

ICurry

S. Antoy¹ M. Hanus² A. Jost¹ S. Libby¹

¹ Computer Science Dept., Portland State University, Oregon, U.S.A.

² Institut für Informatik, Kiel University, D-24098 Kiel, Germany

Abstract. FlatCurry is a well-established intermediate representation of Curry programs used in compilers that translate Curry code into Prolog and Haskell code. Some FlatCurry constructs have no direct translation into imperative code. These constructs must be each handled differently when translating Curry code into C, C++ and Python code. We introduce a new representation of Curry programs, called ICurry, and derive a translation from all FlatCurry constructs into ICurry. We present the syntax of ICurry and the translation from FlatCurry to ICurry. We present a model of functional logic computations as graph rewriting, show how this model can be implemented in a low-level imperative language, and describe the translation from ICurry to this model.

Keywords: Curry, Compilation, Intermediate language, Functional logic computation, Graph rewriting

1 Introduction

Functional logic languages [6] provide fast software prototyping and development, simple elegant solutions to otherwise complicated problems, a tight integration between specifications and code [7], and an ease of provability [8, 16] unmatched by other programming paradigms. Not surprisingly, these advantages place heavy demands on their implementation. Theoretical results must be proven and efficient models of execution must be developed. For these reasons, the efficient implementation of functional logic languages is an active area of research with contributions from many sources. This paper is one such contribution.

Compilers of high-level languages transform a *source* program into a *target* program which is in a lower-level language. This transformation maps constructs available in the source program language into simpler, more primitive, constructs available in the target program language. For example, pattern matching can be translated into a sequence of *switch* and assignment statements available in C, Rust, or Java. We use this idea to map Curry into a C-like language. Our target language is not standard C, but a more abstract language that we call *ICurry*. The “I” in ICurry stands for “imperative”, since a design goal of the language is to be easily mappable into an imperative language.

There are advantages in choosing ICurry over C. ICurry is simpler than C. It has no arrays, *typedef* declarations, types, explicit pointers, or dereferencing

operations. ICurry is more abstract than concrete low level languages. Because of its simplicity and abstraction, it has been mapped with a modest effort to C, C++, and Python.

Section 2 is a brief overview of Curry, with focus on the features relevant to ICurry or to the examples. Section 3 discusses an operational model of execution for functional logic computations. This model can be implemented relatively easily in ICurry or in common imperative languages. Section 4 presents *FlatCurry*, a format of Curry programs similar to ICurry. FlatCurry has been used in the translation of Curry into other, non-imperative, languages, but it is not suitable for the translation of Curry into an imperative language. Section 5 defines ICurry and discusses its use and generation. A detailed algorithm for obtaining ICurry from FlatCurry is presented in Appendix A. Section 6 addresses related work and offers our conclusion.

2 Curry

Curry is a declarative language that joins the most appealing features of functional and logic programming. A *Curry program* declares *data types*, which describe how information is structured, and defines *functions* or *operations*, which describe how information is manipulated. For example:

```
data List a = Nil | Cons a (List a)
```

declares a polymorphic type `List` in which `a` is a type parameter standing for the type of the list elements. The symbols `Nil` and `Cons` are the *constructors* of `List`. The values of a list are either `Nil`, the empty list, or `Cons e l`, a pair in which `e` is an element and `l` is a list.

Since lists are ubiquitous, a special notation eases writing and understanding them. We use `[]` to denote the empty list and `(e : l)` to denote the pair. A finite list is written `[e1, ..., en]`, where `ei` is a list element. For example, `[1,2,3] = 1:2:3:[]`.

Functions are defined by rewrite rules of the form:

$$\begin{array}{l}
 f \bar{p} \mid c_1 = e_1 \\
 \dots \\
 \mid c_n = e_n
 \end{array}
 \tag{1}$$

where `f` is a function symbol, `\bar{p}` stands for a sequence of zero or more expressions made up only of constructor symbols and variables, “`| ci`” is a condition, and `ei` is an expression. Conditions in rules are optional. The expressions in `\bar{p}` are called *patterns*. For example, consider:

$$\begin{array}{l}
 \text{abs } x \mid x < 0 = -x \\
 \quad \mid x \geq 0 = x \\
 \text{length } [] = 0 \\
 \text{length } (_ : xs) = 1 + \text{length } xs
 \end{array}
 \tag{2}$$

where `abs` computes the absolute value of its argument and shows some conditions, and `length` computes the length of its argument and shows some patterns.

In contrast to most other languages, the textual order of the rewrite rules in a program is irrelevant—all the rules that can be applied to an expression are applied. An emblematic example is a function, called *choice*, and denoted by the infix operator “?”, which chooses between two *alternatives*:

```
x ? y = x
x ? y = y
```

Therefore, `0 ? 1` is an expression that produces 0 and 1 non-deterministically. In Curry, there are many other useful syntactic and semantic features, for example, rewrite rules can have nested scopes with local definitions. We omit their description here, since they are largely irrelevant to our discussion, with the exception of *let blocks* and *free variables*.

Let blocks support the definition of circular expressions which allows the construction of cyclic graphs. Fig. 1 shows an example of a let block and the corresponding graph. Expression `oneTwo` evaluates to the infinite list `1:2:1:2:...`

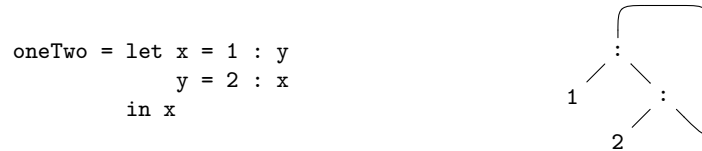


Fig. 1. Example of a let block with mutually recursive variables and the graph it defines.

Free variables abstract unknown information and are “computationally inert” until the information they stand for is required during a computation. When this happens, plausible values for a variable are non-deterministically produced by narrowing [5, 21]. Free variables might occur in initial expressions, conditions, and the right-hand side of rules, and need to be declared by the keyword `free`, unless they are anonymous (denoted by “_”). For instance, the following program defines list concatenation which is exploited to define an operation that returns some element of a list having at least two occurrences:

```
(++) :: [a] -> [a] -> [a]
[]    ++ ys = ys
(x:xs) ++ ys = x : (xs ++ ys)

someDup :: [a] -> a
someDup xs | xs == _ ++ [x] ++ _ ++ [x] ++ _
           = x    where x free
```

3 The Execution Model

A *program* is a graph rewriting system [12, 20] over a *signature*, partitioned into *constructor* and *operation* symbols. We briefly and informally review the underlying theory. A *graph* is a set of *nodes*, where a node is an object with some

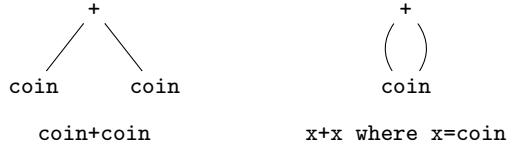


Fig. 2. Graphical and textual representation of expressions. In Curry, all the occurrences of the same variable are shared. Hence, the two occurrences of x stand for the same node. The expression `coin` is conventionally an integer constant with two values, 0 and 1, non-deterministically chosen. The sets of values produced by the two expressions differ.

attributes, and an identity by virtue of being an element in a set. Key attributes of a node are a *label* and a sequence of *successors*. A label is either a symbol of the signature or a variable. A successor is another node, and the sequence of successors may be empty. Exactly one node of a graph is designated as the graph's *root*. Each node of a graph corresponds to an expression in the Curry program.

A *graph rewriting system* is a set of rewrite rules following the *constructor discipline* [19]. A *rule* is a pair of graphs, $l \rightarrow r$, called the left- and right-hand sides respectively. Rules are unconditional without loss of generality [2]. A *rewrite step* of a graph e first identifies both a subgraph t of e , and a rule $l \rightarrow r$ in which t is an instance of l , then replaces t with the corresponding instance of r . The identification of the subgraph t and the rule $l \rightarrow r$ is accomplished by a *strategy* [3]. For example, given the rules (2), a step of `length` [3,4] produces `1+length`[4] where the subgraph reduced in the step is the whole graph, and the rule applied in the step is the second one.

A *computation* of an expression e is a sequence of rewrite steps starting with e , $e = e_0 \rightarrow e_1 \rightarrow \dots$. Expression e is referred to as *top-level*, and each e_i as a *state* of the computation of e . A *value* of a computation is a state in which every node is labeled by a constructor symbol. Such expression is also called a *constructor normal form*. Not every computation has values.

We have modeled a functional logic program as a graph rewriting system [12, 20]. Functional logic computations are executed in this model by rewriting which consists of two relatively simple operations: the construction of graphs and the replacement of subgraphs with other graphs. The most challenging part is selecting the subgraph to be replaced in a way that does not consume computational resources unnecessarily. This is a well-understood problem [3] which is largely separated from the model.

In an implementation of the model, the expressions are objects of a computation and are represented by dynamically linked structures. These structures are similar to those used for computing with lists and trees. The nodes of such a structure are in a bijection with the nodes of the graph they represent. Unless a distinction is relevant, we do not distinguish between a graph and its representation.

The occurrence of a symbol, or variable, in the textual representation of an expression stands for the node labeled by the occurrence. Distinct occurrences may stand for the same node, in which case we say that the occurrences are

shared. The textual representation accommodates this distinction, therefore it is a convenient, linear notation for a graph. Fig. 2 shows two graphs and their corresponding textual expressions.

4 FlatCurry

FlatCurry [13] is an intermediate language used in a variety of applications. These applications include implementing Curry by compiling into other languages, like Prolog [17] or Haskell [10]. FlatCurry is also the basis for specifying the operational semantics of Curry programs [1], building generic analysis tools [18], or verifying properties of Curry programs [15, 16]. The FlatCurry format of a Curry program removes some syntactic constructs, such as nested scopes and infix notation, that make source programs more human readable. This removal still preserves the program’s meaning. We ignore some elements of FlatCurry, such as imported modules or exported symbols, which are not directly related to the execution model presented in Section 3. Instead, we focus on the declaration of data constructors, the definition of functions, and the construction of expressions. These are the elements that play a central role in our execution model.

FlatCurry is a machine representation of Curry programs. As such, it is not intended to be read by human. For example, each variable is identified by an integer, there is only prefix function application, and pattern matching is broken down into a cascade of case distinctions. In the examples that follow, we present a sugared version of FlatCurry in which variables have symbolic names, typically the same as in Curry; the application of familiar infix operators is infix; and indentation, rather than parentheses and commas, show structure and grouping. The intent is to make the examples easier to read without altering the essence of FlatCurry.

In FlatCurry, data constructors are introduced by a type declaration. A type t has attributes such as a name and a visibility, and chief among these attributes is a set of constructors c_1, c_2, \dots, c_n . Each constructor c_i has similar attributes, along with an arity and type of each argument, which are not explicitly used in our discussion. The same information is available for operation symbols, but operations contain their code.

The abstract syntax of FlatCurry operations is summarized in Fig. 3.³ Each operation is defined by a single rule with a linear left-hand side, i.e., the argument variables x_1, \dots, x_n are pairwise different. The right-hand side of the definition consists of (1) variables introduced by the left-hand side or by a *let block* or by a *case* pattern, (2) constructor or function calls, (3) *case* expressions, (4) disjunctions, (5) *let* bindings, or (6) introduction of free variables. The patterns p_i in a *case* expression must be pairwise different constructors applied to variables. Therefore, deep patterns in source programs are represented by nested *case* expressions.

³ In contrast to some other presentations of FlatCurry (e.g., [1, 14]), we omit the difference between rigid and flexible case expressions.

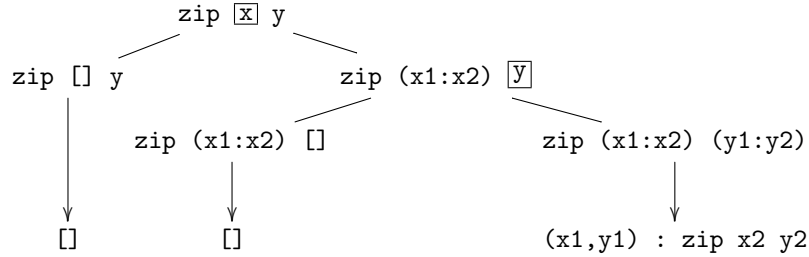
$D ::= f(x_1, \dots, x_n) = e$	(function definition)
$e ::= x$	(variable)
$c(e_1, \dots, e_n)$	(constructor call)
$f(e_1, \dots, e_n)$	(function call)
$\text{case } e \text{ of } \{p_1 \rightarrow e_1; \dots; p_n \rightarrow e_n\}$	(case expression)
$e_1 \text{ or } e_2$	(disjunction)
$\text{let } \{x_1 = e_1; \dots; x_n = e_n\} \text{ in } e$	(let binding)
$\text{let } x_1, \dots, x_n \text{ free in } e$	(free variables)
$p ::= c(x_1, \dots, x_n)$	(pattern)

Fig. 3. Abstract syntax of function definitions in FlatCurry

Case expressions closely resemble definitional trees. We recall that a *definitional tree* of some operation f , of arity n , is a hierarchical structure of expressions of the form $f p_1 \dots p_n$, where p_i is a pattern. Since f is constant and provides no information, except to ease readability, we also call these expressions patterns. The pattern at the root of the tree is $f x_1 \dots x_n$, where the x_i 's are distinct variables. The patterns at the leaves are the left-hand sides of the rules of f , except from the names of the variables. For ease of understanding, in pictorial representations of definitional trees we add the right-hand side of the rules too. If $f p_1 \dots p_n$ is a branch node, β , of the tree, a variable x in some p_j is singled out. We call the variable x *inductive*. The pattern in a child of β is $f p_1 \dots q_j \dots p_n$ where q_j is obtained from p_j by replacing x with $c y_1 \dots y_k$, where c is a constructor of the type of x and each y_i is a fresh variable. For example, consider the usual operation `zip` for zipping two lists:

$$\begin{aligned}
 \text{zip } [] \quad y &= [] \\
 \text{zip } (x1:x2) \quad [] &= [] \\
 \text{zip } (x1:x2) \quad (y1:y2) &= (x1,y1) : \text{zip } x2 \ y2
 \end{aligned} \tag{3}$$

The corresponding definitional tree is shown below where the inductive variable is boxed.



The FlatCurry code of the rules of operation `zip`, closely corresponds to the code below. This would be harder for the programmer to write than (3) and less readable, but is semantically equivalent. Every program can be transformed into an equivalent program in which every operation has a definitional tree [2] which can be obtained from the operation's rules with a relatively simple algorithm [3].

```

zip x y = case x of
  []      → []
  (x1:x2) → case y of
    []      → []
    (y1:y2) → (x1,y1) : zip x2 y2

```

Expressions are the final relevant element of FlatCurry. As the code of `zip` shows, an expression can be a literal, like `[]`; an application of constructors and operations to expressions possibly containing variables, like `(x1,y1) : zip x2 y2`; or a case expression, like `case y of ...`. FlatCurry also has *let blocks* to support the construction of cyclic graphs, as shown in Fig. 1.

FlatCurry programs cannot be directly mapped to code in a C-like target language. There are two problems: case expressions as arguments of a symbol application, and let blocks with shared or mutually recursive variables. A contrived example of the first is:

```

3 + case x of True → 1; False → 2

```

An example of the second is shown in Fig. 1. ICurry proposes a solution to these problems in a language-independent form which is suitable for the imperative paradigm.

5 ICurry

In this section we define ICurry, discuss how to map it to imperative code that implements our earlier model of computation, and show how to obtain it from FlatCurry.

5.1 ICurry Definition

ICurry is a format of Curry programs similar in intent to FlatCurry. The purpose of both is to represent a Curry program into a format with a small number of simple constructs. Properties and manipulations of programs can be more easily investigated and executed in these formats. ICurry is specifically intended for compilation into a low-level language. Each ICurry construct can be translated into a similar construct of languages such as C, Java or Python. This should become apparent once we describe the constructs.

ICurry's data consists of nested applications of symbols represented as graphs. ICurry's key constructs provide the declaration or definition of symbols and variables, construction of graph nodes, assignment, and conditional executions of these constructs. Rewriting steps are implemented in two phases, once the redex and rule are determined. First, the replacement of the redex is constructed. This is defined by the right-hand side of the rule. Then, the successors pointing to the root of the redex are redirected [12, Def. 8], through assignments, to point to the root of the replacement.

The declaration of data constructors in ICurry is identical to that in FlatCurry as described earlier. However, the constructors of a type are in an

D	$::= f = blk$	(function definition)
blk	$::= decl_1 \dots decl_k asgn_1 \dots asgn_n stm$	(block)
$decl$	$::= declare\ x$	(local variable declaration)
	$free\ x$	(free variable declaration)
$asgn$	$::= v = exp$	(variable assignment)
stm	$::= return\ exp$	(return statement)
	$exempt$	(failure statement)
	$case\ x\ of\ \{c_1 \rightarrow blk_1; \dots; c_n \rightarrow blk_n\}$	(case statement)
exp	$::= v$	(variable)
	$NODE(l, exp_1, \dots, exp_n)$	(node construction)
	$exp_1\ or\ exp_2$	(disjunction)
v	$::= x$	(local variable)
	$v[i]$	(node access)
	$ROOT$	(root of function call)
l	$::= c$	(constructor symbol)
	f	(function symbol)
	$LABEL(v)$	(node label symbol)

Fig. 4. Abstract syntax of function definitions in ICurry

arbitrary, but fixed, order. Therefore, we can talk of the first, second, etc., constructor of a type. This index is an attribute of constructor symbols which we call the *tag*. The tag is used to provide efficient pattern matching and to simplify the execution of some strategy. We will return to this topic in Sect. 5.2.

The abstract syntax of operations in ICurry is summarized in Fig. 4. In FlatCurry, the body of a function is an expression. In ICurry, it is a block of statements, the last of which returns an expression. We describe expressions first.

Expressions are nested symbol applications represented as graphs. Therefore, an expression is either a *variable*, which refers to a node in a graph, or a *symbol application*. ICurry makes an application explicit with a *directive*, `NODE`, that constructs a graph node from its label (1st argument) and its successors (remaining arguments), and returns a reference to the node. Accordingly, there are directives to access node components: assuming x is a variable referring to a node, `LABEL(x)` retrieves the label, and $x[k]$ retrieves the k -th successor of the node. A variable is either *free* or *local*. The ICurry format distinguishes between constructor and function application, and between full and partial application. We do not discuss these details in this paper. It is expected that by providing this additional information, processors will be able to generate low-level code more easily, and the generated code should be easier to optimize.

In ICurry, there are only a handful of statement kinds: *declaration* of a variable, and *assignment* to a variable, *return*, and *case* expressions. Following FlatCurry, variables are represented by integers. A declaration introduces a variable. An assignment to a variable is a reference to a graph node. Successors of a node referenced by x are accessed through the $x[...]$ construct. Arguments passed to functions are accessed through local variables. The return statement is intended to return an expression, the result of a function call. Case expressions

in ICurry are structurally similar to those in FlatCurry, but with two differences for algebraically defined types, which have a finite number of data constructors. First, the branches of a case expression are in tag order. We will justify this decision in Sect. 5.2. Second, the set of branches of a case expression is complete, i.e., there is a branch for each constructor of the type.

ICurry code begins with a declaration, and possible assignment, of some variables. It is then followed by either a case statement, or a return statement. Each branch of the case expression may declare and assign variables, and may lead to either another case statement, or a return statement.

Below, we present two examples. The first example is the code of function `oneTwo`, a constant, of Fig. 1:

```
function oneTwo
  declare x
  declare y
  x = NODE(:, NODE(1), y)
  y = NODE(:, NODE(2), x)
  x[2] = y
  return x
```

Symbol application is explicit through `NODE`. The notation `x[2]` stands for the second successor of `x`, which is not yet known when the node referenced by `x` is constructed (since the value of `y` is not yet specified at that point). This allows constructing nodes with a varying number of successors.

The second example is the code of `head`, the usual function returning the head of a non-empty list:

$$\text{head } (x:_) = x \tag{4}$$

The rule of `head` for the argument `[]` is missing in the Curry source code. Consequently, the case branch for the argument `[]` is missing in FlatCurry, too. ICurry has a distinguished expression, `exempt`, to capture the absence of a rule:

```
function head
  declare arg
  arg = ROOT[1]
  case LABEL(arg) of
    [] → exempt
    :  → return arg[1]
```

where `ROOT` is a reference to the root of the expression being evaluated. This expression is rooted by `head`, which is the reason why it is passed to function `head`.

5.2 ICurry Use

The stated goal of ICurry is to be a format of Curry programs suitable for translation into a low-level language. Below, we briefly report our experience in translating ICurry into various target languages. Table 5 shows the size of a Curry program that translates ICurry into a target language. The numerical

values in the table, extended from [22], count the lines of code of the translator. The table is only indicative since “lines of code” is not an accurate measure, and some earlier compilers use older variants of ICurry that have evolved over time. In all these cases, each ICurry construct has a direct translation into the target

C	441
Python	342
Java	790
JavaScript	632
JSON	232

Fig. 5. Number of Curry source lines of code for various translators from ICurry to a target language.

language. The following details refer to the translation into C. Declarations and assignment are the same as in C. The ICurry statements are translated as follows: (1) the ICurry *return* is the same as in C, (2) an ICurry *case statement* is translated into a C *switch statement* where the case selector is the tag of a node, and (3) the ICurry *exempt* statement is translated into code that, when executed, terminates the executing computation without producing any result. This is justified by the facts that the evaluation strategy executes only needed steps, and that failures in non-deterministic programs are natural and expected, therefore they should be silently ignored.

The translation into JSON is used by Sprite [9], a Curry system under development, whose target language is C++. The JSON format is more convenient when the client of ICurry code is not coded in Curry. The ICurry case expressions of a function’s code contain a branch for each constructor in the argument’s type and a branch for each of the following: the choice symbol, the failure symbol, any function symbol [9, Fig. 2], and any free variable. A dispatch table, which is addressed by the argument’s label’s tag, efficiently selects the branch to be executed. The behavior of the additional branches is described below, and is the same across all the functions of a program. A choice symbol in a pattern matched position results in the execution of a pull-tabling step [4, 11]. A failure is propagated to the context. A function symbol triggers the evaluation of the expression rooted by this symbol. Finally, a free variable is instantiated to a choice of *shallow* patterns of the same type as the variable. As an example, the evaluation of `head x where x free`, instantiates `x` to `[] ? (y:ys) where y, ys free`. The alternative, `[]`, will result in failure. This can be determined at compile time and removed during optimizations.

A similar approach is described in [22] where compilers from Curry to Python and to C are developed. As the table above indicates, the compilers (written in Curry) are quite compact. We observe that FlatCurry covers the complete language, since it is the basis for robust Curry implementations, like PAKCS and

KiCS2, and the natural/operational semantics of Curry are defined in FlatCurry [1]. ICurry contains the same information as FlatCurry except type information, since the type correctness of a program has been verified at the point of the compilation process in which ICurry is used.

5.3 ICurry Generation

Current Curry distributions such as PAKCS [17] and KiCS2 [10] provide a package with the definition of FlatCurry and a rich API for its construction and manipulation. Therefore, the ICurry format of a Curry program is conveniently obtained from the FlatCurry format of that program.

A fundamental difference between the two formats concerns expressions. Expressions in FlatCurry may contain *cases* and *lets* as the arguments of a function application. These are banned in ICurry which allows only nested functional application. The reason is that the latter can be directly translated into various imperative languages, where the former cannot. Therefore, any *case* and *let* constructs that are the arguments of a function application are replaced by calls to newly created functions. These new functions execute these constructs at the top level. This transformation takes FlatCurry into ICurry, but it could be executed from FlatCurry into itself, or even from source Curry into itself. Our contrived example below shows the latter for ease of understanding. The code of function *g* is irrelevant, therefore, it is not shown:

```
f x = g x (case x of ...)
```

is transformed into:

```
f x = g x (h x)
h x = case x of ...
```

The offending *case*, as an argument of the application of *g*, has been replaced by a call to a newly created function, *h*. In function *h*, the *case* is no longer an argument of a function application.

The second major difference between FlatCurry and ICurry concerns case expressions. FlatCurry matches a selector against shallow constructor expressions, where ICurry matches against constructor symbols. Furthermore, the set of these symbols is complete and ordered in ICurry. The transformation is relatively simple, except it may require non-local information. A function in a module *M* may pattern match on some instance of a type *t* that is not declared in *M*. Therefore, the constructors of *t* must be accessed in some module different from that being compiled.

A third significant difference between FlatCurry and ICurry concerns *let blocks*. They are banned in ICurry, and replaced by the explicit construction of nodes, and by the assignment of these nodes' references to local variables.

An algorithm to translate FlatCurry into ICurry is shown in Appendix A.

6 Related Work and Conclusion

Our work is centered on the compilation of Curry programs. As in many compilers, our approach is transformational. To compile a Curry program P , we translate P into a language, called *target*, for which a compiler already exists. This is the same route followed by other Curry compilers like PAKCS [17] and KiCS2 [10].

PAKCS translates source Curry code into Prolog, leveraging the existence of native free variables and non-determinism in Prolog. KiCS2 translates source Curry code into Haskell, leveraging the existence of first-class function and their efficient demand-driven execution in Haskell. Both of these compilers use FlatCurry as an intermediate language. They have the same front end which translates Curry into FlatCurry. The use of FlatCurry simplifies the translation process, but is still appropriate to express Curry computations without much effort. FlatCurry has some relatively high-level constructs that can be mapped directly into Prolog and Haskell, because these languages are high-level too.

Before ICurry, a Curry compiler targeting a C-like language would handle certain high-level constructs of FlatCurry in whichever way each programmer would choose. This led to both duplications of code and unnecessary differences. ICurry originates from these efforts. It abstracts the ideas that, over time, proved to be simple and effective in a language-independent way. With ICurry, the effort to produce a Curry compiler targeting an imperative language is both shortened, because more of the front end can be reused, and simplified, because the starting point of the translation is independent of the target and is well understood.

Our work is complementary to, but independent of, other efforts toward the compilation of Curry programs. These efforts include the development of evaluation strategies [5], or the handling of non-determinism [4, 11].

Future work should investigate ICurry to ICurry transformations that are likely to optimize the generated code. For example, different orders of the declaration of variables in a let block lead to different numbers of assignments. Also, cases expressions as arguments of function call can be moved outside the call in some situations rather than be replaced by a call to a new function.

Acknowledgments This material is based in part upon work supported by the National Science Foundation under Grant No. 1317249.

References

1. E. Albert, M. Hanus, F. Huch, J. Oliver, and G. Vidal. Operational semantics for declarative multi-paradigm languages. *Journal of Symbolic Computation*, 40(1):795–829, 2005.
2. S. Antoy. Constructor-based conditional narrowing. In *Proc. of the 3rd International ACM SIGPLAN Conference on Principles and Practice of Declarative Programming (PPDP 2001)*, pages 199–206. ACM Press, 2001.
3. S. Antoy. Evaluation strategies for functional logic programming. *Journal of Symbolic Computation*, 40(1):875–903, 2005.

4. S. Antoy. On the correctness of pull-tabbing. *TPLP*, 11(4-5):713–730, 2011.
5. S. Antoy, R. Echahed, and M. Hanus. A needed narrowing strategy. *Journal of the ACM*, 47(4):776–822, 2000.
6. S. Antoy and M. Hanus. Functional logic programming. *Communications of the ACM*, 53(4):74–85, 2010.
7. S. Antoy and M. Hanus. Contracts and specifications for functional logic programming. In *Proc. of the 14th International Symposium on Practical Aspects of Declarative Languages (PADL 2012)*, pages 33–47. Springer LNCS 7149, 2012.
8. S. Antoy, M. Hanus, and S. Libby. Proving non-deterministic computations in Agda. In *Proc. of the 24th International Workshop on Functional and (Constraint) Logic Programming (WFLP 2016)*, volume 234 of *Electronic Proceedings in Theoretical Computer Science*, pages 180–195. Open Publishing Association, 2017.
9. S. Antoy and A. Jost. A new functional-logic compiler for Curry: Sprite. In *Proceedings of the 26th International Symposium on Logic-Based Program Synthesis and Transformation (LOPSTR 2016)*, pages 97–113. Springer LNCS 10184, 2016.
10. B. Braßel, M. Hanus, B. Peemöller, and F. Reck. KiCS2: A new compiler from Curry to Haskell. In *Proc. 20th International Workshop on Functional and (Constraint) Logic Programming (WFLP 2011)*, pages 1–18. Springer LNCS 6816, 2011.
11. B. Brassel and F. Huch. On a tighter integration of functional and logic programming. In *APLAS'07: Proceedings of the 5th Asian conference on Programming languages and systems*, pages 122–138, Berlin, Heidelberg, 2007. Springer-Verlag.
12. R. Echahed and J.-C. Janodet. On constructor-based graph rewriting systems. Research report IMAG 985-I, IMAG-LSR, CNRS, Grenoble, 1997.
13. M. Hanus. FlatCurry: An intermediate representation for Curry programs, 2008. Available at <http://www.informatik.uni-kiel.de/~curry/flat/>.
14. M. Hanus. Functional logic programming: From theory to Curry. In *Programming Logics - Essays in Memory of Harald Ganzinger*, pages 123–168. Springer LNCS 7797, 2013.
15. M. Hanus. Combining static and dynamic contract checking for Curry. In *Proceedings of the 27th International Symposium on Logic-Based Program Synthesis and Transformation (LOPSTR 2017)*, pages 323–340. Springer LNCS 10855, 2017.
16. M. Hanus. Verifying fail-free declarative programs. In *Proceedings of the 20th International Symposium on Principles and Practice of Declarative Programming (PPDP 2018)*, pages 12:1–12:13. ACM Press, 2018.
17. M. Hanus, S. Antoy, B. Braßel, M. Engelke, K. Höppner, J. Koj, P. Niederau, R. Sadre, F. Steiner, and F. Teegen. PAKCS: The Portland Aachen Kiel Curry System. Available at <http://www.informatik.uni-kiel.de/~pakcs/>, 2018.
18. M. Hanus and F. Skrlac. A modular and generic analysis server system for functional logic programs. In *Proc. ACM SIGPLAN 2014 Workshop on Partial Evaluation and Program Manipulation (PEPM'14)*, pages 181–188. ACM Press, 2014.
19. M. J. O'Donnell. *Equational Logic as a Programming Language*. MIT Press, 1985.
20. D. Plump. Term graph rewriting. In H. Ehrig, G. Engels, H.-J. Kreowski, and G. Rozenberg, editors, *Handbook of Graph Grammars and Computing by Graph Transformation, Volume 2: Applications, Languages and Tools*, pages 3–61. World Scientific, 1999.
21. U.S. Reddy. Narrowing as the operational semantics of functional languages. In *Proc. IEEE Int. Symposium on Logic Programming*, pages 138–151, Boston, 1985.
22. M.A. Wittorf. Generic translation of Curry programs into imperative programs (in German). Master's thesis, Kiel University, 2018.

A Translating FlatCurry to ICurry

A *pure expression* is an expression that only contains literals, variables, constructor applications, and function applications. Any *or* expression and function application may only contain pure expressions. The scrutinee of a case expression must be a variable, literal, or constructor application. An assignment in a *let* expression must be a pure expression or an *or* expression. The branches of a case expression must match all constructors of a data type in an order fixed by the definition of that data type. Branches missing in the original Curry program contain \perp in their right-hand side.

The algorithm is divided into five functions which are described in Fig. 6. \mathcal{F} translates a FlatCurry function into an ICurry function. \mathcal{B} translates a FlatCurry expression into an ICurry block. \mathcal{D} extracts all of the variables declared in a FlatCurry expression. \mathcal{A} generates necessary assignments for ICurry variables. \mathcal{E} translates a FlatCurry expression into an ICurry expression.

The functions \mathcal{F} and \mathcal{E} are straightforward translations. \mathcal{F} simply makes a block, with the root of the block being set to the root of the function. \mathcal{E} is almost entirely a straight translation, but there is one technical point. In a case expression, each branch must be translated into its own block. However, each of the variables in the pattern of a branch need to be related to the scrutinee of the case. This is achieved by setting the root of the block to the scrutinee of the case.

The \mathcal{B} function creates an ICurry block. Blocks are more complicated to construct. Each block will have a root and a list of variables. The root is the root of the expression that created the block. For a function, the root is the root of the function expression. For a case branch, the root is the root of the scrutinee of the case. The variables of a block are the parameters of a function, or the pattern variables of a branch. After declaring variables, all variable in any *let* expressions are declared with the \mathcal{D} function. Then each variable is assigned an expression with the \mathcal{A} function. If either \mathcal{D} or \mathcal{A} is undefined for some expression, it does not generate ICurry code. Finally, we translate the expression into an ICurry statement with the \mathcal{E} function.

The \mathcal{D} function declares variables declared in a *let* or *free* expression. If there is a *case* expression, then a new variable x_e is declared.

The \mathcal{A} function assigns variables in a *let* or *free* expression. The expression for all variables in a *let* is translated with the \mathcal{E} function. Next, if there are any variables declared in the *let* block that are used in one of the expressions, they need to be filled in. Finally, if there is a case expression, we assign x_e to be the root of the scrutinee.

$$\begin{aligned}
\mathcal{F}(f(x_1, \dots, x_n) = e) &:= f = \mathcal{B}(x_1, \dots, x_n, e, \text{ROOT}) \\
\mathcal{B}(x_1, \dots, x_n, \perp, \text{root}) &:= \text{exempt} \\
\mathcal{B}(x_1, \dots, x_n, e, \text{root}) &:= \\
&\quad \text{declare } x_1 \\
&\quad \dots \\
&\quad \text{declare } x_n \\
&\quad \mathcal{D}(e) \\
&\quad x_1 = \text{root}[1] \\
&\quad \dots \\
&\quad x_n = \text{root}[n] \\
&\quad \mathcal{A}(e) \\
&\quad \text{return } \mathcal{E}(e) \quad (\text{omit return if } \mathcal{E}(e) \text{ is a case}) \\
\mathcal{D}(\text{let } x_1, \dots, x_n \text{ free in } e) &:= \\
&\quad \text{free } x_1 \\
&\quad \dots \\
&\quad \text{free } x_n \\
\mathcal{D}(\text{let } \{x_1 = e_1; \dots; x_n = e_n\} \text{ in } e) &:= \\
&\quad \text{declare } x_1 \\
&\quad \dots \\
&\quad \text{declare } x_n \\
\mathcal{D}(\text{case } e \text{ of } \{p_1 \rightarrow e_1; \dots; p_n \rightarrow e_n\}) &:= \text{declare } x_e \\
\mathcal{A}(\text{let } \{x_1 = e_1; \dots; x_n = e_n\} \text{ in } e) &:= \\
&\quad x_1 = \mathcal{E}(e_1) \\
&\quad \dots \\
&\quad x_n = \mathcal{E}(e_n) \\
&\quad x_1[p] = x_i \quad (\text{for each occurrence of } x_i \text{ in } e_1 \text{ at position } p) \\
&\quad \dots \\
&\quad x_1[p] = x_i \quad (\text{for each occurrence of } x_i \text{ in } e_n \text{ at position } p) \\
\mathcal{A}(\text{case } e \text{ of } \{p_1 \rightarrow e_1; \dots; p_n \rightarrow e_n\}) &:= x_e = \mathcal{E}(e) \\
\mathcal{E}(x) &:= x \\
\mathcal{E}(c(e_1, \dots, e_n)) &:= \text{NODE}(c, \mathcal{E}(e_1), \dots, \mathcal{E}(e_n)) \\
\mathcal{E}(f(e_1, \dots, e_n)) &:= \text{NODE}(f, \mathcal{E}(e_1), \dots, \mathcal{E}(e_n)) \\
\mathcal{E}(e_1 \text{ or } e_2) &:= \mathcal{E}(e_1) \text{ or } \mathcal{E}(e_2) \\
\mathcal{E}(\text{let } \{x_1 = e_1; \dots; x_n = e_n\} \text{ in } e) &:= \mathcal{E}(e) \\
\mathcal{E}(\text{let } \{x_1, \dots, x_n\} \text{ free in } e) &:= \mathcal{E}(e) \\
\mathcal{E}(\text{case } e \text{ of } \{c(x_{11}, \dots, x_{1m}) \rightarrow e_1; \dots; c(x_{n1}, \dots, x_{nk}) \rightarrow e_n\}) &:= \\
&\quad \text{case } \mathcal{E}(e) \text{ of } \{ \mathcal{B}(x_{11}, \dots, x_{1m}, e_1, x_e); \\
&\quad \dots; \\
&\quad \mathcal{B}(x_{n1}, \dots, x_{nk}, e_n, x_e); \}
\end{aligned}$$

Fig. 6. Algorithm for translating FlatCurry into ICurry