

Variant Translators for Version 8.10 of Icon

Ralph E. Griswold and Kenneth Walker

Department of Computer Science, The University of Arizona

1. Introduction

A preprocessor, which translates text from source language A to source language B ,

$$A \rightarrow B$$

is a popular and effective means of implementing A , given an implementation of B . B is referred to as the target language. Ratfor [1] is perhaps the best known and most widely used example of this technique, although there are many others.

In some cases A is a variant of B . An example is C_g [2], a variant of C that includes a generator facility similar to Icon's [3]. C_g consists of C and some additional syntax that a preprocessor translates into standard C . A run-time system provides the necessary semantic support for generators. Note that the C_g preprocessor is a source-to-source translator:

$$C_g \rightarrow C$$

where C_g differs from C only in the addition of a few syntactic constructs. This can be viewed as an instance of a more general paradigm:

$$A^+ \rightarrow A$$

The term “translator” is used here in the general sense, and includes both source-to-source translators, such as preprocessors, and source-to-object translators, such as compilers. In practice, the application of a source-to-source translator (preprocessor) may be followed by the application of a source-to-object translator (compiler). The combination is, of course, also a translator.

The term “variant translator” is used here to refer to a translator that differs in its action, in some respect, from a standard one for a language. The applications described in this report relate to source-to-source translators, although the term “preprocessor” is too restrictive to describe all of them.

There are many uses for variant translators. Some of them are:

- the addition of syntactic constructions to produce a superset of a language, as in the case of C_g
- the deletion of features in order to subset a language
- the translation of one source language into another [4]
- the addition of monitoring code, written in the target language
- the insertion of termination code to output monitoring data
- the insertion of initialization code to incorporate additional run-time facilities
- the insertion of code for debugging and checking purposes [5, 6]

Note that in several cases, the translations can be characterized by

$$A \rightarrow A$$

The input text and the output text may be different, but they are both in A . Both the input and the output of the variant translator can be processed by a standard translator for the target language A .

One way to implement a variant translator is to modify a standard source-to-object translator, avoiding the preprocessor. This approach may or may not be easy, depending on the translator. In general, it involves modifying the code generator, which often is tricky and error prone. Furthermore, if the variant is an experiment, the effort involved may be prohibitive.

The standard way to produce a variant translator is the one that is most often used for preprocessors in general, including ones that do not fit the variant translator paradigm — writing a stand-alone program in any convenient language. In the case of Ratfor, the preprocessor is written in Ratfor, providing the advantages of bootstrapping.

This approach presents several problems. In the first place, writing a complete, efficient, and correct preprocessor is a substantial undertaking. In experimental work, this effort may be unwarranted, and it is common to write the preprocessor in a high-level language, handling only the variant portion of the syntax, leaving the detection of errors to the final translator. Such preprocessors have the virtue of being easy to produce, but they often are slow, frequently unfaithful to the source language, and the failure to parse the input language completely may lead to mysterious results when errors are detected, out of context, by the final translator.

Modern tools such as Lex [7] and Yacc [8], that operate on grammatical specifications, have made the production of compilers (and hence translators in general) comparatively easy and have removed many of the sources of error that are commonly found in hand-tailored translators. Nonetheless, the construction of a translator for a large and complicated language is still a substantial undertaking.

If, however, a translator already exists for a language that is based on the use of such tools, it may be easy to produce a variant translator that is efficient and demonstrably correct by modifying grammatical specifications. The key is the use of these tools to produce a source-to-source translator, rather than producing a source-to-object translator. This technique was used in Cg. An existing Yacc specification for the C compiler was modified to generate C source code instead of object code. The idea is a simple one, but it has considerable utility and can be applied to a wide range of situations.

This report describes a system that uses this approach for the construction of variant translators for Icon. This system requires Icon, Yacc, and C.

2. Overview of Variant Translators for Icon

The heart of the system for constructing variant translators for Icon consists of an “identity translator”. The output of this identity translator differs from its input only in the arrangement of nonsemantic “white space” and in the insertion of semicolons between expressions, which are optional in some places in Icon programs.

The identity translator uses the same Yacc grammar as the regular Icon translator, but uses different semantic actions. These semantic actions are cast as macro definitions in the grammar, which are expanded before the grammar is translated by Yacc into a parser. One set of macros is supplied for the regular Icon translator and another set is supplied for the identity translator. The macros used by the identity translator echo the input text, producing source-code output. In addition to the grammar, other code is shared between the two translators, insuring a high degree of consistency between the two systems.

A variant translator is created by first creating an identity translator and then modifying it. There is a shell script for producing identity translators and associated support software to simplify the process of making modifications. This support software allows macro definitions to be changed via specification files, minimizing the clerical work needed to vary the format of the output. There also is a provision for including user functions in the parser, so that more complicated operations can be written in C. Finally, the grammar for the identity translator can be modified in order to make structural changes in the syntax.

The following sections describe this system in more detail and include a number of examples of its use. The material that follows is intended for use in conjunction with listings of source files from the variant translator system.

3. The Grammar for the Icon Identity Translator

The grammar for Icon is in the file `h/grammar.h` in the directory in which variant translators are installed. Many variant translators can be constructed without modifying this grammar, and minor modifications can be made to it without a detailed knowledge of its structure. Knowledge of a few aspects of this grammar are important, however, for understanding the translation process.

There are two types of semantic actions. The semantic action for a declaration outputs text. The semantic action for a component of a declaration, such as an identifier list or an expression, assigns a string to the Yacc attribute for the component. Declarations are parsed by the production:

```

decl      : record {Recdcl($1);} ;
          | proc {Proccl($1);} ;
          | global {Globdcl($1);} ;
          | link {Linkdcl($1);} ;
          | invocable {Invocdcl($1);} ;

```

The non-terminals `record`, `proc`, `global`, `link`, and `invocable` each produce a string and the corresponding macro `Recdcl`, `Proccl`, `Globdcl`, `Linkdcl`, or `Invocdcl` prints the string.

Because the grammar is used for both the regular Icon translator and the variant translator system, the macro calls must be more general than what is required for either one alone. Consider the production for `global`:

```

global    : GLOBAL {Global0($1);} idlist {Global1($1, $2, $3);} ;

```

The macro `Global0` is needed in the regular translator, but performs no operation in the identity translator. The macro `Global1` does the work in the identity translator; it concatenates "global " with the string produced by `idlist`, and this new string becomes the result of this production. The macro `Global1` is passed `$1`, `$2`, and `$3` even though it only uses `$3`. This is done for generality.

The rules and the definitions that construct and output strings are provided as part of the identity translator. When a variant translator is constructed, changes are necessary only in situations in which the input is not to be echoed in the output.

Modifications and additions to the standard grammar require a more thorough understanding of the structure of the grammar.

4. Macro Definitions

The purpose of using macro calls in the semantic actions of the grammar is to separate the structure of the grammar from the format of the output and to allow the output format to be specified without modification of the grammar.

The macro definitions for declarations are all the same. For example the definition of `Global` for the identity translator is:

```

#define Globdcl(x)      if (!nocode) treeprt(x); treefree(x)

```

The variable `nocode` is set when an error is detected during parsing. This helps prevent the variant translator from generating a program with syntax errors. The reason for doing this is that the output of a variant translator is usually piped directly into the regular Icon translator. If syntax errors were propagated, two error messages would result: one from the variant translator and one from the Icon translator. The message from the variant translator is the one that is wanted because it references the line number of the original source whereas the message from the Icon translator references the line number of the generated source.

The function `treeprt()` prints a string and the function `treefree()` reclaims storage. See the Section 5 for details of string representation.

4.1 Specifications for Macros

The macro definitions for expressions produce strings, generally resulting from the concatenation of strings produced by other rules. In order to simplify the definition of macros, a specification format is provided. Specifications are processed by a program that produces the actual definitions. For example, the macro `While1` is used in the rule

```

WHILE expr DO expr {While1($1, $2, $3, $4);} ;

```

A specification for this macro to produce an identity translation is:

```

While1(w, x, y, z) "while " x " do " z

```

Tabs separate the components of the specification. The first component is the prototype for the macro call, which may include optional arguments enclosed in parentheses as illustrated by the example above. The remaining components are the strings to be concatenated with the result being assigned to the Yacc pseudo-variable `$$`.

Specification lines that begin with # or which are empty are treated as comments. A set of lines delineated by %{ and %} is copied unchanged. The “braces” %{ and %} must each occur alone on a separate line; these two delimiting lines are not copied. This feature allows the inclusion of actual macro definitions, as opposed to specifications, and the inclusion of C definitions. The standard macro definitions supplied for the identity translator include examples of these features. These definitions are the file `ident.defs`.

Definitions can be changed by modifying the standard ones or by adding new definitions. In the case of duplicate definitions, the last one holds. Definitions can be provided in several files, so variant definitions can be provided in a separate file that is processed after the standard definitions. See Sec. 8.

Definitions can be deleted by providing a specification that consists only of a prototype for the call. For example, the specification

```
While1()
```

deletes the definition for `While1`. This is a convenient way to insure a macro is undefined. It is usually used along with the copy feature to introduce macro definitions that cannot be generated by the specification system. For example, the following specifications eliminate reclamation of storage, preserving strings between declarations.

```
Globdcl()
Linkdcl()
Procdcl()
Recdcl()
%{
#define Globdcl(x)      if (!ncode) treeprt(x);
#define Linkdcl(x)     if (!ncode) treeprt(x);
#define Procdcl(x)     if (!ncode) treeprt(x);
#define Recdcl(x)      if (!ncode) treeprt(x);
%}
```

4.2 Macros for Icon Operators

There is a distinct macro name for each Icon operator. For example, `Blim(x, y, z)` is the macro for a limitation expression,

```
expr1 \ expr2
```

Note that the parameter `y` is the operator symbol itself. To avoid having to know the names of the macros for the operators, specifications allow the use of operator symbols in prototypes. The symbols are automatically replaced by the appropriate names. Thus

```
\(x, y, z)
```

can be used in a specification in place of

```
Blim(x, y, z)
```

Unary operators are similar. For example, `Uqmark(x, y)`, which is the macro for `?expr`, can be specified as `?(x, by)`. In this case the parameter `x` is the operator symbol.

In most cases, all operators of the same kind are translated in the same way. Since Icon has many operators, a generic form of specification is provided to allow the definition of all operators in a category to be given by a single specification. In a specification, a string of the form `<type>` indicates a category of operators. The categories are:

```
<uop>    unary operators, except as follows
<ucs>    control structures in unary operator format
<bop>    binary operators, except as follows
<aop>    assignment operators
<bcs>    control structures in binary operator format
```

The category `<ucs>` consists only of `|`. The category `<bcs>` consists of `?`, `|`, `!`, and `\`.

For example, the specification for binary operators for identity translations is

```
<bop>(x, y, z)      x      " <bop> " z
```

This specification results in the definition for every binary operator: $+(x, y, z)$, $-(x, y, z)$, and so on. In such a specification, every occurrence of `<bop>` is replaced by the corresponding operator symbol. Note that blanks are necessary to separate the binary operator from its operands. Otherwise,

```
i * *s
```

would be translated into

```
i**s
```

which is equivalent to

```
i ** s
```

The division of operators into categories is based on their semantic properties. For example, a preprocessor may translate all unary operators in the same way, but translate the repeated alternation control structure into a programmer-defined control operation [9].

5. String Handling

Strings are represented as binary trees in which the leaves contain pointers to C strings. The building of these trees can be thought of as doing string concatenation using lazy evaluation. The concatenation operation just creates a new root node with its two operands as subtrees. The real concatenation is only done when the strings are written out. Another view of this is that concatenation builds a list of strings with the list implemented as a binary tree. This view allows “strings” to be treated as a list of tokens. This approach is useful in more complicated situations where there is a need to distinguish more than just syntactic structures. For example, the head of the main procedure can be distinguished from the heads of other procedures by looking at the second string in the list for the procedure declaration.

Strings come from three sources during translation: strings produced by the lexical analyzer, literal strings, and strings produced by semantic actions. The lexical analyzer produces nodes. The cases where the nodes that are produced by the lexical analyzer are of interest occur where strings are recognized for identifiers and literals — the tokens `IDENT`, `STRINGLIT`, `INTLIT`, `REALIT`, and `CSETLIT`. These nodes contain pointers to the strings recognized. (The actual strings are stored in a string space and remain there throughout execution of the translator.) These nodes can be used directly as a tree (of one node) of strings. Other nodes produced by the lexical analyzer, for example those for operators, do not contain strings. However, all of these nodes contain line and column numbers referring to the location of the token in the source text. This line and column information can be useful in variant translators that need to produce output that contains position information from the input.

A literal string must be coerced into a tree of one node. This is done with the C function

```
q(s)
```

This is handled automatically when macros are produced from specifications. For example, the specification

```
Fail(x)      "fail"
```

is translated into the macro

```
#define Fail(x) $$ = q("fail")
```

Most semantic actions concatenate two or more strings and produce a string. They use the C function

```
cat(n, t1, t2, ..., tn)
```

which takes a variable number of arguments and returns a pointer to the concatenated result. The first argument is the number of strings to be concatenated. The other arguments are the strings in tree format. The result is also in tree format.

As an example, the specification

```
While1(w, x, y, z)    "while " x    " do " z
```

produces the definition

```
#define While1(w, x, y, z) $$ = cat(4, q("while "), x, q(" do "), z)
```

Another function, `item(t, n)`, returns the n^{th} node in the “list” `t`. For example, the name of a procedure is contained in the second node in the list for the procedure declaration. Thus, if the procedure heading list is the value of `head`, `item(head, 2)` produces the procedure name.

There are three macros that produce values associated with a node. `Str0()` produces the string. For example, code conditional on the main procedure could be written as follows:

```
if (strcmp(Str0(item(head, 2)), "main") == 0) {  
    :  
    :  
}
```

As this example illustrates, semantic actions may be too complicated to be represented conveniently by macros. In such cases parser functions can be used. A file is provided for such functions. See Section 9 for an example.

The macros `Line` and `Col` produce the source-file line number and column, respectively, of the place where the text for the node begins. The use of these attributes is illustrated in Section 9.

In some sophisticated applications, variant translators may need other capabilities that are available in the translator system. For example, if a function produces a string, it may be necessary place this string in a place that survives the function call. The Icon translator has a string allocation facilities that can be used for this purpose: The macro `AppChar(lex_sbuf, c)` appends a character to the lexical analyzer’s string buffer and the function `str_install(&lex_sbuf)` saves the string in a string table. The use of such facilities requires more knowledge of the translator system than it is practical to provide here. Persons with special needs should study the translator in more detail.

6. Modifying Lexical Components of the Translator

The lexical analyzer for Icon is written in C rather than in Lex in order to make it easier to perform semicolon insertion and other complicated tasks that occur during lexical analysis. Specification files are used to build portions of the lexical analyzer, making it easy to modify. The three kinds of changes that are needed most often are the addition of new keywords, reserved words, and operators.

The identity translator accepts any identifier as a keyword, leaving its resolution to subsequent processing by the Icon translator. Nothing need be done to add a new keyword except for processing it properly in the variant translator.

The specification file `common/tokens.txt` contains a list of all reserved words and operator symbols. Each symbol has associated flags that indicate whether it can begin or end an expression. These flags are used for semicolon insertion.

To add a new reserved word, insert it in proper alphabetical order in the list of reserved words in `tokens.txt` and give it a new token name. To add a new operator, insert it in the list of operators in `tokens.txt` (order there is not important) and give it a new token name. The new token names must be added to the grammar.

The addition of a new operator also requires modifying the specification of a finite-state automaton, `comon/op.txt`. Its structure is straightforward.

7. Building a Variant Translator

The steps for setting up the directory structure for a variant translator are:

- create a directory for the translator
- make that directory the current directory
- execute the shell script `icon_vt` supplied with Version 8 of Icon

For example, if the variant translator is to be in the directory `xtran` and Icon is installed in `/usr/icon/bin`, the following commands will build the variant translator:

```
mkdir xtran
cd xtran
/usr/icon/bin/icon_vt
```

The shell script `icon_vt` creates a number of files in the new directory and in three sub-directories: `common`, `itrans`, and `h`. Unless changes to the lexical analyzer are needed, at most three files need to be modified to produce a new translator:

<code>variant.defs</code>	variant macro definitions (initially empty)
<code>variant.c</code>	parser functions (initially empty)
<code>h/grammar.h</code>	Yacc grammar for Icon

A non-empty `variant.c` usually requires `#include` files to provide needed declarations and definitions. See the example that follows.

The `make` file in the main translator directory just insures that the program `define` has been compiled and then does a `make` in the `itrans` directory. Performing a `make` in the `itrans` directory first combines `variant.defs` with the standard macro definitions (in `ident.defs`) and processes them to produce the definition file, `itrans/ident.h`. The C preprocessor is then used to expand the macros in `h/grammar.h` using these definitions and the result, after some “house keeping”, is put in `itrans/vgram.g`. Next, Yacc uses the grammar in `itrans/vgram.g` to build a new parser, `itrans/tparse.c`. There are over 200 shift/reduce conflicts in the identity translator. All of these conflicts are resolved properly. More conflicts should be expected if additions are made to the grammar. Reduce/reduce conflicts usually indicate errors in the grammar. Finally, all the components of the system are compiled, including `variant.c`, and linked to produce `vitran`, the variant translator.

Most of the errors that may occur in building a variant translator are obvious and easily fixed. Erroneous changes to the grammar, however, may be harder to detect and fix. Error messages from Yacc or from compiling `itrans/tparse.c` refer to line numbers in `itrans/vgram.g`. These errors must be related back to `variant.defs` or `h/grammar.h` by inspection of `itrans/vgram.g`.

8. Using a Variant Translator

The translator, `vitran`, takes an input file on the command line and translates it. The specification – in place of an input file indicates standard input. The output of `vitran` is written to standard output. For example,

```
vitran pre.icn >post.icn
```

translates the file `pre.icn` and produces the output in `post.icn`. The suffix `.icn` on the argument to `vitran` is optional; the example above can be written as:

```
vitran pre >post.icn
```

Assuming the variant translator produces Icon source language, `post.icn` can be translated into object code by

```
icont post.icn
```

where `icont` is the standard Icon command processor.

Variant translators accept the following options:

- `m` process input file with the macro processor `m4` before translation
- `s` suppress informative messages

-P do not generate #line directives

The -P option may be necessary to prevent the insertion of #line directives at places that result in syntactically erroneous output.

9. An Example

As an example of the construction of a variant translator, consider the problem of monitoring string concatenation in Icon programs, writing out the size of each string constructed by concatenation. One way to do this, of course, is to modify Icon itself, adding the necessary monitoring code to the C function that performs concatenation. An alternative approach, which does not require changes to Icon itself, is to produce a variant translator that translates concatenation operations into calls of an Icon procedure, but leaves everything else unchanged:

$$expr_1 \parallel expr_2 \rightarrow \text{Cat}(expr_1, expr_2)$$

The procedure `Cat()` might have the form:

```
procedure Cat(s1, s2)
  write(&errout,"concatenation: ",*s1 + *s2," characters")
  return s1 || s2
end
```

Such a procedure could be added to a preprocessed program (`Cat()` is not preprocessed itself) in order to produce the desired information when the program is run.

A single definition in `variant.defs` suffices:

$$\parallel(x, y, z) \quad \text{"Cat("} \quad x \quad \text{,"} \quad z \quad \text{)"}$$

Note, however, that Icon also has an augmented assignment operator for string concatenation:

$$expr_1 \parallel:= expr_2$$

This operation can be handled by the definition

$$\parallel:=(x, y, z) \quad x \quad \text{"} := \text{Cat("} \quad x \quad \text{,"} \quad z \quad \text{)"}$$

Observe that this definition is not precisely faithful to the semantics of Icon, since it causes `expr1` to be evaluated twice, while `expr2` is evaluated only once in the true augmented assignment operation. This problem cannot be avoided here, since all arguments are passed by value in Icon, but in practice, this discrepancy is unlikely to cause problems.

In the application of such a monitoring facility, it may be useful to have a provision whereby concatenation can be performed without being monitored. This can be accomplished by adding an alternative operator symbol for concatenation, such as

$$expr_1 \sim expr_2 \rightarrow expr_1 \parallel expr_2$$

Adding a new operator to the syntax of Icon requires modifying the grammar in `h/grammar.h`. Since this alternative concatenation operator should have the same precedence and associativity as the regular concatenation operator, it can be added to the definition of `expr5`:

```
expr5      : expr6 ;
           | expr5 CONCAT expr6 {Bcat($1,$2,$3);} ;
           | expr5 TILDE expr6 {Bacat($1,$2,$3);} ;
           | expr5 LCONCAT expr6 {Blcat($1,$2,$3);} ;
```

where `TILDE` is the token name for `~`. Then the definition of `Bacat()` can be added to `variant.defs`:

$$\text{Bacat}(x, y, z) \quad x \quad \text{"} \parallel \text{"} \quad z$$

Such changes to `grammar.h` usually increase the number of shift/reduce conflicts encountered by Yacc.

One difficulty with monitoring concatenation as described above is that the procedure `Cat()` must be added to the translated program. This can be accomplished automatically by arranging to have the code for `Cat()` written out

when the variant translator encounters the main procedure. This is a case where a parser function, as mentioned in Section 5, is more appropriate than a macro definition.

The first step is to change the specifications. The definition for the macro, `Proc1`, that produces procedure declarations is replaced by a call to a parser function. The changes to `variant.defs` are:

```
%{
nodeptr proc();
%}
Proc1(u, v , w, x, y, z)proc(u, w, x, y)
```

The C declaration for `proc()` is included in the file `vgram.g` and subsequently incorporated by Yacc into `tparse.c` where the call to `proc()` is compiled. Note that `proc()` returns a `nodeptr`.

The C function is placed in `variant.c`. It might have the form

```
#include "h/config.h"
#include "itran/tree.h"
#include "itran/tpproto.h"

nodeptr item(), cat(), q();

nodeptr proc(u, w, x, y)
nodeptr u, w, x, y;
{
    static char *catproc = "procedure Cat(s1, s2)\n\
write(&errout,\"concatenation: \", *s1 + *s2,\" characters\")\n\
return s1 || s2\n\
end\n";

    if (strcmp(Str0(item(u, 2)), "main") == 0)
        return cat(7, q(catproc), u, q("; \n"), w, x, y, q("end\n"));
    else
        return cat(6, u, q("; \n"), w, x, y, q("end\n"));
}
```

Thus, when the main procedure is encountered, the text for `Cat()` is written out before the text for the main procedure, but all other procedures are written out as they would be in the absence of this function.

One disadvantage of this way of providing the text for `Cat()` is that the literal string is long, complicated, and difficult to change. In addition, it is necessary to rebuild the variant translator in order to change `Cat()`. Since monitoring of this kind is likely to suggest changes to the format or nature of the data being written, it is useful to be able to change `Cat()` more easily. One solution to this problem is to produce a link declaration for the file containing the translated procedure rather than the text of the procedure. With this change, the parser function might have the form

```
nodeptr proc(u, w, x, y)
nodeptr u, w, x, y;
{
    if (strcmp(Str0(item(u, 2)), "main") == 0)
        return cat(7, q("link cat\n\n"), u, q("; \n"), w, x, y, q("end\n"));
    else
        return cat(6, u, q("; \n"), w, x, y, q("end\n"));
}
```

The monitoring facility described above produces information about all string concatenation operations, but it is not possible to distinguish among them. It might be more useful to know the amount of concatenation performed by each concatenation operation. This can be done if the location of the operator in the source program can be identified. As mentioned in Section 5, tree nodes contain line and column information provided by the lexical analyzer. Thus, the translation for the concatenation operations could provide this additional information as extra arguments to `Cat()`, which then could print out the locations along with information about the amount of

concatenation.

```
procedure Cat(s1, s2, i, j)
  write(&errorout,"concatenation: ", *s1 + *s2, " characters at [", i, ",", j, "]")
  return s1 || s2
end
```

The specifications for the translation of the concatenation operations might be changed to

```
%{
nodeptr proc(), Locargs();
%}
Proc1(u, v, w, x, y, z) proc(u, w, x, y)
||(x, y, z) "Cat(" x "," z Locargs(y) ")"
|:=(x, y, z) x " := Cat(" x "," z Locargs(y)")"
Bacat(x, y, z) x " || " z
```

where `Locargs()` is a parser function that produces a string consisting of the line and column numbers between commas. This function might have the form

```
nodeptr Locargs(x)
nodeptr x;
{
  char buf[25];
  char *s;

  sprintf(buf, "%d,%d", Col(x), Line(x));
  for (s = buf, s != 0, ++s)
    AppChar(lex_sbuf, *s);
  return q(str_install(&lex_sbuf));
}
```

The C function `sprintf()` is used to do the formatting. The resulting string is copied into the translator's string buffer as mentioned in Section 5. The string is installed by `str_install()`, which adds `\0` to null-terminate the string.

10. Conclusions

The system described here for producing variant translators for Icon has been used successfully to provide support for a number of language variants and tools. These include a list scanning facility [10], an animated display of pattern matching [11], An experimental language for manipulating sequences [12, 13], a SNOBOL4-like language with a syntax similar to Icon [4], an Icon program formatter, a tool for monitoring expression evaluation events, and a number of simpler tools.

The value of being able to construct a variant translator quickly and easily is best illustrated by the tool for monitoring expression evaluation events. This translator copies input to output, inserting calls on procedures that tally expression activations, the production of results, and expression resumptions. A similar system was built for Version 2 of Icon [14] and was used to analyze the performance and behavior of generators. In that case, the code generator and run-time system were modified extensively. This involved weeks of tedious and difficult work that required expert knowledge of the internal structure of the Version 2 system. The variant translator for Version 8 was written in a few hours, and required only a knowledge of the format of variant macro specifications and the Icon source language itself. The monitoring of expression evaluation events in Version 8 probably would not have been undertaken if it had been necessary to modify the code generator and the run-time system.

The usefulness of the system described here depends heavily on its support software. The ability to specify macro definitions in a simple format, and particularly to be able to provide a single specification for the translation for all operators in a class, makes it easy to write many variant translators that otherwise would be impractically tedious.

Although the system described in this report is specifically tailored to Icon, the techniques have much broader

applicability. The automatic generation of such systems from grammatical specifications is an interesting project.

Acknowledgements

Tim Budd's Cg preprocessor was the inspiration for the Icon variant translator system described here. Bill Mitchell assisted in adapting the standard Icon translator to its use here. Ken Walker did most of the implementation for the current version, and Gregg Townsend adapted it to Version 8.10 of Icon.

Tim Budd, Dave Hanson, Bill Mitchell, Janalee O'Bagy, and Steve Wampler made a number of helpful suggestions on the variant translator system and the presentation of the material in this report.

References

1. B. W. Kernighan, "Ratfor — A Preprocessor for a Rational Fortran", *Software—Practice & Experience* 5(1975), 395-406.
2. T. A. Budd, "An Implementation of Generators in C", *J. Computer Lang.* 7(1982), 69-87.
3. R. E. Griswold and M. T. Griswold, *The Icon Programming Language*, Prentice-Hall, Inc., Englewood Cliffs, NJ, second edition, 1990.
4. R. E. Griswold, *Rebus — A SNOBOL4/Icon Hybrid*, The Univ. of Arizona Tech. Rep. 84-9, 1984.
5. J. L. Steffen, "Ctrace — A Portable Debugger for C Programs", *UNICOM Conference Proceedings*, Jan. 1983, 187-191. San Diego, California.
6. S. C. Kendall, "Bcc: Runtime Checking for C Programs", *USENIX Software Tools Summer 1983 Toronto Conference Proceedings*, 1983, 5-16.
7. M. E. Lesk and E. Schmidt, *Lex — A Lexical Analyzer Generator*, Bell Laboratories, Murray Hill, New Jersey, 1979.
8. S. C. Johnson, *Yacc: Yet Another Compiler-Compiler*, Bell Laboratories, Murray Hill, New Jersey, 1978.
9. R. E. Griswold and M. T. Griswold, *The Implementation of the Icon Programming Language*, Princeton University Press, 1986.
10. A. J. Anderson and R. E. Griswold, *Unifying List and String Processing in Icon*, The Univ. of Arizona Tech. Rep. 83-4, 1983.
11. K. Walker and R. E. Griswold, *A Pattern-Matching Laboratory; Part I — An Animated Display of String Pattern Matching*, The Univ. of Arizona Tech. Rep. 86-1, 1986.
12. R. E. Griswold and J. O'Bagy, *Seque: A Language for Programming with Streams*, The Univ. of Arizona Tech. Rep. 85-2, 1985.
13. R. E. Griswold and J. O'Bagy, *Reference Manual for the Seque Programming Language*, The Univ. of Arizona Tech. Rep. 85-4, 1985.
14. C. A. Coutant, R. E. Griswold and D. R. Hanson, "Measuring the Performance and Behavior of Icon Programs", *IEEE Trans. on Software Eng.* SE-9, 1 (Jan. 1983), 93-103.