

QA76
U59
fol.
c.2
no.133

SPECIFICATION OF
CONCURRENT EUCLID

(Version 1)

James R. Cordy
Richard C. Holt

Technical Report CSRG-133
August 1981

QA 76
U59
fol.
c.2
no.133

**COMPUTER SYSTEMS RESEARCH GROUP
UNIVERSITY OF TORONTO**



SPECIFICATION OF
CONCURRENT EUCLID

(Version 1)

James R. Cordy
Richard C. Holt

Technical Report CSRG-133
August 1981

Computer Systems Research Group
University of Toronto
Toronto, Canada
M5S 1A1

The Computer Systems Research Group (CSRG) is an interdisciplinary group formed to conduct research and development relevant to computer systems and their application. It is jointly administered by the Department of Electrical Engineering and the Department of Computer Science of the University of Toronto, and is supported in part by the Natural Sciences and Engineering Research Council of Canada.

ABSTRACT

Concurrent Euclid (CE) is a programming language designed for implementing software that is efficient, reliable and portable. It is particularly suited for implementing operating systems, compilers and specialized microprocessor applications. It can serve as the basis for producing verifiable system software.

CE has been designed to allow its compiler to be small, fast and portable. Such a compiler exists, with replaceable high-quality code generators for various target machine architectures including the PDP-11, MC68000 and MC6809.

CONTENTS

INTRODUCTION	1
I. THE SE LANGUAGE	2
Identifiers and Literals	2
Source Program Format	3
Syntactic Notation	3
Programs	3
Modules	4
Declarations	5
Constant Declarations	5
Variable Declarations	6
Types and Type Declarations	6
Type Equivalence and Assignability	9
Variable Bindings	9
Collections	10
Procedures and Functions	11
Type Converters	13
Statements	14
Variables and Constants	16
Expressions	17
Built-in Functions	18
Standard Components	19
Manifest Expressions	19
Precision of Arithmetic	19
Source Inclusion Facility	20
II. CONCURRENCY FEATURES	21
Processes	21
Monitors	22
Conditions	23
The Busy Statement	25
III. SEPARATE COMPILATION	26
External Declarations	26
Compilations	27
Linking of Compilations	28
APPENDIX 1. COLLECTED SYNTAX OF CONCURRENT EUCLID	29
APPENDIX 2. KEYWORDS AND PREDEFINED IDENTIFIERS	41

APPENDIX 3. INPUT/OUTPUT IN CONCURRENT EUCLID	42
IO/1: Terminal Formatted Text I/O	42
IO/2: Sequential Argument File I/O	43
IO/3: Temporary and Non-Argument Files	45
IO/4: Structure Input/Output and Random Access Files	46
Interfacing to Unix	46
 APPENDIX 4. PDP-11 IMPLEMENTATION NOTES	 47
Data Representation	47
Register Usage	47
Calling Conventions	48
External Names	49
Parameter Passing	49
Run-time Checking	50
 REFERENCES	 51
 INDEX	 52

INTRODUCTION

This report defines the programming language Concurrent Euclid, or CE. CE is designed for implementing software, and is particularly suited to implementing operating systems, compilers and specialized microprocessor applications. Because it is based on Euclid [1], it can also serve as the basis for implementing software which is to be formally verified.

CE consists of a subset of the Euclid programming language called Sequential Euclid or SE and a set of concurrency extensions to Euclid based on monitors [2]. The first section of this document defines the SE language independently of Euclid. The second section describes the concurrency features added to form CE. The last section describes CE features that support separate compilation of procedures, functions, modules and monitors. A thorough understanding of the basic concepts of the Pascal family of programming languages is assumed throughout.

I. THE SE LANGUAGE

This section describes the SE subset of Euclid. SE is defined independently of Euclid and no previous knowledge of the Euclid programming language is required. An understanding of the basic concepts of the Pascal family of programming languages is assumed.

IDENTIFIERS AND LITERALS

An identifier consists of any string of at most 50 letters, digits and underscores (_) beginning with a letter. Upper and lower case letters are considered identical in identifiers and keywords, hence aa, aA, Aa and AA all represent the same identifier. Keywords and predefined identifiers of Euclid, SE and CE must not be redeclared. A list of these is given in Appendix 2.

A string literal is any sequence of one or more characters not including a quote (') surrounded by quotes. Within strings, the characters quote, dollar sign, new line and end of file are represented as \$', \$\$, \$N and \$E respectively. As well, \$T, \$S and \$F may be used for tab, space, and form feed respectively.

A character literal is a dollar sign (\$) followed by any single character. The character literals corresponding to quote, dollar sign, space, tab, form feed, new line and end of file are \$\$', \$\$\$, \$\$\$, \$\$T, \$\$F, \$\$N and \$\$E respectively.

In every implementation, the character set for string and character literals will contain at least the upper and lower case letters A-Z and a-z, the digits 0-9 and the special characters ".,;?!?()[]{|}+~*/<=>'\$#^|&*", space, tab, form feed, new line and end of file. Character values are ordered such that A<B<C<...<Z, a<b<c<...<z and 0<1<2<...<9. Ordering of character values is implementation dependent otherwise.

An integer literal is a decimal number, an octal number or a hexadecimal number. A decimal number is any sequence of decimal digits. An octal number is any sequence of octal digits followed by #8. A hexadecimal number is any sequence of hexadecimal digits (represented as the decimal digits plus the capital letters A through F) beginning with a decimal digit and followed by #16. Negative values are obtained using the unary - operator; see "Expressions".

In every implementation, the range of integer literals will include at least 0 through 65535.

SOURCE PROGRAM FORMAT

A comment is any sequence of characters not including comment brackets surrounded by the comment brackets { and }. Comments may cross line boundaries.

A separator is a comment, blank, tab, form feed or source line boundary. Programs are free-format; that is, the identifiers, keywords, literals, operators and special characters which make up a program may have any number of separators between them. Separators cannot be embedded in identifiers, keywords, literals or operators, except that blanks may appear as part of the value of a string literal. Identifiers, keywords and literals must not cross line boundaries. At least one separator must appear between adjacent identifiers, keywords and literals.

SYNTACTIC NOTATION

The following sections define the syntax of SE.

The following notation is used:

{item} means zero or more of the item
[item] means the item is optional

Keywords are given in lower case. Special symbols are enclosed in double quotes (").

The following abbreviations are used:

id for identifier
expn for expression
typeDefn for typeDefinition

Semicolons are not required, but they may optionally appear following statements, declarations and import, export and checked clauses.

PROGRAMS

A main program consists of a module declaration.

A program is:

moduleDeclaration

Execution of a program consists of initializing the main module, see "Modules".

Modules, procedures and functions can be compiled separately; see "Separate Compilation".

MODULES

A moduleDeclaration is:

```
var id ":"
  module
    [imports "(" [var] id {" ," [var] id} ")"]
    [exports "(" id {" ," id} ")"]
    [[not] checked]
    {declarationInModule}
    [initially
      procedureBody]
  end module
```

Execution of a module declaration consists of executing the declarations in the module and then the "initially" procedure of the module. Execution of a program consists of executing the main module declaration in this way.

Module declarations may be nested inside other modules but must not be nested inside procedures and functions.

A module defines a package of variables, constants, types, procedures and functions. The interface of the module to the rest of the program is defined by its imports and exports clauses.

The imports clause lists the global identifiers which are to be visible inside the module. Variable, collection and module identifiers may be imported "var" (or not). Imported variables can be assigned to or passed as var parameters within the module only if they are imported "var". Elements of an imported collection can be allocated, freed, assigned to or passed as var parameters only if the collection is imported "var". Procedures of an imported module can be called only if the module is imported "var". Imported identifiers must not be redeclared inside the module.

The exports clause lists those identifiers defined inside the module which may be accessed outside the module using the "." operator. Exported variables cannot be assigned to or passed as var parameters outside the module. Elements of exported collections cannot be allocated, freed, assigned to or passed as var parameters outside the module. Unexported identifiers cannot be referenced outside the module.

Named types declared inside a module are opaque outside the module, that is, they are not considered equivalent to any other type. Variables and constants whose type is opaque cannot be subscripted, field selected or compared.

Modules may be "checked"; this causes all assert statements, subscripts and case statements in the module to be checked for validity at run-time. In addition, a particular implementation may check other conditions such as ranges in assignments and overflow in expressions. Modules not already nested inside an

unchecked module are checked by default and must be explicitly declared "not checked" to turn off run-time checking.

Even though declared like variables, modules are not variables and cannot be assigned, compared, passed as parameters or exported.

Modules can be separately compiled if desired; see "Separate Compilation".

DECLARATIONS

A declarationInModule is one of the following:

- a. constantDeclaration
- b. variableDeclaration
- c. typeDeclaration
- d. variableBinding
- e. moduleDeclaration
- f. collectionDeclaration
- g. procedureDeclaration
- h. functionDeclaration
- i. converterDeclaration
- j. assert ["(" expn ")"]

Forms (a) through (i) are declarations for new identifiers as explained in the following sections. Form (j) is an assert statement; see "Statements". An identifier must be declared textually preceding any references to it.

CONSTANT DECLARATIONS

A constantDeclaration is one of:

- a. [pervasive] const id ":=" manifestExpn
- b. [pervasive] const id ":" typeDefn ":=" expn
- c. [pervasive] const id ":" typeDefn ":="
"(" manifestExpn {"," manifestExpn} ")"
- d. [pervasive] const id ":=" stringLiteral

A constantDeclaration gives a name to a value which is constant throughout the scope of the declaration. The value of a scalar constant can be manifest or nonmanifest. A manifest expression is one whose value is known at compile-time (see "Manifest Expressions"). A nonmanifest expression must be evaluated at run-time. Non-scalar values are always considered nonmanifest.

Form (a) defines a manifest named constant. The type of the constant is the type of the value expression, which must be manifest. Manifest named constants are not represented at run time since their values are always known at compile time.

Form (b) declares a nonmanifest named constant of the

specified type. The value of the expression may be manifest or nonmanifest, and must be assignable to the constant's type. References to nonmanifest named constants are always considered nonmanifest even if their value is manifest.

Form (c) declares an array constant. The typeDefn must be an array type or named array type whose component type is scalar. The list of expressions gives the values of the elements of the array constant. The element values must be manifest expressions assignable to the element type of the array. The number of element values specified must be exactly the number of elements in the array.

Form (d) allows declaration of an array constant using a string literal value. The type of the constant is "packed array 1..n of Char" where n is the length of the string literal.

Constants declared using "pervasive" are automatically imported into all subsopes of the scope in which they are declared. Such constants need not be explicitly imported.

VARIABLE DECLARATIONS

A variableDeclaration is:

```
[register] var id ["(" at manifestExpn ")"] ":" typeDefn
                [":=" expn]
```

A variableDeclaration declares a variable of the specified type. The "at" clause declares a variable at an absolute machine location. Variables may optionally be declared with an initial value which is assigned to the variable when the declaration is executed. The initial value expression must be assignable to the variable's type.

Local variables in procedures and functions may optionally be declared "register". This is a hint to the compiler that it should attempt to allocate the variable to a register. Register variables cannot be bound to nor passed to a reference parameter. A register variable declaration cannot have an "at" clause.

TYPES AND TYPE DECLARATIONS

A typeDeclaration is:

```
[pervasive] type id "=" typeBody
```

The typeBody is one of:

- a. typeDefn
- b. forward

A typeDeclaration gives a name to a type. The type name can subsequently be used in place of the full type definition. A

named type is equivalent to the type that it names (except when exported, see "Type Equivalence and Assignability").

Named types may optionally be declared "pervasive". Type names declared using "pervasive" are automatically imported into all subscopes of the scope in which they are declared. Such types need not be explicitly imported.

Form (b) declares a forward type. A forward type declares a type name whose definition will be given in a later type declaration in the scope. A forward type can be used only as the element type of a collection until its real type definition is given. This allows the declaration of collections whose elements contain pointers to other elements in the collection.

A typeDefn is one of the following:

- a. standardType
- b. manifestConstant ".." manifestExpn
- c. [packed] array indexType of typeDefn
- d. set of baseType
- e. [packed] recordType
- f. pointerType
- g. namedType

The standardTypes are:

SignedInt	- signed integer, implementation defined range (at least -32768..32767)
UnsignedInt	- unsigned integer, implementation defined range (at least 0..65535)
LongInt	- signed integer, implementation defined range (typically 32 bits)
ShortInt	- unsigned integer, implementation defined range (typically a byte)
Boolean	- values are "true" and "false"
Char	- single character
StorageUnit	- no operations or literals, smallest addressable memory unit (typically a byte)
AddressType	- implementation defined integer range

The standard types and the constants true and false are implicitly declared pervasive in the global scope and need not be imported.

Form (b) is a subrange type. The leading constant must be a (possibly negated) literal or manifest named constant and gives the lower bound of the range of values of the type. The expression, which must be manifest, gives the upper bound of the range. The bounds must be both integer values or both character values. The lower bound must be less than or equal to the upper bound.

A scalar type is a subrange, pointer or one of the standard types.

Form (c) is an array type. The indexType must be a subrange type, Char or a named type which is an indexType. The indexType gives the range of subscripts. The typeDefn gives the type of the elements of the array.

Elements of an array variable are referenced using subscripts (see "Variables and Constants") and themselves used as variables. Array variables and constants may be assigned (but not compared) as a whole.

Arrays can be "packed", which allows the compiler to pack the elements more efficiently. The type of string literals is "packed array 1..n of Char" where n is the length of the string.

Form (d) is a set type. The baseType of the set must be a subrange of integer with lower bound 0 or a namedType which is a baseType. An implementation may limit the upper bound of a set type to insure efficient code; this limit will be at least 15.

A recordType is:

```
record
  var id ":" typeDefn
  {var id ":" typeDefn}
end record
```

Variables declared using a record type have the fields given by the variable declarations in the recordType. Fields of a record variable may be referenced using the "." operator (see "Variables and Constants") and themselves used as variables. Record variables may be assigned (but not compared) as a whole.

The variable declarations in a record type must not have initial values and cannot be declared using "register" or "at" clauses.

Records can be "packed", which allows the compiler to pack the elements more efficiently.

A pointerType is:

```
"^" collectionId
```

Variables declared using a pointerType are pointers to dynamically allocated and freed elements of the specified collection; see "Collections". Pointer variables are used as subscripts of the specified collection to select the element to which they point. The selected element can be used as a variable. Pointer variables may be assigned, compared for equality and passed as parameters.

A namedType is:

```
[moduleId "."] typeId
```

The typeId must be a previously declared type name. Type

names exported from a module are referenced outside the module using the "." operator.

TYPE EQUIVALENCE AND ASSIGNABILITY

Two types are defined to be equivalent if they are

- (a) subranges with equal first and last values
- (b) arrays (both packed or both unpacked) with equivalent index types and equivalent component types
- (c) sets with equivalent base types
- (d) pointers to the same collection

A declared type identifier is equivalent to the type it names, with the following exception. When an exported type identifier is used outside its module, as "moduleId.typeId", it is a unique type, equivalent to no other type.

Each type definition for a record type creates a new type that is not equivalent to any other record type definition.

An array value can be assigned to an array variable, a record value assigned to a record variable, a set value assigned to a set variable and a pointer value assigned to a pointer variable only if the source and target of the assignment have equivalent types.

An expression can be assigned to a scalar variable only if (i) the "root" type of the expression and the "root" type of the variable are equivalent, and (ii) the value of the expression is in the range of the variable's type. The "root" type of Char and character subrange types is Char. The root type of SignedInt, UnsignedInt, LongInt, ShortInt, AddressType and integer subranges is integer. The root type of any other type is the type itself.

A variable can be passed to a reference parameter only if its type is equivalent to the parameter type. An expression can be passed to a value parameter only if it is assignable to the parameter type; see "Procedures and Functions".

VARIABLE BINDINGS

A variableBinding is one of:

- a. bind [register] [var] id to variable
- b. bind "(" [register] [var] id to variable
{" , " [register] [var] id to variable } ")"

A variableBinding declares a new identifier for an arbitrary variable reference which may contain subscripts and "."

operators. The new identifier is subsequently used in place of the variable reference within the scope in which the binding appears. If the bound variable is to be assigned to or passed to a var parameter, the binding must be declared using "var". SE does not allow "aliasing" of variables (i.e., having two names for the same variable in a scope). Hence the "root" variable (the first identifier in the variable reference) becomes inaccessible for the scope of the binding.

Form (b) allows bindings to different elements or fields of the same variable or module. Since SE does not allow aliasing of variables, bindings to the same field, element or variable are not allowed.

Local binds in procedures and functions may optionally be declared "register". This is a hint to the compiler to attempt to allocate the bind to a register.

Elements of packed arrays and fields of packed records cannot be bound to.

COLLECTIONS

A collectionDeclaration is:

```
var id ":" collection of typeDefn
```

A collection is essentially an array whose elements are dynamically allocated and freed at run-time. Elements of a collection are referenced by subscripting the collection name with a variable of the collection's pointer type. This subscripting selects the particular element of the collection located by the pointer variable.

Elements of a collection are allocated and freed dynamically by calls to the built-in operations New and Free. "C.New(p)" allocates a new element in the collection C and sets p to point at it. If no more space is available then p is set to "C.nil". "C.Free(p)" frees the element of C pointed at by p and sets p to "C.nil". In each case p is passed as a var parameter and must be a variable of the pointer type of C. These operations are invoked as statements in procedures, see "Statements". They cannot be used in functions.

The built-in constant "C.nil" is the null pointer value for the collection.

Collections themselves cannot be assigned, compared or passed as parameters.

PROCEDURES AND FUNCTIONS

A procedureDeclaration is:

```
procedure id ["(" [var] id ":" parameterType
             {"," [var] id ":" parameterType} ")"] "="
           procedureBody
```

A functionDeclaration is:

```
function id ["(" id ":" parameterType
            {"," id ":" parameterType} ")"]
           returns id ":" resultType "="
           procedureBody
```

A procedure is invoked by a procedure call statement, with actual parameters if required. A function is invoked by using its name, with actual parameters if required, in an expression.

A procedure may return explicitly by executing a return statement or implicitly by reaching the end of the procedure body. A function must return via "return(expn)".

Procedures and functions may optionally take parameters, the types of which are defined in the header. The parameters can be referred to inside the procedure or function using the names declared in the header. Parameters to a procedure may be declared using "var", which means the parameter may be assigned to or further passed as a var parameter inside the procedure. Parameters declared without using "var" are constants and cannot be assigned to or passed as var parameters. Functions are not allowed to have any side-effects and cannot have var parameters. Only variable references can be passed to var parameters.

A parameter is a reference parameter if it is declared using "var" or if its type is an array or record. Other parameters are value parameters. Hence, a value parameter is a non-var parameter whose type is a scalar or set.

A parameterType is one of:

- a. typeDefn
- b. [packed] array manifestConstant ".." parameter of typeDefn
- c. universal

The type of a variable, record or array passed to a reference parameter must be equivalent to the parameter's type with the following exceptions. (1) The upper bound of the index type of an array parameter can be declared using the keyword "parameter" in which case any array whose element type and index type lower bound are equivalent to the parameter's can be passed to the parameter. (2) The type of a parameter can be specified as "universal", in which case a variable or non-manifest named constant of any type can be passed to the parameter. Inside the procedure, a universal parameter is equivalent to a parameter of

type "array 1..parameter of StorageUnit", where the upper bound is the size of the actual parameter in StorageUnits. Parameters declared using "parameter" or "universal" do not have the ".size" standard component and cannot be assigned or compared as a whole. (Note: Full Euclid does not allow forms (b) and (c).)

The type of an expression passed to a value parameter must be assignable to the parameter's type.

SE does not allow "aliasing" of variables (i.e., having two names for a given variable or part of a given variable in the same scope). Hence a variable or part of a variable which is imported directly or indirectly into a procedure cannot be passed to a reference parameter of the procedure. (A variable is directly imported if it appears in the procedure's import list. It is indirectly imported if an imported module or procedure directly or indirectly imports it.)

Elements of packed arrays and fields of packed records cannot be passed to reference parameters.

The returns clause defines the result type of a function. The return identifier is required for compatibility with full Euclid but cannot be referenced.

A resultType is one of:

- a. standardType
- b. manifestConstant ".." manifestExpn
- c. set of baseType
- d. pointerType
- e. namedType

The result type of a function must be a scalar type or set. The expression in a function's return statement must be assignable to the result type.

A procedureBody is:

```
[imports "(" [var] id {"," [var] id} ")"]
begin
  [[not] checked]
  {declarationInRoutine}
  {statement}
end [id]
```

The identifier following the "end" must be the procedure or function identifier. If the procedure is the initially procedure of a module then the end identifier must not be present.

The imports clause of a procedure or function specifies those global identifiers which are to be visible inside the procedure or function. Only those variables imported into a procedure using "var" may be assigned to or passed to a var parameter inside the procedure. Functions are not allowed to have side-effects and cannot import anything "var". This restriction is

transitive; hence a function cannot import a procedure which imports anything "var". A procedure or function which is recursive must explicitly import itself.

Procedures and functions may be "checked"; this causes assert statements, subscripts and case statements to be checked for validity at run-time. In addition, a particular implementation may check other conditions, such as ranges in assignments and overflow in expressions. Procedures and functions not nested inside an unchecked module are checked by default and must be explicitly declared "not checked" to turn off run-time checking.

A procedure returns when it executes a return statement or reaches the end of the procedure. A function is executed similarly but must return via "return(expn)".

Procedures and functions can be separately compiled; see "Separate Compilation".

A declarationInRoutine is one of:

- a. constantDeclaration
- b. variableDeclaration
- c. typeDeclaration
- d. variableBinding
- e. collectionDeclaration
- f. converterDeclaration
- g. assert ["("expn)"]

Modules, procedures and functions cannot be nested inside a procedure or function. Form (g) allows assert statements to appear in declaration lists.

TYPE CONVERTERS

A converterDeclaration is:

converter id "(" typeId ")" returns typeId

A converterDeclaration declares a type converter. A type converter can be used to convert a variable or nonmanifest named constant to a type other than its declared type. Both the parameter and result type of a type converter must be named or standard types. An implementation is not expected to generate any code for a type conversion.

The type of a converted variable or constant must be equivalent to the converter's parameter type. Expressions, literals, manifest values, elements of packed arrays and fields of packed records cannot be type converted.

If the size of the target type is larger than the size of the source type, or the alignment of the target type is more constrained than the alignment of the source type, then the conversion may be meaningless.

STATEMENTS

A statement is one of:

- a. variable `:=` expn
- b. [`moduleId` "."] `procedureId` ["(" expn { "," expn } ")"]
- c. `assert` ["(" expn ")"]
- d. `return` ["(" expn ")"]
- e. `if` expn `then`
 {statement}
 `elseif` expn `then`
 {statement}
 `else`
 {statement}
 `end if`
- f. `loop`
 {statement}
 `end loop`
- g. `exit` [`when` expn]
- h. `case` expn `of`
 manifestExpn { "," manifestExpn } `=>`
 {statement}
 `end manifestExpn`
 {manifestExpn { "," manifestExpn } `=>`
 {statement}
 `end manifestExpn`
 [`otherwise` `=>`
 {statement}]
 `end case`
- i. `begin`
 {declarationInRoutine}
 {statement}
 `end`
- j. `collectionId` "." `New` "(" variable ")"
- k. `collectionId` "." `Free` "(" variable ")"

Form (a) is an assignment statement. The expression is evaluated and the value assigned to the variable. The expression must be assignable to the variable type; see "Type Equivalence and Assignability".

Form (b) is a procedure call. An exported procedure is called outside the module in which it was declared using the "." operator.

The type of an expression passed to a value parameter must be assignable to the parameter's type. The type of a variable or value passed to a reference parameter must be equivalent to the parameter's type. If the upper bound of the type of an array parameter is declared using "parameter", any array whose element type and index type lower bound are equivalent to the parameter's can be passed to the parameter.

An actual parameter passed to a var parameter must be a variable, a bound variable or a var formal parameter. If it is an imported variable, it must have been imported using "var". Since

SE does not allow aliasing of variables, a variable or part of a variable which is passed to a reference parameter cannot be passed to another reference parameter of the same call.

Form (c) is an assert statement. The parenthesized expression is optional; if it is omitted, it can be replaced by a comment. If present, it must be of type Boolean. The expression is evaluated and checked at run time if it appears in a checked scope. Assert statements may appear in both statement lists and declaration lists. They cannot appear inside records.

Form (d) is a return statement. The return statement causes an immediate return from the procedure or function when executed. The optional parenthesized expression gives the value to be returned from a function. The return expression must be assignable to the function's result type. The return expression is required for function returns. It is forbidden for procedure returns. A function must return via a return statement and not implicitly by reaching the end of the function body. A procedure may return either via a return statement or implicitly by reaching the end of the procedure body.

Form (e) is an if statement. The conditional expression following "if" and each "elseif" is evaluated until one of them is found to be true, in which case the statements following the corresponding "then" are executed. If none of the expressions evaluates to true then the statements following "else" are executed; if no "else" is present then execution continues following the if statement. The conditional expressions must be of type Boolean.

Form (f) is the looping construct. The statements within the loop are repeated until one of its "exit" statements or a "return" statement is executed.

Form (g) is a loop exit. When executed, it causes an immediate exit from the nearest enclosing loop. The optional "when" expression makes the exit conditional. If the expression, which must be Boolean, evaluates to true then the exit is executed, otherwise execution of the loop continues. An exit statement cannot appear outside a loop.

Form (h) is a case statement. The case expression is evaluated and used to select one of the alternative labels. The statements which follow the matching label value are executed. If the case expression value does not match any of the label values then the statements following "otherwise" are executed. If no "otherwise" is present, the case expression must match one of the label values. When execution of the statements following the selected label is completed, execution continues following the case statement.

The root type of the case expression must be integer or Char. All of the label expressions must have the same root type as the case expression. Label expressions must be manifest, i.e., their values must be known at compile time. The values of all label

expressions in a given case statement must be distinct. The value of the manifest expression following the end of an alternative must be equal to the first label expression of the alternative.

An implementation may limit the range of case label expression values to insure efficient code; this range will include at least the ranges of Char and ShortInt.

Form (i) is a begin block. Begin blocks can be used to group local declarations within a procedure or function. In particular, they can be used to make local binds.

Forms (j) and (k) are the built-in collection operations New and Free (see "Collections").

VARIABLES AND CONSTANTS

A variable is:

```
[moduleId "."] id {componentSelector}
```

The syntax for variables includes variable and constant references. An exported variable or constant is referenced outside the module in which it is declared using the "." operator.

A componentSelector is one of:

- a. "(" expn ")"
- b. "." id

Form (a) allows subscripting of variable and constant arrays. The type of the subscript expression must be assignable to the index type of the array. The value of the subscript expression must be in the declared range of the index type of the array. Subscripts which appear in checked scopes are checked for validity at run-time.

Form (a) also allows references to elements of a collection. In this case, the subscript expression must be a pointer to an element of the collection.

Form (b) allows record field selection. Fields of a record variable are referenced using the "." operator.

Form (b) also allows standard component references (see "Standard Components").

EXPRESSIONS

An expn is one of the following:

- a. variable
- b. literalConstant
- c. setTypeId "(" elementList ")"
- d. collectionId "." nil
- e. [moduleId "."] functionId ["(" expn {"," expn} ")"]
- f. [moduleId "."] converterId "(" expn ")"
- g. "(" expn ")"
- h. "-" expn
- i. expn arithmeticOperator expn
- j. expn comparisonOperator expn
- k. not expn
- l. expn booleanOperator expn
- m. expn setOperator expn

The arithmeticOperators are +, -, * (multiply), div (truncating integer divide) and mod (integer remainder). The mod operator is defined by "x mod y = x - y*(x div y)". Operands of the arithmetic operators and unary minus must be integers or expressions having root type integer. The arithmetic operators yield an integer result. (Note: +, - and * are also set operators; see below.)

The comparisonOperators are <, >, =, <=, >= and "not =". Operands of comparison operators must either have equivalent types or the same root type; see "Type Equivalence and Assignability". The comparison operators yield a Boolean result. Arrays and records cannot be compared. Sets and Boolean expressions can be compared for equality only. (Note: <= and >= are also set operators; see below.)

The booleanOperators are "and" (intersection), "or" (union) and -> (implication). The Boolean operators and the "not" operator take Boolean operands and yield a Boolean result. The Boolean operators are conditional; that is, if the result of the operation can be determined from the value of the first operand then the second operand is not evaluated.

The set operators are + (set union), - (set difference), * (set intersection), <= and >= (set inclusion), and "in" and "not in" (element containment). The set operators +, - and * take operands of equivalent set types and yield a set result. The set operators <= and >= take operands of equivalent set types and yield a Boolean result. The operators "in" and "not in" take a set as right operand and an integer expression as left operand. They yield a Boolean result.

The order of precedence is among the following classes of operators (most binding first):

1. unary -
2. *, div, mod
3. +, -

4. <, >, =, <=, >=, not =, in, not in
5. not
6. and
7. or
8. ->

Expression form (a) includes references to constants and variables including elements of arrays and collections, fields of records, and constants and variables exported from a module.

Form (b) includes integer, character and string literal constants.

Form (c) is a set constructor. The setTypeId must be the name of a set type. The set constructor returns a set containing the specified elements.

An elementList is one of:

- a. [expn {" , " expn}]
- b. all

The element list is a (possibly empty) list of expressions of the base type of the set, or "all". If "all" is specified, the constructor returns the complete set. If no elements are specified, the constructor returns the empty set.

Expression form (d) is the null pointer value of the specified collection.

Form (e) is a function call. Functions exported from a module are referenced outside the module using the "." operator. An actual parameter to a function must be an expression assignable to the parameter type.

Form (f) is a type conversion. The type of the actual parameter is changed to the result type of the type converter. The actual parameter must be a variable or nonmanifest named constant whose type is equivalent to the source type of the converter. Type converters exported from a module are referenced outside the module using the "." operator.

BUILT-IN FUNCTIONS

SE has three built-in functions, Chr, Ord and Long. "Chr(i)" returns the character whose machine representation is the positive integer value i. "Ord(c)" returns the positive integer machine representation of the character c. Chr and Ord are defined such that for all characters "c" in the machine character set, Chr(Ord(c)) = c. "Long(i)" forces the integer expression i to be extended to LongInt precision; see "Precision of Arithmetic". (Note: In full Euclid, the Ord built-in function is called "Char.Ord".)

STANDARD COMPONENTS

SE defines two standard components, size and address. "T.size" returns the length in StorageUnits (typically bytes) of the machine representation of the variable or type T. "V.address" returns the AddressType machine address of the variable V. The size and address standard components are not allowed for elements of packed arrays and fields of packed records. The address standard component is not allowed for variables declared "register". (Note: In full Euclid, the address standard component is allowed only for variables of type StorageUnit.)

MANIFEST EXPRESSIONS

A manifest expression is an expression whose value can be computed as a literal constant at compile time. The extent of such compile-time computation is implementation dependent, but every implementation will consider at least the following to be manifest:

1. Integer and Char literal constants
2. The Boolean values "true" and "false"
3. Manifest named constants
4. The arithmetic operations unary -, +, -, *, div and mod when both operands are manifest and both the operands and result lie in the range of SignedInt (at least -32768..32767)
5. The built-in functions Chr and Ord when the actual parameter is manifest

A manifestExpn is an expression whose value is manifest. A manifestConstant is a (possibly negated) literal constant or manifest named constant.

PRECISION OF ARITHMETIC

The precision of an arithmetic operation or comparison is determined by the precision of the operands. Operands have one of three precisions which correspond to the standard types SignedInt, UnsignedInt and LongInt.

The precision of a variable or non-manifest named constant operand is determined by its declared type. If its type is SignedInt, ShortInt or any subrange whose bounds both lie in the range of SignedInt then its operand precision is SignedInt. If its type is UnsignedInt or any subrange whose bounds both lie in the range of UnsignedInt but not in SignedInt then its precision is UnsignedInt. Otherwise, its precision is LongInt.

The precision of a literal or manifest named constant operand is SignedInt if its value lies in the range of SignedInt, UnsignedInt if its value lies in the range of UnsignedInt but not of SignedInt, and LongInt otherwise.

The precision of an arithmetic operation or comparison is LongInt if at least one operand has LongInt precision, UnsignedInt if at least one operand has UnsignedInt precision and neither has LongInt precision, and SignedInt otherwise.

The precision of the result of an arithmetic operation is the precision of the operation. Every implementation will guarantee to obtain the arithmetically correct result if the result of an operation lies within the range of the result precision. If the arithmetically correct result lies outside the range of the result precision then the result may be meaningless.

Note that the precision of an operation or comparison can always be forced to LongInt by extending the precision of one or both of the operands using the Long built-in function (see "Built-in Functions").

SOURCE INCLUSION FACILITY

Other source files may be included as part of a program using the "include" statement.

An includeStatement is:

```
include stringLiteral
```

The stringLiteral gives the name of a source file to be included in the compilation. The include statement is replaced in the program source by the contents of the specified file.

Include statements can appear anywhere in a program and can contain any valid source fragment. Included source files can themselves contain include statements.

II. CONCURRENCY FEATURES

The Concurrent Euclid (CE) language is an extension of SE designed to allow concurrent programming with monitors. SE is a subset of Euclid but CE is not, because concurrency and monitors are not features of Euclid.

The concurrency features of CE will be presented in the following order:

- (1) processes, reentrant procedures and modules;
- (2) monitors, entry procedures and functions;
- (3) conditions, signalling and waiting;
- (4) simulation and the busy statement.

PROCESSES

Each CE module (including the main module) can have any number of concurrent processes in it.

A moduleDeclaration is:

```
var id ":"
  module
    [imports "(" [var] id {""," [var] id} ")"]
    [exports "(" id {""," id} ")"]
    [[not] checked]
    {declarationInModule}
    [initially
      procedureBody]
    {process id ["(" memoryRequirement ")"]
      procedureBody}
  end module
```

Each process is like a parameterless procedure. Concurrent execution of the processes of the module begins following execution of the initially procedure of the module. A process terminates by executing its last statement or by executing a return statement in its body. The process identifier is for documentation only since processes cannot be called.

Processes can communicate with each other by changing and inspecting variables declared in the module or imported into it. Generally, however, processes communicate by means of monitors.

Each process requires a certain amount of memory space for its variables. When the process calls a procedure or function, the requirement increases to provide space for the new local variables. When the procedure or function returns, the requirement decreases to its former amount. The programmer can provide his own estimate of the process's required space as a parenthesized manifest integer expression following the keyword "process". This estimate is in StorageUnits (normally bytes) and can be based on previous program executions. If this estimate is omitted, the implementation provides a default space allocation.

All procedures and functions declared in a CE program are reentrant, meaning that they can be executed simultaneously by more than one process.

Modules, monitors, procedures and functions cannot be nested inside a process.

MONITORS

A monitor is essentially a special kind of module which implements inter-process communication with synchronization.

A declarationInModule is one of the following:

- a. constantDeclaration
- b. variableDeclaration
- c. typeDeclaration
- d. variableBinding
- e. moduleDeclaration
- f. monitorDeclaration
- g. collectionDeclaration
- h. procedureDeclaration
- i. functionDeclaration
- j. converterDeclaration
- k. assert ["(" expn ")"]

Monitors may only be declared inside modules. Monitors cannot be nested inside procedures, functions or other monitors.

A monitorDeclaration is:

```
var id ":"
  monitor
    [imports "(" [var] id {"," [var] id} ")"]
    [exports "(" id {"," id} ")"]
    [[not] checked]
    {declarationInMonitor}
    [initially
      procedureBody]
  end monitor
```

The imports list of a monitor specifies the global identifiers which are accessible inside the monitor, exactly like the imports list in a module.

The exports list of a monitor specifies those identifiers defined inside the monitor which may be accessed outside the monitor using the "." operator. Unlike modules, monitors cannot export variables.

Procedures and functions which are exported from a monitor are called monitor entries. Entry procedures and functions of a monitor cannot be invoked inside the monitor. Outside the monitor, entry procedures and functions can be invoked exactly like the procedures and functions of a module, using the "." operator.

Procedures and functions which are entries of a monitor cannot be separately compiled except as part of the entire monitor.

It is guaranteed that only one process at a time will be executing inside a monitor. As a result, mutually exclusive access to a monitor's variables is implicitly provided, since a monitor cannot export any variables. If a process calls an entry of a monitor while another process is executing in the monitor, the calling process will be blocked and not allowed in the monitor until no other process is executing in the monitor.

A declarationInMonitor is one of the following:

- a. constantDeclaration
- b. variableDeclaration
- c. typeDeclaration
- d. variableBinding
- e. conditionDeclaration
- f. collectionDeclaration
- g. procedureDeclaration
- h. functionDeclaration
- i. converterDeclaration
- j. assert ["(" expn ")"]

Modules and monitors cannot be declared inside a monitor. A monitor cannot contain a nested process.

Monitors can be separately compiled; see "Separate Compilation".

CONDITIONS

A conditionDeclaration is one of:

- a. var id ":" [priority] condition
- b. var id ":" array indexType of [priority] condition

The only place a condition can be declared is as a field of a monitor. The only allowed use of conditions is in the "wait" and "signal" statements and in the "empty" built-in function. Conditions cannot be assigned, compared or passed as parameters. Arrays of conditions are allowed. Conditions may be imported "var" (or not). An imported condition can be used in a wait or signal statement only if it is imported "var".

Two new statements are introduced:

```
wait "(" conditionVar ["," priorityValue] ")"  
signal "(" conditionVar ")"
```

Where a conditionVar is:

```
conditionId ["(" expn ")"]
```

The wait and signal statements each specify a conditionVar.

Each of these must be a conditionId or a subscripted condition array. These statements can appear only in monitors, but not in a monitor's initially procedure.

When a process executes a wait statement for condition C it is blocked and is removed from the monitor. When a process executes a signal statement for condition C, one of the processes (if there are any) waiting for condition C is unblocked and allowed immediately to continue executing the monitor. The signalling process is temporarily removed from the monitor and is not allowed to continue execution until no processes are in the monitor. If no processes were waiting for condition C, the only effect of the signal statement is that the signalling process may be removed from the monitor. The signalling process cannot in general know whether other processes have entered the monitor before the signaller continues in the monitor.

If the condition variable is declared with the "priority" option, the wait statement must specify a priority value; otherwise the priority value is not allowed in wait. The priorityValue is a SignedInt expression that must evaluate to a nonnegative integer value. The processes waiting for a priority condition are ranked in order of their specified priority values, and the process with the smallest priority value is the first to be unblocked by a signal statement.

In the case of processes waiting for non-priority conditions, or waiting with identical priorities for a priority condition, the scheduling is "fair", meaning that a particular waiting process will eventually be unblocked given enough signals on the condition.

A predefined function named "empty" accepts a condition as a parameter. It returns the Boolean value "true" if no processes are waiting for the condition, otherwise "false". Like wait and signal, "empty" can appear only inside a monitor, but not in the initially procedure of a monitor.

The variables in a monitor represent its state. For example, if a monitor allocates a single resource, only one variable inside the monitor is needed and it can be declared as Boolean. When this variable is true, it represents the state in which the resource is available, when false it represents the state of being allocated. When a process enters the monitor and finds that it does not have the desired state, the process leaves the monitor and becomes blocked by executing a wait statement on a condition. The condition corresponds to the state that the process is waiting for. Suppose a process enters a monitor and changes its state to a state that may be waited for by other processes. The process should execute a signal statement for the condition corresponding to the new state. If there are processes waiting for this state transition, then they will be blocked on the condition, and one of them will immediately resume execution in the monitor. Because of this immediate resumption, the signalled process knows the monitor is in the desired state, without testing monitor variables. The signalling process is allowed to

continue executing only when no other processes are in the monitor. If no processes were waiting on the condition, the only effect of the signal statement is to temporarily remove the signallinger from the monitor.

As specified by Hoare, monitors and conditions are intended to be used in the following manner. The programmer should associate with the monitor's variables a consistency criterion. The consistency criterion is a Boolean expression that should be true between monitor activations, or whenever a process enters or leaves a monitor. Hence, the programmer should see that it is made true before each signal or wait statement in the monitor and before each return from an entry of the monitor. The programmer should also associate a Boolean expression, call it E_i , with each condition C_i . The expression E_i should be true whenever a signal is executed for condition C_i . A process that is unblocked after waiting for a condition knows that E_i is true because the signalled process (not the signalling process) executes first. (The consistency criterion and each E_i for a condition do not necessarily appear as executable code in the monitor.) In general, when a process changes the monitor's state so that one of the awaited relations E_i becomes true, the corresponding condition C_i should be signalled.

THE BUSY STATEMENT

A statement is introduced to allow simulation using timing delays:

```
busy "(" time ")"
```

The time must be a nonnegative SignedInt expression. The busy statement can be understood in terms of simulated time recorded by a system clock. This clock is set to zero at the beginning of execution of a program. With the exception of the busy statement (or wait statements causing an indirect delay for a busy statement), statements take negligible simulated time to execute. When the programmer wants to specify that a certain action takes time to complete, the busy statement is used. The process that executes the busy statement is delayed until the system clock ticks (counts off) the specified number of time units.

III. SEPARATE COMPILATION

This section describes the extensions made to CE to allow separate compilation of procedures, functions, modules and monitors.

EXTERNAL DECLARATIONS

Procedures, functions, modules and monitors may be declared "external", which means that they are to be separately compiled and joined with the program at link time. Due to linker restrictions, a particular implementation may be forced to place a limit on the number of significant characters in external module, monitor, procedure and function identifiers.

An externalProcedureDeclaration is:

```
procedure id ["(" [var] id ":" parameterType
             {"," [var] id ":" parameterType} ")"] "="
external
```

An externalFunctionDeclaration is:

```
function id ["(" id ":" parameterType
            {"," id ":" parameterType} ")"]
            returns id ":" resultType "="
external
```

An externalModuleDeclaration is:

```
var id ":"
external module
    [imports "(" [var] id {"," [var] id} ")"]
    [exports "(" id {"," id} ")"]
    {declarationInExternalModule}
end module
```

A declarationInExternalModule is one of:

- a. manifestConstantDeclaration
- b. typeDeclaration
- c. collectionDeclaration
- d. converterDeclaration
- e. externalProcedureDeclaration
- f. externalFunctionDeclaration

An externalMonitorDeclaration is:

```
var id ":"
external monitor
    [imports "(" [var] id {"," [var] id} ")"]
```

```
    [exports "(" id {" ," id} ")"]
    {declarationInExternalMonitor}
end monitor
```

A declarationInExternalMonitor is one of:

- a. manifestConstantDeclaration
- b. typeDeclaration
- c. collectionDeclaration
- d. converterDeclaration
- e. externalProcedureDeclaration
- f. externalFunctionDeclaration

An external declaration can appear in place of the real declaration and specifies that the corresponding procedure, function, module or monitor is to be compiled separately.

Processes and initially procedures of modules cannot be declared external. Procedures and functions which are entries of a monitor cannot be declared external except as part of an external monitor declaration. Nonmanifest and array named constants cannot be declared in an external module or monitor.

COMPILATIONS

A compilation can consist of a main program (see "Programs") or a separate compilation.

A separateCompilation is:

```
{separateDeclaration}
```

Each separateDeclaration is one of the following:

- a. manifestConstantDeclaration
- b. typeDeclaration
- c. collectionDeclaration
- d. converterDeclaration
- e. procedureDeclaration
- f. functionDeclaration
- g. moduleDeclaration
- h. monitorDeclaration

Each separateDeclaration can be a manifest constant declaration, a type declaration, a collection declaration, a procedure or function declared as "external" in another compilation, or a module or monitor declared as "external" in another compilation.

Separately compiled procedures, functions, modules and monitors can be linked to form a complete program. Variables cannot be separately compiled and are not linked across compilations. Consistency of constants, types and collections is not automatically checked across compilations. Consistency of the type and

number of formal parameters and function results between the external declaration and the separate compilation of separately compiled procedures and functions is not automatically checked.

Separately compiled modules and monitors will be initialized at the point of the corresponding "external" declaration. Note that since execution of a program consists of initializing the main module (see "Programs"), only those modules and monitors which are declared in the main module or a module nested within it will be initialized.

LINKING OF COMPILATIONS

A complete program will typically consist of a main module compilation linked together with the separate compilations of any procedures, functions, modules and monitors declared as "external" in it. The compilations must be linked such that the entry point of the program is the beginning of the main module compilation. (Under many systems, this means simply that the main module compilation must be the first in the list of object modules to be linked together.)

APPENDIX 1.
COLLECTED SYNTAX OF CONCURRENT EUCLID

The syntax of SE is given first. Throughout the following, {item} means zero or more of the item, and [item] means the item is optional.

The following abbreviations are used:

id for identifier
expn for expression
typeDefn for typeDefinition

Semicolons are not required, but they may optionally appear following statements, declarations and import, export and checked clauses.

A program is:

moduleDeclaration

A moduleDeclaration is:

```
var id ":"
  module
    [imports "(" [var] id {""," [var] id} ")"]
    [exports "(" id {""," id} ")"]
    [[not] checked]
    {declarationInModule}
    [initially
      procedureBody]
  end module
```

A declarationInModule is one of the following:

- a. constantDeclaration
- b. variableDeclaration
- c. typeDeclaration
- d. variableBinding
- e. moduleDeclaration
- f. collectionDeclaration
- g. procedureDeclaration
- h. functionDeclaration
- i. converterDeclaration
- j. assert ["(" expn ")"]

A constantDeclaration is one of:

- a. [pervasive] const id "!=" manifestExpn
- b. [pervasive] const id ":" typeDefn "!=" expn
- c. [pervasive] const id ":" typeDefn "!="
 "(" manifestExpn {""," manifestExpn} ")"

d. [pervasive] const id " := " stringLiteral

A manifestExpn is:

expn

A variableDeclaration is:

[register] var id ["(" at manifestExpn ")"] ":" typeDefn
[" := " expn]

A typeDeclaration is:

[pervasive] type id "=" typeBody

The typeBody is one of:

- a. typeDefn
- b. forward

A typeDefn is one of the following:

- a. standardType
- b. manifestConstant ".." manifestExpn
- c. [packed] array indexType of typeDefn
- d. set of baseType
- e. [packed] recordType
- f. pointerType
- g. namedType

A standardType is one of:

- a. SignedInt
- b. UnsignedInt
- c. LongInt
- d. ShortInt
- e. Boolean
- f. Char
- g. StorageUnit
- h. AddressType

A manifestConstant is one of:

- a. ["-"] literalConstant
- b. ["-"] [moduleId] "." manifestConstantId

A manifestConstantId is:

id

An indexType is one of:

- a. Char
- b. manifestConstant ".." manifestExpn
- c. namedType

A baseType is one of:

- a. 0 ".." manifestExpn
- b. namedType

A recordType is:

```
record
  var id ":" typeDefn
  {var id ":" typeDefn}
end record
```

A pointerType is:

"^" collectionId

A collectionId is:

id

A namedType is:

[moduleId "."] typeId

A moduleId is:

id

A typeId is:

id

A variableBinding is one of:

- a. bind [register] [var] id to variable
- b. bind "(" [register] [var] id to variable
{"", [register] [var] id to variable} ")"

A collectionDeclaration is:

```
var id ":" collection of typeDefn
```

A procedureDeclaration is:

```
procedure id ["(" [var] id ":" parameterType  
             {"," [var] id ":" parameterType} ")"] "="  
  procedureBody
```

A functionDeclaration is:

```
function id ["(" id ":" parameterType  
            {"," id ":" parameterType} ")"]  
  returns id ":" resultType "="  
  procedureBody
```

A parameterType is one of:

- a. typeDefn
- b. [packed] array manifestConstant ".." parameter of
 typeDefn
- c. universal

A resultType is one of:

- a. standardType
- b. manifestConstant ".." manifestExpn
- c. set of baseType
- d. pointerType
- e. namedType

A procedureBody is:

```
[imports "(" [var] id {"," [var] id} ")"]  
begin  
  [[not] checked]  
  {declarationInRoutine}  
  {statement}  
end [id]
```

A declarationInRoutine is one of:

- a. constantDeclaration
- b. variableDeclaration
- c. typeDeclaration
- d. variableBinding
- e. collectionDeclaration
- f. converterDeclaration
- g. assert ["(" "expn")"]

A converterDeclaration is:

converter id "(" typeId ")" returns typeId

A statement is one of:

- a. variable ":=" expn
- b. [moduleId"."] procedureId ["(" expn {" expn } ")"]
- c. assert ["("expn")"]
- d. return ["("expn")"]
- e. if expn then
 {statement}
elseif expn then
 {statement}}
else
 {statement}}
end if
- f. loop
 {statement}
end loop
- g. exit [when expn]
- h. case expn of
 manifestExpn {" , " manifestExpn "=>"
 {statement}
 end manifestExpn
 {manifestExpn {" , " manifestExpn "=>"
 {statement}
 end manifestExpn}
 [otherwise "=>"
 {statement}]
end case
- i. begin
 {declarationInRoutine}
 {statement}
end
- j. collectionId "." New "(" variable ")"
- k. collectionId "." Free "(" variable ")"

A procedureId is:

id

A variable is:

[moduleId "."] id {componentSelector}

A componentSelector is one of:

- a. "(" expn ")"
- b. "." id
- c. "." size
- d. "." address

An expn is one of the following:

- a. variable
- b. literalConstant
- c. setTypeId "(" elementList ")"
- d. collectionId "." nil
- e. [moduleId "."] functionId ["(" expn {"," expn} ")"]
- f. [moduleId "."] converterId "(" expn ")"
- g. "(" expn ")"
- h. "-" expn
- i. expn arithmeticOperator expn
- j. expn comparisonOperator expn
- k. not expn
- l. expn booleanOperator expn
- m. expn setOperator expn

A setTypeId is:

id

A elementList is one of:

- a. [expn {"," expn}]
- b. all

A functionId is one of:

- a. id
- b. Chr
- c. Ord
- d. Long

A converterId is:

id

An arithmeticOperator is one of:

- a. +
- b. -
- c. *
- d. div
- e. mod

A comparisonOperator is one of:

- a. <
- b. >
- c. =
- d. <=

- e. >=
- f. not =

A booleanOperator is one of:

- a. and
- b. or
- c. ->

A setOperator is one of:

- a. +
- b. -
- c. *
- d. <=
- e. >=
- f. in
- g. not in

Note: The order of precedence is among the following classes of operators (most binding first):

1. unary -
2. *, div, mod
3. +, -
4. <, >, =, <=, >=, not =, in, not in
5. not
6. and
7. or
8. ->

An includeStatement is:

include stringLiteral

Note: Include statements can appear anywhere in a program.

The following changes and additions are made to form CE:

A moduleDeclaration is:

```
var id ":"
  module
    [imports "(" [var] id {"," [var] id} ")"]
    [exports "(" id {"," id} ")"]
    [[not] checked]
    {declarationInModule}
    [initially
      procedureBody]
    {process id [{" memoryRequirement "}]
      procedureBody}
  end module
```

A memoryRequirement is:

```
manifestExpn
```

A declarationInModule is one of the following:

- a. constantDeclaration
- b. variableDeclaration
- c. typeDeclaration
- d. variableBinding
- e. moduleDeclaration
- f. monitorDeclaration
- g. collectionDeclaration
- h. procedureDeclaration
- i. functionDeclaration
- j. converterDeclaration
- k. assert [{" expn "}]

A monitorDeclaration is:

```
var id ":"
  monitor
    [imports "(" [var] id {"," [var] id} ")"]
    [exports "(" id {"," id} ")"]
    [[not] checked]
    {declarationInMonitor}
    [initially
      procedureBody]
  end monitor
```

A declarationInMonitor is one of the following:

- a. constantDeclaration
- b. variableDeclaration
- c. typeDeclaration

- d. variableBinding
- e. conditionDeclaration
- f. collectionDeclaration
- g. procedureDeclaration
- h. functionDeclaration
- i. converterDeclaration
- j. assert ["(" expn ")"]

A conditionDeclaration is one of:

- a. var id ":" [priority] condition
- b. var id ":" array indexType of [priority] condition

A statement is one of:

- a. variable ":=" expn
- b. [moduleId"."] procedureId ["(" expn {" , " expn }")"]
- c. assert ["("expn")"]
- d. return ["("expn")"]
- e. if expn then
 - {statement}
 - {elseif expn then
 - {statement}}
 - [else
 - {statement}]
 - end if
- f. loop
 - {statement}
 - end loop
- g. exit [when expn]
- h. case expn of
 - manifestExpn {" , " manifestExpn} "=>"
 - {statement}
 - end manifestExpn
 - {manifestExpn {" , " manifestExpn} "=>"
 - {statement}
 - end manifestExpn
 - [otherwise "=>"
 - {statement}]
 - end case
- i. begin
 - {declarationInRoutine}
 - {statement}
 - end
- j. collectionId ". New "(" variable ")"
- k. collectionId ". Free "(" variable ")"
- l. wait "(" conditionVar [{" , " priorityValue} ")"
- m. signal "(" conditionVar ")"
- n. busy "(" time ")"

A moduleId is:

moduleOrMonitorId

A moduleOrMonitorId is:

id

A conditionVar is:

conditionId ["(" expn ")"]

A conditionId is:

id

A priorityValue is:

expn

A time is:

expn

A functionId is one of:

- a. id
- b. Chr
- c. Ord
- d. Long
- e. empty

The following extensions allow separate compilation of procedures, functions, modules and monitors:

An externalProcedureDeclaration is:

```
procedure id ["(" [var] id ":" parameterType
             {"," [var] id ":" parameterType} ")"] "="
external
```

An externalFunctionDeclaration is:

```
function id ["(" id ":" parameterType
            {"," id ":" parameterType} ")"]
           returns id ":" resultType "="
external
```

An externalModuleDeclaration is:

```
var id ":"
external module
  [imports "(" [var] id {"," [var] id} ")"]
  [exports "(" id {"," id} ")"]
  {declarationInExternalModule}
end module
```

A declarationInExternalModule is one of:

- a. manifestConstantDeclaration
- b. typeDeclaration
- c. collectionDeclaration
- d. converterDeclaration
- e. externalProcedureDeclaration
- f. externalFunctionDeclaration

An externalMonitorDeclaration is:

```
var id ":"
external monitor
  [imports "(" [var] id {"," [var] id} ")"]
  [exports "(" id {"," id} ")"]
  {declarationInExternalMonitor}
end monitor
```

A declarationInExternalMonitor is one of:

- a. manifestConstantDeclaration
- b. typeDeclaration
- c. collectionDeclaration
- d. converterDeclaration
- e. externalProcedureDeclaration

f. externalFunctionDeclaration

Note: An external declaration can appear in place of the real declaration anywhere in a program.

A manifestConstantDeclaration is:

[pervasive] const id ":@" manifestExpn

A separateCompilation is:

{separateDeclaration}

Each separateDeclaration is one of the following:

- a. manifestConstantDeclaration
- b. typeDeclaration
- c. collectionDeclaration
- d. converterDeclaration
- e. procedureDeclaration
- f. functionDeclaration
- g. moduleDeclaration
- h. monitorDeclaration

APPENDIX 2.
KEYWORDS AND PREDEFINED IDENTIFIERS

The following are reserved words of Euclid. These must not be used as identifiers in SE and CE programs. Those which are not in the SE subset are marked with an *.

*abstraction	*aligned	all	and
*any	array	assert	at
begin	bind	*bits	*bound
case	*checkable	checked	*code
collection	const	converter	*counted
*decreasing	*default	*dependent	div
else	elseif	end	exit
exports	*finally	*for	forward
*from	function	if	imports
in	include	initially	*inline
*invariant	loop	machine	mod
not	of	or	otherwise
packed	parameter	pervasive	*post
*pre	procedure	*readonly	record
return	returns	set	then
*thus	to	type	*unknown
var	when	*with	*xor

The following are additional reserved words of SE and CE. These also must not be used as identifiers in SE and CE programs.

busy	condition	empty	monitor
priority	process	register	signal
universal	wait		

The following are predefined identifiers of Euclid. In general, these are pervasive and must not be redeclared in SE and CE programs. Those which are not in the SE subset are marked with an *.

*Abs	address	AddressType	*alignment
*BaseType	Boolean	Char	Chr
*ComponentType	false	*first	Free
*Index	*IndexType	*Integer	*itsTag
*ItsType	*last	*Max	*Min
New	nil	*ObjectType	*Odd
Ord	*Pred	*refCount	SignedInt
size	*sizeInBits	StorageUnit	*String
*StringIndex	*stringMaxLength		*Succ
*SystemZone	true	UnsignedInt	

The following are additional predefined identifiers of SE and CE. These also must not be redeclared in SE and CE programs.

Long	LongInt	ShortInt
------	---------	----------

APPENDIX 3.
INPUT/OUTPUT IN CONCURRENT EUCLID

This paper presents the standard input/output package for SE and CE. The user can access the I/O facility by including in his program the stub input/output module which corresponds to the level of I/O which his program requires. In this way, the user's compiled and linked program will include code only for the I/O facilities required.

The package provides four levels of sophistication, which are called "IO/1" through "IO/4". Each level includes all the facilities of the previous levels plus certain new features. The levels are as follows:

IO/1: Terminal (standard) input and output; Formatted text input/output of integers, characters and strings (Get and Put).

IO/2: Program argument sequential files; Open and close on argument files; Formatted text input/output of integers, characters and strings to files (FGet and FPut); Internal representation input/output of integers, characters and strings to files (Read and Write); End of file detection (EndFile).

IO/3: Temporary and non-argument sequential files (Assign, Deassign, Delete); Program arguments (FetchArg); Program error exit (SysExit).

IO/4: Record, array and storage input/output (Read and Write); Random access files (Tell and Seek); Error detection (Error).

The procedures and functions of the input/output system are all part of the module "IO" and must be referenced using "IO.". The types and constants which form the interface to the module are global. The user can access the level n facilities of the input/output module by including the statement

```
include '/usr/lib/coneuc/IO'n'
```

as the first declaration in his main module.

We now describe the input/output facilities in detail.

IO/1: Terminal Formatted Text I/O

```
pervasive const newLine := $&N  
pervasive const endOfFile := $&E  
pervasive const maxStringLength :=
```

```
    { Implementation defined; >= 128 }
```

Strings read and written by the input/output routines may be up to maxStringLength characters in length.

IO.PutChar (c: Char)
 Prints the character c on the terminal.

IO.PutInt (i: SignedInt, w: SignedInt)
 Prints the integer i on the terminal, right justified in a field of w characters. Leading blanks are supplied to fill the field. If w is an insufficient width, the value is printed in the minimum possible width with no leading blanks. In particular, if w is 1 then the exact number of characters needed is used. The specified width must be greater than zero and less than maxStringLength.

IO.PutLong (i: LongInt, w: SignedInt)
 Same as IO.PutInt for long integers.

IO.PutString (s: packed array 1..parameter of Char)
 Prints the string s on the terminal. The string must be terminated by an endOfFile character ('\$E'), which is not output. It can contain embedded newLines ('\$N') if desired. (Note: An endOfFile character (\$\$E) can be output using PutChar.)

IO.GetChar (var c:Char)
 Gets a the next input character from the terminal. End of file is indicated by a return of endOfFile (\$\$E).

IO.GetInt (var i: SignedInt)
 Gets an integer from the terminal. The input must consist of any number of optional blanks, tabs and newlines, followed by an optional minus sign, followed by any number of decimal digits.

IO.GetLong (var i: LongInt)
 Same as IO.GetInt for long integers.

IO.GetString (var s: packed array 1..parameter of Char)
 Gets a line of character input from the terminal. The returned string may be up to maxStringLength characters in length. The string returned is ended with the newLine character ('\$N') followed by an endOfFile character ('\$E') if it is a complete line, and by the endOfFile character only if it is a partial line (i.e., if the input line exceeds maxStringLength characters in length). End of file is indicated by returning a string containing endOfFile ('\$E') as the first character.

IO/2: Sequential Argument File I/O

```
pervasive const stdInput := -2
pervasive const stdOutput := -1
pervasive const stdError := 0
pervasive const maxArgs := { Implementation defined; >= 9 }
pervasive const maxFiles :=
  { Implementation defined; >= maxArgs+5 }
type File = stdInput..maxFiles
```

Concurrent Euclid input/output refers to files using a file number. Certain file numbers are preassigned as follows: -2 refers to the terminal input; -1 is the terminal output; 0 is the standard diagnostic output. The file numbers 1..maxArgs refer to the program arguments. The remaining file numbers (maxArgs+1..maxFiles) can be dynamically assigned to files using the "IO.Assign" operation; see "IO/3".

```
pervasive const inFile := 0
pervasive const outFile := 1
pervasive const inOutFile := 2
type FileMode = inFile..inOutFile
Files can be opened for input, output, or input/output using
modes inFile, outFile and inOutFile respectively. (Note:
The input/output mode is not available under Unix V6.)
```

IO.Open (f: File, m: FileMode)

IO.Close (f: File)

With the exception of terminal input/output and the standard diagnostic output, files must be opened before they are used and closed before the program returns. Open opens an existing file for the operations specified by the mode. If the opened file does not exist, it is created. The file number specified must be a preassigned file number or a file number returned from a call to "IO.Assign"; see "IO/3".

IO.FPutChar (f: File, c: Char)

IO.FPutInt (f: File, i: SignedInt, w: SignedInt)

IO.FPutLong (f: File, i: LongInt, w: SignedInt)

IO.FPutString (f: File, s: packed array 1..parameter of Char)

IO.FGetChar (f: File, var c: Char)

IO.FGetInt (f: File, var i: SignedInt)

IO.FGetLong (f: File, var i: LongInt)

IO.FGetString (f: File, var s: packed array 1..parameter of Char)

These operations are identical to the terminal input/output operations of IO/1 except that the put or get is done on the specified file.

IO.WriteChar (f: File, c: Char)

Writes the internal representation of character c to the specified file.

IO.WriteInt (f: File, i: SignedInt)

Writes the internal representation of integer i to the specified file.

IO.WriteLong (f: File, i: LongInt)

Writes the internal representation of long integer i to the specified file.

IO.WriteString (f: File, s: packed array 1..parameter of Char)

Writes the internal representations of the characters in the string s to the specified file.

IO.ReadChar (f: File, var c: Char)

Reads a character in internal representation from the

specified file into c.

- IO.ReadInt (f: File, var i: SignedInt)
Reads an integer in internal representation from the specified file into i.
- IO.ReadLong (f: File, var i: LongInt)
Reads a long integer in internal representation from the specified file into i.
- IO.ReadString (f: File, var s: packed array 1..parameter of Char)
Reads a string of characters terminated by a newLine character ('\$N') in internal representation from the specified file into s. The returned string may be up to maxStringLength characters in length. The string returned is ended with the newLine character ('\$N') followed by an endOfFile character ('\$E') if it is a complete line, and by the endOfFile character only if it is a partial line (i.e., if the input line exceeds maxStringLength characters in length). End of file is indicated by returning a string containing endOfFile ('\$E') as the first character.
- IO.EndFile (f: File)
A function which returns true if the last operation on the specified input file encountered end of file and false otherwise.

IO/3: Temporary and Non-argument Files

- pervasive const maxArgLength :=
 { Implementation defined; >= 32 }
File names and arguments to a program may be up to maxArgLength characters in length.
- IO.Assign (var f: File, s: packed array 1..parameter of Char)
A file number is assigned to the file name supplied in s. The file name is given as a string terminated by the endOfFile character ('\$E'), which is not part of the name. Before the file can be used it must be opened using "IO.Open".
- IO.Deassign (f: File)
The specified file number is freed for assignment to another file name. An open file cannot be deassigned.
- IO.Delete (f: File)
The specified file is destroyed. An open file cannot be deleted. Note that a program can have temporary files using "IO.Assign" and "IO.Delete".
- IO.FetchArg (n: 1..maxArgs, var s: packed array 1..parameter of Char)
The program argument specified by "n" is returned in string s. The returned string is terminated by the endOfFile character ('\$E') and may be up to maxArgLength characters in length.

IO.SysExit (n: SignedInt)
Terminate program execution with the specified return code.
(ConEuc programs return 0 by default.)

IO/4: Structure Input/Output and Random Access Files

IO.Write (f: File, u: universal, n: SignedInt)
The number of StorageUnits specified by "n" are written to the file from u. Write can be used to write out whole arrays and records using a call of the form "IO.Write (f, v, v.size)". The value of n must be positive or zero.

IO.Read (f: File, var u: universal, n: SignedInt)
The number of StorageUnits specified by "n" are read from the file into u. Read can be used to read in whole arrays and records using a call of the form "IO.Read (f, v, v.size)". The value of n must be positive or zero.

type FileIndex = LongInt

IO.Tell (f: File, var x: FileIndex)

IO.Seek (f: File, x: FileIndex)

These operations provide random access input/output by allowing the program to sense a file position, represented as a long integer, and reset the file to a remembered position. Tell returns the current position of the specified file. Seek sets the current position of the specified file to the position specified by the value of x. The representation of file indices is implementation-dependent. (Note: "IO.Tell" and "IO.Seek" are not supported under Unix V6.)

IO.Error (f: File)

A function which returns true if the last operation on the specified file encountered an error and false otherwise.

Interfacing to Unix*

The input/output package is based on standard Unix input/output and is designed to be interfaced to Unix with a minimum of overhead. The Unix implementation is written in C and uses only facilities of the C "stdio" package. This implementation can be compiled unchanged under both V6 and V7 Unix.

Unix* is a trademark of Bell Laboratories.

APPENDIX 4.
PDP-11 IMPLEMENTATION NOTES

This section gives details of the implementation of CE for the PDP-11 under Unix* and provides information necessary for interfacing with CE programs.

DATA REPRESENTATION

The following gives the storage representations of the various CE data types used by the PDP-11 implementation.

<u>Type</u>	<u>Representation</u>
SignedInt and subranges contained in -32768..32767	16-bit signed word
UnsignedInt and subranges contained in 0..65535 but outside -32768..32767	16-bit unsigned word
LongInt and subranges outside the above	32-bit signed doubleword, word aligned; high order word has the lower address
ShortInt and packed subranges in 0..255	8-bit unsigned byte
Boolean	8-bit unsigned byte; true = 1, false = 0
Char	8-bit unsigned byte
StorageUnit	8-bit unsigned byte
AddressType, pointers and binds	16-bit unsigned word
sets of 0..7	8-bit unsigned byte; element 0 is low order bit, element 7 is high order bit
sets of 0..15	16-bit unsigned word; element 0 is low order bit, element 15 is high order bit

REGISTER USAGE

The following register assignments are used by the PDP-11 implementation.

<u>Register</u>	<u>Use</u>
R0, R1	function results, scratch
R2, R3	scratch
R4	line number, register variables and binds
R5	register variables and binds

Since the CE implementation uses the stack pointer register (SP) to address local variables in procedures and functions, there is no local base register.

Function results whose data representation is a byte or word are returned in R0. Doubleword results are returned in R0 and R1, with the high order word in R0.

In order to attain highly efficient code for non-scalar assignments, subscripting and LongInt arithmetic, the CE compiler uses four scratch registers rather than the two used by the C compiler. In particular, CE uses R2 and R3 for scratch and hence does not save and restore them at procedure and function entry and exit. Since the PDP-11 C compiler uses R2 and R3 for register variables, C routines which call CE procedures and functions can use at most one register variable. There is no such restriction on C routines called from CE programs.

Register R5 (and R4 when line numbering is turned off, see below) are used for user variables and binds which are explicitly declared "register".

When run-time line numbering is turned on (which is the default), the CE compiler generates code to maintain the source file and line number in the line number register (R4) during execution. This aids in debugging since the "cedb" program can obtain the source file name and line number from the core dump following a run-time program failure (e.g., assertion failure, subscript or case tag out of range, etc.).

The contents of the line number register is interpreted as a 5 digit unsigned decimal number, the first two digits of which give the source include file number and the last three of which give the source line number within file. Source file numbers are assigned sequentially starting with 1 for the main source file. Source files longer than 999 lines are assigned a new file number for each 1000 lines of source.

Run time line numbering can be turned off using the "-l" compiler toggle.

CALLING CONVENTIONS

CE procedures and functions which are (a) declared "external",

(b) separately compiled, or (c) exported from a separately compiled module or monitor, are called using the C calling convention. A more efficient calling convention is used for calls between CE routines within a single compilation.

Unlike C routines, CE procedures and functions do not save and restore all of the caller's registers, but rather save and restore only those registers which they actually use. Note that since registers R0-R3 are considered scratch registers by the CE compiler, CE routines never save and restore R0-R3. This means that C routines which call CE routines can use at most one register variable. C routines which are called from CE may of course use as many register variables as they wish. Assembly routines called from CE can use R0-R3 as scratch and need not save and restore them. (Exception: the CE built-in routines are called using a special calling convention and must save and restore all registers which they use).

EXTERNAL NAMES

CE procedures and functions which are (a) declared "external", (b) separately compiled, or (c) exported from a separately compiled module or monitor, are assigned external names so that they may be linked with and/or called from other compilations and programs. On the PDP-11 under Unix, these names consist of the routine name preceded by an underscore character. Because of Unix linker restrictions, only the first seven characters of external names are significant and hence care must be taken to avoid conflicts. The "initially" routine of an external module or monitor is given the name of the module/monitor.

PARAMETER PASSING

Like C, CE passes parameters on the PDP-11 stack. Unlike C, however, CE pushes parameters onto the stack in the order in which they appear in the call (C reverses this order). Hence C procedures and functions which are called from CE (and CE procedures and functions which are called from C) must declare their formal parameters in reversed order.

Value parameters as defined in the CE language specification are passed as values on the stack. Byte values are passed in the low order byte of a 16-bit word. Reference parameters are passed as 16-bit word addresses.

A parameter passed to array formal parameter declared using the "parameter" keyword as upper bound is passed with an extra unsigned word parameter following the array address. This extra parameter gives the number of elements in the array minus one. A parameter passed to a "universal" formal parameter is passed as an address only.

RUN-TIME CHECKING

When run-time checking is turned on (which is the default), the CE compiler will generate code to check assert statements, subscript ranges and case selector ranges during execution. It will not generate code to check ranges in assignments and overflow in expressions at run-time. The checking code uses an illegal instruction of the form "jsr r0,rN" to abort the program when a run-time check fails. The second register number in the instruction is an abort code indicating the reason for the abort. The following table gives the abort codes used by the PDP-11 implementation.

<u>Aborting instruction</u>	<u>Reason for abort</u>
jsr r0,r0	assertion failure
jsr r0,r1	subscript out of range
jsr r0,r2	case selector out of range
jsr r0,r3	function failed to return a value

The "cedb" utility will automatically determine the source file name, source line number and reason for abort from the core file produced by a run-time abort.

All run-time checking can be turned off using the "-k" compiler toggle.

Unix* is a trademark of Bell Laboratories.

REFERENCES

1. Lampson, B.W., Horning, J.J., London, R.L., Mitchell, J.G. and Popek, G.J., Report on the Programming Language Euclid. SIGPLAN Notices 12,1 (February 1977).
2. Hoare, C.A.R., Monitors: An Operating System Structuring Concept. Comm. ACM 17,10 (October 1974), 549-557.

INDEX

absolute address,
 variable at 6
 actual parameter 14
 address
 standard component 19
 AddressType 7
 aliasing 10, 12, 15
 all, in set constructor 18
 allocate 10
 alternative label 15
 and operator 17
 arithmetic operation,
 precision of 19
 arithmetic operator 17
 array constant 6
 array type 8
 assert statement 15
 assignability 9
 Assign file utility 45
 assignment statement 14
 at clause 6
 base type 8
 begin block 16
 bind 9
 blocked process 24
 Boolean 7
 boolean operator 17
 built-in function 18
 busy statment 25
 calling conventions,
 of PDP-11
 implementation 48
 case alternative label 15
 case statement 15
 Char 7
 character literal 2
 character set 2
 checked,
 module or monitor 4
 procedure or
 function 13
 Chr built-in function 18
 Close file utility 44
 collection 8, 10
 collection declaration 10
 collection element 10, 16
 comment 3
 comparison operator 17.
 compilation 27
 concurrency 21
 concurrent process 21
 condition 23
 condition declaration 23
 consistency criterion 25
 constant declaration 5
 data representation,
 PDP-11
 implementation 47
 Deassign file utility 45
 declaration,
 external 26
 in external module 26
 in external monitor 27
 in module 5, 22
 in monitor 23
 in procedure
 or function 13
 separate 27
 Delete file utility 45
 difference, set operator 17
 div operator 17
 dynamic allocation,
 of collection elements 10
 empty built-in function 24
 empty set 18
 entry, monitor 22
 exit statement 15
 exports clause,
 of module 4
 of monitor 22
 expression 17
 external declaration 25
 external names,
 PDP-11 implementation 49
 FetchArg,
 program argument
 utility 45
 FGetChar 44
 FGetInt 44
 FGetLong 44
 FGetString 44
 field selection, record 16
 file input/output 43
 file, random access 46
 file, sequential 43
 file, temporary 45
 formal parameter 11
 forward type 7
 FPutChar 44
 FPutInt 44
 FPutLong 44
 FPutString 44
 Free, built-in operation 10
 function call 18

- function declaration 11
- GetChar 43
- GetInt 43
- GetLong 43
- GetString 43
- hexadecimal number 2
- identifier 2
- identifier, predefined 2, 41
- if statement 15
- implication, Boolean 17
- imports clause,
 - of module 4
 - of monitor 22
 - of procedure or function 12
- include statement 20
- index type 8
- initialization,
 - of modules and monitors 4
- initially procedure,
 - of modules and monitors 4
- initial value, of variable 6
- input/output 42
- in, set operator 17
- integer literal 2
- intersection, Boolean 17
- intersection, set 17
- IO/1 42
- IO/2 43
- IO/3 45
- IO/4 46
- IO package 42
- keyword 3, 41
- label, case alternative 15
- line numbering,
 - run-time, PDP-11 implementation 48
- linking,
 - of compilations 27, 28
- literal 2
- Long built-in function 18
- LongInt 7
- loop exit 15
- loop statement 15
- main program 3
- manifest 5
- manifest constant 19
- manifest expression 19
- manifest named constant 5
- memory requirement,
 - of process 21
- mod operator 17
- module 4
- module declaration 4
- monitor 22
- monitor declaration 22
- monitor entry 22
- mutual exclusion,
 - in monitor 23
- named constant 5
- named type 7, 8
- New, built-in operation 10
- nil, collection component 10
- non-aliasing,
 - in binds 10
 - in imports 12
 - in reference actual parameters 15
- nonmanifest 5
- nonmanifest named constant 5
- notation, syntactic 3
- not in, set operator 17
- not = operator 17
- not operator 17
- null pointer 10
- octal number 2
- opaque type 4
- Open file utility 44
- operator,
 - arithmetic 17
 - Boolean 17
 - comparison 17
 - set 17
- operator precedence 17
- Ord built-in function 18
- or operator 17
- packed,
 - array 8
 - record 8
- parameter,
 - actual 14
 - formal 11
 - reference 9, 11
 - value 9, 11
- parameter passing,
 - PDP-11 implementation 49
- parameter type 11
- PDP-11 implementation 47
- pervasive,
 - constant 6
 - type 7
- pointer type,
 - of collection 8
- precedence, operator 17
- precision,
 - of arithmetic 19
- predefined identifier 2, 41
- priority condition 24

priority value 24
 procedure body 12
 procedure call 14
 procedure declaration 11
 process, concurrent 21
 program 3
 PutChar 43
 PutInt 43
 PutLong 43
 PutString 43
 random access files 46
 Read, structure input 46
 ReadChar 44
 ReadInt 45
 ReadLong 45
 ReadString 45
 record field 8, 16
 record type 8
 recursive,
 procedure or function 13
 reentrant procedure
 or function 22
 reference parameter 9, 11
 register,
 bind 10
 variable 6
 register usage,
 PDP-11 implementation 47
 reserved word 41
 result type 12
 return identifier 12
 return,
 procedure and function 13
 returns clause 12
 return statement 15
 root type 9
 run-time checking,
 PDP-11 implementation 50
 run-time line numbering,
 PDP-11 implementation 48
 scalar type 7
 Seek file utility 46
 semicolon 3
 separate compilation 27
 separate declaration 27
 separator 3
 set, complete 18
 set constructor 18
 set difference 17
 set element containment 17
 set, empty 18
 set inclusion 17
 set intersection 17
 set operator 17
 set type 8
 set union 17
 ShortInt 7
 side-effects, function 12
 signal statement 23
 SignedInt 7
 simulation 25
 size standard component 19
 source file inclusion 20
 special symbol 3
 standard component 16, 18
 standard type 7
 statement 14
 StorageUnit 7
 string constant 6
 string literal 2
 structured types 8, 16
 structure input/output 46
 subrange type 7
 subscript, array 16
 subscript, collection 16
 synchronization 22
 SysExit,
 program return
 code utility 46
 Tell file utility 46
 temporary files 45
 terminal input/output 42
 type conversion 18
 type converter 13
 type declaration 6
 type definition 7
 type equivalence 9
 union, Boolean 17
 union, set operator 17
 universal 11
 UnsignedInt 7
 value parameter 9, 11
 variable declaration 6
 variable reference 16
 wait statement 23
 Write, structure output 46
 WriteChar 44
 WriteInt 44
 WriteLong 44
 WriteString 44

University of Toronto
Computer Systems Research Group

BIBLIOGRAPHY OF CSRG TECHNICAL REPORTS+

- * CSRG-1 EMPIRICAL COMPARISON OF LR(k) AND PRECEDENCE PARSERS
J.J. Horning and W.R. Lalonde, September 1970
[ACM SIGPLAN Notices, November 1970]
- * CSRG-2 AN EFFICIENT LALR PARSER GENERATOR
W.R. Lalonde, February 1971
[M.A.Sc. Thesis, EE 1971]
- * CSRG-3 A PROCESSOR GENERATOR SYSTEM
J.D. Gorrie, February 1971
[M.A.Sc. Thesis, EE 1971]
- * CSRG-4 DYLAN USER'S MANUAL
P.E. Bonzon, March 1971
- CSRG-5 DIAL - A PROGRAMMING SYSTEM FOR INTERACTIVE ALGEBRAIC MANIPULATION
Alan C.M. Brown and J.J. Horning, March 1971
- *CSRG-6 ON DEADLOCK IN COMPUTER SYSTEMS
Richard C. Holt, April 1971
[Ph.D. Thesis, Dept. of Computer Science,
Cornell University, 1971]
- CSRG-7 THE STAR-RING SYSTEM OF LOOSELY COUPLED DIGITAL DEVICES
John Neill Thomas Pötvin, August 1971
[M.A.Sc. Thesis, EE 1971]
- * CSRG-8 FILE ORGANIZATION AND STRUCTURE
G.M. Stacey, August 1971
- CSRG-9 DESIGN STUDY FOR A TWO-DIMENSIONAL COMPUTER-ASSISTED
ANIMATION SYSTEM
Kenneth B. Evans, January 1972
[M.Sc. Thesis, DCS, 1972]
- * CSRG-10 HOW A PROGRAMMING LANGUAGE IS USED
William Gregg Alexander, February 1972
[M.Sc. Thesis, DCS 1971; Computer, v.8, n.11, November 1975]
- * CSRG-11 PROJECT SUE STATUS REPORT
J.W. Atwood (ed.), April 1972

+ Abbreviations:

DCS - Department of Computer Science, University of Toronto
EE - Department of Electrical Engineering, University of
Toronto

* - Out of print

- * CSRG-12 THREE DIMENSIONAL DATA DISPLAY WITH HIDDEN LINE REMOVAL
Rupert Bramall, April 1972
[M.Sc. Thesis, DCS, 1971]

- * CSRG-13 A SYNTAX DIRECTED ERROR RECOVERY METHOD
Lewis R. James, May 1972
[M.Sc. Thesis, DCS, 1972]

- CSRG-14 THE USE OF SERVICE TIME DISTRIBUTIONS IN SCHEDULING
Kenneth C. Sevcik, May 1972
[Ph.D. Thesis, Committee on Information Sciences,
University of Chicago, 1971; JACM, January 1974]

- CSRG-15 PROCESS STRUCTURING
J.J. Horning and B. Randell, June 1972
[ACM Computing Surveys, March 1972]

- *CSRG-16 OPTIMAL PROCESSOR SCHEDULING WHEN SERVICE TIMES ARE
HYPEREXPONENTIALLY DISTRIBUTED AND PREEMPTION OVERHEAD
IS NOT NEGLIGIBLE
Kenneth C. Sevcik, June 1972
[Proceedings of the Symposium on Computer-Communication,
Networks and Teletraffic, Polytechnic Institute of Brooklyn, 1972]

- * CSRG-17 PROGRAMMING LANGUAGE TRANSLATION TECHNIQUES
W.M. McKeeman, July 1972

- CSRG-18 A COMPARATIVE ANALYSIS OF SEVERAL DISK SCHEDULING ALGORITHMS
C.J.M. Turnbull, September 1972

- CSRG-19 PROJECT SUE AS A LEARNING EXPERIENCE
K.C. Sevcik *et al*, September 1972
[Proceedings AFIPS Fall Joint Computer Conference,
v. 41, December 1972]

- * CSRG-20 A STUDY OF LANGUAGE DIRECTED COMPUTER DESIGN
David B. Wortman, December 1972
[Ph.D. Thesis, Computer Science Department,
Stanford University, 1972]

- CSRG-21 AN APL TERMINAL APPROACH TO COMPUTER MAPPING
R. Kvaternik, December 1972
[M.Sc. Thesis, DCS, 1972]

- * CSRG-22 AN IMPLEMENTATION LANGUAGE FOR MINICOMPUTERS
G.G. Kalmar, January 1973
[M.Sc. Thesis, DCS, 1972]

- CSRG-23 COMPILER STRUCTURE
W.M. McKeeman, January 1973
[Proceedings of the USA-Japan Computer Conference, 1972]

- * CSRG-24 AN ANNOTATED BIBLIOGRAPHY ON COMPUTER PROGRAM ENGINEERING
J.D. Gannon (ed.), March 1973
- CSRG-25 THE INVESTIGATION OF SERVICE TIME DISTRIBUTIONS
Eleanor A. Lester, April 1973
[M.Sc. Thesis, DCS, 1973]
- * CSRG-26 PSYCHOLOGICAL COMPLEXITY OF COMPUTER PROGRAMS:
AN INITIAL EXPERIMENT
Larry Weissman, August 1973
- * CSRG-27 STRUCTURED SUBSETS OF THE PL/I LANGUAGE
Richard C. Holt and David B. Wortman, October 1973
- * CSRG-28 ON REDUCED MATRIX REPRESENTATION OF LR(k)
PARSER TABLES
Marc Louis Joliat, October 1973
[Ph.D. Thesis, EE 1973]
- * CSRG-29 A STUDENT PROJECT FOR AN OPERATING SYSTEMS COURSE
B. Czarnik and D. Tsichritzis (eds.), November 1973
- * CSRG-30 A PSEUDO-MACHINE FOR CODE GENERATION
Henry John Pasko, December 1973
[M.Sc. Thesis, DCS 1973]
- * CSRG-31 AN ANNOTATED BIBLIOGRAPHY ON COMPUTER PROGRAM ENGINEERING
J.D. Gannon (ed.), Second Edition, March 1974
- * CSRG-32 SCHEDULING MULTIPLE RESOURCE COMPUTER SYSTEMS
E.D. Lazowska, May 1974
[M.Sc. Thesis, DCS, 1974]
- * CSRG-33 AN EDUCATIONAL DATA BASE MANAGEMENT SYSTEM
F. Lochovsky and D. Tsichritzis, May 1974
[INFOR, 14 (3), pp.270-273, 1974]
- * CSRG-34 ALLOCATING STORAGE IN HIERARCHICAL DATA BASES
P. Bernstein and D. Tsichritzis, May 1974
[Information Systems Journal, v.1, pp.133-140]
- * CSRG-35 ON IMPLEMENTATION OF RELATIONS
D. Tsichritzis, May 1974
- * CSRG-36 SIX PL/I COMPILERS
D.B. Wortman, P.J. Khaiat, and D.M. Lasker, August 1974
[Software Practice and Experience, v.6, n.3,
July-Sept. 1976]
- * CSRG-37 A METHODOLOGY FOR STUDYING THE PSYCHOLOGICAL COMPLEXITY
OF COMPUTER PROGRAMS
Laurence M. Weissman, August 1974
[Ph.D. Thesis, DCS, 1974]

- * CSRG-38 AN INVESTIGATION OF A NEW METHOD OF CONSTRUCTING SOFTWARE
David M. Lasker, September 1974
[M.Sc. Thesis, DCS, 1974]

- CSRG-39 AN ALGEBRAIC MODEL FOR STRING PATTERNS
Glenn F. Stewart, September 1974
[M.Sc. Thesis, DCS, 1974]

- * CSRG-40 EDUCATIONAL DATA BASE SYSTEM USER'S MANUAL
J. Klebanoff, F. Lochovsky, A. Rozitis, and
D. Tsichritzis, September 1974

- * CSRG-41 NOTES FROM A WORKSHOP ON THE ATTAINMENT OF
RELIABLE SOFTWARE
David B. Wortman (ed.), September 1974

- * CSRG-42 THE PROJECT SUE SYSTEM LANGUAGE REFERENCE MANUAL
B.L. Clark and F.J.B. Ham, September 1974

- * CSRG-43 A DATA BASE PROCESSOR
E.A. Ozkarahan, S.A. Schuster and K.C. Smith,
November 1974 [Proceedings National Computer
Conference 1975, v.44, pp.379-388]

- * CSRG-44 MATCHING PROGRAM AND DATA REPRESENTATION TO A
COMPUTING ENVIRONMENT
Eric C.R. Hehner, November 1974
[Ph.D. Thesis, DCS, 1974]
See Computer, Vol.9, No.9, August 1976, pp.65-70.

- * CSRG-45 THREE APPROACHES TO RELIABLE SOFTWARE: LANGUAGE DESIGN,
DYADIC SPECIFICATIONS, COMPLEMENTARY SEMANTICS
J.E. Donahue, J.D. Gannon, J.V. Guttag and
J.J. Horning, December 1974

- CSRG-46 THE SYNTHESIS OF OPTIMAL DECISION TREES FROM
DECISION TABLES
Helmut Schumacher, December 1974
[M.Sc. Thesis, DCS, 1974; CACM, v.19, n.6, June 1976]

- * CSRG-47 LANGUAGE DESIGN TO ENHANCE PROGRAMMING RELIABILITY
John D. Gannon, January 1975
[Ph.D. Thesis, DCS, 1975]

- CSRG-48 DETERMINISTIC LEFT TO RIGHT PARSING
Christopher J.M. Turnbull, January 1975
[Ph.D. Thesis, EE, 1974]

- * CSRG-49 A NETWORK FRAMEWORK FOR RELATIONAL IMPLEMENTATION
D. Tsichritzis, February 1975 [in Data Base Description,
Dongue and Nijssen (eds.), North Holland Publishing Co.]

- * CSRG-50 A UNIFIED APPROACH TO FUNCTIONAL DEPENDENCIES AND RELATIONS
P.A. Bernstein, J.R. Swenson and D.C. Tsichritzis
February 1975 [Proceedings of the ACM SIGMOD Conference, 1975]
- * CSRG-51 ZETA: A PROTOTYPE RELATIONAL DATA BASE MANAGEMENT SYSTEM
M. Brodie (ed). February 1975 [Proceedings Pacific ACM Conference, 1975]
- CSRG-52 AUTOMATIC GENERATION OF SYNTAX-REPAIRING AND PARAGRAPHING PARSERS
David T. Barnard, March 1975
[M.Sc. Thesis, DCS, 1975]
- * CSRG-53 QUERY EXECUTION AND INDEX SELECTION FOR RELATIONAL DATA BASES
J.H. Gilles Farley and Stewart A. Schuster, March 1975
- CSRG-54 AN ANNOTATED BIBLIOGRAPHY ON COMPUTER PROGRAM ENGINEERING
J.V. Guttag (ed.), Third Edition, April 1975
- CSRG-55 STRUCTURED SUBSETS OF THE PL/1 LANGUAGE
Richard C. Holt and David B. Wortman, May 1975
- * CSRG-56 FEATURES OF A CONCEPTUAL SCHEMA
D. Tsichritzis, June 1975 [Proceedings Very Large Data Base Conference, 1975]
- * CSRG-57 MERLIN: TOWARDS AN IDEAL PROGRAMMING LANGUAGE
Eric C.R. Hehner, July 1975
see Acta Informatica Col.10, No.3, pp.229-243, 1978
- CSRG-58 ON THE SEMANTICS OF THE RELATIONAL DATA MODEL
Hans Albrecht Schmid and J. Richard Swenson,
July 1975 [Proceedings of the ACM SIGMOD Conference, 1975]
- * CSRG-59 THE SPECIFICATION AND APPLICATION TO PROGRAMMING OF ABSTRACT DATA TYPES
John V. Guttag, September 1975
[Ph.D. Thesis, DCS, 1975]
- * CSRG-60 NORMALIZATION AND FUNCTIONAL DEPENDENCIES IN THE RELATIONAL DATA BASE MODEL
Phillip Alan Bernstein, October 1975
[Ph.D. Thesis, DCS, 1975]
- * CSRG-61 LSL: A LINK AND SELECTION LANGUAGE
D. Tsichritzis, November 1975 [Proceedings ACM SIGMOD Conference, 1975]

- * CSRG-62 COMPLEMENTARY DEFINITIONS OF PROGRAMMING LANGUAGE SEMANTICS
James E. Donahue, November 1975
[Ph.D. Thesis, DCS, 1975]

- CSRG-63 AN EXPERIMENTAL EVALUATION OF CHESS PLAYING HEURISTICS
Lazlo Sugar, December 1975
[M.Sc. Thesis, DCS, 1975]

- CSRG-64 A VIRTUAL MEMORY SYSTEM FOR A RELATIONAL ASSOCIATIVE PROCESSOR
S.A. Schuster, E.A. Ozkarahan, and K.C. Smith,
February 1976 [Proceedings National Computer Conference 1976, v.45, pp.855-862]

- CSRG-65 PERFORMANCE EVALUATION OF A RELATIONAL ASSOCIATIVE PROCESSOR
E.A. Ozkarahan, S.A. Schuster, and K.C. Sevcik,
February 1976 [ACM Transactions on Database Systems, v.1, n.4, December 1976]

- CSRG-66 EDITING COMPUTER ANIMATED FILM
Michael D. Tilson, February 1976
[M.Sc. Thesis, DCS, 1975]

- CSRG-67 A DIAGRAMMATIC APPROACH TO PROGRAMMING LANGUAGE SEMANTICS
James R. Cordy, March 1976
[M.Sc. Thesis, DCS, 1976]

- * CSRG-68 A SYNTHETIC ENGLISH QUERY LANGUAGE FOR A RELATIONAL ASSOCIATIVE PROCESSOR
L. Kerschberg, E.A. Ozkarahan, and J.E.S. Pacheco,
April 1976

- CSRG-69 AN ANNOTATED BIBLIOGRAPHY ON COMPUTER PROGRAM ENGINEERING
D. Barnard and D. Thompson (eds.), Fourth Edition,
May 1976

- * CSRG-70 A TAXONOMY OF DATA MODELS
L. Kerschberg, A. Klug, and D. Tschritzis, May 1976
[Proceedings Very Large Data Base Conference, 1976]

- * CSRG-71 OPTIMIZATION FEATURES FOR THE ARCHITECTURE OF A DATA BASE MACHINE
E.A. Ozkarahan and K.C. Sevcik, May 1976
[ACM Transactions of Database Systems, v.2, n.4, December 1977]

- CSRG-72 THE RELATIONAL DATA BASE SYSTEM OMEGA - PROGRESS REPORT
H.A. Schmid (ed.), P.A. Bernstein (ed.), B. Arlow,
R. Baker and S. Pozgaj, July 1976

- * CSRG-73 AN ALGORITHMIC APPROACH TO NORMALIZATION OF
RELATIONAL DATA BASE SCHEMAS
P.A. Bernstein and C. Beeri, September 1976

- CSRG-74 A HIGH-LEVEL MACHINE-ORIENTED ASSEMBLER LANGUAGE
FOR A DATA BASE MACHINE
E.A. Ozkarahan and S.A. Schuster, October 1976

- * CSRG-75 DO CONSIDERED OD: A CONTRIBUTION TO THE PROGRAMMING
CALCULUS
Eric C.R. Hehner, November 1978
Acta Informatica to appear 1979

- CSRG-76 SOFTWARE HUT: A COMPUTER PROGRAM ENGINEERING
PROJECT IN THE FORM OF A GAME
J.J. Horning and D.B. Wortman, November 1976
[IEEE Transactions on Software Engineering, v.SE-3, n.4, July 1977]

- CSRG-77 A SHORT STUDY OF PROGRAM AND MEMORY POLICY BEHAVIOUR
G. Scott Graham, January 1977

- * CSRG-78 A PANACHE OF DBMS IDEAS
D. Tsichritzis (ed.), February 1977

- CSRG-79 THE DESIGN AND IMPLEMENTATION OF AN ADVANCED LALR
PARSE TABLE CONSTRUCTOR
David H. Thompson, April 1977
[M.Sc. Thesis, DCS, 1976]

- CSRG-80 AN ANNOTATED BIBLIOGRAPHY ON COMPUTER PROGRAM
ENGINEERING
D. Barnard (ed.), Fifth Edition, May 1977

- * CSRG-81 PROGRAMMING METHODOLOGY: AN ANNOTATED BIBLIOGRAPHY
FOR IFIP WORKING GROUP 2.3
Sol J. Greenspan and J.J. Horning (eds.), First Edition, May 1977

- CSRG-82 NOTES ON EUCLID
edited by W. David Elliott and David T. Barnard, August 1977

- CSRG-83 TOPICS IN QUEUEING NETWORK MODELING
edited by G. Scott Graham, July 1977

- CSRG-84 TOWARD PROGRAM ILLUSTRATION
Edward Yarwood, September 1977
[M.Sc. Thesis, DCS, 1974]

- CSRG-85 CHARACTERIZING SERVICE TIME AND RESPONSE TIME
DISTRIBUTIONS IN QUEUEING NETWORK MODELS OF COMPUTER
SYSTEMS
Edward D. Lazowska, September 1977
[Ph.D. Thesis, DCS, 1977]

- CSRG-86 MEASUREMENTS OF COMPUTER SYSTEMS FOR QUEUEING
NETWORK MODELS
Martin G. Kienzie, October 1977
[M.Sc. Thesis, DCS, 1977; Proc. Int. Symp. on Modelling and Performance
Evaluation of Computer Systems, Vienna, 1979]
- CSRG-87 'OLGA' LANGUAGE REFERENCE MANUAL
B. Abourbih, H. Trickey, D.M. Lewis, E.S. Lee,
P.I.P. Boulton, November 1977
- * CSRG-88 USING A GRAMMATICAL FORMALISM AS A PROGRAMMING LANGUAGE
Brad A. Silverberg, January 1978
[M.Sc. Thesis, DCS, 1978]
- CSRG-89 ON THE IMPLEMENTATION OF RELATIONS: A KEY TO EFFICIENCY
Joachim W. Schmidt, January 1978
- CSRG-90 DATA BASE MANAGEMENT SYSTEM USER PERFORMANCE
Frederick H. Lochovsky, April 1978
[Ph.D. Thesis, DCS, 1978]
- CSRG-91 SPECIFICATION AND VERIFICATION OF DATA BASE
SEMANTIC INTEGRITY
Michael Lawrence Brodie, April 1978
[Ph.D. Thesis, DCS, 1978]
- CSRG-92 STRUCTURED SOUND SYNTHESIS PROJECT (SSSP):
AN INTRODUCTION
by William Buxton, Guy Fedorkow, with Ronald Baecker,
Gustav Ciarnaga, Leslie Mezei and K.C. Smith, June 1978
- * CSRG-93 A DEVICE-INDEPENDENT, GENERAL-PURPOSE GRAPHICS SYSTEM
IN A MINICOMPUTER TIME-SHARING ENVIRONMENT
William T. Reeves, August 1978
[M.Sc. Thesis, DCS, 1978]
- * CSRG-94 ON THE AXIOMATIC VERIFICATION OF
CONCURRENT ALGORITHMS
Christian Lengauer, August 1978
[M.Sc. Thesis, DCS, 1978]
- CSRG-95 PISA: A PROGRAMMING SYSTEM FOR INTERACTIVE
PRODUCTION OF APPLICATION SOFTWARE
Rudolf Marty, August 1978
- CSRG-96 ADAPTIVE MICROPROGRAMMING AND PROCESSOR MODELING
Walter G. Rosocha
[Ph.D. Thesis, EE, August 1978]
- * CSRG-97 DESIGN ISSUES IN THE FOUNDATION OF A COMPUTER-BASED
TOOL FOR MUSIC COMPOSITION
William Buxton
[M.Sc. Thesis, CSRG, October 1978]

- CSRG-98 THEORY OF DATABASE MAPPINGS
Anthony C. Klug
[Ph.D. Thesis, DCS, December 1978]
- CSRG-99 HIERARCHICAL COROUTINES: A MECHANISM FOR IMPROVED
PROGRAM STRUCTURE
Leonard I. Vanek, February 1979
- CSRG-100 TOPICS IN PERFORMANCE EVALUATION
G. Scott Graham (ed.), July 1979
- * CSRG-101 A PANACHE OF DBMS IDEAS II
F.H. Lochovsky (ed.), May 1979
- CSRG-102 A SIMPLE SET THEORY FOR COMPUTING SCIENCE
Eric C.R. Hehner, May 1979
- CSRG-103 THE CENTRALIZED ALGORITHM IN DISTRIBUTED SYSTEMS
Ernest J.H. Chang
[Ph.D. Thesis, DCS, July 1979]
- CSRG-104 ELIMINATING THE VARIABLE FROM DIJKSTRA'S
MINI-LANGUAGE
D. Hugh Redelmeier, July 1979
- CSRG-105 A LANGUAGE FACILITY FOR DESIGNING INTERACTIVE
DATABASE-INTENSIVE APPLICATIONS
John Mylopoulos, Philip A. Bernstein, Harry K.T. Wong,
July 1979
- CSRG-106 ON APPROXIMATE SOLUTION TECHNIQUES FOR
QUEUEING NETWORK MODELS OF COMPUTER SYSTEMS
Satish Kumar Tripathi, July 1979
- CSRG-107 A FRAMEWORK FOR VISUAL MOTION UNDERSTANDING
John K. Tsotsos, John Mylopoulos, H. Dominic Cowvey
Steven W. Zucker, DCS, June 1979
- * CSRG-108 DIALOGUE ORGANIZATION AND STRUCTURE FOR
INTERACTIVE INFORMATION SYSTEMS
John Leonard Barron
[M.Sc. Thesis, DCS, 1980]
- * CSRG-109 A UNIFYING MODEL OF PHYSICAL DATABASES
D.S. Batory, C.C. Gotlieb, April 1980
- * CSRG-110 OPTIMAL FILE DESIGNS AND REORGANIZATION POINTS
D.S. Batory, April 1980
- * CSRG-111 A PANACHE OF DBMS IDEAS III
D. Tsichritzis (ed.), April 1980

- CSRG-112 TOPICS IN PSN - II: EXCEPTIONAL CONDITION
HANDLING IN PSN; REPRESENTING PROGRAMS IN PSN;
CONTENTS IN PSN
Yves Lesperance, Byran M. Kramer, Peter F. Schneider
April, 1980
- CSRG-113 SYSTEM-ORIENTED MACRO-SCHEDULING
C.C. Gotlieb and A. Schonbach
May 1980
- CSRG-114 A FRAMEWORK FOR VISUAL MOTION UNDERSTANDING
John Konstantine Tsotsos
[Ph.D. Thesis, DCS, June 1980]
- CSRG-115 SPECIFICATION OF CONCURRENT EUCLID
James R. Cordy and Richard C. Holt
July 1980
- CSRG-116 THE REPRESENTATION OF PROGRAMS IN THE
PROCEDURAL SEMANTIC NETWORK FORMALISM
Bryan M. Kramer
[M.Sc. Thesis, DCS, 1980]
- CSRG-117 CONTEXT-FREE GRAMMARS AND DERIVATION TREES AS
PROGRAMMING TOOLS
Volker Linnemann
September 1980
- CSRG-118 S/SL: SYNTAX/SEMANTIC LANGUAGE
INTRODUCTION AND SPECIFICATION
R.C. Holt, J.R. Cordy, D.B. Wortman
CSRG, September 1980
- CSRG-119 PT: A PASCAL SUBSET
Alan Rosselet
[M.Sc. Thesis, DCS, October 1980]
- CSRG-120 PTED: A STANDARD PASCAL TEXT EDITOR BASED ON
THE KERNIGHAN AND PLAUGER DESIGN
Ken Newman, DCS
October 1980
- CSRG-121 TERMINAL CONTEXT GRAMMARS
Howard W. Trickey
[M.Sc. Thesis, EE, September 1980]
- CSRG-122 THE APPROXIMATE SOLUTION OF LARGE QUEUEING
NETWORK MODELS
John Zaherjan
[Ph.D. Thesis, DCS, August 1980]

CSRG-123 A FORMAL TREATMENT OF IMPERFECT INFORMATION
IN DATABASE MANAGEMENT

Yannis Vassiliou
[Ph.D. Thesis, DCS, September 1980]

CSRG-124 AN ANALYTIC MODEL OF PHYSICAL DATABASES

Don S. Batory
[Ph.D. Thesis, DCS, January 1981]

CSRG-125 MACHINE-INDEPENDENT CODE GENERATION

Richard H. Kozlak
[M.Sc. Thesis, DCS, January 1981]

CSRG-126 COMPUTER MACRO-SCHEDULING FOR HIGH PRODUCTIVITY

Abraham Schonbach
[Ph.D. Thesis, DCS, March 1981]

CSRG-127 OMEGA ALPHA

D. Tschritzis (ed.), March 1981

CSRG-128 DIALOGUE AND PROCESS DESIGN FOR INTERACTIVE
INFORMATION SYSTEMS USING TAXIS

John Barron, April 1981

CSRG-129 DESIGN AND VERIFICATION OF INTERACTIVE INFORMATION
SYSTEMS USING TAXIS

Harry K.T. Wong
[Ph.D. Thesis, DCS, to be submitted]

CSRG-130 DYNAMIC PROTECTION OF OBJECTS IN A COMPUTER UTILITY

Leslie H. Goldsmith, April, 1981

CSRG-131 INTEGRITY ANALYSIS: A METHODOLOGY FOR EDP AUDIT
AND DATA QUALITY CONTROL

Maija Irene Svanks
[Ph.D. Thesis, DCS, February 1981]

CSRG-132 A PROTOTYPE KNOWLEDGE-BASED SYSTEM
FOR COMPUTER-ASSISTED MEDICAL DIAGNOSIS

Stephen A. Ho-Tai
[M.Sc.Thesis, DCS, January 1981]

CSRG-133 SPECIFICATION OF CONCURRENT EUCLID

James R. Cordy, Richard C. Holt
August 1981 (Version 1)