

The Rexx Language (Background)

13th February 2012

Mike Cowlishaw

`mfc@speleotrove.com`

`http://speleotrove.com/mfc/`

Version 4.00

Copyright © Mike Cowlshaw 1979, 2012.
Parts Copyright © Prentice Hall 1985, 1990.
All rights reserved.

Table of Contents

Background	5
What Kind of a Language is REXX?	6
Summary of the REXX Language	8
Fundamental Language Concepts	12
Design Principles	17
History	19
Index	21

Background

This introductory part of the book is in five sections. The first two sections introduce the REXX language, the next two sections describe the concepts and design principles that shaped it, and the final section reviews the history of the language.

What Kind of a Language is REXX?

REXX is a procedural language that allows programs and algorithms to be written in a clear and structured way. The primary design goal has been that it should be genuinely easy to use both by computer professionals and by “casual” general users. A language that is designed to be easy to use must be effective at manipulating the kinds of symbolic objects that people normally deal with: words, numbers, names, and so on. Most of the features in REXX are included to make this kind of symbolic manipulation easy. REXX is also designed to be independent of its supporting system software, but with the capability of issuing both commands and conventional inter-language calls to its host environment.

The REXX language covers several application areas that traditionally have been served by fundamentally different types of programming language.

Personal programming REXX is a language that provides powerful character and arithmetical abilities in a simple framework. You may write short programs with a minimum of overhead, yet facilities exist to allow the writing of robust large programs. The language is well adapted to interpretation, and is therefore rather suitable for many of the applications for which languages such as BASIC are currently used.

REXX has proved to be an easy language to learn and to teach. As a first language for students, it has the advantage of being a practical and structured programming language which is also easy to use and to debug.

Tailoring user commands Command program interpreters are an important component of modern operating systems. Nearly all operating systems include some form of Executive, Shell, or Batch language. In many cases the language is so embedded into the operating system that it is unlikely to be of use outside its primary environment, but there is a clear trend towards providing command programming languages that are both powerful and capable of more general usage. REXX carries this principle further by being a language that is designed primarily for generality but also for suitability as a command programming language.

Over the years, many REXX programs for tailoring operating systems have been written – originally for the Conversational Monitor System component of the IBM Virtual Machine/System Product, and later for other operating systems. Many of these programs run to hundreds or thousands of lines, and some are in the tens of thousands. One laboratory that uses REXX has over four million lines of code written in REXX, with more than ten percent of the files on its main computer system being REXX programs.

Macros Many applications are programmable by means of macros. In the data processing world there is a different macro language for almost every type of application. There are macro languages for editors, assemblers, interactive systems, text processors, spreadsheets, databases, and of course for other languages. The work of Stephenson¹ and others has highlighted the similarities between these applications, and the need for a common language. Since REXX is essentially a character manipulation language, it can provide the macro language for all these applications.

¹ Stephenson, C. J. **On the structure and control of commands.** *ACM Operating Systems Review (SIGOPS)*, Vol 7, No 4, pp22-26 and 127-136 (1973).

Macro languages often have unusual qualities and syntax that restrict their use to skilled programmers. REXX has a more conventional syntax and is a flexible language, and so makes it possible for the same jobs to be done in less time by less skilled personnel.

*Prototype
development*

Interpreter implementations of REXX can be highly interactive, and permit rapid program development. This productivity advantage, together with the ease of interfacing REXX to system utilities for display and for data input and output, makes the language very suitable for modelling applications and products. It has also proved useful for setting up experimental systems for human factors studies.

The design of REXX is such that the same language can effectively and efficiently be used for many different applications that previously required the learning of several languages.

Summary of the REXX Language

REXX is a language that is superficially similar to earlier languages. However, every aspect of the language has been critically reviewed and usually differs from other languages in ways that make REXX more suited to general users. It was possible to make these improvements because REXX was designed as an entirely new language, without the requirement that it be compatible with any earlier design.

The structure of a REXX program is extremely simple. This sample program, TOAST, is complete, documented, and executable as it stands.

```
TOAST
/* This wishes you the best of health. */
say 'Cheers!'
```

TOAST consists of two lines: the first is a comment that describes the purpose of the program, and the second is an instance of the SAY instruction. SAY simply displays the result of the expression following it – in this case a literal string.

Of course, REXX can do more than just display a character string. Although the language is composed of a small number of instructions and options, it is powerful. Where a function is not built-in it can be added by using one of the defined mechanisms for external interfaces.

The rest of this section introduces most of the features of REXX. It is intended as a brief introduction to the language to serve as a background for the rest of the book. Since many of the subtleties of REXX are best appreciated with use, you are urged to use the language yourself.

REXX provides a conventional selection of *control constructs*. These include IF... THEN... ELSE for simple conditional processing, SELECT... WHEN... OTHERWISE... END for selecting from a number of alternatives, and several varieties of DO... END for grouping and repetition. These constructs are similar to those of PL/I, but with several enhancements and simplifications. The DO (looping) construct can be used to step a variable TO some limit, FOR a specified number of iterations, and WHILE or UNTIL some condition is satisfied. DO FOREVER is also provided. Loop execution may be modified by LEAVE and ITERATE instructions that significantly reduce the complexity of many programs. No GOTO instruction is included, but a SIGNAL instruction is provided for abnormal transfer of control, such as error exits and computed branching.

REXX *expressions* are general, in that any operator combinations may be used (provided, of course, that the data values are valid for those operations). There are 9 arithmetic operators (including integer division, remainder, and power operators), 3 concatenation operators, 12 comparative operators, and 4 logical operators. All the operators act upon strings of characters, which may be of any length (typically limited only by the amount of storage available).

This sample program shows both expressions and a conditional instruction:

```
GREET
/* A short program to greet you. */
/* First display a prompt: */
say 'Please type your name and then press ENTER:'
parse pull answer /* Get the reply into ANSWER */

/* If nothing was typed, then use a fixed greeting, */
/* otherwise echo the name politely. */
if answer='' then say 'Hello Stranger!'
else say 'Hello' answer''
```

The expression on the last SAY (display) instruction concatenates the string 'Hello' to the value of variable ANSWER with a blank in between them (the blank is here a valid operator, meaning “concatenate with blank”). The string '!' is then directly concatenated to the result built up so far. These simple and unobtrusive concatenation operators make it very easy to build up strings and commands, and may be freely mixed with the other operators.

The layout of control constructs is very flexible. In the GREET example, for instance, the IF construct could be laid out in a number of ways, according to personal preference. Line breaks can be added at either side of the THEN (or following the ELSE), or multiple instructions can be placed on one line with the aid of the semicolon separator.

In REXX, any string or symbol may be a *number*. Numbers are all “real” and may be specified in exponential notation if desired. (An implementation may use appropriately efficient internal representations, of course.) The arithmetic operations in REXX are designed for people rather than for the machine, so are decimal rather than binary and have a number of user-oriented features. The operations are completely defined so that different implementations will always give the same results.

The NUMERIC instruction may be used to select the *arbitrary precision* of calculations (you may calculate with one thousand significant digits, for example). The same instruction may also be used to set the *fuzz* to be used for comparisons (that is, the number of significant digits of error permitted when making a numerical comparison) and the exponential notation (scientific or engineering) that REXX will use to present results.

Variables all hold strings of characters, and cannot have aliases under any circumstances. The simple *compound variable* mechanism allows the use of arrays (many-dimensional) that have the property of being indexed by arbitrary character strings. These are in effect content-addressable data structures, which can also be used for building lists and trees. Groups of variables (arrays) with a common stem to their name can be set, reset, or manipulated by references to that stem alone.

This example is a routine that removes all duplicate words from a string of words:

```
JUSTONE
/* This removes duplicate words from a string, and */
/* shows the use of a compound variable (HADWORD) */
/* which is indexed by arbitrary data (words). */
Justone: procedure /* make all variables private */
  parse arg wordlist /* get the list of words */
  hadword.=0 /* show all possible words as new */
  outlist='' /* initialize the output list */
  do while wordlist~='' /* loop while we have data */
    /* split WORDLIST into first word and remainder */
    parse var wordlist word wordlist
    if hadword.word then iterate /* loop if had word */
    hadword.word=1 /* remember we have had this word */
    outlist=outlist word /* add word to output list */
  end
  return outlist /* finally return the result */
```

This example also shows some of the built-in *string parsing* available with the PARSE instruction. This provides a fast and simple way of decomposing strings of characters using a primitive form of pattern matching. A string may be split into parts using various forms of patterns, and then assigned to variables by words or as a whole.

A variety of internal and external calling mechanisms are defined. The most primitive is the *command*

(which is quite similar to a *message* in the Smalltalk-80² system and in other object-oriented systems), in which a clause that consists of just an expression is evaluated. The resulting string of characters is passed to the currently selected external environment, which might be an operating system, an editor, or any other functional object. This ability to send commands to different environments is a primary concept of the language and is especially important when REXX is used as a “macro” language for extending applications.

The REXX programmer can also invoke *functions* and *subroutines*. These may be internal to the program, built-in (part of the language), or external. Within an internal routine, variables may be shared with the caller, or protected by the PROCEDURE instruction (that is, be made local to the routine). If protected, selected variables or groups of variables belonging to the caller may be exposed to the routine for read or write access.

Certain types of *exception handling* are supported. A simple mechanism (associated with the CALL and SIGNAL instructions) allows the trapping of run-time errors, halt conditions (external interrupts), command errors (errors resulting from external commands), stream (input and output) errors, and the use of uninitialized variables. Where appropriate it is possible to call a subroutine to handle the exception, and error handling is supported by a useful set of built-in functions.

The INTERPRET instruction (expected to be supported by interpreters only) allows any string of REXX instructions to be interpreted dynamically. It is useful for some kinds of interactive or interpretive environments, and can be used to build the following SHOWME program – an almost trivial “instant calculator”:

```
SHOWME
/* Simple calculator that evaluates REXX expressions. */
numeric digits 20          /* Work to 20 digits          */
parse arg input           /* Get expression into INPUT */
interpret 'Say' input     /* Build and execute SAY    */
```

This program first sets REXX arithmetic to work to 20 digits. It then assigns the first argument string (perhaps typed by a user) to the variable INPUT. The final instruction evaluates the expression following the keyword INTERPRET to build a SAY instruction which is then executed. If you were to call this program with the argument “22/7” then the instruction “Say 22/7” would be built and executed. This would therefore display the result

```
3.1428571428571428571
```

Input and *output* functions in REXX are defined only for simple character-based operations. Included in the language are the concepts of named character streams (whose actual source or destination are determined externally). These streams may be accessed on a character basis or on a line-by-line basis. One input stream is linked with the concept of an *external data queue* that provides for limited formal communication with external programs.

A rich set of built-in functions is included. These provide extensive string and word manipulations, date and time extraction (in a variety of formats), conversions, bit manipulations, number manipulation and formatting, state and error handling, input and output, and random number generation.

The language defines an extensive *tracing* (debugging) mechanism, though it is recognized that some implementations may be unable to support the whole package or may prefer to provide an alternative process. The tracing options allow various subsets of instructions to be traced (Commands, Labels, All, and so on), and also control the tracing of various levels of expression evaluation results

2 See, for example: Xerox Learning Research Group, **The Smalltalk-80 system**, *Byte* 6, No. 8, pp36-47 (August 1981).

(intermediate calculation results, or just the final results). Furthermore, for a suitable implementation, the language describes an *interactive tracing* option, in which the execution of the program may be halted selectively. Once execution has paused, you may then type in any REXX instructions (to display or alter variables, and so on), step to the next pause, or re-execute the last clause traced.

An example, longer than those shown above, of a REXX program is included as an appendix to this book.

Fundamental Language Concepts

Language design is always subtly affected by unconscious biases and by historical precedent. To minimize these effects a number of concepts were chosen and used as guidelines for the design of the REXX language. The following list includes the major concepts that were consciously followed during the design of REXX.

A complete treatment of some of these topics would fill another book, so unfortunately these paragraphs can only be summaries of the extensive discussions that led to the current design.

Readability

If there is one concept that has dominated the evolution of REXX syntax, it is *readability* (used here in the sense of perceived legibility). Readability in this sense is a rather subjective quality, but the general principle followed in REXX is that the tokens which form a program can be written much as one might write them in Western European languages (English, French, and so forth). Although the semantics of REXX is, of course, more formal than that of a natural language, REXX is lexically similar to normal text.

The structure of the syntax means that the language readily adapts itself to a variety of programming styles and layouts. This helps satisfy user preferences and allows a lexical familiarity that also increases readability. Good readability leads to enhanced understandability, thus yielding fewer errors both while writing a program and while reading it for information, debugging, or maintenance.

Important factors here are:

1. There is deliberate support throughout the language for upper and lower case letters, both for processing data and for the program itself.
2. The essentially free format of the language (and the way blanks are treated around tokens and so on) lets you lay out the program in the style that you feel is the most readable.
3. Punctuation is required only when absolutely necessary to remove ambiguity (though it may often be added according to personal preference, so long as it is syntactically correct). This relatively tolerant syntax has proved to be less frustrating than the syntax of languages such as Pascal.
4. Modern concepts of structured programming are available in REXX, and can undoubtedly lead to programs that are easier to read than they might otherwise be. The structured programming constructs also make REXX a good language for teaching the concepts of good structure.
5. Loose binding between lines and program source ensure that even though programs are affected by line ends, they are not irrevocably so. You may spread a clause over several lines or put it on just one line. Clause separators are optional (except where more than one clause is put on a line), again letting you adjust the language to your own preferred style.

Natural data typing

“Strong typing”, in which the values that a variable may take are tightly constrained, first became a fashionable attribute for languages in the 1970s. I believe that the greatest advantage of strong typing is for the interfaces between program modules, where errors may be difficult to catch. Errors *within* modules that would be detected by strong typing (and would not be detected from context) are much rarer, certainly when compared with design errors, and in the majority of

cases do not justify the added program complexity.

REXX, therefore, treats types as naturally as possible. The meaning of data depends entirely on their usage. All values are defined in the form of the symbolic notation (strings of characters) that a user would normally write to represent that data. Since no internal or machine representation is exposed in the language, the need for many data types is reduced. There are, for example, no fundamentally different concepts of *integer* and *real*; there is just the single concept of *number*. The results of all operations have a defined symbolic representation, so you can always inspect values (for example, the intermediate results of an expression evaluation). Numeric computations and all other operations are precisely defined, and will therefore act consistently and predictably for every correct implementation.

This language definition does not exclude the future addition of a data typing mechanism for those applications that require it, though there seems to be little call for this. The mechanism could perhaps be in the form of ASSERT-like instructions that assign data type checking to variables during execution flow. An optional restriction, similar to the existing trap for uninitialized variables, could be defined to provide enforced assertion for all variables.

*Emphasis on
symbolic
manipulation*

The values that REXX manipulates are (from the user's point of view, at least) in the form of strings of characters. It is extremely desirable to be able to manage this data as naturally as you would manipulate words on a page or in a text editor. The language therefore has a rich set of character manipulation operators and functions.

Concatenation, the most common string operation, is treated specially in REXX. In addition to a conventional concatenate operator (“||”), there is a novel *blank operator* that concatenates two data strings together with a blank in between. Furthermore, if two syntactically distinct terms (such as a string and a variable name) are abutted, then the data strings are concatenated directly. These operators make it especially easy to build up complex character strings, and may at any time be combined with the other operators available.

For example, the SAY instruction consists of the keyword SAY followed by any expression. In this instance of the instruction, if the variable N has the value '6' then

```
say n*100/50'% ' ARE REJECTS
```

would display the string

```
12% ARE REJECTS
```

Concatenation has a lower priority than the arithmetic operators. The order of evaluation of the expression is therefore first the multiplication, then the division, then the direct concatenation, and finally the two “concatenate with blank” operations.

Since the concatenation operators are distinct from the arithmetic operators, very natural coercion (automatic conversion) between numbers and character strings is possible and has become a highly valued feature of the language.

Dynamic scoping Most languages (especially those designed to be compiled) rely on static scoping,

where the physical position of an instruction in the program source may alter its meaning. Languages that are interpreted (or that have advanced compilers) generally have *dynamic scoping*. Here, the meaning of an instruction is only affected by the instructions that have already been executed (rather than those that precede or follow it in the program source).

REXX scoping is purely dynamic. This implies that it may be efficiently interpreted because only minimal look-ahead is needed. It also implies that a compiler is harder to implement, so the semantics includes restrictions that ease the task of the compiler writer. Most importantly, though, it implies that in general a person reading the program need only be aware of the program *above* the point which is being studied. Not only does this aid comprehension, but it also makes programming and maintenance easier when only a computer display terminal is being used.

The GOTO instruction is a necessary casualty of dynamic scoping. In a truly dynamically scoped language, a GOTO cannot be used as an error exit from a loop. If it were, the loop would never become inactive.³ REXX instead provides an “abnormal transfer of control” instruction, SIGNAL, that terminates all active control structures when it is executed. Note that it is not just a synonym for GOTO since it cannot be used to transfer control within a loop (for which alternative instructions are provided).

Nothing to declare

Consistent with the philosophy of simplicity, REXX provides no mechanism for declaring variables. Variables may of course be documented and initialized at the start of a program, and this covers the primary advantages of declarations. The other, data typing, is discussed above.

Implicit declarations do take place during execution, but the only true declarations in the REXX language are the markers (*labels*) that identify points in the program that may be used as the targets of SIGNAL instructions or internal routine calls.

System independence

The REXX language is independent of both system and hardware. REXX programs, though, must be able to interact with their environment. Such interactions necessarily have system dependent attributes. However, these system dependencies are clearly bounded and the rest of the language has no such dependencies. In some cases this leads to added expense in implementation (and in language usage), but the advantages are obvious and well worth the penalties.

As an example, string-of-characters comparison is normally independent of leading and trailing blanks. (The string “ Yes ” *means* the same as “Yes” in most applications.) However, the influence of underlying hardware has subtly affected this kind of decision, so that many languages only allow trailing blanks but not leading blanks. By contrast, REXX permits both leading and trailing blanks during general comparisons.

Limited span syntactic units

The fundamental unit of syntax in the REXX language is the clause, which is a piece of program text terminated by a semicolon (usually implied by the end of a line). The span of syntactic units is therefore small, usually one line or less. This means that the syntax parser in the language processor can rapidly detect and locate errors, which in turn means that error messages can be both precise and

³ Some interpreted languages detect control jumping outside the body of the loop and terminate the loop if this occurs. These languages are therefore relying on static scoping.

concise.

It is difficult to provide good diagnostics for languages (such as Pascal and its derivatives) that have large fundamental syntactic units. For these languages, a small error can often have a major or distributed effect on the parser, which can lead to multiple error messages or even misleading error messages.

Dealing with reality

A computer language is a tool for use by real people to do real work. Any tool must, above all, be reliable. In the case of a language this means that it should do what the user expects. User expectations are generally based on prior experience, including the use of various programming and natural languages, and on the human ability to abstract and generalize.

It is difficult to define exactly how to meet user expectations, but it helps to ask the question “Could there be a high *astonishment factor* associated with this feature?”. If a feature, accidentally misused, gives apparently unpredictable results, then it has a high *astonishment factor* and is therefore undesirable.

Another important attribute of a reliable software tool is *consistency*. A consistent language is by definition predictable and is often elegant. The danger here is to assume that because a rule is consistent and easily described, it is therefore simple to understand. Unfortunately, some of the most elegant rules can lead to effects that are completely alien to the intuition and expectations of a user who, after all, is human.

Consistency applied for its own sake can easily lead to rules that are either too restrictive or too powerful for general human use. During the design process, I found that simple rules for REXX syntax quite often had to be rethought to make the language a more usable tool.

REXX originally allowed almost all options on instructions to be variable (and even the names of functions were variable), but many users fell into the pitfalls that were the side-effects of this powerful generality. For example, the TRACE instruction allows its options to be abbreviated to a single letter (as it needs to be typed often during debugging sessions). Users therefore often used the instruction “TRACE I”, but when I had been used as a variable (perhaps as a loop counter) then this instruction could become “TRACE 10” – a correct but unexpected action. The TRACE instruction was therefore changed to treat the symbol as a constant (and the language became more complex as a consequence) to protect users against such happenings; a VALUE option on TRACE allows variability for the experienced user. There is a fine line to tread between concise (terse) syntax and usability.

Be adaptable

Wherever possible the language allows for extension of instructions and other language constructs. For example, there is a useful set of common characters available for future extensions, since only a restricted set is allowed for the names of variables (symbols). Similarly, the rules for keyword recognition allow instructions to be added whenever required without compromising the integrity of existing programs that are written in the appropriate style. There are no globally reserved words (though a few are reserved within the local context of a single clause).

A language needs to be adaptable because *it certainly will be used for applications*

not foreseen by the designer. Although proven effective as a command programming and personal language, REXX may (indeed, probably will) prove inadequate in certain future applications. Room for expansion and change is included to make the language more adaptable.

Keep the language small

Every suggested addition to the language has been considered only if it would be of use to a significant number of users. The intention has been to keep the language as small as possible, so that users can rapidly grasp most of the language. This means that:

- The language appears less formidable to the new user.
- Documentation is smaller and simpler.
- The experienced user can be aware of all the abilities of the language, and so has the whole tool at his or her disposal to achieve results.
- There are few exceptions, special cases, or rarely used embellishments.
- The language is easier to implement.

No defined size or shape limits

The language does not define limits on the size or shape of any of its tokens or data (although there may be implementation restrictions). It does, however, define the *minimum* requirements that must be satisfied by an implementation. Wherever an implementation restriction has to be applied, it is recommended that it should be of such a magnitude that few (if any) users will be affected.

Where implementation limits are necessary, the language encourages the implementer to use familiar and memorable values for the limits. For example 250 is preferred to 255, 500 to 512, and so on. There is no longer any excuse for forcing the artifacts of the binary system onto a population that uses only the decimal system. Only a tiny minority of future programmers will need to deal with base-two-derived number systems.

Design Principles

A good philosophy for a language is of little use if there is not an effective process for testing the resulting design and tuning it to the needs of its users. As REXX evolved, so too did a certain design ethic; these principles are still followed today for REXX – other projects, too, are using similar techniques.

The design process started rather conventionally – the language was first designed and documented; this initial informal specification was then circulated to a number of appropriate reviewers. The revised initial description then became the basis for the first specification and implementation.

From then on, other less common design principles were followed. The most significant was the intense use of a communications network, but all three items in this list have had a considerable influence on the evolution of REXX.

Communications Once an initial implementation was complete, the most important factor in the development of REXX began to take effect. IBM has an internal network, known as VNET, that at the time linked nearly 1000 mainframe computers in 40 countries. REXX rapidly spread throughout this network, so from the start many hundreds of people were using the language. All the users, from temporary staff to professional programmers, were able to provide immediate feedback to the designer on their preferences, needs, and suggestions for changes. (At times it seemed as though most of them did – at peak periods I was replying to an average of 350 pieces of electronic mail each day.)

An informal language committee soon appeared spontaneously, communicating entirely electronically, and the language discussions grew to be hundreds of thousands of lines.

On occasions it became clear as time passed that incompatible changes to the language were needed. Here the network was both a hindrance and a help. It was a hindrance as its size meant that REXX was enjoying very wide usage and hence many people had a heavy investment in existing programs. It was a help because it was possible to communicate directly with the users to explain why the change was necessary, and to provide aids to help and persuade people to change to the new version of the language. The decision to make an incompatible change was never taken lightly, but because changes could be made relatively easily the language was able to evolve much further than would have been the case if only upwards compatible extensions were considered.

Documentation before implementation Every major section of the REXX language was documented (and circulated for review) before implementation. The documentation was not in the form of a functional specification, but was instead complete reference documentation that in due course became part of this language definition. At the same time (before implementation) sample programs were written to explore the usability of any proposed new feature. This approach resulted in the following benefits:

- The majority of usability problems were discovered before they became embedded in the language and before any implementation included them.
- Writing the documentation was found to be the most effective way of spotting inconsistencies, ambiguities, or incompleteness in a design. (But the documentation must itself be complete, to “final draft” standard.)

- I deliberately did not consider the implementation details until the documentation was complete. This minimized the implementation's influence upon the language.
- Reference documentation written after implementation is likely to be inaccurate or incomplete, since at that stage the author will know the implementation too well to write an objective description.

*The language
user is usually
right*

User feedback was fundamental to the process of evolution of the REXX language. Although users can occasionally be naïve in their suggestions, even those suggestions which appeared to be shallow were considered carefully since they often acted as pointers to deficiencies in the language or documentation. The language has often been tuned to meet user expectations; some of the desirable quirks of the language are a direct result of this necessary tuning. Much would have remained unimproved if users had had to go through a formal suggestions procedure instead of simply sending a piece of electronic mail directly to me. All of this mail was reviewed some time after the initial correspondence in an effort to perceive trends and generalities that might not have been apparent on a day-to-day basis.

Many (if not most) of the good ideas embodied in the language came directly or indirectly from suggestions made by users. It is impossible to overestimate the value of the direct feedback from users that was available while REXX was being designed.

History

The REXX language (originally called “REX”) borrows from many earlier languages; PL/I, Algol, and even APL have had their influences, as have several unpublished languages that I developed during the 1970s.

The language has developed in two distinct phases: the first being the rapid evolution of the language in an essentially experimental environment, and the second being a more cautious series of enhancements following the commercial availability of implementations of the language.

The first phase took place as a personal project of about four thousand hours during the years 1979 through 1982, at the IBM UK Laboratories near Winchester (England) and at the IBM T. J. Watson Research Center in New York (USA). With this background REXX has an international flavour, with roots in both the European and North American programming cultures.

In 1983, my own System/370 implementation became part of the Virtual Machine/System Product, as the System Product Interpreter for the Conversational Monitor System (CMS). This implementation of the language is described in the Reference Manual for that product.⁴ In 1985 the first edition of this book was published, and soon after that the pioneer non-IBM implementation of REXX was announced by the Mansfield Software Group: this implementation runs under the MS-DOS and PC-DOS operating systems for Personal Computers. A number of other implementations have followed from a variety of suppliers: one which perhaps best demonstrates the suitability of REXX for different environments is a version for the Commodore Amiga computer.

The next milestone for REXX was its choice by IBM as the Procedures Language for the Systems Application Architecture (SAA).⁵ This 1987 announcement implies a common REXX language across all the SAA operating systems: VM, MVS, OS/400, and OS/2. ⁶ The language interpreter development for all these environments is coordinated at the IBM Endicott Programming Laboratory, New York.

All the first implementations of REXX were interpreters: notable, then, was the announcement in 1989 of IBM’s CMS REXX Compiler, developed at the IBM Vienna Software Development Laboratory in Austria with help from the IBM Scientific Centre at Haifa in Israel.

Inevitably the commercial exploitation of the language has required a stable language definition – the radical changes in the language that were characteristic of its first years are no longer possible. Fortunately, those early years of heavy use and rapid evolution probably mean that such radical changes are no longer necessary: rather one would expect to see incremental changes and adjustments consistent with the philosophy of keeping the language small and approachable. Even so it is not impossible that major enhancements could be added to the base REXX language: over the years there have been research proposals for both a REXX “systems programming language” and an object-oriented REXX. REXX will doubtless continue to evolve as software technology itself evolves. I hope, and expect, that even as it changes it will always remain true to its original goal.

4 **IBM Virtual Machine/System Product: System Product Interpreter Reference.** *IBM Reference Manual*, Order No. SC24-5239, IBM (1983).

5 The Procedures Language for SAA comprises the REXX language, Double Byte Character Set support, and a series of common interfaces to the language.

6 More formally: CMS in the VM/System Product or VM/Extended Architecture, TSO/E in the Enterprise Systems Architecture/370, Operating System/400 for the Application System/400 (AS/400), and Operating System/2 Extended Edition.

Index

A

Adaptability 15
Algol language 19
Amiga implementation 19
APL language 19
Application areas for REXX 6
AS/400 implementation 19
Astonishment factor 15

B

Background 5
BASIC language 6
Batch languages 6

C

CMS 6
 implementation 19
Communications 17
Compiler
 first for REXX 19
Consistency 15

D

Data
 type checking 12
Datatyping 12
Dealing with reality 15
Declarations, why none in REXX 14
Design principles for REXX 17
Documentation before implementation 17
Dynamic scoping 13

E

Electronic mail 17
Examples
 of programs 8-10
Executive languages 6

F

Feedback from users, value of 18

G

GOTO instruction, why not in REXX 14
GREET example program 8

H

History of REXX 19

I

Implementations of REXX 19
Incompatible changes 17
Introduction to REXX 8

J

JUSTONE example program 9

L

Language committee 17
Language concepts 12
Legibility, perceived 12
Limits of size 16

M

Macros 6
Mail, electronic 17
MS-DOS implementation 19
MVS implementation 19

N

Natural data typing 12
Network, electronic 17
Nothing to declare 14

O

OS/2 implementation 19
OS/400 implementation 19

P

PC-DOS implementation 19
Perceived legibility 12
Personal programming 6

Procedures Language 19
Programming style 12
Programs
 examples 8-10
Prototype development 7
Prototyping
 see Prototype development 7

Q

quot.PL/I language 19

R

Readability, of programs 12
Reality, dealing with 15
Reliability, of a language 15

REXX

 background 5
 compiler 19
 design principles 17
 history 19
 implementations 19
 language concepts 12
 object-oriented 19
 summary of the language 8

S

SAA
 see Systems Application Architecture 19

Shell languages 6
SHOWME example program 10
Size

 of language 16
 see Length 16

Sparse arrays

 see Compound variables 9

Strong typing 12

Structured programming concepts 12

Style, programming 12

Summary of the REXX language 8

Symbolic manipulation 13

Syntactic units 14

System independence 14

Systems Application Architecture 19

T

Tailoring user commands 6

TOAST example program 8

Tools, reliability of 15

TSO/E implementation 19

U

User is usually right 18

V

VM/CMS implementation 19

VNET 17