi^\times_1\ h)\ @^{\{6\}}_8\ (\pi^\times_2\ h)\ @^{\{2\}}_4\ 7$$

where $\sigma_1 = \times\left[(\times[]\ \tfrac{\{5\}}{\{8\}}\text{int}\ \tfrac{\{1\}}{\{4\}}\text{int}),\times[]\right]$ and $\sigma_2 = \times\left[(\text{int}\ \tfrac{\{6\}}{\{8\}}\text{int}\ \tfrac{\{2\}}{\{4\}}\text{int}),\text{int}\right]$

---

[5] Notation: $\wedge(M_1,\dots,M_n)$ constructs a term of intersection type $\wedge[\tau_1,\dots,\tau_n]$ whose components are extracted via $\pi^\wedge_i$. $(\mathbf{in}^\vee_i\ M)^{\tau_i}$ constructs a term of union type $\vee[\tau_1,\dots,\tau_n]$ which is analyzed by $\mathbf{case}^\vee$.

Another option is to use only one representation for the closed function, but to *tag* it to distinguish it from the open function representation. In this case, the virtual product stays virtual, but a case analysis will distinguish between the call protocols of the real variants at run-time:

$$\boxed{\lambda^1_{\{3,4\}} x^{\text{int}}.x * 2} \qquad \boxed{\times\left((\lambda^6_{\{8\}} e^{\text{int}}.\lambda^2_{\{4\}} y^{\text{int}}.y + e), a^{\text{int}}\right)}$$

$$\boxed{\square\ @^{\{1\}}_3\ 5}$$

$$\boxed{\begin{aligned} &\textbf{case}^+\left(\textbf{if}\ b^{\text{bool}}\ \textbf{then}\ (\textbf{in}^+_1 \square)^{+\left[\text{int}\xrightarrow{\{1\}}{\{4\}}\text{int},\sigma_2\right]}\right.\\ &\qquad\qquad\qquad\qquad \left.\textbf{else}\ (\textbf{in}^+_2 \square)^{+\left[\text{int}\xrightarrow{\{1\}}{\{4\}}\text{int},\sigma_2\right]}\right)\\ &\textbf{bind}\ h\ \textbf{in}\ (\text{int}\xrightarrow{\{1\}}{\{4\}}\text{int}) \Rightarrow h\ @^{\{1\}}_4\ 7,\\ &\qquad\qquad \sigma_2 \Rightarrow (\pi^\times_1\ h)\ @^{\{6\}}_8\ (\pi^\times_2\ h)\ @^{\{2\}}_4\ 7 \end{aligned}}$$

λ^{CIL} can also handle *inlining*, a vital compiler optimization, as another function representation choice. In our example, the code of the open function can be inlined at its single call site, and the open function can be represented by the value of its sole free variable a rather than as a closure.

$$\boxed{\lambda^1_{\{3,4\}} x^{\text{int}}.x * 2} \qquad \boxed{a^{\text{int}}}$$

$$\boxed{\square\ @^{\{1\}}_3\ 5}$$

$$\boxed{\begin{aligned} &\textbf{case}^+\left(\textbf{if}\ b^{\text{bool}}\ \textbf{then}\ (\textbf{in}^+_1 \square)^{+\left[\text{int}\xrightarrow{\{1\}}{\{4\}}\text{int},\text{int}\right]}\right.\\ &\qquad\qquad\qquad\qquad \left.\textbf{else}\ (\textbf{in}^+_2 \square)^{+\left[\text{int}\xrightarrow{\{1\}}{\{4\}}\text{int},\text{int}\right]}\right)\\ &\textbf{bind}\ h\ \textbf{in}\ (\text{int}\xrightarrow{\{1\}}{\{4\}}\text{int}) \Rightarrow h\ @^{\{1\}}_4\ 7,\\ &\qquad\qquad \text{int} \Rightarrow 7 + h \end{aligned}}$$

Not only does λ^{CIL} support the inlining of open functions, but the flow annotations in λ^{CIL} expose flow-based inlining opportunities that may not be apparent from the program text.

Every change from $\wedge$ to $\times$ or $\vee$ to $+$ may lead to a cascade of changes necesary to preserve well typedness and meaning. Our calculus aids in automatically handling these changes. Space does not permit us to specify the closure conversion transformations here; for details, see [9].

3   Design Issues

This section discusses some of the goals that guided the design of our language λ^{CIL} and some of the technical challenges that had to be overcome.

Finitary Types and Typings: A central theme of our work is the desire for types and typings containing detailed information on the uses of functions and data representations. Some type system designs conflict with these goals. For example, although universal and existential quantifiers are capable of representing

strong behavioral guarantees, they tend to inhibit access by the compiler to information on implementation decisions. As a result, a standard implementation method for languages with universal quantification is *boxing*, i.e., accessing every value that can not fit in a register through a pointer. Boxing is expensive due to run-time overhead and compile-time inhibition of optimization. The dynamic dispatch problem of object-oriented languages is essentially the same as boxing.

Implicit or deep subtyping can cause similar problems. Implicit subtyping fails to record decisions on the placement of coercions. A use of deep subtyping represents a potential coercion which modifies a value *at some other location in the program which may not even exist yet*. This interferes with optimization.

As an alternative to the approaches mentioned above, we have deliberately formulated our language to increase the *concrete* type information available to the compiler and to make typing decisions explicit instead of implicit. Thus, for handling code *polymorphism* and *abstractness*, we use the *finitary* intersection and union types instead of the *infinitary* universal and existential types. Finitary types allow typing as many or more terms as infinitary types.

Encoding Type Annotations: Intersection types are ordinarily implicitly typed using the following typing rule for introducing an intersection type:

$$\frac{A \vdash M : \sigma; \ A \vdash M : \tau}{A \vdash M : \sigma \wedge \tau} \ (\wedge \ \text{intro})$$

As a result, for any subterm M in a typing, there may be multiple typing derivations. Thus, formulating explicit intersection types requires deciding (1) how to annotate the types of bound variables, (2) how to combine different typing annotations for the same term, and (3) how to determine if two different type annotations are for the same term. The new $\wedge$-introduction rule will look something like this:

$$\frac{\begin{array}{l} A \vdash M_1 : \sigma; \ A \vdash M_2 : \tau; \\ M_1 \text{ and } M_2 \text{ are "the same modulo type annotations";} \\ M_3 \text{ is the "combination" of } M_1 \text{ and } M_2 \end{array}}{A \vdash M_3 : \sigma \wedge \tau}$$

The approach used by Reynolds in the language Forsythe [24] annotates the binding of an abstraction $(\lambda x.M)$ with a list of types as in $(\lambda x{:}\sigma_1|\cdots|\sigma_n.M)$, requires the body M of the abstraction to be typable with the same type τ for each possible type σ_i of the bound variable x, and then assigns the abstraction the type $(\sigma_1 \to \tau) \wedge \cdots \wedge (\sigma_n \to \tau)$. Unfortunately, this method is not sufficient to represent dependencies between the types of nested variable bindings. Pierce gives a more general approach using a special term-level construct to bind a type variable to some set of types [22]. For example, using this method the term $(\lambda x.\lambda y.x)$ could be annotated as $(\textbf{for } \alpha \in \{\sigma, \tau\}.\lambda x{:}\alpha.\lambda y{:}\alpha.x)$ to have the type $(\sigma \to \sigma \to \sigma) \wedge (\tau \to \tau \to \tau)$. However, this method is insufficient to represent some typings, e.g., giving the term $(\lambda x.\lambda y.\lambda z.(xy, xz))$ the type $(((\alpha \to \alpha) \wedge (\beta \to \beta)) \to \alpha \to \beta \to (\alpha \times \beta)) \wedge ((\gamma \to \gamma) \to \gamma \to \gamma \to (\gamma \times \gamma))$.

To provide a place for multiple conflicting type annotations, we altered the standard typing rule to "combine" the multiple type-annotated versions of a term by simply keeping both versions:

$$A \vdash M_1 : \sigma; \;\; A \vdash M_2 : \tau;$$
$$\underline{M_1 \text{ and } M_2 \text{ are "the same modulo type annotations"}}$$
$$A \vdash \wedge(M_1, M_2) : \sigma \wedge \tau$$

We call the term $\wedge(M_1, M_2)$ a *virtual tuple* and prefix it with the "$\wedge$" symbol to distinguish it from an ordinary tuple, which we now prefix with "$\times$". The intended meaning is that M_1, M_2, and $\wedge(M_1, M_2)$ are merely different type-annotated versions of the *same* term. Given this choice, we can then use ordinary type annotations on variable bindings. For example, to give the term $\lambda x.x$ the type $(\sigma \to \sigma) \wedge (\tau \to \tau)$, we annotate it as $\wedge(\lambda x^\sigma.x^\sigma, \lambda x^\tau.x^\tau)$.

One implication of our choice is that the tree structure of an explicitly typed term follows the tree structure of its typing proof instead of the tree structure of the untyped term which it represents. A difficulty this introduces is that reduction must essentially work on typing derivations, which is non-trivial to formulate. Wells [31] has developed an alternative formulation where typed and untyped terms have essentially the same tree structure, but the reduction rules are quite complex.

Difficulties with Union Types: It is difficult to formulate an implicitly typed calculus with union types which has the subject-reduction property. For an explicitly typed calculus, this problem manifests itself as a difficulty in guaranteeing that any computation that can be performed on an untyped program can be duplicated on a typed version of the same program. In an implicitly typed calculus, the $\vee$-elimination rule is usually formulated as:

$$\frac{A, x{:}\sigma \vdash M : \rho; \;\; A, x{:}\tau \vdash M : \rho; \;\; A \vdash N : \sigma \vee \tau}{A \vdash M[x{:=}N] : \rho} \;\; (\vee \text{ elim})$$

With this formulation, the subject-reduction property is lost. Barbanera and Dezani-Ciancaglini give as an example the term $(\lambda x.\lambda y.\lambda z.x((\lambda t.t)yz)((\lambda t.t)yz))$, which can be given the type $((\sigma \to \sigma \to \tau) \wedge (\rho \to \rho \to \tau)) \to (\pi \to (\sigma \vee \rho)) \to \pi \to \tau$, but the term $(\lambda x.\lambda y.\lambda z.x(yz)((\lambda t.t)yz))$ to which it reduces can not.

Since the $\vee$-elimination rule given above also causes other difficulties in formulating explicitly typed terms, it seems a solution to this might be to change the elimination rule to:

$$\frac{A, x{:}\sigma \vdash M : \rho; \;\; A, x{:}\tau \vdash M : \rho; \;\; A \vdash N : \sigma \vee \tau}{A \vdash (\lambda x.M)N : \rho} \;\; (\vee \text{ elim})$$

The same example above would still have a problem with this because one could just perform an extra β-reduction step. To solve the problem, it is sufficient to additionally require call-by-value reduction, if a variable is not considered a value. The base values are constants and abstractions and the set of values is closed under tuple and variant formation. This ensures that every reduction at the untyped level will have a corresponding reduction at the typed level.

4 Formal Language Definition

4.1 General Notation and Terminology

A *context* is a term containing holes. However, in this paper, it is simpler to view *terms* as contexts without holes. The expression $C[M_1, \dots, M_n]$ denotes the result of placing M_1, $\dots$, M_n in the n holes of the context C from left to right, possibly capturing free variables. For terms, $M \equiv N$ denotes that M and N are the same term after renaming bound variables. For contexts, $C_1 \equiv C_2$ is similar but only allows renaming bound variables whose scopes do not include a hole. The statement $X \lhd Y$ means that the syntactic entity X occurs properly within the syntactic entity Y; $X \unlhd Y$ has the same meaning except X and Y may be the same. The expression $M[x{:=}N]$ denotes the result of replacing all free occurrences of x in M by N after first renaming the bound variables of M to be distinct from the free variables of N. The expression $\mathrm{FV}(M)$ denotes the set of free variables of M.

Our presentation generalizes *notions of reduction* (n.o.r.). A *simple* n.o.r. R is a pair $(\leadsto_R, \mathbf{C}_R)$ of a redex/contractum relation $\leadsto_R$ and a set of reduction contexts $\mathbf{C}_R$. For a simple n.o.r., $M \longrightarrow_R N$ means M is transformed into N by contracting R-redexes in positions in M specified by an *R-reduction context*, i.e., there are a context $C \in \mathbf{C}_R$ with k holes and terms M_i and N_i for $i \in \{1, \dots, k\}$ such that $M \equiv C[M_1, \dots, M_k]$ and $N \equiv C[N_1, \dots, N_k]$ and $M_i \leadsto_R N_i$ for $i \in \{1, \dots, k\}$. A *composite* n.o.r. R is a rule composing reduction steps of simple n.o.r.'s; in this case $M \longrightarrow_R N$ means M and N are related by the rule. Writing "$\longrightarrow\!\!\!\twoheadrightarrow_R$" denotes the transitive and reflexive closure of "$\longrightarrow_R$". A term M is in *normal form* with respect to R, written $R\text{-nf}(M)$, when there is no term N such that $M \longrightarrow_R N$. The statement $M \xrightarrow{\text{nf}}_R N$ means $M \longrightarrow\!\!\!\twoheadrightarrow_R N$ and $R\text{-nf}(N)$.

4.2 Untyped Language $\lambda_{\mathrm{ut}}^{\mathbf{CIL}}$

Figure 1 shows the syntax and semantics of the untyped language $\lambda_{\mathrm{ut}}^{\mathrm{CIL}}$.

Theorem 1 (Confluence of Untyped Reduction). *If* $\hat{M} \longrightarrow\!\!\!\twoheadrightarrow_{\mathrm{ut}} \hat{N}_1$ *and* $\hat{M} \longrightarrow\!\!\!\twoheadrightarrow_{\mathrm{ut}} \hat{N}_2$, *then there exists* $\hat{M}'$ *such that* $\hat{N}_1 \longrightarrow\!\!\!\twoheadrightarrow_{\mathrm{ut}} \hat{M}'$ *and* $\hat{N}_2 \longrightarrow\!\!\!\twoheadrightarrow_{\mathrm{ut}} \hat{M}'$.

4.3 Explicitly Typed Language $\lambda^{\mathbf{CIL}}$

Figure 2 shows the syntax of our explicitly typed language λ^{CIL}.

Although this presentation omits recursive types, they can be added by extending the types to regular trees. This causes no difficulties with the theorems given in this paper. Of course, a finite representation must be chosen, e.g., the usual $\mu\alpha.\tau$ syntax.

The type erasure $|C|$ of a type-annotated context C (defined in figure 2) is the corresponding untyped and unlabelled context. Some contexts do not have

Untyped Syntax

$$\hat{C} \in \textbf{UntContext} \ ::= \ \Box \ | \ c \ | \ x \ | \ \textbf{rec } x.\hat{C} \ | \ \lambda x.\hat{C} \ | \ \hat{C}_1 @ \hat{C}_2$$
$$| \ \times\left(\hat{C}_1,\ldots,\hat{C}_n\right) \ | \ \pi_i^\times \hat{C}$$
$$| \ \textbf{in}_i^+\hat{C} \ | \ \textbf{case}^+ \hat{C} \ \textbf{bind } x \ \textbf{in} \ \hat{C}_1,\ldots,\hat{C}_n$$

$$\hat{M},\hat{N} \in \textbf{UntTerm} \ = \ \{\,\hat{C} \mid \Box \not\trianglelefteq \hat{C}\,\}$$

$$\hat{V} \in \textbf{UntValue} \quad ::= c \ | \ \lambda x.\hat{M} \ | \ \times\left(\hat{V}_1,\ldots,\hat{V}_n\right) \ | \ \textbf{in}_i^+\hat{V}$$

Untyped Reduction

$$
\begin{array}{lll}
(\lambda x.\hat{M})@\hat{V} & \rightsquigarrow_\text{ut} \ \hat{M}[x{:=}\hat{V}] & \\
\pi_i^\times \times\left(\hat{V}_1,\ldots,\hat{V}_n\right) & \rightsquigarrow_\text{ut} \ \hat{V}_i & \text{if } 1 \le i \le n \\
\textbf{case}^+ (\textbf{in}_i^+\hat{V}) \ \textbf{bind } x \ \textbf{in} \ \hat{M}_1,\ldots,\hat{M}_n & \rightsquigarrow_\text{ut} \ \hat{M}_i[x{:=}\hat{V}] & \text{if } 1 \le i \le n \\
\textbf{rec } x.\hat{V} & \rightsquigarrow_\text{ut} \ \hat{V}[x{:=}(\textbf{rec } x.\hat{V})] &
\end{array}
$$

Reduction contexts: $\mathbf{C}_\text{ut} = \{\,\hat{C} \mid \hat{C} \in \textbf{UntContext} \text{ and } \hat{C} \text{ has exactly one hole}\,\}$

Fig. 1. Untyped language $\lambda_\text{ut}^\text{CIL}$.

a type erasure, i.e., those containing virtual tuples like $\wedge(C_1,\ldots,C_n)$ or virtual case expressions like

$$\textbf{case}^\vee C \ \textbf{bind } x \ \textbf{in} \ \tau_1 \Rightarrow C_1,\ldots,\tau_1 \Rightarrow C_1$$

where the type erasures of $C_1, \ldots, C_n$ are not identical.

Figure 3 gives the typing rules of λ^CIL. A *type environment* is a finite mapping from term variables to types, i.e., a set of variable/type pairs. If A is a type environment, then $A, x{:}\tau$ denotes A extended to map x to type τ. The domain of definition of A is $\text{DomDef}(A)$. A triple $A \vdash M : \tau$ is a *judgement*. A *derivation* $\mathcal{D}$ in language X is a sequence of judgements, each obtained from the previous ones by the typing rules of X. We write "$A \vdash_X M : \tau$ via $\mathcal{D}$" to mean derivation $\mathcal{D}$ is valid in language X and $\mathcal{D}$ ends with $A \vdash M : \tau$. In this case, $\mathcal{D}$ is a *typing* for M in X and M is *well typed* in X. The statement $A \vdash_{\lambda^\text{CIL}} M : \tau$ means there exists some $\mathcal{D}$ such that $A \vdash_{\lambda^\text{CIL}} M : \tau$ via $\mathcal{D}$.

The ($\wedge$ intro) rule requires the equivalence of the type erasure of all components of the virtual tuple, while the ($\vee$ elim) rule requires the equivalence of the type erasures of all clause bodies of a $\textbf{case}^\vee$ expression. These two rules formalize the restrictions on virtual tuples and virtual variants mentioned earlier.

Theorem 2 (Uniqueness of Typings in λ^CIL). *For $M \in \textbf{Term}$, there is at most one type environment A and type τ such that $\text{DomDef}(A) = \text{FV}(M)$ and $A \vdash_{\lambda^\text{CIL}} M : \tau$.*

The call-by-value reduction rules for our typed language λ^CIL are in figure 4. The main notion of reduction, r-reduction, is divided into three steps: simplifying type annotations, performing a computation step, and then simplifying type

Syntax Shared between Types and Terms

$$Q ::= P \mid S \qquad\qquad S ::= \vee \mid + \qquad\qquad P ::= \wedge \mid \times$$

$$l, k \in \mathbf{Label} = \mathbb{N} \qquad\qquad \varnothing \neq \phi, \psi \subset \mathbf{Label}$$

Types

$$\rho, \sigma, \tau \ ::= \ o \ \mid \ \sigma \xrightarrow[\psi]{\phi} \tau \ \mid \ Q[\tau_1, \dots, \tau_n]$$

Type-Annotated Contexts

$$
\begin{aligned}
C \in \mathbf{Context} \ ::= \quad & \square \ \mid \ c \ \mid \ x^\tau \ \mid \ \mathbf{rec}\ x^\tau.C \ \mid \ \lambda_\psi^l x^\tau.C \ \mid \ C_1 \ @_k^\phi\ C_2 \\
& \mid \ P(C_1, \dots, C_n) \ \mid \ \pi_i^P\ C \ \mid \ \mathbf{coerce}\,(\sigma, \tau)\,C \\
& \mid \ \left(\mathbf{in}_i^S\ C\right)^\tau \ \mid \ \mathbf{case}^S\ C\ \mathbf{bind}\ x\ \mathbf{in}\ \tau_1 \Rightarrow C_1, \dots, \tau_n \Rightarrow C_n
\end{aligned}
$$

Type Erasure (a partial function)

$$
\begin{array}{llll}
|\square| & \equiv \square & |c| & \equiv c \\
|x^\tau| & \equiv x & |\mathbf{rec}\ x^\tau.C| & \equiv \mathbf{rec}\ x.|C| \\
\left|\lambda_\psi^l x^\tau.C\right| & \equiv \lambda x.|C| & \left|C_1\ @_k^\phi\ C_2\right| & \equiv |C_1|\ @\ |C_2| \\
|\times(C_1, \dots, C_n)| & \equiv \times(|C_1|, \dots, |C_n|) & |\mathbf{coerce}\,(\sigma, \tau)\,C| & \equiv |C| \\
|\pi_i^\times\ C| & \equiv \pi_i^\times\ |C| & |\pi_i^\wedge\ C| & \equiv |C| \\
\left|\left(\mathbf{in}_i^+\ C\right)^\tau\right| & \equiv \mathbf{in}_i^+ |C| & \left|\left(\mathbf{in}_i^\vee\ C\right)^\tau\right| & \equiv |C|
\end{array}
$$

$$\left|\mathbf{case}^+\ C\ \mathbf{bind}\ x\ \mathbf{in}\ \tau_1 \Rightarrow C_1, \dots, \tau_n \Rightarrow C_n\right| \equiv \mathbf{case}^+\ |C|\ \mathbf{bind}\ x\ \mathbf{in}\ |C_1|, \dots, |C_n|$$

$$\left|\mathbf{case}^\vee\ C\ \mathbf{bind}\ x\ \mathbf{in}\ \tau_1 \Rightarrow C_1, \dots, \tau_n \Rightarrow C_n\right| \equiv \begin{cases} (\lambda x.|C_1|)\ @\ |C| & \text{if } |C_1| \equiv \cdots \equiv |C_n|, \\ \textit{undefined} & \text{otherwise.} \end{cases}$$

$$|\wedge(C_1, \dots, C_n)| \equiv \begin{cases} |C_1| & \text{if } |C_1| \equiv \cdots \equiv |C_n|, \\ \textit{undefined} & \text{otherwise.} \end{cases}$$

Type-Annotated Terms, Values, Parallel Contexts

$$
\begin{aligned}
M, N \in \mathbf{Term} & = \{\, C \mid \text{the type erasure } |C| \in \mathbf{UntTerm} \,\} \\
V \in \mathbf{Value} & = \{\, C \mid \text{the type erasure } |C| \in \mathbf{UntValue} \,\} \\
Cp \in \mathbf{ParallelContext} & = \{\, C \mid \text{the type erasure } |C| \text{ has exactly one hole} \,\}
\end{aligned}
$$

Syntactic Sugar for Examples

$$\mathrm{bool} = +[\times[\,], \times[\,]] \qquad \mathrm{true} \equiv \left(\mathbf{in}_1^+ \times()\right)^{\mathrm{bool}} \qquad \mathrm{false} \equiv \left(\mathbf{in}_2^+ \times()\right)^{\mathrm{bool}}$$

$$(\mathbf{if}\ M_1\ \mathbf{then}\ M_2\ \mathbf{else}\ M_3) \equiv \mathbf{case}^+\ M_1\ \mathbf{bind}\ x\ \mathbf{in}\ \times[\,] \Rightarrow M_2, \times[\,] \Rightarrow M_3 \quad (\text{fresh } x)$$

$$(\mathbf{let}\ x^\tau = N\ \mathbf{in}\ M) \equiv ((\lambda_{\{k\}}^l x^\tau.M)\ @_k^{\{l\}}\ N) \quad (\text{fresh } l, k)$$

Fig. 2. Syntax of explicitly typed language λ^{CIL}.

$$(\text{const}) \quad \overline{A \vdash c : o} \qquad\qquad (\text{var}) \quad \overline{A, x{:}\tau \vdash x^{\tau} : \tau}$$

$$(\rightarrow \text{elim}) \quad \frac{A \vdash M : \sigma \xrightarrow[\{k\}]{\phi} \tau; \ A \vdash N : \sigma}{A \vdash M \, @_k^{\phi} \, N : \tau} \qquad (\rightarrow \text{intro}) \quad \frac{A, x{:}\sigma \vdash M : \tau}{A \vdash \lambda_\psi^l x^\sigma.M : \sigma \xrightarrow[\psi]{\{l\}} \tau}$$

$$(\times \text{intro}) \quad \frac{\forall_{i=1}^n.\ A \vdash M_i : \tau_i}{A \vdash \times(M_1, \ldots, M_n) : \times[\tau_1, \ldots, \tau_n]} \qquad (\text{coerce}) \quad \frac{A \vdash M : \sigma; \ \sigma \leq \tau}{A \vdash \mathbf{coerce}\,(\sigma, \tau)\,M : \tau}$$

$$(\wedge \text{intro}) \quad \frac{\forall_{i=1}^n.\ A \vdash M_i : \tau_i; \ |M_1| \equiv \cdots \equiv |M_n|}{A \vdash \wedge(M_1, \ldots, M_n) : \wedge[\tau_1, \ldots, \tau_n]} \qquad (\text{rec}) \quad \frac{A, x{:}\tau \vdash M : \tau}{A \vdash \mathbf{rec}\ x^\tau.M : \tau}$$

$$(\times,\wedge \text{ elim}) \quad \frac{A \vdash M : P[\tau_1, \ldots, \tau_n]; \ 1 \leq i \leq n}{A \vdash \pi_i^P \, M : \tau_i} \qquad (\text{arrow-}\leq) \quad \frac{\phi \subseteq \phi'; \ \psi' \subseteq \psi}{\sigma \xrightarrow[\psi]{\phi} \tau \leq \sigma \xrightarrow[\psi']{\phi'} \tau}$$

$$(+,\vee \text{ intro}) \quad \frac{A \vdash M : \tau_i; \ 1 \leq i \leq n}{A \vdash \left(\mathbf{in}_i^S \, M\right)^{S[\tau_1, \ldots, \tau_n]} : S[\tau_1, \ldots, \tau_n]}$$

$$(+ \text{ elim}) \quad \frac{A \vdash M : +[\tau_1, \ldots, \tau_n]; \ \forall_{i=1}^n.\ A, x{:}\tau_i \vdash M_i : \tau}{A \vdash \mathbf{case}^+ M \ \mathbf{bind} \ x \ \mathbf{in} \ \tau_1 \Rightarrow M_1, \ldots, \tau_n \Rightarrow M_n : \tau}$$

$$(\vee \text{ elim}) \quad \frac{A \vdash M : \vee[\tau_1, \ldots, \tau_n]; \ \forall_{i=1}^n.\ A, x{:}\tau_i \vdash M_i : \tau; \ |M_1| \equiv \cdots \equiv |M_n|}{A \vdash \mathbf{case}^\vee M \ \mathbf{bind} \ x \ \mathbf{in} \ \tau_1 \Rightarrow M_1, \ldots, \tau_n \Rightarrow M_n : \tau}$$

Fig. 3. Typing rules of explicitly typed language λ^{CIL}.

```
┌─────────────────────────────────────────────────────────────────────────────┐
│ Main Notion of Reduction for Type-Annotated Terms                           │
│                                                                              │
│         M ⟶ᵣ N   iff   ∃M′, N′. (M ━ⁿᶠ→ₜ M′ ⟶꜀ N′ ━ⁿᶠ→ₜ N)                │
│                                                                              │
│ Computation Reduction                                                        │
└─────────────────────────────────────────────────────────────────────────────┘
```

Main Notion of Reduction for Type-Annotated Terms

$$M \longrightarrow_r N \quad \text{iff} \quad \exists M', N'. \ (M \xrightarrow{\text{nf}}_t M' \leadsto_c N' \xrightarrow{\text{nf}}_t N)$$

Computation Reduction

$$
\begin{array}{lll}
(\lambda^l_\psi x^\tau.M) \ @^\phi_k V & \leadsto_c M[x:=V] & \\
\pi^\times_i \times(V_1,\ldots,V_n) & \leadsto_c V_i & \text{if } 1 \le i \le n \\
\mathbf{case}^+ \left(\mathbf{in}^+_i V\right)^\tau \mathbf{bind}\ x\ \mathbf{in}\ \tau_1 \Rightarrow M_1,\ldots,\tau_n \Rightarrow M_n \leadsto_c M_i[x:=V] & & \text{if } 1 \le i \le n \\
\mathbf{rec}\ x^\tau.V & \leadsto_c V[x:=(\mathbf{rec}\ x^\tau.V)] &
\end{array}
$$

Reduction contexts: $\mathbf{C}_c = \mathbf{ParallelContext}$

Type-Annotation-Simplification Reduction

$$
\begin{array}{lll}
\pi^\wedge_i \wedge(M_1,\ldots,M_n) & \leadsto_t M_i & \text{if } 1 \le i \le n \\
(\mathbf{case}^\vee \left(\mathbf{in}^\vee_i N\right)^\tau \mathbf{bind}\ x\ \mathbf{in} & \leadsto_t (\lambda^1_{\{1\}} x^{\tau_i}.M_i) \ @^{\{1\}}_1 N & \text{if } 1 \le i \le n \\
\quad \tau_1 \Rightarrow M_1,\ldots,\tau_n \Rightarrow M_n) & & \\
(\mathbf{coerce}\,(\sigma,\tau)\,(\lambda^l_\psi x^\rho.M)) \ @^\phi_k N \leadsto_t (\lambda^l_{\{k\}} x^\rho.M) \ @^{\{l\}}_k N & & \\
\mathbf{coerce}\,(\sigma_1,\tau)\ \mathbf{coerce}\,(\rho,\sigma_2)\,M \leadsto_t \mathbf{coerce}\,(\rho,\tau)\,M & &
\end{array}
$$

Reduction contexts: $\mathbf{C}_t = \{\, C \mid C \in \mathbf{Context}\ \text{and}\ C\ \text{has exactly one hole}\,\}$

Fig. 4. Reduction rules of explicitly typed language λ^{CIL}.

annotations again. Type annotations that might block a computation step are removed by t-reduction. Since t-reduction is terminating, it is convenient to go to t-normal form before and after computation steps. The notion of c-reduction performs real computation steps. In our term formulation, *parallel* c-redexes (i.e., different type-annotated versions of the *same* program phrase) must be contracted simultaneously. This is formalized using *parallel contexts* (members of **ParallelContext**), which require parallel c-redexes to fill holes that map to the same hole in the type-erased program.

Theorem 3 (Subject Reduction for λ^{CIL}). *If $M \longrightarrow_r N$ and $A \vdash_{\lambda^{\text{CIL}}} M : \tau$, then $A \vdash_{\lambda^{\text{CIL}}} N : \tau$.*

Theorem 4 (Typed/Untyped Reduction Correspondence).
If $A \vdash_{\lambda^{\text{CIL}}} M : \tau$, then

1. *If $M \longrightarrow_r N$, then $|M| \longrightarrow_{ut} |N|$.*
2. *If $|M| \longrightarrow_{ut} \hat{N}$, then there exists a term N where $M \longrightarrow_r N$ and $|N| \equiv \hat{N}$.*

Theorem 5 (Confluence of Typed Reduction). *If $M \longrightarrow\!\!\!\twoheadrightarrow_r N_1$ and $M \longrightarrow\!\!\!\twoheadrightarrow_r N_2$, then there exist M'_1 and M'_2 such that $|M'_1| \equiv |M'_2|$ and $N_1 \longrightarrow\!\!\!\twoheadrightarrow_r M'_1$ and $N_2 \longrightarrow\!\!\!\twoheadrightarrow_r M'_2$.*

4.4 Implicitly Typed Language λ^{CIL}_i

The implicitly typed language λ^{CIL}_i is automatically obtained from λ^{CIL} and $\lambda^{\text{CIL}}_{ut}$. The syntax and semantics of λ^{CIL}_i are the same as $\lambda^{\text{CIL}}_{ut}$ as given in figure 1.

The typing rules of $\lambda_{\mathrm{i}}^{\mathrm{CIL}}$ are the rules of figure 3 modified by replacing every judgement $A \vdash M : \tau$ mentioned in a rule by $A \vdash |M| : \tau$, using the type erasure rules from figure 2.

Theorem 6 (Subject Reduction for $\lambda_{\mathrm{i}}^{\mathrm{CIL}}$). *If $\hat{M} \longrightarrow_{\mathrm{ut}} \hat{N}$ and $A \vdash_{\lambda_{\mathrm{i}}^{\mathrm{CIL}}} \hat{M} : \tau$, then $A \vdash_{\lambda^{\mathrm{CIL}}} \hat{N} : \tau$.*

5   Related Work

Typed intermediate languages are used in several experimental compilers. Most typed intermediate languages for polymorphic programming languages can be seen as variants of the Girard/Reynolds λ-calculus, System F [10, 23].

Recent versions of the Standard ML of New Jersey (SML/NJ) compiler [5, 26] use a variant of system F as the representation in the front-end of the compiler. In SML/NJ, type inference annotates polymorphic functions with universally quantified types and annotates function applications with the simple types to which the polymorphic types are instantiated. The compiler uses the type information to select efficient data representations and to minimize boxing coercions [17]. The SML/NJ compiler also uses *minimal typing derivations* [8] to reduce boxing coercions for let-polymorphic definitions. The compiler uses a simply typed representation in later stages of the compiler.

The Glasgow Haskell Compiler (GHC) [20] also uses a variant of System F. In GHC, type inference annotates polymorphic functions with type abstractions and uses of polymorphic functions with type arguments. This allows the compiler to preserve the well-typedness of the intermediate representation across program transformations. The type information is used in the later stages of the compiler to improve code generation.

System F can also be seen as the basis of the typed intermediate language λ_i^{ML} of the TIL compiler for Standard ML [19, 18]. The calculus λ_i^{ML} is a predicative variant of System F extended with intensional polymorphism [12]. The key feature is the support for dynamic type dispatch at run-time. This aids in efficient compilation of polymorphism without sacrificing separate compilation. A use of a polymorphic function can dispatch on a type argument to yield a monomorphic routine suitable for the type. This approach to compiling polymorphism yields excellent results [27] since many type dispatch redexes can be eliminated at compile-time and the compiler can then gain the resulting benefits of type specialization including in-lining and common subexpression elimination.

Our intermediate language λ^{CIL} was inspired by the earlier work on rank-2 intersection types of Jim [15]. As we have shown in this paper, intersection types naturally lead to a flow-directed approach to compilation. Our flow labels encode information about the operational behavior of the program that cannot be obtained from types without flow labels. At the same time, intersection and union types support a natural encoding of polyvariant flow information [6]. While it is clearly possible to compute, record, and use the flow and type information separately, we believe that a single representation is more natural for compilation.

General research into intersection types that has influenced our thinking includes the work of Van Bakel [29] and Jim [15]. Research on both intersection and union types that we have consulted includes the work by Pierce [22], Aiken, Wimmers, and Lakshman [3, 4], Barbanera and Dezani-Ciancaglini [7], and Trifonov and Smith [28]. Of the above, only Pierce considers intersection and union types in an explicitly typed language. Even that is somewhat distant from our work because Pierce includes a general subtyping relation on intersection and union types which we have deliberately avoided.

6   Conclusions and Future Work

We have presented λ^{CIL}, a typed intermediate language suitable for optimizing compilers for higher-order polymorphic programming languages such as ML. The intermediate language is designed to facilitate verifiable flow-directed compiling. Based on λ^{CIL}, we have developed a framework for typed-directed flow-based representation transformations, and have illustrated this framework in a closure conversion application that supports multiple function representations, including the inlining of open functions [9]. This application (informally sketched in section 2) is an example of how λ^{CIL} supplies the compiler writer both important information and great flexibility in making optimization decisions.

Below, we outline some of the work ahead.

Labelling All Terms: In this presentation, only abstractions, applications, and function types were given flow labels. In order to track the flows of non-function values, it will be necessary to to annotate *all* terms and types in the language.

Compiling Polymorphism by Specialization: The λ^{CIL}-calculus suggests an approach to compiling polymorphism of flow-directed specialization. The number of specializations required for a given definition can be minimized if they are determined by representation types rather than source types. We are currently studying the issue of representation types.

Separate Compilation: If a program is compiled as a single unit, it is possible to express all instances of polymorphism and data abstraction in terms of intersection and union types. However, if a program is decomposed into separately compiled modules, universal and existential types may be necessary to model the module interfaces. λ^{CIL} will need to be extended in order to support separate compilation. Additionally, flow-directed specialization is difficult to extend to separately compiled modules. We are currently studying link-time specialization in which the linker determines whether new specializations of a definition are required.

Flow Analysis: The typed control flow analyses alluded to in this paper are limited by our shallow subtyping relation. We would like to weaken this restriction to permit more powerful control flow analysis algorithms.

Term Duplication: An important practical consideration in compiling with types is controlling the size of the intermediate representations. Our current language duplicates terms when it duplicates types. While this language is concep-

tually convenient for specification, for implementation purposes a considerable size savings can be obtained by using a typed calculus with intersection and union types in the style of [31].

References

1. ACM. *Conf. Rec. 21st Ann. ACM Symp. Princ. of Prog. Langs.*, 1994.
2. ACM. *Conf. Rec. POPL '96: 23rd ACM Symp. Princ. of Prog. Langs.*, 1996.
3. A. S. Aiken and E. L. Wimmers. Type inclusion constraints and type inference. In *FPCA '93, Conf. Funct. Program. Lang. Comput. Arch.*, pp. 31–41. ACM, 1993.
4. A. S. Aiken, E. L. Wimmers, and T. K. Lakshman. Soft typing with conditional types. In POPL '94 [1], pp. 163–173.
5. A. W. Appel. *Compiling with Continuations*. Cambridge University Press, 1992.
6. A. Banerjee. A modular, polyvariant, and type-based closure analysis. In ICFP '97 [14].
7. F. Barbanera, M. Dezani-Ciancaglini, and U. de'Liguoro. Intersection and union types: Syntax and semantics. *Inform. & Comput.*, 119:202–230, 1995.
8. N. S. Bjørner. Minimal typing derivations. In *ACM SIGPLAN Workshop on ML and its Applications*, pp. 120–126, 1994.
9. A. Dimock, R. Muller, F. Turbak, and J. B. Wells. Strongly typed flow-directed representation transformations. In ICFP '97 [14], pp. 11–24.
10. J.-Y. Girard. *Interprétation Fonctionnelle et Elimination des Coupures de l'Arithmétique d'Ordre Supérieur*. Thèse d'Etat, Université de Paris VII, 1972.
11. J. Hannan. Type systems for closure conversion. In *Workshop on Types for Program Analysis*, pp. 48–62, 1995. The TPA '95 proceedings are DAIMI PB-493.
12. R. Harper and G. Morrisett. Compiling polymorphism using intensional type analysis. In *Conf. Rec. 22nd Ann. ACM Symp. Princ. of Prog. Langs.* ACM, 1995.
13. N. Heintze. Control-flow analysis and type systems. In *Proc. 2nd Int'l Static Analysis Symp.*, vol. 983 of *LNCS*, pp. 189–206, 1995.
14. *Proc. 1997 Int'l Conf. Functional Programming*. ACM Press, 1997.
15. T. Jim. What are principal typings and what are they good for? In POPL '96 [2].
16. M. P. Jones. Dictionary-free overloading by partial evaluation. In *PEPM '94 — ACM SIGPLAN Workshop Partial Eval. & Semantics-Based Prog. Manipulation*, 1994.
17. X. Leroy. Unboxed objects and polymorphic typing. In *Conf. Rec. 19th Ann. ACM Symp. Princ. of Prog. Langs.*, pp. 177–188. ACM, 1992.
18. Y. Minamide, G. Morrisett, and R. Harper. Typed closure conversion. In POPL '96 [2].
19. G. Morrisett. *Compiling with Types*. Ph.D. thesis, Carnegie Mellon University, 1995.
20. S. L. Peyton Jones. Compiling Haskell by program transformation: A report from the trenches. In *Proc. European Symp. on Programming*, 1996.
21. S. L. Peyton Jones and E. Meijer. Henk: A typed intermediate language. In *Proc. First Int'l Workshop on Types in Compilation*, June 1997.
22. B. C. Pierce. Programming with intersection types, union types, and polymorphism. Technical Report CMU-CS-91-106, Carnegie Mellon University, Feb. 1991.
23. J. C. Reynolds. Towards a theory of type structure. In *Colloque sur la Programmation*, vol. 19 of *LNCS*, pp. 408–425, Paris, France, 1974. Springer-Verlag.

24. J. C. Reynolds. Design of the programming language Forsythe. In P. O'Hearn and R. D. Tennent, eds., *Algol-like Languages*. Birkhauser, 1996.

25. Z. Shao. *Compiling Standard ML for Efficient Execution on Modern Machines*. Ph.D. thesis, Princeton University, 1994.

26. Z. Shao and A. Appel. A type-based compiler for Standard ML. In *Proc. ACM SIGPLAN '95 Conf. Prog. Lang. Design & Impl.*, 1995.

27. D. Tarditi, G. Morrisett, P. Cheng, C. Stone, R. Harper, and P. Lee. TIL: A type-directed optimizing compiler for ML. In *Proc. ACM SIGPLAN '96 Conf. Prog. Lang. Design & Impl.*, 1996.

28. V. Trifonov and S. Smith. Subtyping constrained types. In *Proc. 3rd Int'l Static Analysis Symp.*, pp. 349–365, 1996.

29. S. J. van Bakel. *Intersection Type Disciplines in Lambda Calculus and Applicative Term Rewriting Systems*. Ph.D. thesis, Catholic University of Nijmegen, 1993.

30. M. Wand and P. Steckler. Selective and lightweight closure conversion. In POPL '94 [1], pp. 435–445.

31. J. B. Wells. Intersection types revisited in the Church style. Manuscript, June 1996.

32. J. B. Wells, A. Dimock, R. Muller, and F. Turbak. A typed intermediate language for flow-directed compilation. In *Proc. 7th Int'l Joint Conf. Theory & Practice of Software Development*, pp. 757–771, 1997. Superseded by [33].

33. J. B. Wells, A. Dimock, R. Muller, and F. Turbak. A calculus with polymorphic and polyvariant flow types. *J. Funct. Programming*, 200X. To appear. Supersedes [32].