Transporting Version 7 of Icon*

Ralph E. Griswold

TR 88-9a

January 29, 1988; revised March 5, 1988

Department of Computer Science

The University of Arizona

Tucson, Arizona 85721

*This work was supported by the National Science Foundation under Grant DCR-8502015.

Transporting Version 7 of Icon

1. Background

The implementation of the Icon programming language is large and complex [1]. It is, however, written almost entirely in C, and it is designed to be portable to a wide range of computers and operating systems.

The implementation was developed on a UNIX* system. It has been installed on a wide range of UNIX systems, from mainframes to personal computers. Putting Icon on a new UNIX system is more a matter of installation than porting [2]. There presently also are installations of Icon for VAX/VMS, the Amiga, the Atari ST, the Macintosh under MPW, and MS-DOS. This document addresses the problems and procedures for porting Icon to other operating systems and computers.

The current version of Icon is Version 7 [3]. All installations of Version 7 of Icon are obtained from common source code, using conditional compilation to select system-dependent code. Consequently, transporting Icon to a new system is largely a matter of selecting appropriate values for configuration parameters, deciding among alternative definitions, and possibly adding some code that is computer- or operating-system-dependent.

A small amount of assembly-language code is needed for a complete installation. See Section 7. This code is optional and only affects co-expressions and checking for arithmetic overflow. A running version of the language can be obtained by working only in C.

Transporting Icon to a new system is a fairly complex task, although there are many aids to simplify mechanical tasks. Read this report carefully before beginning a port. Understanding the Icon programming language is helpful during the debugging phase of a port. See [3-5]

2. Requirements

C Data Sizes

Icon places the following requirements on C data sizes:

- *char*s must be 8 bits.
- *int*s must be 16, 32, or 64 bits.
- *long*s and pointers must be 32 or 64 bits.
- All pointers must be the same length.
- *long*s and pointers must be the same length.

If your C data sizes do not meet these requirements, do not attempt to transport Icon. Call the Icon Project for advice. *Note:* Icon has not yet been ported to a computer with 64-bit C data sizes; such a port may encounter problems.

The C Compiler

The main requirement for implementing Icon is a production-quality C compiler that supports at least the *de facto* ‘‘K&R’’ standard [6]. The term ‘‘production quality’’ implies robustness, correctness, the ability to handle large files and complicated expressions, and a comprehensive run-time library.

C preprocessor should conform either to the ANSI C standard [7] or to the *de facto* standard for UNIX C preprocessors. In particular, Icon uses the C preprocessor to concatenate strings and substitute arguments within quotation marks. For the ANSI standard, the following definitions are used:

*UNIX is a trademark of AT&T Bell Laboratories.

```
#define Cat(x,y) x##y
#define Lit(x) #x
```

For the UNIX *de facto* standard, the following definitions are used:

```
#define Ident(x) x
#define Cat(x,y) Ident(x)y
#define Lit(x) "x"
```

The following program can be used to test these preprocessor facilities:

```
Cat(ma,in)()
{
   printf(Lit(Hello world\n));
}
```

If this program does not compile and print Hello world using one of the sets of definitions above, there is no point in proceeding. Contact the Icon Project as described in Section 8 for alternative approaches.

Memory

The Icon programming language requires a substantial amount of memory. The practical minimum is 512kb.

File Space

The source code for Icon is large — about 850kb. Compilation and testing require considerably more space. While the implementation can be divided into components that can be transported separately, this approach may be painful.

3. Organization of the Implementation

Icon was developed on a hierarchical file system. To facilitate file transfer between different operating systems and to simplify porting to systems that do not support file hierarchies, the source code for Icon is provided both in hierarchical form and in a ‘‘flat’’ form in which all files reside in the same area. This document applies to both the hierarchical and flat forms. Some of the supplementary documentation on Icon refers to file hierarchies. In interpreting this documentation on flat systems, simply ignore the directories in path specifications; the file names themselves are the same in the hierarchical and flat version.

3.1 Source Code

There are four components of Icon:

- itran — a translator that converts source-language programs to *ucode*, an assembly language for an abstract ‘‘Icon machine’’.
- ilink — a linker that combines one or more ucode files into a single binary *icode* file in executable format for the Icon machine.
- iconx — an executor for icode, including a run-time system that supports the operations of the Icon language.
- icont — a command processor that provides a user interface for running itran, ilink, and iconx.

itran

The translator is relatively straightforward. It contains a lexical analyzer, a parser, a code generator, and support routines. The translator is independent of the architecture of the computer on which it runs and its ucode output consists of printable text.

ilink

The linker is somewhat more complex than the translator. It reads ucode files and outputs binary code and data structures that are needed during execution. Because these data structures depend to some extent on the architecture of the computer on which Icon runs, the linker contains some machine-dependent parameters.

iconx

The executor is the largest and most complex of the components of Icon. Its data structures depend on the architecture of the computer on which Icon runs, and it includes code for all the operations in the Icon language. In addition, it manages storage dynamically.

icont

The command processor is a small program that executes the other components of Icon. Its main function is to analyze the user's command line and take appropriate actions. The command processor is a convenience, not a necessity. A command-language script can be used in its place.

The files related to the source are packaged in five groups:

itran	files for itran
ilink	files for ilink
iconx	files for iconx
icont	files for icont
h	header files

The header files are in a separate package, since some are used in several components of Icon. In some forms of the distribution, iconx comes in two parts, since it is is too large to fit on a single diskette.

Appendix A lists the files of each component of Icon. Some header files are used in several components; these are identified in the appendix. The files itran.bat, ilink.bat, iconx.bat, and icont.bat are scripts that indicate what files are to be compiled and loaded to produce the respective components. These scripts were derived from a UNIX implementation, but they can be adapted easily to other systems.

3.2 Tests

Test programs are divided into two parts. The first part, referred to as suite1, contains test programs and the expected output for itran and ilink. The second part, referred to as suite2, contains test programs and expected output for iconx.

See Section 6 for more information about the test programs.

4. An Overview of the Porting Process

The first step in the porting process is to configure the source code for the new system. This process is described in Section 5.1. After this is done, the components need to be ported, one by one.

The porting process for each component of Icon is essentially the same:

- provide code and definitions that are system-dependent
- compile the source files and link them to produce executable binary files
- test the result
- debug, iterating over the previous steps as necessary

The components should be ported in the order given above: itran, ilink, iconx, and icont. Of course, bugs in previously ported components may not show up until subsequent components are tested.

In addition to this obvious sequence of steps, some aspects of the implementation may be deferred until the entire system is running, or they may be implemented in a preliminary manner and subsequently refined. For example, the assembly-language portions of Icon are best left unimplemented until the rest of the system is running. Considerable frustration can be avoided if problems that come up can be circumvented with temporary expedients until the majority of the implementation is working properly. Similarly, conservative choices should be made during the initial phases of the implementation.

5. Conditional Compilation

Conditional compilation is used extensively in Icon to select code that is appropriate to a particular installation. Conceptually, conditional compilation can be divided into two categories:

(1) Matters related to the details of computer architecture, run-time system idiosyncrasies, specific C compilers, and operating-system variants.

(2) Matters that are specific to operating systems that are distinctly different, such as MS-DOS, UNIX, and VMS.

5.1 Parameters and Definitions

There are many constants in the source code for Icon that vary from system to system. Default values are provided so that the usual cases are handled automatically. The file define.h contains C preprocessor definitions for parameters that differ from the defaults or that must be provided on an individual basis. This file initially contains values for a ''vanilla'' 32-bit system. If your system closely approximates a ''vanilla'' system, you will have few changes to make to define.h. Over the range of possible systems, there are many possibilities as described below.

The definitions are grouped into categories so that any necessary changes to define.h can be approached in a logical way.

C preprocessor: Icon preprocessor directives use string concatenation and substitution of arguments within quotation marks as described in Section 2. By default, the original UNIX method of accomplishing this is used. A different method is specified in the C draft standard [6]. If your C compiler supports the draft standard, add

```
#define Standard
```

to define.h.

C sizing and alignment: There are four constants that relate to the size of C data and alignment:

IntBits	(default: 32)
WordBits	(default: 32)
Double	(default: undefined)

IntBits is the number of bits in a C *int*. It may be 16, 32, or 64. WordBits is the number of bits in a C *long* and pointer (Icon's ''word''). It may be 32 or 64. If your C library expects *double*s to be aligned at double-word boundaries, add

```
#define Double
```

to define.h.

Floating-point arithmetic: There are four optional definitions related to floating-point arithmetic:

Big	(default: 9007199254740092.)
LogHuge	(default: 309)
Precision	(default: 10)
ZeroDivide	(default: undefined)

The values of Big, LogHuge, and Precision give, respectively, the largest floating-point number that does not loose precision, the maximum base-10 exponent + 1 of a floating-point number, and the number of digits provided in the string representation of a floating-point number. If the default values given above do not suit the floating-point arithmetic on your system, add appropriate definitions to define.h. If your system needs a software check for division by floating-point zero, add

```
#define ZeroDivide
```

to define.h.

Include file location: The location of the include file time.h varies from system to system. Its default location is <time.h>. If it resides at a different location on your system (usually <sys/time.h>), add an appropriate definition of SysTime to define.h, as in

```
#define SysTime <sys/time.h>
```

If the location is incorrect, a fatal error will occur during the compilation of src/iconx/lmisc.c. The use of this definition also depends on your C preprocessor making macro substitutions in #include directives. Most preprocessors do, but if yours does not, edit src/iconx/lmisc.c and replace SysTime there by the appropriate value. If you have to do this, make a note to come back later and place the definition under the control of conditional compilation.

Run-time routines: The support for some run-time routines varies from system to system. The related constants are:

Fork	(default: fork)
IconGcvt	(default: undefined)
IconQsort	(default: undefined)
NoAtof	(default: undefined)
SysMem	(default: undefined)
index	(default: undefined)
rindex	(default: undefined)

If your system supports *vfork(2)*, add

```
#define Fork vfork
```

to define.h. This definition is used only in icont.c.

If IconGcvt and IconQsort are defined, versions of *gcvt(3)* and *qsort(3)* in the Icon system are used in place of the routines normally provided in the C run-time system. These constants only need to be defined if the versions of these routines in your run-time system are defective or missing.

The C run-time routine *atof(3)* normally is used in the Icon linker to convert strings for real literals to corresponding floating-point numbers. If the version of *atof* on your system does not work properly, add

```
#define NoAtof
```

to define.h. This replaces the use of *atof* by in-line conversion code.

If your run-time system includes *memcpy(3)* and *memset(3)*, add

```
#define SysMem
```

to define.h. Otherwise, versions of these routines in the Icon system are used.

Different run-time systems use different names for the routines for locating substrings within strings. The source code for Icon uses index and rindex. The other possibilities are strchr and strrchr. If your system uses the latter names, add

```
#define index strchr
#define rindex strrchr
```

to define.h.

Host identification: The identification of the host computer as given by the Icon keyword &host needs to be specified in define.h. The definition

```
#define HostStr "unknown host"
```

is provided in define.h initially. This definition should be changed to an appropriate value for your system.

Storage management: Storage management is one of the most complex components of Icon and the one that is most likely to cause problems in porting. There are two alternative forms of memory management: fixed-regions and expandable-regions. In fixed-regions memory management, Icon's storage regions are allocated at the beginning of execution and do not increase in size. If they are not large enough, program execution terminates with an error if more space is needed. In expandable-regions memory management, Icon's storage regions also are allocation at the beginning of execution, but it they are not big enough, they are increased using *sbrk*.

Many systems do not support *sbrk*, so the fixed-regions form of memory management is necessary, although it is more limiting. This form of memory management is recommended for an initial port, even if *sbrk* is supported. If the port is otherwise successful, expandable-regions management can be attempted later.

Initially, define.h contains

```
#define FixedRegions
```

which enables fixed-regions management. Remove this definition if you want expandable-regions management.

Icon includes its own versions of *malloc, calloc*, *realloc*, and *free* so that it can manage its storage regions without interference from allocation by the operating system. Normally, Icon's versions of these routines are loaded instead of the system library routines. Leave things are they are in the initial configuration, but if your system insists on loading its own library routines, multiple definitions will occur as a result of the *ld* in src/iconx. If multiple definitions occur, go back and add

```
#define IconAlloc
```

to define.h. This definition causes Icon's routines to be named differently to avoid collision with the system routine names.

One possible effect of this definition is to interfere with Icon's expansion of its memory region in expandable-regions memory management.

Environment variables: If your system does not support environment variables (via the run-time library routine getenv), add the following line to define.h:

```
#define NoEnvVars
```

This disables Icon's ability to change internal parameters to accommodate special user needs (such as using memory region sizes different from the defaults), but does not otherwise interfere with the use of Icon.

Character set:

The implementation of Icon was designed on the assumption it would run on computers using the ASCII character set. Some work has been done toward supporting the EBCDIC character set as an alternative. If you are porting Icon to a computer that uses the EBCDIC character set, add

```
#define EBCDIC
```

to define.h.

Support for the EBCDIC character set is not yet complete and some places in the source code need modification for it. Examine code under #ifdef EBCDIC in fstr.c, fstranl.c, imain.c, lex.c, rconv.c, and toktab.c.

There also is an option that allows alternative representations of some characters commonly used in Icon programs that are not supported by many EBCDIC terminals and printers. The standard characters and their alternative forms are:

standard	*alternative*
`{`	`$(`
`}`	`$)`
`[`	`$<`
`]`	`$>`

To enable this option, add

```
#define ExtChars
```

to define.h.

Miscellaneous: There are two other definitions that may be needed in some cases:

Hz	(default: 60)
UpStack	(default: undefined)

If you are running in a 50-hz environment, add

```
#define Hz 50
```

to define.h.

Most computers have downward-growing C stacks, for which stack addresses decrease as values are pushed. If you have an upward-growing stack, for which stack addresses increase as values are pushed, add

```
#define UpStack
```

to define.h.

Optional features: The implementation of co-expressions and arithmetic overflow checking require assembly language routines. Initially, define.h contains

```
#define NoCoexpr
#define NoOver
```

These definitions disable co-expressions and arithmetic overflow checks. Leave these definitions in for the first round, although you may want to remove them later and implement these features (see Section 7).

5.2 Operating System Differences

Conditional compilation for operating systems usually is due to differences in run-time library routines, differences in file naming, the handling of input and output, and environmental factors.

The presently supported operating system are AmigaDos, Atari ST TOS, the Macintosh under MPW, MS-DOS, UNIX, and VAX/VMS. There hooks for transporting to MVS and VM and for an unspecified system (a new port). The associated defined symbols are

AMIGA	AmigaDos
ATARI_ST	Atari ST TOS
MACINTOSH	Macintosh
MSDOS	MS-DOS
MVS	MVS
PORT	new port
UNIX	UNIX
VM	VM
VMS	VMS

Conditional compilation uses logical expressions composed from these symbols. An example is:

```
                .
                .
                .
#if MSDOS
                .
                .          /* code for MS–DOS */
                .
#endif

#if UNIX || VMS
                .
                .          /* code for UNIX and VMS */
                .
#endif
                .
                .
                .
```

Each symbol must be defined to be either 1 (for the target operating system) or 0 (for all other operating systems). This is accomplished by defining the symbol for the target operating system to be 1 in define.h. In config.h, which includes define.h, all other operating-system symbols are automatically defined to be 0.

Logical conditionals with #if are used instead of defined or undefined names with #ifdef to avoid nested conditionals, which become very complicated and difficult to understand when there are several alternative operating systems. Note that it is important not to use #ifdef accidentally in place of #if, since all the names are defined.

The file define.h initially contains

```
#define PORT
```

It you are porting to MVS or VM, change the symbol accordingly. Otherwise, leave it as is; later you should come back and change PORT to some more appropriate name.

To make it easy to locate all the places where there is code that may be dependent on the operating system, such code is bracketed by unique comments of the following form:

```
/*
 * The following code is operating–system dependent.
 */
                        .
                        .
                        .
/*
 * End of operating–system specific code.
 */
```

Between these beginning and ending comments, the code for different operating systems is provided using conditional expressions such as those indicated above.

There presently are a total of 76 segments that contain such code. The files that contain operating-system-dependent code are listed in Appendix B. Look through some of the files that contain such segments to get an idea of what is involved. Each segment contains comments that describe the purpose of the code. In most cases, the most likely code or a suggestion is given in the conditional code under PORT. In some cases, no code will be needed. In others, code for an existing system may suffice for the new system.

In any event, code for the new operating system name must be added to each such segment, either by adding it to a logical disjunction, as in

```
#if MSDOS || UNIX || PORT
        .
        .
        .
#endif

#if VMS
        .
        .
        .
#endif
```

or by filling in the segment with the appropriate code, as in

```
#if PORT
            .
            .           /* code for the the port */
            .
#endif
```

If no code for the target operating system, a comment should be added so that it is clear that the situation has been considered.

You may find need for code that is operating-system dependent at a place where no such dependency presently exists. If this happens, add a new segment similar to the other ones, being sure to provide something appropriate for all operating systems. Do not simply add code like

```
#if PORT
        :
        :
#endif
```

without empty code for the other systems, since this will interfere with transportation to other systems in the future.

Do not use #else constructions; this encourages errors and obscures the mutually exclusive nature of operating system differences.

6. Building and Testing

6.1 The Translator

Start by compiling all the C programs listed in tran.bat. Load the resulting object files to produce itran. With any luck, this will go without problems, since itran is largely machine-independent. If you encounter problems, check the portions of code containing operating system dependencies.

Once you have a version of itran, try it on the Icon programs in suite1. For example, to translate bitops.icn in suite1, do

```
itran bitops.icn
```

This should yield two ucode files, bitops.u1 and bitops.u2. The .u1 file contains procedure declarations and code for the Icon machine; the .u2 file contains global declaration information.

These files both consist of printable text. They should be identical to the corresponding files in suite1. Be careful to run itran in a way that does not overwrite the ucode files in suite1.

More than likely, if you get any ucode files at all, they will be correct, since the translator is machine-independent and portable and no significant problems have been encountered with it in other ports.

6.2 The Linker

Compile all the C files listed in ilink.bat and load them to get ilink. You may encounter more problems here, since the linker is somewhat dependent on the sizes of C data objects and does more sophisticated input and output than the translator.

Once you have ilink, test it on ucode files as follows:

```
ilink –L bitops.u1
```

Only the .u1 file is named for the linker, but both the .u1 and .u2 files must be present for ilink. The –L option causes ilink to produce a debugging output file with the suffix .ux — in this case, bitops.ux. Compare this file to the corresponding one in suite1. The two files may not be exactly the same, since the debugging output depends on C sizes (the distributed files are for 32-bit ints) and idiosyncrasies of fprintf. However, the files should be similar. Do not worry about small differences.

6.3 The Run-Time System

If you get this far without apparent problems, you are ready for the next — and most difficult — part of the transporting process: iconx. Compile all the C programs listed in iconx.bat and load them to form iconx.

As a first test, try iconx on hello.icn in suite1 as follows:

```
itran hello.icn
ilink hello.u1
iconx hello
```

If all is well, the last step should print out "hello world" and some identifying information.

Once this test has been passed, more rigorous testing should follow. At this point, you probably will want to devise a way of testing programs, since there are a large number of tests. This is done for the MS-DOS implementation using the following .bat file:

```
echo off
itran –s %1.icn
ilink %1.u1
echo Executing
if not exist %1.dat goto skip1
iconx  %1 <%1.dat >%1.out
goto comp
:skip1
iconx –e – %1 >%1.out
:comp
compare %1.out stand\%1.out –t –w –l %1.dif
type %1.dif
del %1
del %1.u1
del %1.u2
erase local\%1.out
erase local\%1.dif
movefile %1.out local
movefile %1.dif local
echo on
```

On UNIX systems, the following script is used to test all the files in a list:

```
for i in ‘cat $1.lst‘
do
    rm –f local/$i.out
    echo Running $i
    itran $i.icn
    ilink $i.u1
    if test –r $i.dat
    then
        iconx $i <$i.dat >local/$i.out 2>&1
    else
        iconx $i >local/$i.out 2>&1
    fi
    echo Checking $i
    diff local/$i.out stand/$i.out
    rm –f $i
done
```

While these methods are adapted to specific operating-system capabilities, something similar can be concocted for other systems. Making such a facility as easy to use as possible is worth the effort.

In suite2 there many Icon programs for testing different aspects of iconx. These range from simple tests to ‘‘grinders’’. The names of the test programs are listed in the following files:

icon.lst	short programs∗
expr.lst	programs that contain a wide variety of expressions
check.lst	programs that may produce different results on different systems
extra.lst	programs that test additional features
work.lst	long-running programs∗
version7.lst	programs that test new features in Version 7∗

The lists flagged with a ∗ contains tests that require data files that are included in suite2 with the ending .dat. For example, the Icon program meander.icn, listed in icon.lst, takes data from meander.dat. suite2 also contains files whose names end in .out that contain the expected output of each test program. For example, the expected output of meander.icn is contained in meander.out.

The programs listed in icon.lst should produce output identical to that in their corresponding .out files. Any discrepancies should be checked carefully and corrections made before continuing.

The programs listed in expr.lst execute a wide variety of individual expressions. Ideally, there should be no discrepancies between their output and the expected output. If there are many discrepancies, something serious probably is wrong. If there are only a few discrepancies, they may be noted while other testing is conducted.

The programs listed in check.lst certainly will show some differences, since they test features whose results are time- and environment-dependent. Other differences may show up also. These do not necessarily indicate problems. For examples, minor differences in the results of floating-point arithmetic are common in these tests.

The programs listed in extra.lst test some features that are not tested elsewhere. They should be treated like the programs listed in icon.lst.

In work.lst are listed several programs that run for a very long time and exercise some parts of the implementation extensively.

Since storage management is one of the parts of Icon that is likely to give trouble, there are special storage-management tests in gc.lst. These programs run for a long period of time. One program may show a difference in output if the fixed-regions version of memory management is used, since it may run out of space.

Not much general advice can be given about locating and correcting problems that may show up in testing iconx. It has to be done the hard way and may involve learning more about the Icon language [4] and how it is implemented [1]. A good debugger can be very helpful.

6.4 The Command Processor

See the remarks about icont in Section 3.1. If icont is to be implemented, compile icont.c and load it to form icont.

Test icont on hello.icn in suite1 as follows:

```
icont hello.icn –x
```

If all is well, hello.icn should be translated, linked, and executed to produce the output given in the preceding section.

icont supports several options. See [8] for a complete list. Various combinations of options should be tested to confirm that icont is working properly.

7. Assembly-Language Code

Once Icon is running satisfactorily, you may wish to implement the features that require assembly language: co-expressions and arithmetic overflow checking.

7.1 Co-Expressions

Note: If your system does not allow the C stack to be at an arbitrary place in memory, there is probably little hope of implementing co-expressions. If you do not implement co-expressions, the only effect will be that Icon programs that attempt to use a co-expression will terminate with an error message.

All aspects of co-expression creation and activation are written in C in Version 7 except for a routine, coswitch, that is needed for context switching. This routine requires assembly language, since it must manipulate hardware registers. It either can be written as a C routine with asm directives or as an assembly language routine.

Calls to the context switch have the form coswitch(old_cs,new_cs,first), where old_cs is a pointer to an array of words that contain C state information for the current co-expression, new_cs is a pointer to an array of words that hold C state information for a co-expression to be activated, and first is 1 or 0, depending on whether or not the new co-expression has or has not been activated before. The zeroth element of a C state array always contains the hardware stack pointer (*sp*) for that co-expression. The other elements can be used to save any C frame pointers and any other registers your C compiler expects to be preserved across calls.

The default number of elements for saving the C state is 15. This number may be changed by adding

```
#define CStateSize n
```

to define.h, where *n* is the number of elements needed.

The first thing coswitch does is to save the current pointers and registers in the old_cs array. Then it tests first. If first is zero, coswitch sets *sp* from new_cs[0], clears the C frame pointers, and *calls* interp. If first is not zero, it loads the (previously saved) *sp*, C frame pointers, and registers from new_cs and returns.

Written in C, coswitch has the form:

```
/*
 * coswitch
 */
coswitch(old_cs, new_cs, first)
int *old_cs, *new_cs;
int first;
{
                              .
                              .
                              .
          /* save sp, frame pointers, and other registers in old_cs */
                              .
                              .
                              .
   if (first == 0) {/* this is first activation */
                              .
                              .
                              .
          /* load sp from new_cs[0] and clear frame pointers */
                              .
                              .
                              .
      interp(0, 0);
      syserr("interp() returned in coswitch");
      }
   else {
                              .
                              .
                              .
          /* load sp, frame pointers, and other registers from new_cs */
                              .
                              .
                              .
      }
}
```

After you implement coswitch, remove the #define NoCoexpr from define.h.

To test your context switch, run the programs in coexpr.lst. Ideally, there should be no differences in the comparison of outputs.

If you have trouble with your context switch, the first thing to do is double-check the registers that your C compiler expects to be preserved across calls — different C compilers on the same computer may have different requirements.

Another possible source of problems is built-in stack checking. Co-expressions rely on being able to specify an arbitrary region of memory for the C stack. If your C compiler generates code for stack probes that expects the C stack to be at a specific location, you may need to disable this code or replace it with something more appropriate.

7.2 Overflow Checking

C does not provide overflow checking for integer addition, subtraction, or multiplication. Icon, on the other hand, is supposed to check for overflow. This usually requires assembly-language code.

Initially, define.h contains the definition

```
#define NoOver
```

which causes overflow checking to be bypassed.

If you do not want to implement overflow checking, you need do nothing. The only effect will be that overflow will not be detected.

If you want to implement overflow checking, remove the definition of NoOver from your define.h and write routines ckadd, cksub, and ckmul that call fatalerr(–203,0) in the case of overflow.

Note: It often is harder to test for overflow for multiplication than for addition and subtraction. A dummy routine that simply returns can be provided for multiplication if this is the case on your system.

To test overflow checking, run the programs in over.lst. There should be no differences in the comparison of outputs if overflow checking is working properly.

You should also rerun previous tests at this point to make sure that arithmetic still works properly.

8. Trouble Reports and Feedback

If you run into problems, contact us at the Icon Project:

Icon Project
Department of Computer Science
Gould-Simpson Building
The University of Arizona
Tucson, AZ 85721
U.S.A.

(602) 621-2018

icon-project@arizona.edu (Internet)
… {uunet, allegra, cmcl2, noao}!arizona!icon-project (uucp)

Please also let us know of any suggestions for improvements to the porting process.

Once you have completed your port, please send us copies of any files that you modified so that we can make corresponding changes in the central version of the source code. Once this is done, you can get a new copy of the source code whenever changes or extensions are made to the implementation.

Also be sure to include documentation on any features that are not implemented in your port or any changes that would affect users.

Acknowledgements

Many persons have been involved in the implementation of Icon. Contributions to its portability have been made by Bill Mitchell, Kelvin Nilsen, Gregg Townsend, and Cheyenne Wills.

References

1. R. E. Griswold and M. T. Griswold, *The Implementation of the Icon Programming Language*, Princeton University Press, 1986.
2. R. E. Griswold, *Installation Guide for Version 7 of Icon on UNIX Systems*, The Univ. of Arizona Tech. Rep. 88-6, 1988.
3. R. E. Griswold, G. M. Townsend and K. Walker, *Version 7 of Icon*, The Univ. of Arizona Tech. Rep. 88-5, 1988.
4. R. E. Griswold and M. T. Griswold, *The Icon Programming Language*, Prentice-Hall, Inc., Englewood Cliffs, NJ, 1983.
5. R. E. Griswold, *An Overview of the Icon Programming Language*, The Univ. of Arizona Tech. Rep. 83-3, 1983.
6. B. W. Kernighan and D. M. Ritchie, *The C Programming Language*, Prentice-Hall, Inc., Englewood Cliffs, NJ, 1978.

7. Technical Committee X3J11, *Draft Proposed American National Standard for Information Systems — Programming Language C*, Oct. 1986.

8. R. E. Griswold, *Version 7.5 of Icon for MS-DOS*, The Univ. of Arizona Icon Project Document IPD47c, 1988.

Appendix A — Files Used by Components of Icon

Files marked by ∗ are used in more than one component.

Files Used by itran

config.h∗	general configuration information
define.h∗	system-dependent definitions
keyword.h∗	keyword definitions
itran.h	general heading information of itran
lex.h	header for lexical analysis
lfile.h	header for link declarations
sym.h	header for symbol tables
token.h	token definitions
tree.h	code tree information
version.h∗	version information
code.c	code generator
err.c	error messages
itran.c	main program
keyword.c	keyword structure
lex.c	lexical analyzer
lnklist.c	file linking
mem.c	memory management
optab.c	state tables for operator recognition
parse.c	parser
sym.c	symbol table management
toktab.c	token table
tree.c	code tree constructor

Files Used by ilink

config.h∗	general configuration information
cpuconf.h∗	computer architecture information
define.h∗	system-dependent definitions
fdefs.h∗	function definitions
header.h∗	icode header structure
ilink.h	general heading information for ilink
keyword.h∗	keyword definitions
opcode.h	opcode structure
opdefs.h∗	icode instruction definitions
paths.h∗	file paths
rt.h∗	header for run-time system
version.h∗	version information
glob.c	processor for global information
ilink.c	main program
lcode.c	code generator
llex.c	lexical analyzer
lmem.c	memory management
lsym.c	symbol table management
opcode.c	opcode table

Files Used by iconx

config.h*	general configuration information
cpuconf.h*	computer configuration information
define.h*	system-dependent definitions
fdefs.h*	function definitions
gc.h	garbage collection definitions
header.h*	icode header
keyword.h*	keyword definitions
memsize.h*	memory sizing
opdefs.h*	icode definitions
rt.h*	run-time definitions
version.h*	version information
fconv.c	conversion functions
fmisc.c	miscellaneous functions
fscan.c	scanning functions
fstr.c	string construction functions
fstranl.c	string analysis functions
fstruct.c	data structure functions
fsys.c	system functions
fxtra.c	extra functions
idata.c	data
imain.c	main program
interp.c	icode interpreter
invoke.c	function and procedure invocation
lmisc.c	miscellaneous library routines
lrec.c	library routines for record
lscan.c	scanning routines
oarith.c	arithmetic operations
oasgn.c	assignment operations
ocat.c	concatenation operations
ocomp.c	comparison operations
omisc.c	miscellaneous operations
oref.c	referencing operations
oset.c	set operations
ovalue.c	value operations
rcomp.c	comparison routines
rconv.c	conversion routines
rdefault.c	default value routines
rdoasgn.c	assignment routines
rlocal.c	locally needed routines
rmemexp.c	memory management routines for expandable regions
rmemfix.c	memory management routines for fixed regions
rmemmgt.c	general memory management routines
rmisc.c	miscellaneous routines
rstruct.c	structure routines
rsys.c	system routines

Files Used by icont

config.h*	general configuration information
define.h*	system-dependent definitions
paths.h*	path definitions
icont.c	main program

Appendix B — System-Dependent Code

The following source files contain code that is operating-system dependent. The number of places where such code occurs in each file is given in parentheses.

Translator:

itran.c (5)

Linker

ilink.c (7)
lcode.c (2)
llex.c (1)
lmem.c (6)

Executor:

fmisc.c (1)
fsys.c (14)
imain.c (10)
interp.c (6)
lmisc.c (6)
rconv.c (2)
rmemexp.c (2)
rmemmgt.c (1)
rmisc.c (1)
rsys.c (2)

Command Processor

icont.c (10)