

UC Irvine

ICS Technical Reports

Title

RGA users manual

Permalink

<https://escholarship.org/uc/item/56w3v2xm>

Author

Morgan, E. Timothy

Publication Date

1984-12-04

Peer reviewed

RGA Users Manual

by

E. Timothy Morgan

ABSTRACT

RGA is an interpreter for a special language designed for the analysis of reachability graphs, or control flow graphs, generated from Petri nets. Although in some cases the reachability graph can become too large to be tractable, or can even be infinite, many interesting problems exist whose reachability graphs of reasonable size. In RGA, the user has access to the names of the places in the net, and to the states of the reachability graph. The structure of the graph is also available through functions which return the sets of successor or predecessor states of a state, and the transition-firings connecting the states. The RGA language allows dynamic typing of identifiers, recursion, and function and operator overloading. Rather than providing a number of predefined analysis functions, RGA provides primitive functions which allow the user to conduct complex analyses with little programming effort. RGA is part of a suite of tools intended to make the analysis of concurrent systems described by Petri nets easier.

Technical Report #243

**Department of Information and Computer Science
University of California, Irvine
Irvine, CA 92717**

December 4, 1984

© Copyright - 1984

Contents

	Page
Introduction	1
1. Execution Environment	2
2. Lexical Issues	3
3. Expression Types and Execution Semantics.	4
3.1. Arithmetic Expressions.	6
3.2. Boolean Expressions	7
3.3. State, Place, and Transition Expressions	9
3.4. Set and Sequence Expressions.	9
4. User-Defined Functions	12
5. Some Examples	13
6. Running RGA	16
7. Implementation and Performance	17
8. Conclusions and Future Work	18
Appendix A: BNF Grammer for RGA Language	19
Appendix B: Some Performance Measurements	22
References	23

RGA Users Manual

Introduction

RGA is an interpreter for a special language designed for the analysis of reachability graphs, or control flow graphs, generated from Petri nets [PETE77]. Although in some cases the reachability graph can become too large to be tractable, or can even be infinite, many interesting problems exist whose reachability graphs of reasonable size. In RGA, the user has access to the names of the places in the net, and to the states of the reachability graph. The structure of the graph is also available through functions which return the sets of successor or predecessor states of a state, and the transition-firings connecting the states. The RGA language allows dynamic typing of identifiers, recursion, and function and operator overloading. Rather than providing a number of predefined analysis functions, RGA provides primitive functions which allow the user to conduct complex analyses with little programming effort. RGA is part of a suite of tools intended to make the analysis of concurrent systems described by Petri nets easier.

In RGA, the user merely types an expression, and the interpreter evaluates it and prints the resulting value. For example, using the function **nsucc** which returns the number of successor states of a state, and the set of all states **S**, the user can write

```
forall s in S [nsucc(s) > 0]
```

This expression will return **true** if for each state in the set **S**, the number of successors is greater than zero. Thus this expression is a test for deadlock-freeness of the Petri net [AGER79].

Another test might be to determine if the net is conservative, that is, that tokens are never gained or lost [AGER79]. The function **tokens(s)** returns the sum

of the tokens on all places in a state s . The first state in the graph is written #0, so the expression for net conservation might be

$$\text{forall } s \text{ in } S [\text{tokens}(\#0) = \text{tokens}(s)]$$

The following sections describe the properties of the interpreter for the language, the data types and expressions which exist in the language, and how the user may define functions using the primitive functions provided by the interpreter. Then some examples are given to show how the system may be used to answer more complex questions than those shown above. Finally, some implementation issues are discussed, and some conclusions are drawn.

1. Execution Environment

RGA is an interpreter, and thus its operation is similar to that of most LISP interpreters. Any expression which the user types is immediately evaluated, and that value is printed on the standard output. The expression is then thrown away, and the user is prompted again for another command. In addition to typing expressions, the user may define expressions to be evaluated later as functions. Expressions and function definitions may be read from a file as well as from the standard input.

Unlike LISP, RGA has a number of distinct data types which it uses. But there is no explicit way to declare variables. In fact, all variables in the RGA language are dynamically typed when they are assigned values: an identifier or expression always represents a $\langle \text{value}, \text{type} \rangle$ pair. The user never explicitly deals with the *type* component, however. During execution, an identifier may have more than one value (and therefore, type) associated with it simultaneously. These values are stored on an execution stack, and only the most recently bound value may be accessed at any time.

Because identifiers need not be declared before their use, it is very easy to define functions. However, it also means that much of the type checking which needs to be performed must be delayed until execution time, since the types and values of identifiers used in a function definition will not be known at the time the function is defined.

Three types of errors are possible using RGA. The first error type is a syntax error in an expression or command. This type of error results in the message "Command ignored." The second type of error is a run-time error, such as a type conflict or a division by zero. A run-time error usually results in an appropriate message's being printed, followed by a prompt. The execution stack is *not* "cleaned up" so that variables will have the values they had at the time of the error; this facilitates debugging of defined functions. If user-defined functions were being executed at the time of the error, a stack-back trace of function calls is printed. The final type of error is an internal error in the RGA interpreter, which should not happen under normal circumstances.

2. Lexical Issues

RGA is case sensitive. All command keywords and predefined function names are written in lower case. All identifiers which the user defines may be written in lower, upper, or mixed case. The user may not redefine a reserved language keyword, but predefined identifiers may be redefined, although that is not recommended. In addition to the three predefined identifiers **S**, **P**, and **T** (defined later), the reachability graph which is loaded during initialization will typically define a number of identifiers to be places; these identifiers must follow the requirements for identifiers described below.

An identifier is represented as an upper or lower case alphabetic letter, followed by zero or more letters, digits, single-quote characters, periods, and underscores.

A number is represented as an optional minus sign followed by one or more digits; only integer values may be represented.

A command to RGA normally terminates with a newline character. Receiving this character will cause RGA to perform the appropriate function. For very long expressions or function definitions, a line may be terminated with a backslash (\) followed by a newline. This combination of two characters is treated as a single space character, so its only affect is to delimit other tokens. Multiple space and tab characters and comments are treated as a single space. Comments may be inserted using the conventions of the C and PL/I languages: */* comment text */*.

3. Expression Types and Execution Semantics

This section describes the syntax and semantics of the expressions available in RGA, and it describes the built-in primitive functions which are available. It is divided into subsections which describe each of the different data types which the language supports. Section 4 describes the commands which may be used to define new functions. Although the syntax of the language is described in section 3, a formal BNF description is given in Appendix A. All expressions in the RGA language evaluate to a value whose type is either a state, an integer, a boolean, a transition, a transition-firing (TF), a set, or a sequence. Identifiers can be assigned values of any of these types, and they will automatically take on the appropriate type. Normally, evaluating a place is interpreted to mean the integer number of tokens on that place within the context of a particular state. Only in enumerating the elements of a set of places may a place value be assigned directly to an identifier, and that value is popped from the execution stack when the enumeration finishes unless there is a runtime error.

Calls to functions, both those which are predefined, and to those defined by the user, have the same syntax:

id (list-of-expressions)

where a *list-of-expressions* is a single expression, or multiple expressions separated with commas. If the function takes no arguments then the parentheses are omitted. When invoking a function, the expressions are evaluated from right to left in the current context, and then all the values are bound to the formal parameters, from left to right. Thus all parameters are passed by value, so they cannot be changed in any way by the called function unless it accesses them globally. Variable access is dynamically scoped.

When evaluating a place identifier, as mentioned above, its value is the integer number of tokens on that place in a certain state of the control flow graph. The state may be specified explicitly by the user, or it will default to the "current state." To specify the state to use explicitly, the place identifier should be followed by a state-valued expression in parentheses. The current state is a global value implicitly set by the **forall** and **exists** boolean expressions, and in the subset construction, which are described below.

There are four special expression operators whose type depends on their arguments, and which operate on expressions of all types. (1) The **print** function takes an arbitrary expression in parentheses, and it returns the value of that argument. It has a side-effect of printing the value returned on the standard output. (2) The infix assignment function ":@" assigns the identifier on its left the *<value, type>* pair which results from evaluating the expression on its right. Like **print**, the assignment operation also returns the value which has been evaluated. (3) The semicolon infix operator ";" takes two expressions of arbitrary type. It evaluates the left expression and discards its value, if any, and then it returns the value of the right hand expression. The semicolon operator is left-recursive in its evaluation.

(4) Finally, the *if* expression allows for conditional execution of expressions. There are two forms of the *if* command:

if boolean-expression then expression fi

if boolean-expression then expression else expression fi

The type and value returned by the *if* expression depends on what expression, if any, is executed. It is unique, however, in that it may not return a value at all if the *boolean-expression* in the first form evaluates to **false**. The only time that this form of the *if* expression can be used is as the left argument to a semicolon operator, which would discard any value returned. The **else-less** form is not allowed in any other situation, when a value is required.

3.1. Arithmetic Expressions

Arithmetic expressions follow the conventions of most modern programming languages. An arithmetic value may be an integer constant, an identifier whose value is an integer, a place (which is evaluated as the number of tokens on that place, as explained above), or an integer-valued function. A number of arithmetic functions are written in conventional infix notation: addition (+), subtraction (-), multiplication (*), division (/), modulo (%), and exponentiation (^). Unary negation is recognized. Parentheses can be used to control the evaluation of an expression; conventional precedence and left-to-right evaluation order otherwise hold.

The following are the predefined integer-valued primitive functions. The argument types (states and sets) are described in later sections.

- | | |
|-----------------------------|---|
| tokens(<i>state</i>) | The total number of tokens on all places in a specified state. The state is given as an argument to the function, as tokens(#0) . An alternate way of writing this function is to put the state within vertical bars, as an absolute value. For example, #1 . |
| marked(<i>state</i>) | Returns the number of places in the argument state which have at least one token on them. If marked(<i>s</i>)=tokens(<i>s</i>) then the state <i>s</i> is a safe state [AGER79]. |

- nsucc**(state) Returns the number of successor states of the argument state. If there are two or more transition-firings which lead to the same state, then **nsucc** will actually return the cardinality of the **tfout** set rather than that of the **succ** set.
- npred**(state) Returns the cardinality of the **tfin** set, as **nsucc** returns the cardinality of the **tfout** set.
- card**(s) Returns the number of elements of a set or sequence *s*, which is the single argument.

3.2. Boolean Expressions

As with other expressions, boolean expressions are built up from constants, infix operators, and predefined and user-defined function calls. The boolean constants are the reserved words **true** and **false**.

The infix boolean operators are the conventional arithmetic comparison tests: **<**, **<=**, **>**, **>=**, **=**, and **!=**; not equal may also be written as **<>**. The equal and not equal tests may be applied to any data types (places, states, sets, sequences, and booleans, as well as integer-valued expressions), while the other operators are restricted to integer expressions. Of course, the **=** operation applied to boolean expressions is a logical equivalence test. Other infix boolean operators which apply to boolean expressions are **implies**, **iff**, **and**, and **or**. For convenience, the **and** and **or** operators may also be written as **&** and **|**. Both the **and** and **or** operators are "short-circuit" operators which evaluate the lefthand operand, and then only evaluate the righthand operand if necessary. Their precedence, from lowest to highest, is the logical relational operators **implies** and **iff**, which have neither left nor right associativity, **or**, **and**, which are left associative, and the arithmetic relational operators, which also have no associativity. The prefix unary operator **not** may be used to negate a logical expression. It has the precedence of arithmetic relational operators, so it is higher than the other logical operators.

As a special case, places may be used as boolean values if they contain at most one token. This is for convenience when working with safe nets. When a boolean

expression is expected, and a place name is found instead, then the number of tokens on that place is evaluated, and **false** is returned if it is zero, and **true** if it is one.

Two of the language's most important operators are **forall** and **exists**, the universal and existential quantifiers. Their syntax is the same, so only that of **forall** will be given:

forall *id* in *set-expression* [*boolean-expression*]

This expression is evaluated as follows. First, the current value of *id* is pushed on the execution stack, to be popped off when the **forall** expression is finished being evaluating. The global current state is also pushed and popped at the same time if the set is a set of states. Next, the *set-expression* is evaluated once and only once. The *id* is then looped through the elements of the set one at a time. If the *id* is a state, then the current state is set to be that state also. For each value of the *id*, the *boolean-expression* is evaluated. If for all values, the expression evaluates to **true**, then the whole expression returns that value. But if the expression ever evaluates to **false**, then execution of the loop is halted immediately and the **forall** expression returns **false**.

The **exists** expression is similar to **forall**, but with the logical tests reversed. It continues to evaluate the boolean expression until it exhausts all the elements of the set or until the expression evaluates to **true**. If the set is exhausted, then **exists** returns **false**, and otherwise, **true**.

There is only one primitive function which returns a boolean value:

in (<i>item</i> , <i>s</i>)	The in function takes two arguments, an <i>item</i> of any type, and a set or sequence of items <i>s</i> . It returns true if the item is an element of the indicated set or sequence, and false otherwise.
--------------------------------------	--

3.3. State, Place, and Transition Expressions

State constants may be written as a pound sign (#) followed by an integer. The first state in the reachability graph is #0. Places can only be referred to through the identifiers defined in the original Petri net from which the reachability graph is derived, and through loop control identifiers in the **forall** and **exists** expressions, and the *subset* construct described in the following subsection.

Transition constants are written as a dollar sign (\$) followed by an integer, with the first transition written as \$0. Transition-firing constants (TFs) are written as a triple of the source and destination states of the firing, and the transition which is fired. The components of the triple are written separated by commas, between square brackets. For example, [#0, #10, \$10] is a firing of transition \$10 taking the net from state #0 to state #10. It is a syntax error to write a transition constant which does not exist in the reachability graph.

Three primitive functions exist which return state or transition values. All three take a transition-firing as their single argument, and return the separate components of the TF. A fourth function is used to display the marked places in a state. As a side effect, it returns the state which is its argument:

src (<i>tf</i>)	Returns the source state of <i>tf</i> .
dest (<i>tf</i>)	Returns the destination state of <i>tf</i> .
trans (<i>tf</i>)	Returns the transition involved in the firing <i>tf</i> .
showstate (<i>state</i>)	Returns the state argument, and prints its marked places as a side effect. If more than one token is on a place, the token count is shown in parentheses.

3.4. Set and Sequence Expressions

The set operations are probably the single most powerful feature of the language. Sets and sequences are composed of elements which must be of the same type. Any legal type is acceptable, including other sets and sequences; all the elements of a set or sequence must be of the same type. Although sets should be

considered to be unordered, they are always maintained in ascending numerical order for convenience in reading and comparing. Sets do not contain duplicate elements, while sequences are ordered and may contain duplicates. A single set is either a set variable, a set constant, or a set function. Three predefined set variables exist: **S**, **P**, and **T**. The **S** set is the set of all states in the reachability graph, and **P** is similarly the set of all places in the original Petri net. **T** is the set of all transitions in the Petri net potentially-firable from the initial state.

A set constant is written as a list of expressions within curly braces `{}`. The list is written with the elements separated with commas. For convenience, a constant range of states may also be entered by giving the first state, `".."`, and the final state of the range. For example, the set consisting of states 1, 5 through 10, and 12 could be written

$$\{\#1, \#5..\#10, \#12\}$$

The list of elements may be empty, resulting in the empty set. As a special case, an empty set may be used in the context of a set of any type without a type-conflict error.

Another powerful way of specifying a set constant is the *subset* construct. It allows elements to be selected from a set using any boolean expression as the selection criterion. This construct is similar to a **forall** command, but it always loops through the entire set evaluating the *boolean-expression* for each element. Like the **forall** and **exists** statements, the *id*'s value is pushed at the beginning of the loop and restored when the subset has been constructed. If the *set-expression* is a set of states, then the "current state" will also be pushed and popped, and set to each value along with the *id*. The *set-expression* is evaluated only once. The subset construct is written

$$\{id \text{ in } set-expression \mid boolean-expression\}$$

Sequence constants are written just like set constants, except that their elements are ordered, and two slashes (//) are used to indicate the beginning and ending of the sequence, instead of open and close braces.

It is recommended that set and sequence constants be used only in contexts where they will not be repeatedly evaluated, as within a loop or a recursive function, because they are relatively expensive to compute. If a constant is to be used in these situations, it should be evaluated once and the value assigned to some identifier. Then that identifier may be used in place of the constant.

There are several predefined functions which return sets as their values:

<code>tfin(state)</code>	Returns the set of transition-firings whose destination state is the indicated <i>state</i> .
<code>tfout(state)</code>	Returns the set of TFs whose source state is the indicated <i>state</i> .
<code>succ(state)</code>	The succ function takes a state expression as its argument. It returns the (possibly empty) set of immediate successor states in the reachability graph of the specified state.
<code>pred(state)</code>	The pred function is similar to the succ function, but it returns the set of immediate predecessor states instead of the successor set.
<code>allsucc(state)</code>	Returns the set of all the successors of the indicated state, and recursively, all their successors.
<code>allpred(state)</code>	Returns the set of all the predecessors of the indicated state, and recursively, all their predecessors.
<code>union(s1, s2)</code>	The union function takes two sets or sequences <i>s1</i> and <i>s2</i> as its arguments. Both set/sequences must consist of elements of the same type, or at least one must have zero cardinality. This function returns the set union of the two sets, or the concatenation of the two sequences in the order given. The infix plus operator (+) may be written in place of the union function.
<code>intersection(s1, s2)</code>	This function is similar to the union function, but it returns the set intersection of its two arguments which must both be sets. The two arguments must both be sets of the same type, or at least one must be the empty set.

- setdiff(*s1*, *s2*)** The **setdiff** command takes two arguments with the same restrictions as the **intersection** function. It returns a copy of *s1* minus any elements it has in common with *s2*. Elements of *s2* which do not appear in *s1* are ignored. The **setdiff** function may be written using the infix minus (-) operator.
- setop(*func*, *set*)** The **setop** operator applies the function *func*, which must be a function of one argument, to each element of the *set*. The results of the function executions are **unioned** into the resulting set, which is returned as the value of the **setop** function. The function *func* may return values which are either individual elements or sets of elements; it may be either a user-defined function or one of the predefined functions **SUCC**, **PRED**, **TFIN**, **TFOUT**, **CARD**, **MARKED**, **NSUCC**, **NPRED**, **SRC**, **DEST**, **TRANS**, **ALLPRED**, **SHOWSTATE**, and **ALLSUCC**.

4. User-Defined Functions

Defining a function is similar to assigning a value to a variable. The primary difference is that the expression is not immediately evaluated. Instead, the parse tree which represents the expression is stored as the value of the identifier, with a special type indicating that the "value" of the identifier is an unevaluated expression tree. Whenever that identifier is subsequently evaluated, the expression tree is retrieved and evaluated, with its value being returned as the value of the identifier. Functions defined in this way may be written recursively; as in pure LISP, recursion is the primary mechanism of looping and flow control.

To define a function, the "::<=" operator is used. **Functions may be defined only at the top command level.** If RGA is being used interactively, then defining a function will cause the message "ok" to be printed on the terminal. The list of formal parameters for the function, if any, is enclosed in parentheses after the identifier and before the :::= operator. As with function calls, the parentheses are omitted if there are no parameters. Local variables, if any, of the function are listed within square brackets following the formal parameters. If there are no local

variables, the brackets are omitted. The expression which defines the function is given to the right of the operator. At the top command level, the special command **show id** will print the definition of the *id* if it is a function. The full syntax of a function definition is given in Appendix A.

Like other identifiers in the language, the arguments are given their types at the time they are bound to values, when the function is invoked. For example, assume the following definition:

$$\text{setx}(v) ::= x := v$$

Then one may type **setx(1)** and the identifier *x* will be defined as having the value of the integer constant 1. The **setx** function will also return the integer value 1, since it expands to an assignment expression which has that value. Subsequently typing the command

$$\text{setx}(\{s \text{ in } S \mid \text{nsucc}(s)=0\})$$

causes *x* to be assigned to the set of all deadlocked states. The previous value of *x* is thrown away. Note that the formal parameter *v* takes on values of different types dynamically. The existing value of *v*, if any, will still be valid when **setx** has returned. Incidentally, it is the dynamically scoped "global" value of *x* which is assigned in the above examples.

5. Some Examples

In the introduction, some expressions for net deadlock-freeness and net conservation are given. In this section, some more complex examples are given to illustrate the full power of the system.

If the Petri net is *safe*, then each place will have at most one token on it (ie, each place is 1-bounded) [PETE77]. One might test for this condition with the

```

/*
 * Findmax returns the maximum value attained by marked() over all states.
 */
findmax[max] ::= \
    max:=0; \
    forall s in S [if marked(s) > max then max := marked(s) fi; true]; \
    max

```

Figure 1

Function to find the maximum value of **marked**

expression

```
forall s in S [forall p in P [p <= 1]]
```

Note that the value of **p** in **p <= 1** is **p(s)**. There is a faster way to test for net safety, however:

```
forall s in S [marked(s) = tokens(s)]
```

This function works because the **marked** function returns the number of places which have at least one token on them for the state **s**. If any of these places has more than one token on it, then **tokens(s)** will be greater than **marked(s)**. It is faster because it avoids the doubly-nested loop of the first approach.

Suppose one wishes to know the maximum value of the **marked** function over all the states in the graph. This could be obtained with the function shown in Figure 1. Notice the use of the semicolon operator to make the expression in the **forall** statement return a **true** value, thus guaranteeing that each state in **S** will be tested. The entire expression returns the maximum value found, **max**, which is a local variable of the function.

Now suppose that we wish to define a boolean function which returns **true** if a particular state can be reached from another state in the graph (see Figure 2). The definition given here is breadth first: it always has a set of states that it knows have already been checked, and one whose successors have not yet been checked. The

```

/*
 * Breadth-first search version of reachable
 */
reachable(s, final)[morestates, tried] := \
  s = final | \
  (tried := emptyset := {}); \
  try({s}))

/*
 * For each s in nextset, test if final is in succ(s).
 * If not, iterate on all those successors.
 */
try(nextset) := \
  morestates := emptyset; \
  exists s in nextset [in(final, succ(s)) | \
    card(morestates := union(morestates, succ(s))) < 0] | \
  card(morestates:=setdiff(morestates, tried:=union(tried, nextset))) > 0 & \
  try(morestates)

```

Figure 2

Recursive Breadth-First Reachability Function

function is actually divided into two parts, **reachable**(*s*, *final*) and **try**(*nextset*). The **reachable** function is the top-level definition. It checks if the starting state, *s*, is the same as the desired state, *final*. If not, it initializes the set constant *emptyset*, which is used in the **try** function only for speed and clarity, and calls **try**.

This pair of functions takes advantage of the the dynamic binding of RGA. The **reachable** function has two local variables *morestates* and *tried*. They are both actually locals of the pair of functions, since they are shared by both. If they were not locals of **reachable**, then any global value with the same name would be lost by executing the **reachable** function. *Morestates* contains the next set of states to be tested by the **try** function; *tried* is the set of states whose successors have already been tested.

The **try** function takes one argument: the set of states which have recently been tested against *final* (*nextset*) whose successors now need to be checked. For each element of *nextset*, **try** compares its successors to *final*. This test can be made quickly since it uses only the built-in functions **exists**, **in**, and **succ**. If a

match is found, **try** returns **true**. If all the matches fail, *morestates* will have been assigned the set of all the successors just tested. Any states in *morestates* which are in the set *tried* are removed, and the cardinality of the resulting set is compared to zero. If it is zero, then **try** will return **false** since there are no more states whose successors have yet to be compared to *final*. Otherwise, **try** is invoked recursively to try the elements in *morestates*, with *tried* augmented with the set of states just tried.

The **allsucc** function could have been used to determine the same function. If the matching successor is near the end of those tested by the above function, or if there is no match, it would be significantly faster than the **reachable** function given. But if the match occurs early in the search, then the above code could be significantly faster, since it would not bother to generate further successors.

As a final example, suppose that all the states in the reachability graph have been partitioned in to two sets, **good** and **bad**. One might then wish to know the transition(s) which lead from the **good** to the **bad** states. The set of transition-firings between the two sets is expressed as

```
tfs := intersection(setop(tfout, good), setop(tfin, bad))
```

The set of "critical" transitions (those which take the net from the "good" states to the "bad" ones, [RAZO80]) is then given by the expression

```
setop(trans, tfs)
```

6. Running RGA

Typically, the user will enter a Petri net representation of the system to be analyzed in a symbolic notation. This symbolic representation is translated into the reachability graph via other programs which are described elsewhere. Their output may then be analyzed by RGA. RGA is executed like an editor, taking the

graph from a file rather than from the standard input. If no file is specified on the command line, the file `main.input` will be read. RGA then prompts the user for commands with a ">" character reading the commands from `stdin`. RGA exits when it receives an end of file from its primary input.

It is possible to have RGA read its input from a disk file instead of from `stdin`. When the end of the file is reached, it resumes reading commands from the previous input source. The command to read from a file is the "@" character, followed by the name of the file to read. If the name of the file does not follow the lexical rules for an identifier in RGA, then it must be quoted in double quotation marks. One command file may recursively read another file by using the @ command. If a command file defines global values, it is convenient to end such commands with the ; operator so that the values will not be printed as the file is interpreted.

As mentioned previously, the user may have the definition of a function printed by issuing the `show` command followed by the name of the function. The formal parameters and local variables, if any, will be displayed in addition to the function's definition. The other command which may be used only at the prompting level is defining a function with the `::=` command.

7. Implementation and Performance

The RGA program is written in the C programming language running on the 4.2BSD UNIX[‡] operating system. The current implementation was written on a VAX-11/750 computer.[†]

The RGA system implements a minimum number of primitive operations to allow all the different operations which were desired in the initial design of the system. It has been designed also to be extensible, since the user may define new functions in terms of the primitives. In particular, the ability to pass values to

[‡] UNIX is a trademark of AT&T Bell Labs

[†] VAX is a trademark of Digital Equipment Corporation

parameters of functions, and the existence of the semicolon operator and the `if` expression, were included expressly for extensibility purposes.

But RGA was also written to be fairly efficient. Efficiency is necessary if large reachability graphs are to be handled, and the system will only really be useful if realistically-large graphs can be analyzed in a reasonable period of time. In the interests of efficiency, some non-primitive operations which could be implemented as user-defined functions have been coded as primitive routines instead. For example, as a test, the primitive function `nsucc` was defined as `ns(s) ::= card(tfout(s))`. By executing the primitive and then the defined function in a loop 75,705 times, the `nsucc` function was measured to be 4.17 times faster than the user-defined equivalent. More performance measurements are given in Appendix B.

8. Conclusions and Future Work

A powerful extension to the language would be the inclusion of temporal logic expressions. The current system allows for the definition of some limited temporal logic expressions, but they could be made more general as well as more efficient if included directly in the language. There is no way currently to pass a function as an argument, or to pass identifiers by-reference.

Perhaps future versions of RGA will overcome these problems if they prove to be serious. Other commonly-used functions may become primitives as they are identified. New primitives to make sequences more powerful also need to be identified and implemented.

Appendix A

BNF Grammer for RGA Language

```

<formals> :: ( <list_of_exprs> )
           |

<locals> :: [ <list_of_exprs> ]
           |

<definition> :: <ident> <formals> <locals> ::= <expr>

<list_of_exprs> :: <expr>
                 | <list_of_exprs> , <expr>

<transition> :: $ <number>

seqstart :: //

addop :: +
        | -

mulop :: *
        | /
        | %

lrelop :: iff
        | implies

relop :: =
        | >
        | >=
        | <
        | <=
        | <>
        | !=

<seqconst> :: <seqstart> <list_of_elems> <seqstart>

<expr> :: <seqconst>

```

```

<identifier> := <expr>
true
false
[ <state> , <state> , <transition> ]
<transition>
if <expr> then <expr> else <expr> fi
if <expr> then <expr> fi
<subset>
print ( <expr> )
<expr> and <expr>
<expr> or <expr>
not <expr>
forall <identifier> in <set> [ <expr> ]
exists <identifier> in <set> [ <expr> ]
( <expr> )
in ( <expr> , <set> )
in ( <expr> , <seqconst> )
<expr> <relop> <expr>
<expr> ; <expr>
<expr> <lrelop> <expr>
<expr> <addop> <expr>
<expr> <mulop> <expr>
<addop> <expr>
<identifier> ( <list_of_exprs> )
tokens ( <state> )
| <state> |
card ( <set> )
card ( <seqconst> )
marked ( <state> )
<expr> ^ <expr>
nsucc ( <state> )
npred ( <state> )
src ( <expr> )
dest ( <expr> )
trans ( <expr> )
# <number>
<identifier>
<number>

<state> ::
<identifier>
<identifier> ( <list_of_exprs> )
Print ( <state> )
# <number>

```

```

|           <identifier> Becomes <state>

<state_range> :: # <number> .. # <number>

<list_of_elems> :: <expr>
|               <state_range>
|               <list_of_elems> , <expr>
|               <list_of_elems> , <state_range>

<setopfunc> :: Succ
|             Pred
|             Tfin
|             Tfout
|             Tokens
|             Marked
|             Nsucc
|             Npred
|             Src
|             Dest
|             Trans

<subset> :: Succ ( <state> )
|             Pred ( <state> )
|             Allpred ( <state> )
|             Allsucc ( <state> )
|             Tfin ( <state> )
|             Tfout ( <state> )
|             Setop ( <identifier> , <set> )
|             Setop ( <setopfunc> , <set> )
|             Union ( <set> , <set> )
|             Union ( <seqconst> , <seqconst> )
|             Setdiff ( <set> , <set> )
|             Intersection ( <set> , <set> )
|             { <list_of_elems> }
|             { }
|             { <identifier> In <set> Or <expr> }

<set> :: <identifier>
|       <identifier> ( <list_of_exprs> )
|       <identifier> Becomes <set>
|       Print ( <set> )
|       <subset>

```


Appendix B

Some Performance Measurements

The table below contains some time and space measurements of the current implementation of RGA. The problem measured was the dining philosophers problem for a varying number of philosophers, between two and eight. For each number of philosophers, the number of states in the reachability graph and the time to load the reachability graph were measured. The time is divided into user and system CPU time, in seconds. In addition, the size of the interpreter in kilobytes was measured after loading the graph but before executing any tests. Finally, two typical problems were executed, testing for the safety of the net, and determining the set of states which can reach state zero (#0) after zero or more transition firings. The `canreach` function is a user-coded version of the `allpred` primitive. On the average, it is about 37% as fast as the primitive function. For each of these tests, the execution time minus the load time is given, in CPU seconds.

Some Performance Measurements					
Dining n	States	Load Time	RGA Size (Kb)	Safe	Canreach(#0)
2	21	0.1 + 0.4	59	0.1	0.6
3	26	0.1 + 0.4	59	0.1	0.8
4	80	0.5 + 0.5	73	0.3	1.3
5	242	1.8 + 0.7	121	0.7	3.5
6	728	6.1 + 1.5	269	2.2	15.4
7	2186	21.3 + 5.3	967	7.2	95.2
8	6560	76.7 + 17.2	2913	23.7	750.6

Figure 3

Some Performance Measurements of RGA
