

NetREXX Language Reference

Mike Cowlshaw and REXXLA

Version 4.02-GA of January 24, 2022

THE REXX LANGUAGE ASSOCIATION
NetRexx Programming Series
ISBN 978-90-819090-1-3

Publication Data

©Copyright The Rexx Language Association, 2011- 2022

All original material in this publication is published under the Creative Commons - Share Alike 3.0 License as stated at <http://creativecommons.org/licenses/by-nc-sa/3.0/us/legalcode>.

The responsible publisher of this edition is identified as *IBizz IT Services and Consultancy*, Amsteldijk 14, 1074 HR Amsterdam, a registered company governed by the laws of the Kingdom of The Netherlands.

This edition is registered under ISBN 978-90-819090-1-3

ISBN 978-90-819090-1-3



9 789081 909013 >

Contents

List of Tables	VII
The NetREXX Programming Series	i
1 Introduction	1
1.1 Language Objectives	1
1.2 Language Concepts	4
1.3 Acknowledgements	9
2 Introduction to the current edition	11
3 NetREXX Language Definition	12
4 Notations	13
5 Characters and Encodings	14
5.1 Character Sets	14
6 Structure and General Syntax	16
6.1 Blanks and White Space	16
6.2 Comments	16
6.3 Tokens	17
6.4 Implied semicolons and continuations	20
6.5 The case of names and symbols	21
6.6 Hexadecimal and binary numeric symbols	21
7 Types and Classes	23
7.1 Primitive types	23
7.2 Dimensioned types	24
7.3 Minor and Dependent classes	24

8	Terms	26
8.1	Simple terms	26
8.2	Compound terms	27
8.3	Evaluation of terms	28
8.4	Simple term evaluation	28
8.5	Stub evaluation	28
8.6	Continuation evaluation	30
8.7	Arrays in terms	31
9	Methods and Constructors	32
9.1	Method call instructions	33
9.2	Method resolution (search order)	33
9.3	Method overriding	35
9.4	Return Types	36
9.5	Constructor methods	36
10	Type conversions	38
10.1	Permitted automatic conversions	38
10.2	Permitted explicit conversions	40
10.3	Costs of conversions	40
11	Expressions and Operators	41
11.1	Operators	41
11.2	Numbers	45
11.3	Parentheses and operator precedence	46
12	Clauses and Instructions	48
13	Assignments and Variables	49
13.1	The use and scope of variables	50
13.2	Terms on the left of assignments	52
14	Indexed strings and Arrays	53
14.1	Arrays	54
15	Keyword Instructions	57
16	Annotation instruction	58
16.1	Example	58
17	Address instruction	60

18 Class instruction	62
18.1 Visibility	63
18.2 Modifier	63
18.3 Binary	64
18.4 Deprecated	64
18.5 Extends	65
18.6 Uses	65
18.7 Implements	65
19 Do instruction	66
20 Exit instruction	68
21 If instruction	69
22 Import instruction	71
23 Iterate instruction	73
24 Leave instruction	74
25 Loop instruction	75
25.1 Syntax notes:	76
25.2 Indefinite loops	76
25.3 Bounded loops	76
25.4 Label phrase	79
25.5 Protect phrase	80
25.6 Exceptions in loops	80
25.7 Programmer's model - how a typical loop is executed	80
26 Method instruction	82
26.1 Arguments	83
26.2 Visibility	83
26.3 Modifier	84
26.4 Protect	85
26.5 Binary	85
26.6 Deprecated	86
26.7 Returns	86
26.8 Signals	86
26.9 Duplicate methods	86

27	Nop instruction	88
28	Numeric instruction	89
29	Options instruction	91
30	Package instruction	96
31	Parse instruction	97
32	Properties instruction	99
32.1	Visibility	100
32.2	Modifier	100
32.3	Deprecated	101
32.4	Unused	101
32.5	Properties in interface classes	101
33	Return instruction	102
34	Say instruction	103
35	Select instruction	104
35.1	Label phrase	105
35.2	Protect phrase	105
35.3	Case phrase	106
35.4	Exceptions in select constructs	107
36	Signal instruction	108
37	Trace instruction	109
37.1	Tracing clauses	110
37.2	Tracing variables	110
37.3	The format of trace output	111
38	Program structure	114
38.1	Program defaults	115
39	Minor and Dependent classes	117
39.1	Minor classes	117
39.2	Dependent classes	118
39.3	Restrictions	120

40	Special names and methods	121
40.1	Special names	121
40.2	Special methods	124
41	JavaBean Support	125
41.1	Indirect properties	125
42	Parsing templates	129
42.1	Introduction to parsing	129
42.2	Parsing definition	131
43	Numbers and Arithmetic	136
43.1	Introduction	136
43.2	Definition	137
44	Binary values and operations	146
44.1	Operations in binary classes and methods	146
44.2	Binary constructors	148
45	Exceptions	149
45.1	Syntax and example	150
45.2	Exceptions after catch and finally clauses	151
45.3	Checked exceptions	151
46	Thread Pool Support	153
46.1	Examples	153
47	Structured Lists Interface	154
47.1	Essential List Processing Methods	154
47.2	Convenience Methods	155
47.3	NetRexx Structured List Format	155
47.4	Structured List Ruleset Description	156
48	Methods for NetREXX strings	159
48.1	General notes on the built-in methods:	159
48.2	The built-in methods	160
49	Appendix A - A Sample NetREXX Program	177

50	Appendix B - The netrex.lang Package	180
50.1	Exception classes	180
50.2	The REXX class	181
50.3	REXX constructors	181
50.4	REXX arithmetic methods	182
50.5	REXX miscellaneous methods	185
50.6	The REXXIO class	186
50.7	The REXXDate class	187
50.8	The REXXTime class	189
50.9	The REXXOperators interface class	191
50.10	The REXXSet class	191
51	Appendix C - Translator Options	193
	Index	199

List of Tables

1	Escape sequences	18
2	Concatenation operators	42
3	Arithmetic operators	42
4	Normal comparative operators	43
5	Strict comparative operators	44
6	Boolean operators	44
7	Operator precedence	47
8	Trace identifier tags	112
9	Options	193

The NetREXX Programming Series

This book is part of a library, the *NetREXX Programming Series*, documenting the NetREXX programming language and its use and applications. This section lists the other publications in this series, and their roles. These books can be ordered in convenient hardcopy and electronic formats from the REXX Language Association.

Quick Start Guide

This guide is meant for an audience that has done some programming and wants to start quickly. It starts with a quick tour of the language, and a section on installing the NetREXX translator and how to run it. It also contains help for troubleshooting if anything in the installation does not work as designed, and states current limits and restrictions of the open source reference implementation.

Programming Guide

The Programming Guide is the one manual that at the same time teaches programming, shows lots of examples as they occur in the real world, and explains about the internals of the translator and how to interface with it.

Language Reference

Referred to as the NRL, this is meant as the formal definition for the language, documenting its syntax and semantics, and prescribing minimal functionality for language implementors.

Pipelines Guide & Reference

The Data Flow oriented companion to NetREXX, with its CMS Pipelines compatible syntax, is documented in this manual. It discusses running Pipes for NetREXX in the command shell and the Workspace, and has ample examples of defining your own stages in NetREXX.

Introduction

NetREXX is a general-purpose programming language inspired by two very different programming languages, REXX[™] and Java[™]. It is designed for people, not computers. In this respect it follows REXX closely, with many of the concepts and most of the syntax taken directly from REXX or its object-oriented version, Object REXX. From Java it derives static typing, binary arithmetic, the object model, and exception handling. The resulting language not only provides the scripting capabilities and decimal arithmetic of REXX, but also seamlessly extends to large application development with fast binary arithmetic.

The open source reference implementation (version 3 and later) of NetREXX produces classes for the Java Virtual Machine, and in so doing demonstrates the value of that concrete interface between language and machine: NetREXX classes and Java classes are entirely equivalent – NetREXX can use any Java class (and vice versa) and inherits the portability and robustness of the Java environment.

This document is in three parts:

1. The objectives of the NetREXX language and the concepts underlying its design, and acknowledgements.
2. An overview and introduction to the NetREXX language.
3. The definition of the language.

Appendices include a sample NetREXX program, a description of an experimental feature, and some details of the contents of the `netrexx.lang` package.

1.1 Language Objectives

This document describes a programming language, called NetREXX, which is derived from both REXX and Java. NetREXX is intended as a dialect of REXX that can be as efficient and portable as languages such as Java, while preserving the low threshold to learning and the ease of use of the original REXX language.

1.1.1 Features of REXX

The REXX programming language¹ was designed with just one objective: to make programming easier than it was before. The design achieved this by emphasizing readability and usability, with a minimum of special notations and restric-

¹Cowlshaw, M. F., *The REXX Language* (second edition), ISBN 0-13-780651-5, Prentice-Hall, 1990.

tions. It was consciously designed to make life easier for its users, rather than for its implementers. One important feature of REXX syntax is *keyword safety*. Programming languages invariably need to evolve over time as the needs and expectations of their users change, so this is an essential requirement for languages that are intended to be executed from source.

Keywords in REXX are not globally reserved but are recognized only in context. This language attribute has allowed the language to be extended substantially over the years without invalidating existing programs. Even so, some areas of REXX have proved difficult to extend – for example, keywords are reserved within instructions such as **do**. Therefore, the design for NetREXX takes the concept of keyword safety even further than in REXX, and also improves extensibility in other areas.

The great strengths of REXX are its human-oriented features, including

- simplicity
- coherent and uncluttered syntax
- comprehensive stringhandling
- case-insensitivity
- arbitrary precision decimal arithmetic.

Care has been taken to preserve these. Conversely, its interpretive nature has always entailed a lack of efficiency: excellent REXX compilers do exist, from IBM and other companies, but cannot offer the full speed of statically-scoped languages such as C² or Java³.

1.1.2 Influence of Java

The system-independent design of REXX makes it an obvious and natural fit to a system-independent execution environment such as that provided by the Java Virtual Machine (JVM). The JVM, especially when enhanced with “just-in-time” bytecode compilers that compile bytecodes into native code just before execution, offers an effective and attractive target environment for a language like REXX.

Choosing the JVM as a target environment does, however, place significant constraints on the design of a language suitable for that environment. For example, the semantics of method invocation are in several ways determined by the environment rather than by the source language, and, to a large extent, the object model (class structure, *etc.*) of the Java environment is imposed on languages that use it.

Also, Java maintains the C concept of primitive datatypes; types (such as `int`, a 32-bit signed integer) which allow efficient use of the underlying hardware yet do not describe true objects. These types are pervasive in classes and interfaces written in the Java language; any language intending to use Java classes effectively must provide access to these types.

²Kernighan, B. W., and Ritchie, D. M., *The C Programming Language* (second edition), ISBN 0-13-110362-8, Prentice-Hall, 1988.

³Gosling, J. A., *et al.* *The Java Language Specification*, ISBN 0-201-63451-1, Addison-Wesley, 1996.

Equally, the *exception* (error handling) model of Java is pervasive, to the extent that methods must check certain exceptions and declare those that are not handled within the method. This makes it difficult to fit an alternative exception model.

The constraints of safety, efficiency, and environment necessitated that NetREXX would have to differ in some details of syntax and semantics from REXX; unlike Object REXX, it could not be a fully upwards-compatible extension of the language⁴. The need for changes, however, offered the opportunity to make some significant simplifications and enhancements to the language, both to improve its keyword safety and to strengthen other features of the original REXX design⁵. Some additions from Object REXX and ANSI REXX⁶ are also included.

Similarly, the concepts and philosophy of the REXX design can profitably be applied to avoid many of the minor irregularities that characterize the C and Java language family, by providing suitable simplifications in the programming model. For example, the NetREXX looping construct has only one form, rather than three, and exception handling can be applied to all blocks rather than requiring an extra construct. Also, as in REXX, all NetREXX storage allocation and de-allocation is implicit – an explicit new operator is not required.

Further, the human-oriented design features of REXX (case-insensitivity for identifiers, type deduction from context, automatic conversions where safe, tracing, and a strong emphasis on string representations of common values and numbers) make programming for the Java environment especially easy in NetREXX.

1.1.3 A hybrid or a whole?

As in other mixtures, not all blends are a success; when first designing NetREXX, it was not at all obvious whether the new language would be an improvement on its parents, or would simply reflect the worst features of both.

The fulcrum of the design is perhaps the way in which datatyping is automated without losing the static typing supported by Java. Typing in NetREXX is most apparent at interfaces – where it provides most value – but within methods it is subservient and does not obscure algorithms. A simple concept, *binary classes*, also lets the programmer choose between robust decimal arithmetic and less safe (but faster) binary arithmetic for advanced programming where performance is a primary consideration.

The “seamless” integration of types into what was previously an essentially typeless language does seem to have been a success, offering the advantages of strong typing while preserving the ease of use and speed of development that REXX programmers have enjoyed.

⁴Nash, S. C., *Object-Oriented REXX* in Goldberg, G, and Smith, P. H. III, *The REXX Handbook*, pp115-125, ISBN 0-07-023682-8, McGraw-Hill, Inc., New York, 1992.

⁵See Cowlshaw, M. F., *The Early History of REXX*, IEEE Annals of the History of Computing, ISSN 1058-6180, Vol 16, No. 4, Winter 1994, pp15-24, and Cowlshaw, M. F., *The Future of REXX*, Proceedings of Winter 1993 Meeting/SHARE 80, Volume II, p.2709, SHARE Inc., Chicago, 1993.

⁶See *American National Standard for Information Technology – Programming Language REXX*, X3.274-1996, American National Standards Institute, New York, 1996.

The end result of adding Java typing capabilities to the REXX language is a single language that has both the REXX strengths for scripting and for writing macros for applications and the Java strengths of robustness, good efficiency, portability, and security for application development.

1.2 Language Concepts

As described in the last section, NetREXX was created by applying the philosophy of the REXX language to the semantics required for programming the Java Virtual Machine (JVM). Despite the assumption that the JVM is a “target environment” for NetREXX, it is intended that the language not be environment-dependent; the essentials of the language do not depend on the JVM. Environment-dependent details, such as the primitive types supported, are not part of the language specification.

The primary concepts of REXX have been described before, in *The REXX Language*, but it is worth repeating them and also indicating where modifications and additions have been necessary to support the concepts of statically-typed and object-oriented environments. The remainder of this section is therefore a summary of the principal concepts of NetREXX.

1.2.1 Readability

One concept was central to the evolution of REXX syntax, and hence NetREXX syntax: *readability* (used here in the sense of perceived legibility). Readability in this sense is a somewhat subjective quality, but the general principle followed 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 NetREXX is more formal than a natural language, its syntax is lexically similar to everyday text.

The structure of the syntax means that the language is readily adapted 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. Punctuation and other special notations are required only when absolutely necessary to remove ambiguity (though punctuation may often be added according to personal preference, so long as it is syntactically correct). Where notations are used, they follow established conventions.
2. The language is essentially case-insensitive. A NetREXX programmer may choose a style of use of uppercase and lowercase letters that he or she finds most helpful (rather than a style chosen by some other programmer).
3. The classical constructs of structured and object-oriented programming are available in NetREXX, and can undoubtedly lead to programs that are easier

to read than they might otherwise be. The simplicity and small number of constructs also make NetREXX an excellent language for teaching the concepts of good structure.

4. Loose binding between the physical lines in a program and the syntax of the language ensures that even though programs are affected by line ends, they are not irrevocably so. A clause may be spread over several lines or put on just one line; this flexibility helps a programmer lay out the program in the style felt to be most readable.

1.2.2 Natural data typing and decimal arithmetic

“Strong typing”, in which the values that a variable may take are tightly constrained, has been an attribute of some languages for many years. The greatest advantage of strong typing is for the interfaces between program modules, where errors are easy to introduce and difficult to catch. Errors *within* modules that would be detected by strong typing (and which 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.

NetREXX, therefore, treats types as unobtrusively as possible, with a simple syntax for type description which makes it easy to make types explicit at interfaces (for example, when describing the arguments to methods).

By default, common values (identifiers, numbers, and so on) are described in the form of the symbolic notation (strings of characters) that a user would normally write to represent those values. This natural datatype for values also supports decimal arithmetic for numbers, so, from the user’s perspective, numbers look like and are manipulated as strings, just as they would be in everyday use on paper.

When dealing with values in this way, no internal or machine representation of characters or numbers is exposed in the language, and so 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, and will therefore act consistently and predictably for every correct implementation.

This concept also underlies the BASIC⁷ language; indeed, Kemeny and Kurtz’s vision for BASIC included many of the fundamental principles that inspired REXX. For example, Thomas E. Kurtz wrote:

“Regarding variable types, we felt that a distinction between ‘fixed’ and ‘floating’ was less justified in 1964 than earlier ... to our potential audience the distinction between an integer number and a non-integer number would seem esoteric. A number is a number is a number.”⁸

For REXX, intended as a scripting language, this approach was ideal; symbolic operations were all that were necessary.

For NetREXX, however, it is recognized that for some applications it is necessary

⁷Kemeny, J. G. and Kurtz, T. E., **BASIC programming**, John Wiley & Sons Inc., New York, 1967.

⁸Kurtz, T. E., **BASIC** in Wexelblat, R. L. (Ed), **History of Programming Languages**, ISBN 0-12-745040-8, Academic Press, New York 1981.

to take full advantage of the performance of the underlying environment, and so the language allows for the use and specification of binary arithmetic and types, if available. A very simple mechanism (declaring a class or method to be *binary*) is provided to indicate to the language processor that binary arithmetic and types are to be used where applicable. In this case, as in other languages, extra care has to be taken by the programmer to avoid exceeding limits of number size and so on.

1.2.3 Emphasis on symbolic manipulation

Many values that NetREXX manipulates are (from the user's point of view, at least) in the form of strings of characters. Productivity is greatly enhanced if these strings can be handled as easily as manipulating words on a page or in a text editor. NetREXX therefore has a rich set of character manipulation operators and methods, which operate on values of type `Rexx` (the name of the class of NetREXX strings).

Concatenation, the most common string operation, is treated specially in NetREXX. In addition to a conventional concatenate operator (“||”), the novel *blank operator* from REXX 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.

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 'Sorry,' n*100/50'% were rejected'
```

would display the string

```
Sorry, 12% were rejected
```

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 concatenate-with-blank, and finally the direct concatenation. Since the concatenation operators are distinct from the arithmetic operators, very natural coercion (automatic conversion) between numbers and character strings is possible. Further, explicit type-casting (conversion of types) is effected by the same operators, at the same priority, making for a very natural and consistent syntax for changing the types of results. For example,

```
i=int 100/7
```

would calculate the result of 100 divided by 7, convert that result to an integer (assuming `int` describes an integer type) and then assign it to the variable `i`.

1.2.4 Nothing to declare

Consistent with the philosophy of simplicity, NetREXX does not require that variables within methods be declared before use. Only the *properties*⁹ of classes – which may form part of their interface to other classes – need be listed formally. Within methods, the type of variables is deduced statically from context, which saves the programmer the menial task of stating the type explicitly. Of course, if preferred, variables may be listed and assigned a type at the start of each method.

1.2.5 Environment independence

The core NetREXX language is independent of both operating systems and hardware. NetREXX programs, though, must be able to interact with their environment, which implies some dependence on that environment (for example, binary representations of numbers may be required). Certain areas of the language are therefore described as being defined by the environment.

Where environment-independence is defined, however, there may be a loss of efficiency – though this can usually be justified in view of the simplicity and portability gained.

As an example, character string comparison in NetREXX is normally independent of case and of leading and trailing blanks. (The string “ Yes ” *means* the same as “yes” in most applications.) However, the influence of underlying hardware has often subtly affected this kind of design decision, so that many languages only allow trailing blanks but not leading blanks, and insist on exact case matching. By contrast, NetREXX provides the human-oriented relaxed comparison for strings as default, with optional “strict comparison” operators.

1.2.6 Limited span syntactic units

The fundamental unit of syntax in the NetREXX 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 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.

⁹Class variables and instance variables.

1.2.7 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.

These constraints make programming language design more of an art than a science, if the usability of the language is a primary goal. The problems are further compounded for NetREXX because the language is suitable for both scripting (where rapid development and ease of use are paramount) and for application development (where some programmers prefer extensive checking and redundant coding). These conflicting goals are balanced in NetREXX by providing automatic handling of many tasks (such as conversions between different representations of strings and numbers) yet allowing for “strict” options which, for example, may require that all types be explicit, identifiers be identical in case as well as spelling, and so on.

1.2.8 Be adaptable

Wherever possible NetREXX allows for the extension of instructions and other language constructs, building on the experience gained with REXX. For example, there is a useful set of common characters available for future use, since only small set is used for the few special notations in the language.

Similarly, the rules for keyword recognition allow instructions to be added whenever required without compromising the integrity of existing programs. There are no reserved keywords in NetREXX; variable names chosen by a programmer always take precedence over recognition of keywords. This ensures that NetREXX programs may safely be executed, from source, at a time or place remote from their original writing – even if in the meantime new keywords have been added to the language.

A language needs to be adaptable because *it certainly will be used for applications not foreseen by the designer*. Like all programming languages, NetREXX may (indeed, probably will) prove inadequate for certain future applications; room for expansion and change is included to make the language more adaptable and

robust.

1.2.9 Keep the language small

NetREXX is designed as a small language. It is not the sum of all the features of REXX and of Java; rather, unnecessary features have been omitted. 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
- there are few exceptions, special cases, or rarely used embellishments
- the language is easier to implement.

Many languages have accreted “neat” features which make certain algorithms easier to express; analysis shows that many of these are rarely used. As a rough rule-of-thumb, features that simply provided alternative ways of writing code were added to REXX and NetREXX only if they were likely to be used more often than once in five thousand clauses.

1.2.10 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 a few *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 arbitrary implementation limits are necessary, the language requires that the implementer use familiar and memorable decimal values for the limits. For example 250 would be used in preference to 255, 500 to 512, and so on.

1.3 Acknowledgements

Much of NetREXX is based on earlier work, and I am indebted to the hundreds of people who contributed to the development of REXX, Object REXX, and Java.

In the 1990s I gained many insights from the deliberations of the members of the X3J18 technical committee, which, under the remarkable chairmanship of Brian Marks, led to the 1996 ANSI Standard for REXX. Many of the committee’s suggestions are incorporated in NetREXX.

Equally important have been the comments and feedback from the pioneering users of NetREXX, and all those people who sent me comments on the language either directly or in the NetREXX mailing list or forum. I would especially like to thank Ian Brackenbury, Barry Feigenbaum, Davis Foulger, Norio Furukawa,

Dion Gillard, Martin Lafaix, Max Marsiglietti, and Trevor Turton for their insightful comments and encouragement.

I also thank IBM; my appointment as an IBM Fellow made it possible to make the implementation of NetREXX a reality in months rather than years. IBM has also donated the NetREXX implementation to the REXX Language Association, with special thanks due to Matthew Emmons for piloting NetREXX through the convoluted legal and other processes, and to René Jansen for massaging the NetREXX reference implementation into shape for its Open Source release.

Finally, this document has relied on old but trusted technology for its creation: its GML markup was processed using macros originally written by Bob O'Hara, and formatted using SCRIPT/VS, the IBM Document Composition Facility. Geoff Bartlett provided critical advice on character sets and fonts for the NetREXX book. This version uses a set of REXX programs to translate that same GML markup to OpenOffice Document Text format (XML files).

Mike Cowlshaw, 1997 and 2009

Introduction to the current edition

After the open sourcing of the NetREXX reference implementation in 2011 the REXXLA NetREXX ARB (Architecture Review Board), in which Mike Cowlshaw takes part as Language Architect, took responsibility for the definition of the language. Starting from version 3.00, changes in the language definition¹⁰ in this publication will be marked with the introducing release number, in the form of a margin note. 3.00

For this version of the NetREXX Language Reference, a NetREXX program was used to translate the original GML markup to \LaTeX . This edition describes the 4.02-GA version of the language and supercedes all earlier versions. The previously included chapter “A Quick Tour of the NetRexx Language” can now be exclusively found in the *NetRexx Quickstart Guide*. 3.07

Version 4, a new major version number of the language translator, signifies a major milestone in the development of this implementation. Thanks to Marc Remes, NetREXX now supports the Java Platform Module System (JPMS), which enables it to compile (or interpret) programs on current JDK versions. 4.01

NetREXX 4 depends on JSR 203 (NIO.2) and thus requires a minimum JDK level of Java 7, whereas NetREXX 3 runs on Java 6. NetREXX 4 compiles and runs on Java 7/8 (without JPMS) and on Java 9+¹¹ (with JPMS) including the most recent versions.

René Vincent Jansen, January 24, 2022

¹⁰This publication is traditionally known as NRL, short for NetRexx Language Reference. This title however, has (for reasons of clarity for new users) been changed in the filename of the PDF version of the book in favour of a longer and more descriptive name.

¹¹JDK 9,11,13,15 and 17 have been tested.

NetREXX Language Definition

This part of the document describes the NetREXX language, version 4.02-GA. This version includes the original NetREXX language reference¹² together with additions made from 1997 through 2000 and previously published in the *NetREXX Language Supplement*.

The language is described first in terms of the characters from which it is composed and its low-level syntax, and then progressively through more complex constructions. Finally, special sections describe the semantics of the more complicated areas.

Some features of the language, such as options keywords and binary arithmetic, are implementation-dependent. Rather than leaving these important aspects entirely abstract, this description includes summaries of the treatment of such items in the *reference implementation* of NetREXX. The reference implementation is based on the Java environment and class libraries.

Paragraphs that refer to the reference implementation, and are therefore not strictly part of the language definition, are shown in italics, like this one.

¹²The NetREXX Language, M. F. Cowlishaw, ISBN 0-13-806332-X, Prentice-Hall, 1997

Notations

In this part of the book, various notations such as changes of font are used for clarity. Within the text, a sans-serif bold font is used to indicate keywords, and an italic font is used to indicate *technical terms*. An italic font is also used to indicate a reference to a *technical term defined elsewhere* or a *word* in a syntax diagram that names a segment of syntax.

Similarly, in the syntax diagrams in this book, words (symbols) in the sans-serif bold font also denote keywords or sub-keywords, and words (such as *expression*) in italics denote a token or collection of tokens defined elsewhere. The brackets [and] delimit optional (and possibly alternative) parts of the instructions, whereas the braces { and } indicate that one of a number of alternatives must be selected. An ellipsis (...) following a bracket indicates that the bracketed part of the clause may optionally be repeated.

Occasionally in syntax diagrams (*e.g.*, for indexed references) brackets are "real" (that is, a bracket is required in the syntax; it is not marking an optional part). These brackets are enclosed in single quotes, thus: '[' or ']'.

Note that the keywords and sub-keywords in the syntax diagrams are not case-sensitive: the symbols "IF" "If" and "iF" would all match the keyword shown in a syntax diagram as *if*.

Note also that most of the clause delimiters (";") shown can usually be omitted as they will be implied by the end of a line.

Characters and Encodings

In the definition of a programming language it is important to emphasize the distinction between a *character* and the *coded representation*¹³ (encoding) of a character. The character "A", for example, is the first letter of the English (Roman) alphabet, and this meaning is independent of any specific coded representation of that character. Different coded character sets (such as, for example, the ASCII¹⁴ and EBCDIC¹⁵ codes) use quite different encodings for this character (decimal values 65 and 193, respectively). Except where stated otherwise, this book uses characters to convey meaning and not to imply a specific character code (the exceptions are certain operations that specifically convert between characters and their representations). At no time is NetREXX concerned with the glyph (actual appearance) of a character.

5.1 Character Sets

Programming in the NetREXX language can be considered to involve the use of two character sets. The first is used for expressing the NetREXX program itself, and is the relatively small set of characters described in the next section. The second character set is the set of characters that can be used as character data by a particular implementation of a NetREXX language processor. This character set may be limited in size (sometimes to a limit of 256 different characters, which have a convenient 8-bit representation), or it may be much larger. The *Unicode*¹⁶ character set, for example, allows for 1,114,112 code points, of which currently 128,000 are defined as characters. These are represented, depending on the serialization format, in one to four bytes.

Usually, most or all of the characters in the second (data) character set are also allowed within a NetREXX program, but only within commentary or immediate (literal) data. The NetREXX language explicitly defines the first character set, in order that programs will be portable and understandable; at the same time it avoids restrictions due to the language itself on the character set used for data. However, where the language itself manipulates or inspects the data (as when carrying out arithmetic operations), there may be requirements on the data character set (for example, numbers can only be expressed if there are

¹³These terms have the meanings as defined by the International Organization for Standardization, in ISO 2382 :cit.Data processing - Vocabulary:ecit..

¹⁴ American Standard Code for Information Interchange.

¹⁵ Extended Binary Coded Decimal Interchange Code.

¹⁶*The Unicode Standard, version 6.0.*, The Unicode Consortium, Mountain View, 2011, ISBN 09781936213016.

digit characters in the set).

Structure and General Syntax

A NetREXX program is built up out of a series of *clauses* that are composed of: zero or more blanks (which are ignored); a sequence of tokens (described in this section); zero or more blanks (again ignored); and the delimiter ";" (semicolon) which may be implied by line-ends or certain keywords. Conceptually, each clause is scanned from left to right before execution and the tokens composing it are resolved.

Identifiers (known as symbols) and numbers are recognized at this stage, comments (described below) are removed, and multiple blanks (except within literal strings) are reduced to single blanks. Blanks adjacent to operator characters (see page 20) and special characters (see page 20) are also removed.

6.1 Blanks and White Space

Blanks (spaces) may be freely used in a program to improve appearance and layout, and most are ignored. Blanks, however, are usually significant

- within literal strings (see below)
- between two tokens that are not special characters (for example, between two symbols or keywords)
- between the two characters forming a comment delimiter
- immediately outside parentheses ("(" and ")") or brackets ("[" and "]").

For implementations that support tabulation (tab) and form feed characters, these characters outside of literal strings are treated as if they were a single blank; similarly, if the last character in a NetREXX program is the End-of-file character (EOF, encoded in ASCII as decimal 26), that character is ignored.

6.2 Comments

Commentary is included in a NetREXX program by means of *comments*. Two forms of comment notation are provided: *line comments* are ended by the end of the line on which they start, and *block comments* are typically used for more extensive commentary.

Line comments A line comment is started by a sequence of two adjacent hyphens ("--"); all characters following that sequence up to the end of the line are then ignored by the NetREXX language processor.

Example:

```
i=j+7 -- this line comment follows an assignment
```

Block comments A block comment is started by the sequence of characters `"/**"`, and is ended by the same sequence reversed, `**/"`. Within these delimiters any characters are allowed (including quotes, which need not be paired). Block comments may be nested, which is to say that `"/**"` and `**/"` must pair correctly. Block comments may be anywhere, and may be of any length. When a block comment is found, it is treated as though it were a blank (which may then be removed, if adjacent to a special character).

Example:

```
/* This is a valid block comment */
```

The two characters forming a comment delimiter (`"/**"` or `**/"`) must be adjacent (that is, they may not be separated by blanks or a line-end).

Note: It is recommended that NetREXX programs start with a block comment that describes the program. Not only is this good programming practice, but some implementations may use this to distinguish NetREXX programs from other languages. **Implementation minimum:** Implementations should support nested block comments to a depth of at least 10. The length of a comment should not be restricted, in that it should be possible to "comment out" an entire program.

6.3 Tokens

The essential components of clauses are called *tokens*. These may be of any length, unless limited by implementation restrictions,¹⁷ and are separated by blanks, comments, ends of lines, or by the nature of the tokens themselves.

The tokens are:

Literal strings A sequence including any characters that can safely be represented in an implementation¹⁸ and delimited by the single quote character (`'`) or the double-quote (`"`). Use `""` to include a `"` in a literal string delimited by `"`, and similarly use two single quotes to include a single quote in a literal string delimited by single quotes. A literal string is a constant and its contents will never be modified by NetREXX. Literal strings must be complete on a single line (this means that unmatched quotes may be detected on the line that they occur). Any string with no characters (*i.e.*, a string of length 0) is called a *null string*.

Examples:

¹⁷Wherever arbitrary implementation restrictions are applied, the size of the restriction should be a number that is readily memorable in the decimal system; that is, one of 1, 25, or 5 multiplied by a power of ten. 500 is preferred to 512, the number 250 is more "natural" than 256, and so on. Limits expressed in digits should be a multiple of three.

¹⁸Some implementations may not allow certain "control characters" in literal strings. These characters may be included in literal strings by using one of the escape sequences provided.

TABLE 1: Escape sequences

<code>\t</code>	the escape sequence represents a tabulation (tab) character
<code>\n</code>	the escape sequence represents a new-line (line feed) character
<code>\r</code>	the escape sequence represents a return (carriage return) character
<code>\f</code>	the escape sequence represents a form-feed character
<code>\"</code>	the escape sequence represents a double-quote character
<code>\'</code>	the escape sequence represents a single-quote character
<code>\</code>	the escape sequence represents a backslash character
<code>\-</code>	the escape sequence represents a "null" character (the character whose encoding equals zero), used to indicate continuation in a say instruction
<code>\0(zero)</code>	the escape sequence represents a "null" character (the character whose encoding equals zero); an alternative to <code>\-</code>
<code>\xhh</code>	the escape sequence represents a character whose encoding is given by the two hexadecimal digits (" hh ") following the " x ". If the character encoding for the implementation requires more than two hexadecimal digits, they are padded with zero digits on the left.
<code>\uhhhh</code>	the escape sequence represents a character whose encoding is given by the four hexadecimal digits (" hhhh ") following the " u ". It is an error to use this escape if the character encoding for the implementation requires fewer than four hexadecimal digits.

```
'Fred'
[]Ay'
"Don't Panic!"
":x"
'You shouldn't'    /* Same as "You shouldn't" */
''                /* A null string */
```

Within literal strings, characters that cannot safely or easily be represented (for example "control characters") may be introduced using an *escape sequence*. An escape sequence starts with a *backslash* ("`\`"), which must then be followed immediately by one of the following (letters may be in either uppercase or lowercase) - see table 1.

Hexadecimal digits for use in the escape sequences above may be any decimal digit (0-9) or any of the first six alphabetic characters (a-f), in either lowercase or uppercase. **Examples:**

```
'You shouldn't' /* Same as "You shouldn't" */
```

```
'\x6d\u0066\x63' /* In Unicode: 'mfc' */
'\\u005C'         /* In Unicode, two backslashes */
```

Implementation minimum: Implementations should support literal strings of at least 100 characters. (But note that the length of string expression results, *etc.*, should have a much larger minimum, normally only limited by the amount of storage available.)

Symbols Symbols are groups of characters selected from the Roman alphabet in uppercase or lowercase (A-Z, a-z), the Arabic numerals (0-9), or the

characters underscore, dollar, and euro¹⁹ ("_\$ €") Implementations may also allow other alphabetic and numeric characters in symbols to improve the readability of programs in languages other than English. These additional characters are known as *extra letters* and *extra digits*.²⁰

Examples:

```
DanYr0gof
minx
Élan
$Virtual3D
```

A symbol may include other characters only when the first character of the symbol is a digit (0-9 or an extra digit). In this case, it is a *numeric symbol* - it may include a period (".") and it must have the syntax of a number. This may be *simple number*, which is a sequence of digits with at most one period (which may not be the final character of the sequence), or it may be a *hexadecimal or binary numeric symbol* (see page 21) , or it may be a number expressed in *exponential notation*.

A number in exponential notation is a simple number followed immediately by the sequence "E" (or "e"), followed immediately by a sign ("+" or "-"),²¹ followed immediately by one or more digits (which may not be followed by any other symbol characters).

Examples:

```
1
1.3
12.007
17.3E-12
3e+12
0.03E+9
```

When *extra digits* are used in numeric symbols, they must represent values in the range zero through nine. When numeric symbols are used as numbers, any extra digits are first converted to the corresponding character in the range 0-9, and then the symbol follows the usual rules for numbers in NetREXX (that is, the most significant digit is on the left, etc.).

In the reference implementation, the extra letters are those characters (excluding A-Z, a-z, and underscore) that result in 1 when tested with `java.lang.Character.isJavaIdentifierPart`. Similarly, the extra digits are those characters (excluding 0-9) that result in 1 when tested with `java.lang.Character.isDigit`. 3.03

The meaning of a symbol depends on the context in which it is used. For example, a symbol may be a constant (if a number), a keyword, the name of a variable, or identify some algorithm.

It is recommended that the dollar and euro only be used in symbols in mechanically generated programs or where otherwise essential. **Implementation minimum:** Implementations should support symbols of at least 50

¹⁹Note that only UTF8-encoded source files can currently use the euro character.

²⁰It is expected that implementations of NetREXX will be based on Unicode or a similarly rich character set. However, portability to implementations with smaller character sets may be compromised when extra letters or extra digits are used in a program.

²¹The sign in this context is part of the symbol; it is not an operator.

characters. (But note that the length of its value, if it is a string variable, should have a much larger limit.)

Operator characters The characters + - * % | / & \ = < > are used (sometimes in combination) to indicate operations (see page 41) in expressions. A few of these are also used in parsing templates, and the equals sign is also used to indicate assignment. Blanks adjacent to operator characters are removed, so, for example, the sequences:

```
345>=123
345 >=123
345 >= 123
345 > = 123
```

are identical in meaning. Some of these characters may not be available in all character sets, and if this is the case appropriate translations may be used. **Note:** The sequences "--", "/**", and "**/" are comment delimiters, as described earlier. The sequences "++" and "\\\" are not valid in NetREXX programs.

Special characters The characters . , ;) (] [together with the operator characters have special significance when found outside of literal strings, and constitute the set of *special characters*. They all act as token delimiters, and blanks adjacent to any of these (except the period) are removed, except that a blank adjacent to the outside of a parenthesis or bracket is only deleted if it is also adjacent to another special character (unless this is a parenthesis or bracket and the blank is outside it, too). Some of these characters may not be available in all character sets, and if this is the case appropriate translations may be used.

To illustrate how a clause is composed out of tokens, consider this example:

```
'REPEAT' B + 3;
```

This is composed of six tokens: a literal string, a blank operator (described later), a symbol (which is probably the name of a variable), an operator, a second symbol (a number), and a semicolon. The blanks between the "B" and the "+" and between the "+" and the "3" are removed. However one of the blanks between the 'REPEAT' and the "B" remains as an operator. Thus the clause is treated as though written:

```
'REPEAT' B+3;
```

6.4 Implied semicolons and continuations

A semicolon (clause end) is implied at the end of each line, except if:

1. The line ends in the middle of a block comment, in which case the clause continues at the end of the block comment.
2. The last token was a hyphen. In this case the hyphen is functionally replaced by a blank, and hence acts as a *continuation character*.

This means that semicolons need only be included to separate multiple clauses on a single line.

Notes:

1. A comment is not a token, so therefore a comment may follow the continuation character on a line.
2. Semicolons are added automatically by NetREXX after certain instruction keywords when in the correct context. The keywords that may have this effect are `else`, `finally`, `otherwise`, `then`; they become complete clauses in their own right when this occurs. These special cases reduce program entry errors significantly.

6.5 The case of names and symbols

In general, NetREXX is a *case-insensitive* language. That is, the names of keywords, variables, and so on, will be recognized independently of the case used for each letter in a name; the name "**Swildon**" would match the name "**swilDon**".

NetREXX, however, uses names that may be visible outside the NetREXX program, and these may well be referenced by case-sensitive languages. Therefore, any name that has an external use (such as the name of a property, method, constructor, or class) has a defined spelling, in which each letter of the name has the case used for that letter when the name was first defined or used.

Similarly, the lookup of external names is both case-preserving and case-insensitive. If a class, method, or property is referenced by the name "**Foo**", for example, an exact-case match will first be tried at each point that a search is made. If this succeeds, the search for a matching name is complete. If it does not succeed, a case-insensitive search in the same context is carried out, and if one item is found, then the search is complete. If more than one item matches then the reference is ambiguous, and an error is reported.

Implementations are encouraged to offer an option that requires that all name matches are exact (case-sensitive), for programmers or house-styles that prefer that approach to name matching.

6.6 Hexadecimal and binary numeric symbols

A *hexadecimal numeric symbol* describes a whole number, and is of the form *nXstring*. Here, *n* is a simple number with no decimal part (and optional leading insignificant zeros) which describes the effective length of the hexadecimal string, the **X** (which may be in lowercase) indicates that the notation is hexadec-

imal, and *string* is a string of one or more hexadecimal characters (characters from the ranges "a-f", "A-F", and the digits "0-9").

The *string* is taken as a signed number expressed in *n* hexadecimal characters. If necessary, *string* is padded on the left with "0" characters (note, not "sign-extended") to length *n* characters.

If the most significant (left-most) bit of the resulting string is zero then the number is positive; otherwise it is a negative number in twos-complement form. In both cases it is converted to a NetRexx number which may, therefore, be negative. The result of the conversion is a number comprised of the Arabic digits 0-9, with no insignificant leading zeros but possibly with a leading "-".

The value *n* may not be less than the number of characters in *string*, with the single exception that it may be zero, which indicates that the number is always positive (as though *n* were greater than the the length of *string*).

Examples:

```
1x8    == -8
2x8    == 8
2x08   == 8
0x08   == 8
0x10   == 16
0x81   == 129
2x81   == -127
3x81   == 129
4x81   == 129
04x81  == 129
16x81  == 129
4xF081 == -3967
8xF081 == 61569
0Xf081 == 61569
```

A *binary numeric symbol* describes a whole number using the same rules, except that the identifying character is **B** or **b**, and the digits of *string* must be either **0** or **1**, each representing a single bit.

Examples:

```
1b0    == 0
1b1    == -1
0b10   == 2
0b100  == 4
4b1000 == -8
8B1000 == 8
```

Note: Hexadecimal and binary numeric symbols are a purely syntactic device for representing decimal whole numbers. That is, they are recognized only within the source of a NetRexx program, and are not equivalent to a literal string with the same characters within quotes.

Types and Classes

Programs written in the NetREXX language manipulate values, such as names, numbers, and other representations of data. All such values have an associated *type* (also known as a *signature*).

The type of a value is a descriptor which identifies the nature of the value and the operations that may be carried out on that value.

A type is normally defined by a *class*, which is a named collection of values (called *properties*) and procedures (called *methods*) for carrying out operations on the properties.

For example, a character string in NetREXX is usually of type **REXX**, which will be implemented by the class called **REXX**. The class **REXX** defines the properties of the string (a sequence of characters) and the methods that work on strings. This type of string may be the subject of arithmetic operations as well as more conventional string operations such as concatenation, and so the methods implement string arithmetic as well as other string operations.

The names of types can further be qualified by the name of a *package* where the class is held. See the package instruction for details of packages; in summary, a package name is a sequence of one or more non-numeric symbols, separated by periods. Thus, if the **REXX** class was part of the **netrexx.lang** package,²² then its *qualified type* would be **netrexx.lang.REXX**.

In general, only the class name need be specified to refer to a type. However, if a class of the same name exists in more than one known (imported) package, then the name should be qualified by the package name. That is, if the use of just the class name does not uniquely identify the class then the reference is ambiguous and an error is reported.

7.1 Primitive types

Implementations may optionally provide *primitive types* for efficiency. Primitive types are "built-in" types that are not necessarily implemented as classes. They typically represent concepts native to the underlying environment (such as 32-bit binary integer numbers) and may exhibit semantics that are different from other types. NetREXX, however, makes no syntax distinction in the names of primitive types, and assumes *binary constructors* (see page 148) exist for primitive values.

²²This is in fact where it may be found in the reference implementation.

Primitive types are necessary when performance is an overriding consideration, and so this definition will assume that primitive types corresponding to the common binary number formats are available and will describe how they differ from other types, where appropriate.

In the reference implementation, the names of the primitive types are:

```
boolean char byte short int long float double
```

*where the first two describe a single-bit value and Unicode character respectively, and the remainder describe signed numbers of various formats. The main difference between these and other types is that the primitive types are not a subclass of **Object**, so they cannot be assigned to a variable of type **Object** or passed to methods "by reference". To use them in this way, an object must be created to "wrap" them; Java provides classes for this (for example, the class **Long**).*

7.2 Dimensioned types

Another feature that is provided for efficiency is the concept of *dimensioned types*, which are types (normal or primitive) that have an associated dimension (in the sense of the dimensions of an array). Dimensioned values are described in detail in the section on *Arrays* (see page 54) .

The dimension of a dimensioned type is represented in NetREXX programs by square brackets enclosing zero or more commas, where the dimension is given by the number of commas, plus one. A dimensioned type is distinct from the type of the same name but with no dimensions.

Examples:

```
Rexx  
int  
Rexx[]  
int[,,]
```

The examples show a normal type, a primitive type, and a dimensioned version of each (of dimension 1 and 3 respectively). The latter type would result from constructing an array thus:

```
myarray=int[10,10,10]
```

That is, the dimension of the type matches the count of indexes defined for the array.

7.3 Minor and Dependent classes

A *minor class* in NetREXX is a class whose name is qualified by the name of another class, called its *parent*. This qualification is indicated by the form of the name of the class: the short name of the minor class is prefixed by the name of

its parent class (separated by a period). For example, if the parent is called Foo then the full name of a minor class Bar would be written Foo.Bar.

A *dependent class* is a minor class that has a link to its parent class that allows a child object simplified access to its parent object and its properties.

These refinements of classes and are described in the section *Minor and Dependent classes* (see page 117) .

Terms

A *term* in NetREXX is a syntactic unit which describes some value (such as a literal string, a variable, or the result of some computation) that can be manipulated in a NetREXX program.

Terms may be either *simple* (consisting of a single element) or *compound* (consisting of more than one element, with a period and no other characters between each element).

8.1 Simple terms

A simple term may be:

- A *literal string* (see page 17) - a character string delimited by quotes, which is a constant.
- A *symbol* (see page 18) . A symbol that does not begin with a digit identifies a variable, a value, or a type. One that does begin with a digit is a *numeric symbol*, which is a constant.
- A *method call* (see page 32) , which is of the form
`symbol([expression[,expression]...])`
- An *indexed reference* (see page 53) , which is of the form ²³
`symbol' '['expression[,expression]...]']'`
- An *array initializer* (see page 55) , which is of the form
`' ['expression[,expression]...]']'`
- A *sub-expression* (see page 46) , which consists of any expression enclosed within a left and a right parenthesis.

Blanks are not permitted between the symbol in a method call and the "(" , or between the symbol in an indexed reference and the "[" .

Within simple terms, method calls with no arguments (that is, with no expressions between the parentheses) may be expressed without the parentheses provided that they refer to a method in the current class (or to a static method in a class *used* by the current class) and do not refer to a constructor (see page 36) . An implementation may optionally provide a mechanism that disallows this parenthesis omission.

²³The notations '[' and ']' indicate square brackets appearing in the NetREXX program.

8.2 Compound terms

Compound terms may start with any simple term, and, in addition, may start with a qualified class name (see page 96) or a qualified constructor (see page 32). These last two both start with a package name (a sequence of non-numeric symbols separated by periods and ending in a period).

This first part of a compound term is known as the *stub* of the term. **Example stubs:**

```
"A string"
Arca
12.10
paint(g)
indexedVar[i+1]
("A" "string")
java.lang.Math          -- qualified class name
netrexx.lang.Rexx(1)   -- qualified constructor
```

All stubs are syntactically valid terms (either simple or compound) and may optionally be followed by a *continuation*, which is one or more additional non-numeric symbols, method calls, or indexed references, separated from each other and from the stub by *connectors* which are periods. **Example compound terms:**

```
"A rabbit".word(2).pos('b')
Fluffy.left(3)
12.10.max(j)
paint(g).picture
indexedVar[i+1].length
("A" "string").word(1)
java.lang.Math.PI
java.lang.Math.log(10)
```

Within compound terms, method calls with no arguments (that is, with no expressions between the parentheses) may be expressed without the parentheses provided that they do not refer to a constructor (see page 36). For example, the term:

```
Thread.currentThread().suspend()
```

could be written:

```
Thread.currentThread.suspend
```

An implementation may optionally provide a mechanism that disallows this parenthesis omission.

8.3 Evaluation of terms

Simple terms are evaluated as a whole, as described below. Compound terms are evaluated from left to right. First the stub is evaluated according to the rules detailed below. The type of the value of the stub then qualifies the next element of the term (if any) which is then evaluated (again, the exact rules are detailed below). This process is then repeated for each element in the term.

For instance, for the example above:

```
"A rabbit".word(2).pos('b')
```

the evaluation proceeds as follows:

1. The stub ("A rabbit") is evaluated, resulting in a string of type **REXX**.
2. Because that string is of type **REXX**, the **REXX** class is then searched for the **word** method. This is then invoked on the string, with argument **2**. In other words, the **word** method is invoked with the string "A rabbit" as its current context (the properties of the **REXX** class when the method is invoked reflect that value).
This returns a new string of type **REXX**, "rabbit" (the second word in the original string).
3. In the same way as before, the **pos** method of the **REXX** class is then invoked on the new string, with argument "b".
This returns a new string of type **REXX**, "3" (the position of the first "b" in the previous result).

This value, "3", is the final value of the term.

The remainder of this section gives the details of term evaluation; it is best skipped on first reading.

8.4 Simple term evaluation

All simple terms may also be used as stubs, and are evaluated in precisely the same way as stubs, as described below. For example, numeric symbols are evaluated as though they were enclosed in quotes; their value is a string of type **REXX**.

In binary classes (see page 64), however, simple terms that are strings or numeric symbols are given an implementation-defined string or primitive type respectively, as described in the section on *Binary values and operations* (see page 146)

8.5 Stub evaluation

A term's stub is evaluated according to the following rules:

- If the stub is a literal string, its value is the string, of type **REXX**, constructed from that literal.
- If the stub is a numeric symbol, its value is the string, of type **REXX**, constructed from that literal (as though the literal were enclosed in quotes).
- If the stub is an unqualified method or constructor call, or a qualified constructor call, then its value and type is the result of invoking the method identified by the stub, as described in *Methods and Constructors* (see page 32) .
- If the stub is a (non-numeric) symbol, then its value and type will be determined by whichever of the following is first found:
 1. A local variable or method argument within the current method, or a property in the current class.
 2. A method whose name matches the symbol, and takes no arguments, and that is not a constructor, in the current class. ²⁴ If the stub is part of a compound symbol, then it may also be in a superclass, searching upwards from the current class.
 3. A static or constant property, or a static method, ²⁵ whose name matches the symbol (and that takes no arguments, if a method) in a class listed in the uses phrase of the class instruction. Each class from the list is searched for a matching property or method, and then its superclasses are searched upwards from the class in the same way; this process is repeated for each of the classes, in the order specified in the list.
 4. One of the allowed special words described in *Special words and methods* (see page 121) , such as **this** or **version**.
 5. The short name of a known class or primitive type (in which case the stub has no value, just a type).
- If the stub is an indexed reference, then its value and type will be determined by whichever of the following is first found:
 1. The symbol naming the reference is an undimensioned local variable or method argument within the current method, or a property in the current class, which has type **REXX**. In this case, the reference is to an *indexed string* (see page 53) ; the expressions within the brackets must be convertible to type **REXX**, and the type of the result will be **REXX**.
 2. The symbol naming the reference is a dimensioned local variable or method argument within the current method, or a property in the current class. In this case, the reference is to an *array* (see page 54) , and the expressions within the brackets must be convertible to whole numbers allowed for array indexes. The type of the result will have the type of the array, with dimensions reduced by the number of dimensions specified in the stub. For example, if the array **foo** was of type **Baa**[,,,] and the stub referred to **foo**[1,2], then the result would be of type **Baa**[,]. It would have been an error for the stub to have specified more than four dimensions.
 3. The symbol naming the reference is the name of a static or constant property in a class listed in the uses phrase of the class instruction.

²⁴Unless parenthesis omission is disallowed by an implementation option, such as options strictargs.

²⁵Unless parenthesis omission is disallowed by an implementation option, such as options strictargs.

Each class from the list is searched in the same way as for symbols, above. The reference may be to an *indexed string* or an *array*, as for properties in the current class.

4. The symbol naming the reference is the name of a primitive type or the short name of a known class, and there are no expressions within the brackets (just optional commas). In this case, the stub describes a *dimensioned type* (see page 24).
 5. The symbol naming the reference is the name of a primitive type or is the short name of a known class, and there are one or more expressions within the brackets. In this case, the reference is to an *array constructor* (see page 54); the expressions within the brackets must be convertible to non-negative whole numbers allowed for array indexes, and the value is an array of the specified type, dimensions, and size.
- If the stub is a sub-expression, then its value and type will be determined by evaluating the *expression* (see page 41) within the parentheses.
 - If the stub starts with the name of a package, then it will either describe a qualified type (see page 23) or a qualified constructor (see page 36). Its type will be same in either case, and if a constructor then its value will be the value returned by the constructor.

If the stub is not followed by further segments, the final value and type of the term is the value and type of the stub.

8.6 Continuation evaluation

Each segment of a term's continuation is evaluated from left to right, according to the type of the evaluation of the term so far (the *continuation type*) and the syntax of the new segment:

- If the segment is a method call, then its value and type is the result of invoking the matching method in the class defining the continuation type (or a superclass or subclass of that class), as described in *Methods and Constructors* (see page 32). Note that method calls in term continuations cannot be constructors.
- If the stub is an indexed reference, then it will refer to a property in the class defining the continuation type (or a superclass of that class). That property will either be an undimensioned NetREXX string (type **REXX**) or a dimensioned array. In either case, it is evaluated in the same way as an indexed reference found in the stub, except that it is not necessarily in the current class, cannot be an array constructor, and cannot result in a dimensioned type.
- If the segment is a symbol, then it refers to either a property or a method in the class defining the continuation type (or a superclass of that class).²⁶ The search for the property or method starts with the class defining the continuation type. If a property name matches, it is used; if not, a method

²⁶Unless parenthesis omission is disallowed by an implementation option, such as options `strictargs`, in which case it can only be a property.

of the same name and having no arguments (or only optional arguments) will match. If neither property nor method is found, then the same search is applied to each of the continuation class's superclasses in turn, starting with the class that it extends.

As a convenient special case, if the property or method is not found, then if the segment is the final segment of the term and is the simple symbol **length** and the continuation type is a single-dimensioned array, then the segment evaluates to the size of the array. This will be a non-negative whole number of an appropriate primitive type (or of type **REXX** if there is no appropriate type).

The final value and type of the term is the value and type determined by the evaluation of the last segment of the continuation.

8.7 Arrays in terms

If a partially-evaluated term results in a dimensioned array (see page 54) , its type is treated as type **Object** so that evaluation of the term can continue. For example, in

```
ca=char[] "tosh"  
say ca.toString()
```

the variable **ca** is an array of characters; in the expression on the second line the method **toString()** of the class **Object** will be found. ²⁷

²⁷In the reference implementation, this would return an identifier for the object.

Methods and Constructors

Instructions in NetRexx are grouped into *methods*, which are named routines that always belong to (are part of) a *class*.

Methods are invoked by being referenced in a term (see page 26), which may be part of an expression or be a clause in its own right (a method call instruction). In either case, the syntax used for a method invocation is:

```
symbol( [expression[, expression]... ] )
```

The *symbol*, which must be non-numeric, is called the *name* of the method. It is important to note that the name of the method must be followed immediately by the "(" , with **no** blank in between, or the construct will not be recognized as a method call (a *blank operator* would be assumed at that point instead). The *expressions* (separated by commas) between the parentheses are called the *arguments* to the method. Each argument expression may include further method calls. The argument expressions are evaluated in turn from left to right and the resulting values are then passed to the method (the procedure for locating the method is described below). The method then executes some algorithm (usually dependent on any arguments passed, though arguments are not mandatory) and will eventually return a value. This value is then included in the original expression just as though the entire method reference had been replaced by the name of a variable whose value is that returned data.

For example, the **substr** method is provided for strings of type **REXX** and could be used as:

```
c='abcdefghijk'  
a=c.substr(3,7)  
/* would set A to "cdefghi" */
```

Here, the value of the variable **c** is a string (of type **REXX**). The **substr** (substring) method of the **REXX** class is then invoked, with arguments **3** and **7**, on the value referred to by **c**. That is, the the properties available to (the context of) the **substr** method are the properties constructed from the literal string **'abcdefghijk'**. The method returns the substring of the value, starting at the third character and of length seven characters.

A method may have a variable number of arguments: only those required need be specified. For example, **'ABCDEF'.substr(4)** would return the string **'DEF'**, as the **substr** method will assume that the remainder of the string is to be returned if no length is provided.

Method invocations that take no arguments may omit the (empty) parentheses in circumstances where this would not be ambiguous. See the section on *Terms* (see page 26) for details.

Implementation minimum: At least 10 argument expressions should be allowed in a method call.

9.1 Method call instructions

When a clause in a method consists of just a term, and the final part of the term is a method invocation, the clause is a *method call instruction*:

```
symbol( [expression[, expression]... ] );
```

The method is being called as a subroutine of the current method, and any returned value is discarded. In this case (and in this case only), the method invoked need not return a value (that is, the `return` instruction that ends it need not specify an expression).²⁸

A method call instruction that is the first instruction in a constructor (see below) can only invoke the special constructors **this** and **super**.

9.2 Method resolution (search order)

Method resolution in NetREXX proceeds as follows:

- If the method invocation is the first part (stub) of a term, then:
 1. The current class is searched for the method (see below for details of searching).
 2. If not found in the current class, then the superclasses of the current class are searched, starting with the class that the current class extends.
 3. If still not found, then the classes listed in the `uses` phrase of the `class` instruction are searched for the method, which in this case must be a static method (see page 84). Each class from the list is searched for the method, and then its superclasses are searched upwards from the class; this process is repeated for each of the classes, in the order specified in the list.
 4. If still not found, the method invocation must be a constructor (see below) and so the method name, which may be qualified by a package name, should match the name of a primitive type or a known class (type). The specified class is then searched for a constructor that matches the method invocation.

²⁸A method call instruction is equivalent to the `call` instruction of other languages, except that no keyword is required.

- If the method invocation is not the first part of the term, then the evaluation of the parts of the term to the left of the method invocation will have resulted in a value (or just a type), which will have a known type (the continuation type). Then:
 1. The class that defines the continuation type is searched for the method (see below for details of searching).
 2. If not found in that class, then the superclasses of that class are searched, starting with the class that that class extends.

If the search did not find a method, an error is reported.

If the search did find a method, that is the method which is invoked, except in one case:

- If the evaluation so far has resulted in a value (an object), then that value may have a type which is a subclass of the continuation type. If, within that subclass, there is a method that exactly overrides (see page 35) the method that was found in the search, then the method in the subclass is invoked.

This case occurs when an object is earlier assigned to a variable of a type which is a superclass of the type of the object. This type simplification hides the real type of the object from the language processor, though it can be determined when the program is executed.

Searching for a method in a class proceeds as follows:

1. Candidate methods in the class are selected. To be a candidate method:
 - the method must have the same name as the method invocation (independent of the case (see page 21) of the letters of the name)
 - the method must have the same number of arguments as the method invocation (or more arguments, provided that the remainder are shown as optional in the method definition)
 - it must be possible to assign the result of each argument expression to the type of the corresponding argument in the method definition (if strict type checking is in effect, the types must match exactly).
2. If there are no candidate methods then the search is complete; the method was not found.
3. If there is just one candidate method, that method is used; the search is complete.
4. If there is more than one candidate method, the sum of the costs of the conversions (see page 40) from the type of each argument expression to the type of the corresponding argument defined for the method is computed for each candidate method.
5. The costs of those candidates (if any) whose names match the method invocation exactly, including in case, are compared; if one has a lower cost than all others, that method is used and the search is complete.
6. The costs of all the candidates are compared; if one has a lower cost than all others, that method is used and the search is complete.
7. If there remain two or more candidates with the same minimum cost, the method invocation is ambiguous, and an error is reported.

Note: When a method that is not an exact match to a call is found in a class, superclasses of that class are also searched for methods which may have a lower-cost of conversion and the method with the lowest cost, hence the closest match, is used to resolve the search.

The current method of method resolution has been chosen to maximize interoperability with Java-language programs.²⁹

9.3 Method overriding

A method is said to *exactly override* a method in another class if

1. the method in the other class has the same name as the current method
2. the method in the other class is not private
3. the other class is a superclass of the current class, or is a class that the current class implements (or is a superclass of one of those classes).
4. the number and type of the arguments of the method in the other class exactly match the number and type of the arguments of the current method (where subsets are also checked, if either method has optional arguments).

For example, the **Rexx** class includes a **substr** (see page 171) method, which takes from one to three strings of type **Rexx**. In the class:

```
class mystring extends Rexx
  method substr(n=Rexx, length=Rexx)
    return this.reverse.substr(n, length)

  method substr(n=int, length=int)
    return this.reverse.substr(Rexx n, Rexx length)
```

the first method exactly overrides the **substr** method in the **Rexx** class, but the second does not, because the types of the arguments do not match.

A method that exactly overrides a method is assumed to be an extension of the overridden method, to be used in the same way. For such a method, the following rules apply:

- It must return a value of the same type³⁰ as the overridden method (or none, if the overridden method returns none).
- It must be at least as visible as the overridden routine. For example, if the overridden routine is public then it must also be public.
- If the overridden method is static then it must also be static.
- If the overridden method is not static then it must not be static.
- If the underlying implementation checks exceptions (see page 149), only those checked exceptions that are signalled by the overridden method may be left uncaught in the current method.

²⁹This in contrast to all versions before 3.02, where this rule was: *When a method is found in a class, superclasses of that class are not searched for methods, even though a lower-cost method may exist in a superclass.* The latter was chosen to guard the program optimally against changes in superclasses.

³⁰observing what is stated in the next paragraph

9.4 Return Types

- 3.02 NetREXX allows covariant return types such as have been allowed in Java since the version 1.5 release. Prior to Java 1.5, in order for a method to override or implement a method from another class, the return type of the methods had to be an exact match. Since the Java 1.5 release, methods which override a superclass method or implement an interface class method are allowed to have a return type which is a subclass of the return type of the method replaced or implemented. An exact match is no longer required.

9.5 Constructor methods

As described above, methods are usually invoked in the context of an existing value or type. A special kind of method, called a constructor method, is used to actually create a value of a given type (an object).

Constructor methods always have the same short name as the class in which they are found, and construct and return a value of the type defined by that class (sometimes known as an *instance* of that class). If the class is part of a package, then the constructor call may be qualified by the package name.

Example constructors:

```
File('Dan.yr.Ogof')
java.io.File('Speleogroup.letter')
Rexx('some words')
netrexx.lang.Rexx(1)
```

There will always be at least one constructor if values can be created for a class. NetREXX will add a default public constructor that takes no arguments if no constructors are provided, unless the components of the class are all static or constant, or the class is an interface class.

All constructors follow the same rules as other methods, and in addition:

1. Constructor calls always include parentheses in the syntax, even if no arguments are supplied. This distinguishes them from a reference to the type of the same name.
2. Constructors must call a constructor of their superclass (the class they extend) before they carry out any initialization of their own. This is so any initialization carried out by the superclass takes place, and at the appropriate moment. Only after this call is complete can they make any reference to the special words **this** or **super** (see page 121) .

Therefore, the first instruction in a constructor must be either a call to the superclass, using the special constructor **super()** (with optional arguments), or a call to to another constructor in the same class, using the special constructor **this()** (with optional arguments). In the latter case, eventually a constructor that explicitly calls **super()** will be invoked and the chain of local constructor calls ends.

As a convenience, NetREXX will add a default call to **super()**, with no arguments, if the first instruction in a constructor is not a call to **this()** or **super()**.

3. The properties of a constructed value are initialized, in the order given in the program, after the call to **super()** (whether implicit or explicit).
4. By definition, constructors create a value (object) whose type is defined by the current class, and then return that value for use. Therefore, the **return** keyword on the **method** instruction (see page 82) that introduces the constructor is optional (if given, the type specified must be that of the class). Similarly, the only possible forms of the **return** instruction used in a constructor are either "**return this;**", which returns the value that has just been constructed, or just "**return;**", in which case, the "**this**" is assumed (this form will be assumed at the end of a method, as usual, if necessary).

Here is an example of a class with two constructors, showing the use of **this()** and **super()**, and taking advantage of some of the assumptions:

```
class MyChars extends SomeClass

  properties private
    /* the data 'in' the object */
    value=char[]

  /* construct the object from a char array */
  method MyChars(array=char[])
    /* initialize superclass */
    super()
    value=array          -- save the value

  /* construct the object from a String */
  method MyChars(s=String)
    /* convert to char[] and use the above */
    this(s.toCharArray())
```

Objects of type **MyChars** could then be created thus:

```
myvar=MyChars("From a string")
```

or by using an argument that has type **char[]**.

Type conversions

As described in the section on *Types and classes* (see page 23), all values that are manipulated in NetREXX have an associated type. On occasion, a value involved in some operation may have a different type than is needed, and in this case conversion of a value from one type to another is necessary.

NetREXX considerably simplifies the task of programming, without losing robustness, by making many such conversions automatic. It will automatically convert values providing that there is no loss of information caused by the automatic conversion (or if there is, an exception would be raised).

Conversions can also be made explicit by concatenating a type to a value (see page 44) and in this case less robust conversions (sometimes known as *casts*) may be effected. See below for details of the permitted automatic and explicit conversions.

Almost all conversions carry some risk of failure, or have a performance cost, and so it is expected that implementations will provide options to either report costly conversions or require that programmers make all conversions explicit.³¹ Such options might be recommended for certain critical programming tasks, but are not necessary for general programming.

10.1 Permitted automatic conversions

In general, the semantics of a type is unknown, and so conversion (from a *source type* to a *target type*) is only possible in the following cases:

- The target type and the source type are identical (the trivial case).
- The target type is a superclass of the source type, or is an interface class implemented by the source type. This is called *type simplification*, and in this case the value is not altered, no information is lost, and an explicit conversion may be used later to revert the value to its original type.
- The source type has a dimension, and the target type is **Object**.
- The source type is **null** and the target type is not primitive.
- The target and source types have known semantics (that is, they are "well-known" to the implementation) and the conversion can be effected without loss of information (other than knowledge of the original type). These are called *well-known conversions*.

³¹In the reference implementation, options `strictassign` may be used to disallow automatic conversions.

Necessarily, the well-known conversions are implementation-dependent, in that they depend on the standard classes for the implementation and on the primitive types supported (if any). Equally, it is this automatic conversion between strings and the primitive types of an implementation that offer the greatest simplifications of NetREXX programming.

For example, if the implementation supported an **int** binary type (perhaps a 32-bit integer) then this can safely be converted to a NetREXX string (of type **REXX**). Hence a value of type **int** can be added to a number which is a NetREXX string (resulting in a NetREXX string) without concern over the difference in the types of the two terms used in the operation.

Conversely, converting a long integer to a shorter one without checking for truncation of significant digits could cause a loss of information and would not be permitted.

In the reference implementation, the semantics of each of the following types is known to the language processor (the first four are all string types, and the remainder are known as binary numbers):

- *netrexx.lang.REXX* - the NetREXX string class
- *java.lang.String* - the Java string class
- *char* - the Java primitive which represents a single character
- *char[]* - an array of *chars*
- *boolean* - a single-bit primitive
- *byte, short, int, long*, - signed integer primitives (8, 16, 32, or 64 bits)
- *float, double* - floating-point primitives (32 or 64 bits)

Under the rules described above, the following well-known conversions are permitted:

- *REXX* to *binary number, char[], String, or char*
- *String* to *binary number, char[], REXX, or char*
- *char* to *binary number, char[], String, or REXX*
- *char[]* to *binary number, REXX, String, or char*
- *binary number* to *REXX, String, char[], or char*
- *binary number* to *binary number* (if no loss of information can take place - no sign, high order digits, decimal part, or exponent information would be lost)

Notes:

1. *Some of the conversions can cause a run-time error (exception), as when a string represents a number that is too large for an **int** and a conversion to **int** is attempted, or when a string that does not contain exactly one character is converted to a **char**.*
2. *The **boolean** primitive is treated as a binary number that may only take the values 0 or 1. A **boolean** may therefore be converted to any binary number type, as well as any of the string (or **char**) types, as the character "0" or "1". Similarly, any binary number or string can be converted to **boolean** (and must have a value of 0 or 1 for the conversion to succeed).*
3. *The **char** type is a single-character string (it is not a number that represents the encoding of the character).*

10.2 Permitted explicit conversions

Explicit conversions are permitted for all permitted automatic conversions and, in addition, when:

- The target type is a subclass of the source type, or implements the source type. This conversion will fail if the value being converted was not originally of the target type (or a subclass of the target type).
- Both the source and target types are primitive and (depending on the implementation) the conversion may fail or truncate information.
- The target type is **Rexx** or a well-known string type (all values have an explicit string representation).

10.3 Costs of conversions

All conversions are considered to have a cost, and, for permitted automatic conversions, these costs are used to select a method for execution when several possibilities arise, using the algorithm described in *Methods and Constructors* (see page 33) .

For permitted automatic conversions, the cost of a conversion from a *source type* to a *target type* will range from zero through some arbitrary positive value, constrained by the following rules:

- The cost is zero only if the source and target types are the same, or if the source type is **null** and the target type is not primitive.
- Conversions from a given primitive source type to different primitive target types should have different costs. For example, conversion of an 8-bit number to a 64-bit number might cost more than conversion of a 8-bit number to a 32-bit number.
- Conversions that may require the creation of a new object cost more than those that can never require the creation of a new object.
- Conversions that may fail (raise an exception) cost more than those that may require the creation of an object but can never fail.

Within these constraints, exact costs are arbitrary, and (because they mostly involve implementation-dependent primitive types) are necessarily implementation-dependent. The intent is that the "best performance" method be selected when there is a possibility of more than one.

Expressions and Operators

Many clauses can include *expressions*. Expressions in NetREXX are a general mechanism for combining one or more data items in various ways to produce a result, usually different from the original data. Expressions consist of one or more terms (see page 26) , such as literal strings, symbols, method calls, or sub-expressions, and zero or more *operators* that denote operations to be carried out on terms. Most operators act on two terms, and there will be at least one of these *dyadic* operators between every pair of terms.³² There are also *prefix* (monadic) operators, that act on the term that is immediately to the right of the operator. There may be one or more prefix operators to the left of any term, provided that adjacent prefix operators are different.

Evaluation of an expression is left to right, modified by parentheses and by operator precedence (see page 46) in the usual "algebraic" manner. Expressions are wholly evaluated, except when an error occurs during evaluation.

As each term is used in an expression, it is evaluated as appropriate and its value (and the type of that value) are determined.

The result of any operation is also a value, which may be a character string, a data object of some other type, or (in special circumstances) a binary representation of a character or number. The type of the result is well-defined, and depends on the types of any terms involved in an operation and the operation carried out. Consequently, the result of evaluating any expression is a value which has a known type.

Note that the NetREXX language imposes no restriction on the maximum size of results, but there will usually be some practical limitation dependent upon the amount of storage and other resources available during execution.

11.1 Operators

The operators in NetREXX are constructed from one or more operator characters (see page 20). Blanks (and comments) adjacent to operator characters have no effect on the operator, and so the operators constructed from more than one character may have embedded blanks and comments. In addition, blank characters, where they occur between tokens within expressions but are not adjacent to another operator, also act as an operator. The operators may be subdivided into five groups: concatenation, arithmetic, comparative, logical, and type oper-

³²One operator, direct concatenation, is implied if two terms abut (see page 42) .

TABLE 2: Concatenation operators

(blank)	Concatenate terms with one blank in between.
 	Concatenate without an intervening blank.
(abuttal)	Concatenate without an intervening blank.

TABLE 3: Arithmetic operators

+	Add
-	Subtract
*	Multiply
/	Divide
%	Integer divide. Divide and return the integer part of the result.
//	Remainder. Divide and return the remainder (this is not modulo, as the result may be negative).
**	Power. Raise a number to a whole number power.
Prefix -	Same as the subtraction: "0-number".
Prefix +	Same as the addition: "0+number".

ators. The first four groups work with terms whose type is "well-known" (that is, strings, or known types that can be converted to strings without information loss). The operations in these groups are defined in terms of their operations on strings.

11.1.1 Concatenation

The concatenation operators are used to combine two strings to form one string by appending the second string to the right-hand end of the first string. The concatenation may occur with or without an intervening blank: Concatenation without a blank may be forced by using the `||` operator, but it is useful to remember that when two terms are adjacent and are not separated by an operator,³³ they will be concatenated in the same way. This is the *abuttal* operation. For example, if the variable `Total` had the value `'37.4'`, then `Total '%'` would evaluate to `'37.4%'`. Values that are not strings are first converted to strings before concatenation. The concatenation operators are listed in table 2.

11.1.2 Arithmetic

Character strings that are numbers (see page 45) may be combined using the arithmetic operators listed in table 3. The section on *Numbers and Arithmetic* (see page 136) describes numeric precision, the format of valid numbers, and the operation rules for arithmetic. Note that if an arithmetic result is shown in exponential notation, then it is likely that rounding has occurred.

³³This can occur when the terms are syntactically distinct (such as a literal string and a symbol).

TABLE 4: **Normal** comparative operators

=	Equal (numerically or when padded, <i>etc.</i>).
\=	Not equal (inverse of =).
>	Greater than.
<	Less than.
><, <>	Greater than or less than (same as "Not equal").
>=, \<	Greater than or equal to, not less than.
<=, \>	Less than or equal to, not greater than.

In binary classes (see page 64) , the arithmetic operators will use binary arithmetic if both terms involved have values which are binary numbers. The section on *Binary values and operations* (see page 146) describes binary arithmetic.

11.1.3 Comparative

The comparative operators compare two terms and return the value '1' if the result of the comparison is true, or '0' otherwise. Two sets of operators are defined: the *strict* comparisons (listed in table 5) and the *normal* comparisons (listed in table 4). The strict comparative operators all have one of the characters defining the operator doubled. The "==" and "\==" operators test for strict equality or inequality between two strings. Two strings must be identical to be considered strictly equal. Similarly, the other strict comparative operators (such as ">>" or "<<") carry out a simple left-to-right character-by-character comparison, with no padding of either of the strings being compared. If one string is shorter than, and is a leading sub-string of, another then it is smaller (less than) the other. Strict comparison operations are case sensitive, and the exact collating order may depend on the character set used for the implementation.³⁴ For all the other comparative operators, if **both** the terms involved are numeric,³⁵ a numeric comparison (in which leading zeros are ignored, *etc.*) is effected; otherwise, both terms are treated as character strings. For this character string comparison, leading and trailing blanks are ignored, and then the shorter string is padded with blanks on the right. The character comparison operation takes place from left to right, and is **not** case sensitive (that is, "Yes" compares equal to "yes"). As for strict comparisons, the exact collating order may depend on the character set used for the implementation.

The equal and not equal operators ("=", "==", "\=", and "\==") may be used to compare two objects which are not strings for equality, if the implementation allows them to be compared (usually they will need to be of the same type). The strict operators test whether the two objects are in fact the same object,³⁶

³⁴For example, in an ASCII or Unicode environment, the digits 0-9 are lower than the alphabets, and lowercase alphabets are higher than uppercase alphabets. In an EBCDIC environment, lowercase alphabets precede uppercase, but the digits are higher than all the alphabets.

³⁵That is, if they can be compared numerically without error.

³⁶Note that two distinct objects compared in this way may contain values (properties) that are identical, yet they will not compare equal as they are not the same object.

TABLE 5: **Strict** comparative operators

<code>==</code>	Strictly equal (identical).
<code>\==</code>	Strictly not equal (inverse of <code>==</code>).
<code>>></code>	Strictly greater than.
<code><<</code>	Strictly less than.
<code>>>=, \<<</code>	Strictly greater than or equal to, strictly not less than.
<code><<=, \>></code>	Strictly less than or equal to, strictly not greater than.

TABLE 6: Boolean operators

<code>&</code>	And. Returns 1 if both terms are true.
<code> </code>	Inclusive or. Returns 1 if either term is true.
<code>&&</code>	Exclusive or. Returns 1 if either (but not both) is true.
Prefix <code>\</code>	Logical not. Negates; 1 becomes 0 and <i>vice versa</i> .

and the normal operators may provide a more relaxed comparison, if available to the implementation.³⁷

In binary classes (see page 64), all the comparative operators will use binary arithmetic to effect the comparison if both terms involved have values which are binary numbers. In this case, there is no distinction between the strict and the normal comparative operators. The section on *Binary values and operations* (see page 146) describes the binary arithmetic used for comparisons.

11.1.4 Boolean

A character string is taken to have the value "false" if it is '0', and "true" if it is '1'. The logical operators take one or two such values (values other than '0' or '1' are not allowed) and return '0' or '1' as appropriate. The Boolean operators are listed in table 6. In binary classes (see page 64), the logical operators will act on all bits in the values if both terms involved have values which are boolean or integers. The section on *Binary values and operations* (see page 146) describes this in more detail.

11.1.5 Type

Several of the operators already described can be used to carry out operations related to types. Specifically:

- Any of the concatenation operators may be used for *type concatenation*, which concatenates a type to a value. All three operators (blank, "`||`", and abuttal) have the same effect, which is to convert (see page 38)³⁸ the value

³⁷In the reference implementation, the *equals* method is used for normal comparisons. Where not provided by a type, this is implemented by the *Object* class as a strict comparison.

³⁸This is sometimes known as *casting*

to the type specified (if the conversion is not possible, an error is reported or an exception is signalled). The type must be on the left-hand side of the operator. **Examples:**

```
String "abc"  
int (a+1)  
long 1  
Exception e  
InputStream myfile
```

- A type on the left hand side of an operator that could be a prefix (+, - or \) type concatenation after the prefix operator is applied to the right-hand operand, as though an explicit concatenation operator were placed before the prefix operator.

For example:

```
x=int -y
```

means that `-y` is evaluated, and then the result is converted to `int` before being assigned to `x`.³⁹ The "less than or equal" and the "greater than or equal" operators ("`<=`" and "`>=`") may be used with a type on either side of the operator, or on both sides. In this case, they test whether a value or type is a subclass of, or is the same as, a type, or vice versa. **Examples:**

```
if i<=Object then say 'I is an Object'  
if String>=i then say 'I is a String'  
if String<=Object then say 'String is an Object'
```

The precedence of these operators is not affected by their being used with types as operands.

11.2 Numbers

The arithmetic operators above require that both terms involved be numbers; similarly some of the comparative operators carry out a numeric comparison if both terms are numbers. Numbers are introduced and defined in detail in the section on *Numbers and arithmetic* (see page 136) . In summary, *numbers* are character strings consisting of one or more decimal digits optionally prefixed by a plus or minus sign, and optionally including a single period (".") which then represents a decimal point. A number may also have a power of ten suffixed in conventional exponential notation: an "E" (uppercase or lowercase) followed by a plus or minus sign then followed by one or more decimal digits defining the power of ten. Numbers may have leading blanks (before and/or after the sign, if any) and may have trailing blanks. Blanks may not be embedded among the digits of a number or in the exponential part. **Examples:**

```
'12'  
'-17.9'  
'127.0650'
```

³⁹This could also have been written `x=int (-y)`.


```
'73e+128'
' + 7.9E-5 '
'00E+000'
```

Note that the sequence **-17.9** (without quotes) in an expression is not simply a number. It is a minus operator (which may be prefix minus if there is no term to the left of it) followed by a positive number. The result of the operation will be a number. A *whole number* (see page 144) in NetREXX is a number that has a zero (or no) decimal part.

Implementation minimum: All implementations must support 9-digit arithmetic. In unavoidable cases this may be limited to integers only, and in this case the divide operator (" $/$ ") must not be supported. If exponents are supported in an implementation, then they must be supported for exponents whose absolute value is at least as large as the largest number that can be expressed as an exact integer in default precision, *i.e.*, 999999999.

11.3 Parentheses and operator precedence

Expression evaluation is from left to right; this is modified by parentheses and by operator precedence:

- When parentheses are encountered, other than those that identify method calls (see page 32), the entire *sub-expression* delimited by the parentheses is evaluated immediately when the term is required.
- When the sequence

```
term1 operator1 term2 operator2 term3
```

is encountered, and **operator₂** has a higher precedence than **operator₁**, then the operation (**term₂ operator₂ term₃**) is evaluated first. The same rule is applied repeatedly as necessary. Note, however, that individual terms are evaluated from left to right in the expression (that is, as soon as they are encountered). It is only the order of **operations** that is affected by the precedence rules.

For example, " $*$ " (multiply) has a higher precedence than " $+$ " (add), so $3+2*5$ will evaluate to 13 (rather than the 25 that would result if strict left to right evaluation occurred). To force the addition to be performed before the multiplication the expression would be written $(3+2)*5$, where the first three tokens have been formed into a sub-expression by the addition of parentheses. The order of precedence of the operators is (highest at the top) is listed in table 7.

If, for example, the symbol **a** is a variable whose value is '3', and **day** is a variable with the value 'Monday', then:

```
a+5           == '8'
a-4*2        == '-5'
a/2          == '1.5'
a\%2         == '1'
0.5**2       == '0.25'
(a+1)>7      == '0'           /* that is, False */
```

TABLE 7: Operator precedence

Prefix operators	+ - \
Power operator	**
Multiplication and division	* and /
Addition and subtraction	+ -
Concatenation	(blank) (abuttal)
Comparative operators	== > < <= >= << \>> etc.
And	&
Or, exclusive or	&&

```

' '=' == '1' /* that is, True */
' ==' == '0' /* that is, False */
' \ == '1' /* that is, True */
(a+1)*3=12 == '1' /* that is, True */
'077'>'11' == '1' /* that is, True */
'077'>>'11' == '0' /* that is, False */
'abc'>>'ab' == '1' /* that is, True */
'If it is' day == 'If it is Monday'
day.substr(2,3) == 'ond'
'!'day'!' == '!Monday!'

```

Note: The NetREXX order of precedence usually causes no difficulty, as it is the same as in conventional algebra and other computer languages. There are two differences from some common notations; the prefix minus operator always has a higher priority than the power operator, and power operators (like other operators) are evaluated left-to-right. Thus

```

-3**2 == 9 /* not -9 */
-(2+1)**2 == 9 /* not -9 */
2**2**3 == 64 /* not 256 */

```

These rules were found to match the expectations of the majority of users when the REXX language was first designed, and NetREXX follows the same rules.

Clauses and Instructions

Clauses (see page 16) are recognized, and can usefully be classified, in the following order:

Null clauses A clause that is empty or comprises only blanks, comments, and continuations is a *null clause* and is completely ignored by NetREXX (except that if it includes a comment it will be traced, if reached during execution).

Note: A null clause is not an instruction, so (for example) putting an extra semicolon after the *then* or *else* in an *if* instruction is not equivalent to putting a dummy instruction (as it would be in C or PL/I). The *nop* instruction is provided for this purpose.

Assignments Single clauses within a class and of the form *term=expression*; are instructions known as *assignments* (see page 49) . An assignment gives a variable, identified by the *term*, a type or a new value.

In just one context, where property assignments are expected (before the first method in a class), the "=" and the expression may be omitted; in this case, the term (and hence the entire clause) will always be a simple non-numeric symbol which names the property

Method call instructions A method call instruction (see page 33) is a clause within a method that comprises a single term that is, or ends in, a method invocation.

Keyword instructions A *keyword instruction* consists of one or more clauses, the first of which starts with a non-numeric symbol which is not the name of a variable or property in the current class (if any) and is immediately followed by a blank, a semicolon (which may be implied by the end of a line), a literal string, or an operator (other than "=", which would imply an assignment). This symbol, the *keyword*, identifies the instruction.

Keyword instructions control the external interfaces, the flow of control, and so on. Some keyword instructions (see page 57) (*do*, *if*, *loop*, or *select*) can include nested instructions.

Assignments and Variables

A *variable* is a named item whose value may be changed during the course of execution of a NetREXX program. The process of changing the value of a variable is called *assigning* a new value to it.

Each variable has an associated type, which cannot change during the execution of a program; therefore, the values assigned to a given variable must always have a type that can safely be assigned to that variable.

Variables may be assigned a new value by the `method` or `parse` instructions, but the most common way of changing the value of a variable is by using an *assignment instruction*. Any clause within a class and of the form:

```
assignment;  
  
where assignment is:  
  
term=expression
```

is taken to be an assignment instruction. The result of the *expression* becomes the new value of the variable named by the *term* to the left of the equals sign. When the term is simply a symbol, this is called the *name* of the variable. **Example:**

```
/* Next line gives FRED the value 'Frederic' */  
fred='Frederic'
```

The symbol naming the variable cannot begin with a digit (0-9).⁴⁰

Within a NetREXX program, variable names are not case-sensitive (for example, the names `fred`, `Fred`, and `FRED` refer to the same variable). Where public names are exposed (for example, the names of properties, classes, and methods, and in cross-reference listings) the case used for the name will be that used when the name was first introduced ("first" is determined statically by position in a program rather than dynamically).

Similarly, the type of a NetREXX variable is determined by the type of the value of the expression that is first assigned to it.⁴¹ For subsequent assignments, it is an error to assign a value to a variable with a type mismatch unless the language

⁴⁰Without this restriction on the first character of a variable name, it would be possible to redefine a number, in that for example the assignment "`3=4`;" would give a variable called "`3`" the value '`4`'.

⁴¹Since NetREXX infers the type of a variable from usage, substantial programs can be written without introducing explicit type declarations, although these are allowed.

processor can determine that the value can be assigned safely to the type of the variable.

In practice, this means that the types must match exactly, be a simplification, or both be "well-known" types such as **REXX**, **String**, **int**, *etc.*, for which safe conversions are defined. The possibilities are described in the section on *Conversions* (see page 38).⁴²

For example, if there are types (classes) called **ibm.util.hex**, **RunKnown**, and **Window**, then:

```
hexy=ibm.util.hex(3) -- 'hexy' has type 'ibm.util.hex'  
rk=RunKnown()      -- 'rk' has type 'RunKnown'  
fred=Window(10, 20) -- 'fred' has type 'Window'  
s="Los Lagos"      -- 's' has type 'R\textsc{exx}'  
j=5                -- 'j' has type 'R\textsc{exx}'
```

The first three examples invoke the *constructor* method for the type to construct a value (an object). A constructor method always has the same name as the class to which it belongs, and returns a new value of that type. Constructor methods are described in detail in *Methods and Constructors* (see page 32).

The last two examples above illustrate that, by default, the types of literal strings and numbers are NetREXX strings (type **REXX**) and so variables tend to be of type **REXX**. This simplifies the language and makes it easy to learn, as many useful programs can be written solely using the powerful **REXX** type. Potentially more efficient (though less human-oriented) primitive or built-in types for literals will be used in binary classes (see page 64). *If the examples above were in a binary class, then, in the reference implementation, the types of s and j would have been java.lang.String and int respectively.*

A variable may be introduced ("declared") without giving it an initial value by simply assigning a type to it:

```
i=int  
r=R\textsc{exx}  
f=java.io.File
```

Here, the expression to the right of the "=" simply evaluates to a type with no value.

13.1 The use and scope of variables

NetREXX variables all follow the same rules of assignment, but are used in different contexts. These are:

Properties Variables which name the values (the data) owned by an object of the type defined by the class are called *properties*. When an object is constructed by the class, its properties are created and are initialized to either a default value (**null** or, for variables of primitive type, an implementation-defined value, typically 0) or to a value provided by the programmer.

⁴²Implementations may provide for a stricter rule for assignment (where the types must be identical), controlled by the `options` instruction.

The attributes of properties can be changed by the `properties` instruction (see page 99). For example, properties may also be *constant*, which means that they are initialized when the class is first loaded and do not change thereafter.

Method arguments When a method is invoked, arguments may be passed to it. These *method arguments* are assigned to the variables named on the method instruction (see page 82) that introduces the method.

Local variables Variables that are known only within a method are called *local variables*; each time a method is invoked a distinct set of local variables is available. Local variables are normally given an initial value by the programmer. If they are not, they are initialized to a default value (**null** or, for variables of primitive type, an implementation-defined value, typically 0).

In order for types to be determined and type-checking to be possible at “compile-time”, and easily determined by inspection, the use and type of every variable is determined by its position in the program, not by the order in which assignments are executed. That is, variable typing is static.

The visibility of a variable depends on its use. Properties are visible to all methods in a class; method arguments and local variables are only visible within the method in which they appear. In particular:

- Within a class, properties have unique names (they cannot be overridden by method arguments or by local variables within methods); this avoids error-prone ambiguity.
- Within a method, a method argument acts like a local variable (that is, it is in the same name-space as local variables, and can be assigned new values); it can be considered to be a local variable that is assigned a value just before the body of the method is executed. There cannot be both a method argument and a local variable in a method with the same name.
- Within methods, variables can take only one type, the type assigned to them when first encountered in the method (in a strict “physical” sense, that is, as parsed from top to bottom of the program and from left to right on each line). Since methods tend to be small, there is no local scoping of variables inside the constructs within a method.⁴³

Thus, in this example:

```
method iszero(x)
  if x=0 then qualifier='is zero'
           else qualifier='is not zero'
  say 'The argument' qualifier'.'
```

the variable **qualifier** is known throughout the method and hence has a known type and value when the `say` instruction is executed.

To summarize: a symbol that names a variable in the current class either refers to a property (and in any use of it within the class refers to that property), or it refers to a variable that is unique within a method (and any use of the name within that method refers to the same variable).

⁴³Unlike the block scoping of PL/I, C, or Java.

Note: A variable is just a name, or “handle” for a value. It is possible for more than one variable to refer to the same value, as in the program:

```
first='A string'  
second=first
```

Here, both variables refer to the same value. If that value is changeable then a change to the value referred to by one of the variable names would also be seen if the value is referred to by the other. For example, sub-values of a NetREXX string can be changed, using *Indexed references* (see page 53) , so a change to a sub-value of **first** would also be seen in an identical indexed reference to **second**.

13.2 Terms on the left of assignments

In an assignment instruction, the *term* to the left of the equals sign is most commonly a simple non-numeric symbol, which always names a variable in the current class. The other possibilities, as seen in the example below, are:

1. The term is an *indexed reference* (see page 53) , to an existing variable that refers to a string of type **REXX** or an array (see page 54) . The variable may be in the current class, or be a property in a class named in the *uses* phrase of the class instruction for the current class.
2. The term is a compound term (see page 27) that ultimately refers to a property (see above) in some class (which may be the current class). This property cannot be a constant.

Examples:

```
r=Rexx ''  
r['foo']='?'      -- indexed string assignment  
s=String[3]  
s[0]='test'      -- array assignment  
Sample.value=1   -- property assignment  
this.value=1     -- property assignment  
super.value=1    -- property assignment
```

The last two examples show assignments to a property in the current class or in a superclass of the current class, respectively. Note that references to properties in other classes must always be qualified in some way (for example, by the prefix **super**.) The use of the prefix **this** for properties in the current class is optional.

Indexed strings and Arrays

Any NetREXX string (that is, a value of type **REXX**), has the ability to have *sub-values*, values (also of type **REXX**) which are associated with the original string and are indexed by an *index string* which identifies the sub-value. Any string with such sub-values is known as an *indexed string*.

The sub-values of a NetREXX string are accessed using *indexed references*, where the name of a variable of type **REXX** is followed immediately by square brackets enclosing one or more expressions separated by commas:⁴⁴

```
symbol '['[expression[, expression]...]' ]'
```

It is important to note that the *symbol* that names the variable must be followed immediately by the "[", with **no** blank in between, or the construct will not be recognized as an indexed reference. The *expressions* (separated by commas) between the brackets are called the *indexes* to the string. These index expressions are evaluated in turn from left to right, and each must evaluate to a value is of type **REXX** or that can be converted to type **REXX**.

The resulting index strings are taken "as-is" - that is, they must match exactly in content, case, and length for a reference to find a previously-set item. They may have any length (including the null string) and value (they are not constrained to be just those strings which are numbers, for example).

If a reference does not find a sub-value, then a copy of the non-indexed value of the variable is used. **Example:**

```
surname='Unknown'      -- default value
surname['Fred']='Bloggs'
surname['Davy']='Jones'
try='Fred'
say surname[try] surname['Bert']
```

would say "**Bloggs Unknown**".

When multiple indexes are used, they indicate accessing a hierarchy of strings. A single NetREXX string has a single set of indexes and subvalues associated with it. The sub-values, however, are also NetREXX strings, and so may in turn have indexes and sub-values. When more than one index is specified in an indexed reference, the indexes are applied in turn from left to right to each retrieved sub-value.

⁴⁴The notations '[' and ']' indicate square brackets appearing in the NetREXX program.

For example, in the sequence:

```
x='?'  
x['foo', 'bar']='OK'  
say x['foo', 'bar']  
y=x['foo']  
say y['bar']
```

both `say` instructions would display the string "OK". Indexed strings may be used to set up "associative arrays", or dictionaries, in which the subscript is not necessarily numeric, and thus offer great scope for the creative programmer. A useful application is to set up a variable in which the subscripts are taken from the value of one or more variables, so effecting a form of associative (content addressable) memory. **Notes:**

1. A variable of type **REXX** must have been assigned a value before indexing is used on it. This is the value that is used as the default value whenever an indexed reference finds no sub-value.
2. The indexes, and hence the sub-values, of a **REXX** object can be retrieved in turn using the `over` (see page 78) keyword of the `loop` instruction.
3. The `exists` method (see page 166) of the **REXX** class may be used to test whether an indexed reference has an explicitly-set value.
4. Assigning `null` to an indexed reference (for example, the assignment `switch[7]=null;`) drops the sub-value; until set to a new value, any reference to the sub-value (including use of the `exists` method) will return the same result as when it had never been set.

14.1 Arrays

In addition to indexed strings, NetREXX also includes the concept of fixed-size *arrays*, which may be used for indexing values of any type (including strings).

Arrays are used with the same syntax and in the same manner as indexed strings, but with important differences that allow for compact implementations and access to equivalent data structures constructed using other programming languages:

1. The indexes for arrays must be whole numbers that are zero or positive. There will usually be an implementation restriction on the maximum value of the index (typically 999999999 or higher).
2. The elements of an array are considered to be *ordered*; the first element has index **0**, the second **1**, and so on.
3. An array is of fixed size; it must be constructed before use.
4. Variables that are assigned arrays can only be assigned arrays (of the same dimension, see below) in the future. That is, being an array changes the type of a value; it becomes a *dimensioned type* (see page 24) .

Array references use the NetREXX *indexed reference* syntax defined above. The same syntax is used for constructing arrays, except that the symbol before the

left bracket describes a type (and hence may be qualified by a package name). The expression or expressions between the brackets indicate the size of the array in each dimension, and must be a positive whole number or zero:

```
arg=String[4]      -- makes an array for four Strings
arg=java.io.File[4] -- makes an array for four Files
i=int[3]           -- makes an array for three 'int's
```

(Another way of describing this is that array constructors look just like other object constructors, except that brackets are used instead of parentheses.)

Once an array has been constructed, its elements can be referred to using brackets and expressions, as before:

```
i[2]=3 -- sets the '2'-indexed value of 'i'
j=i[2] -- sets 'j' to the '2'-indexed value of 'i'
```

Regular multiple-dimensional arrays may be constructed and referenced by using multiple expressions within the brackets:

```
i=int[2,3] -- makes a 2x3 array of 'int' type objects
i[1,2]=3   -- sets the '1,2'-indexed value of 'i'
j=i[1,2]   -- sets 'j' to the '1,2'-indexed value of 'i'
```

As with indexed strings, when multiple indexes are used, they indicate accessing a hierarchy of arrays (the underlying model is therefore of a hierarchy of single-dimensional arrays). When more than one index is specified in an indexed reference to an array, the indexes are applied in turn from left to right to each array.

As described in the section on *Types* (see page 23) , the type of a variable that refers to an array can be set (declared) by assignment of the type with array notation that indicates the dimension of an array without any sizes:

```
k=int[]      -- one-dimensional array of 'int' objects
m=float[, ,] -- 3-dimensional array of 'float' objects
```

The same syntax is also used when describing an array type in the arguments of a method instruction or when converting types. For example, after:

```
gg=char[] "Horse"
```

the variable **gg** has as its value an array of type **char[]** containing the five characters **H**, **o**, **r**, **s**, and **e**.

14.1.1 Array initializers

An *array initializer* is a *simple term* which is recognized if it does not immediately follow (abut) a symbol, and has the form ⁴⁵

```
' [ 'expression [, expression] ... ' ] '
```

An array initializer therefore comprises a list of one or more expressions, separated by commas, within brackets. When an array initializer is evaluated, the

⁴⁵The notations '[' and ']' indicate square brackets appearing in the NetRexx program.

expressions are evaluated in turn from left to right, and all must result in a value. An array is then constructed, with a number of elements equal to the number of expressions in the list, with each element initialized by being assigned the result of the corresponding expression.

The type of the array is derived by adding one dimension to the type of the result of the first expression in the list, where the type of that expression is determined using the same rules as are used to select the type of a variable when it is first assigned a value(see page 49). All the other expressions in the list must have types that could be assigned to the chosen type without error.

For example, in

```
var1=['aa', 'bb', 'cc']
var2=[char 'a', 'b', 'c']
var3=[String 'a', 'bb', 'c']
var4=[1, 2, 3, 4, 5, 6]
var5=[[1,2], [3,4]]
```

the types of the variables would be **REXX[]**, **char[]**, **String[]**, **REXX[]**, and **REXX[,]** respectively. In a binary class in the reference implementation, the types would be **String[]**, **char[]**, **String[]**, **int[]**, and **int[,]**.

Array initializers are most useful for initializing properties and variables, but like other simple terms, they may start a compound term.

So, for example

```
say [1,1,1,1].length
```

would display **4**. Note that an array of length zero cannot be constructed with an array initializer, as its type would be undefined. An explicitly typed array constructor (for example, **int[0]**) must be used.

Keyword Instructions

A *keyword instruction* is one or more clauses, the first of which starts with a keyword that identifies the instruction. Some keyword instructions affect the flow of control; the remainder just provide services to the programmer. Some keyword instructions (`do`, `if`, `loop`, or `select`) can include nested instructions. Appendix A (see page 177) includes an example of a NetREXX program using many of the instructions available. As can be deduced from the syntax rules described earlier, a keyword instruction is recognized **only** if its keyword is the first token in a clause, and if the second token is not an "=" character (implying an assignment). It would also not be recognized if the second token started with "(", "[", or "." (implying that the first token starts a term). Further, if a current local variable, method argument, or property has the same name as a keyword then the keyword will not be recognized. This important rule allows NetREXX to be extended with new keywords in the future without invalidating existing programs.

Thus, for example, this sequence in a program with no `say` variable:

```
say 'Hello'  
say('1')  
say=3  
say 'Hello'
```

would be a `say` instruction, a call to some `say` method, an assignment to a `say` variable, and an error. In NetREXX, therefore, keywords are not reserved; they may be used as the names of variables (though this is not recommended, where known in advance). Certain other keywords, known as *sub-keywords*, may be known within the clauses of individual instructions - for example, the symbols `to` and `while` in the `loop` instruction. Again, these are not reserved; if they had been used as names of variables, they would not be recognized as sub-keywords. Blanks adjacent to keywords have no effect other than that of separating the keyword from the subsequent token. For example, this applies to the blanks next to the sub-keyword `while` in

```
loop while a=3
```

Here at least one blank was required to separate the symbols forming the keywords and the variable name, `a`. However the blank following the `while` is not necessary in

```
loop while 'Me'=a
```

though it does aid readability.

Annotation instruction

@

An *annotation* starts with an @ (commercial at sign) and is passed through unchanged⁴⁶ To interpret a program with an annotation is an error.

16.1 Example

```
/* standard annotations like @Override and @Deprecated are */
/* used, as are some custom ones                               */
/* (those need to be compiled first to be used)                */
options binary
@Author(name="Class Author")
class AnnotateTest
properties private unused
propz
a = ArrayList()
test = TreeMap()

@SuppressWarnings("unchecked")
method main(args=String[]) static
  say 'hello annotations'
  t=AnnotateTest()
  t.old()

@Override
method toString() returns String
  return 'Annotations'

@Deprecated
method old() /* a comment with an @ in it */
  say 'do no use anymore'

@Author(name = "Jane Doe")
@Author(name = "John Doe")
method repeating()
  say 'repeating annotations'

@Author( name = "Fifi the Cat", date = "2016-01-01" )
```

⁴⁶dependent on the setting of option -annotations, which is the default. When option -noannotations is in effect, no annotations are passed through. In this case, no @SuppressWarnings("unchecked") annotations are generated on methods, which might lead to (harmless) javac warnings.

```
method parameters()  
  say 'parameters are possible, but all on one line'
```

Address instruction

```
address [environment] [expression]
```

where *environment* is a native executable or script on the path

The keyword *address* temporarily or permanently changes the destination of commands. Commands are strings sent to an external environment. You can send commands by specifying clauses consisting of only an expression or by using the Address instruction.

To send a single command to a specified environment, code an environment, a literal string or a single symbol, which is taken to be a constant, followed by an expression.

The environment name is the name of an external procedure or process that can process commands.

The expression is evaluated to produce a character string value, and this string is routed to the environment to be processed as a command.

After execution of the command, environment is set back to its original state, thus temporarily changing the destination for a single command. The default environment is SYSTEM, which is the shell on most operating systems.

After execution, the most recent return code is in the variable RC.

Moreover, any expression which is not intercepted by the translator as a Net-Rexx clause, is sent to the address environment. The environment sends the expression to STDIN of the external process or procedure for execution. Any output received on STDOUT is printed on the console.

Address processing can be switched off by the `noaddress` option.

Example:

```
'echo "Hello world"'  
say RC  
address bash 'echo "Hello world"'  
  
address cat  
'Hello world'  
  
address 'bash'  
'echo "Hello world"'
```

```
exitcode=4  
'exit '||exitcode  
say RC
```

The program greets the world 4 times, and shows usage of special variable RC.

Class instruction

```
class name [visibility] [modifier] [binary] [deprecated]
    [extends classname]
    [uses useslist]
    [implements interfacelist];
```

where *visibility* is one of:

```
private
public
shared
```

and *modifier* is one of:

```
abstract
adapter
final
interface
```

and *useslist* and *interfacelist* are lists of one or more *classnames*, separated by commas.

The class instruction is used to introduce a class, as described in the sections *Types and Classes* (see page 23) and *Program structure* (see page 114), and define its attributes. The class must be given a *name*, which must be different from the name of any other classes in the program. The *name*, which must be a non-numeric symbol, is known as the *short name* of the class.

A *classname* can be either the short name of a class (if that is unambiguous in the context in which it is used), or the qualified name of the class - the name of the class prefixed by a package name and a period, as described under the package instruction (see page 96).

The *body* of the class consists of all clauses following the class instruction (if any) until the next class instruction or the end of the program.

The *visibility*, *modifier*, and binary keywords, and the extends, uses, and implements phrases, may appear in any order.

18.1 Visibility

Classes may be public, private, or shared:

- A *public class* is visible to (that is, may be used by) all other classes.
- A *private class* is visible only within same program and to classes in the same package (see page 96) .
- A *shared class* is also visible only within same program and to classes in the same package.⁴⁷

A program may have only one public class, and if no class is marked public then the first is assumed to be public (unless it is explicitly marked private).

18.2 Modifier

Most classes are collections of data (properties) and the procedures that can act on that data (methods); they completely implement a datatype (type), and are permitted to be subclassed. These are called *standard classes*. The *modifier* keywords indicate that the class is not a standard class - it is special in some way. Only one of the following modifier keywords is allowed:

abstract An *abstract class* does not completely implement a datatype; one or more of the methods that it defines (or which it inherits from classes it extends or implements) is abstract - that is, the name of the method and the types of its arguments are defined, but no instructions to implement the method are provided.

Since some methods are not provided, an object cannot be constructed from an abstract class. Instead, the class must be extended and any missing methods provided. Such a subclass can then be used to construct an object.

Abstract classes are useful where many subclasses can share common data or methods, but each will have some unique attribute or attributes (data and/or methods). For example, some set of geometric objects might share dimensions in X and Y, yet need unique methods for calculating the area of the object.

adapter An *adapter class* is a class that is guaranteed to implement all unimplemented abstract methods of its superclasses and interface classes that it inherits or lists as implemented on the `class` instruction.

If any unimplemented methods are found, they will be automatically generated by the language processor. Methods generated in this way will have the same visibility and signature as the abstract method they implement, and if a return value is expected then a default value is returned (as for the initial value of variables of the same type: that is, **null** or, for values of primitive type, an implementation-defined value, typically 0). Other than

⁴⁷The `shared` keyword on the `class` instruction means exactly the same as the keyword `private`, and is accepted for consistency with the other meanings of `shared`.

possibly returning a value, these methods are empty; that is, they have no side-effects.

An adapter class provides a concrete representation of its superclasses and the interface classes it implements. As such, it is especially useful for implementing event handlers and the like, where only a small number of event-handling methods are needed but many more might be specified in the interface class that describes the event model.⁴⁸

An adapter class cannot have any abstract methods.

final A *final class* is considered to be complete; it cannot be subclassed (extended), and all its methods are considered complete.⁴⁹

interface An *interface class* is an abstract class that contains only abstract method definitions and/or constants. That is, it defines neither instructions that implement methods nor modifiable properties, and hence cannot be used to construct an object.

Interface classes are used by classes that claim to *implement* them (see the `implements` keyword, described below). The difference between abstract and interface classes is that the former may have methods which are not abstract, and hence can only be subclassed (extended), whereas the latter are wholly abstract and may only be implemented.

18.3 Binary

The keyword `binary` indicates that the class is a *binary class*. In binary classes, literal strings and numeric symbols are assigned native string or binary (primitive) types, rather than NetRExx types, and native binary operations are used to implement operators where possible. When `binary` is not in effect (the default), terms in expressions are converted to NetRExx types before use by operators. The section *Binary values and operations* (see page 146) describes the implications of binary classes in detail.

Individual methods in a class which is not binary can be made into *binary methods* using the `binary` keyword on the method instruction (see page 82).

18.4 Deprecated

The keyword `deprecated` indicates that the class is *deprecated*, which implies that a better alternative is available and documented. A compiler can use this information to warn of out-of-date or other use that is not recommended.

⁴⁸For example, see the "Scribble" sample in the NetRExx package.

⁴⁹This modifier is provided for consistency with other languages, and may allow compilers to improve the performance of classes that refer to the final class. In many cases it will reduce the reusability of the class, and hence should be avoided.

18.5 Extends

Classes form a hierarchy, with all classes (except the top of the tree, the **Object**⁵⁰ class) being a *subclass* of some other class. The `extends` keyword identifies the *classname* of the immediate *superclass* of the new class - that is, the class immediately above it in the hierarchy. If no `extends` phrase is given, the superclass is assumed to be **Object** (or **null**, in the case where the current class is **Object**).

18.6 Uses

The `uses` keyword introduces a list of the names of one or more classes that will be used as a source of constant (or static) properties and/or methods.

When a term (see page 26) starts with a symbol, method call, or indexed reference that is not known in the current context, each class in the *useslist* and its superclasses are searched (in the order specified in the *useslist*) for a constant or static method or property that matches the item. If found, the method or property is used just as though explicitly qualified by the name of the class in which it was found.

The `uses` mechanism affects only the syntax of terms in the current class; it is not inherited by subclasses of the current class.

18.7 Implements

The `implements` keyword introduces a list of the names of one or more interface classes (see above). These interface classes are then known to (inherited by) the current class, in the order specified in the *interfacelist*. Their methods (which are all abstract) and constant properties act as though part of the current class, unless they are overridden (hidden) by a method or constant of the same name in the current class.

If the current class is not an interface class then it must implement (provide non-abstract methods for) all the methods inherited from the interface classes in the `implements` list.

Interface classes, therefore, can be used to:

1. Define a common set of methods (possibly with associated constants) that will be implemented by other classes.
2. Conveniently package collections of constants for use by other classes.

The `implements` list may not include the superclass of the current class.

⁵⁰In the reference implementation, `java.lang.Object`.

Do instruction

```
do [label name] [protect term] [binary];
    instructionlist
    [catch [vare =] exception;
        instructionlist]...
    [finally[;]
        instructionlist]
end [name];
```

where *name* is a non-numeric *symbol*

and *instructionlist* is zero or more *instructions*

The do instruction is used to group instructions together for execution; these are executed once. The group may optionally be given a label, and may protect an object while the instructions in the group are executed; exceptional conditions can be handled with catch and finally.

The most common use of do is simply for treating a number of instructions as group.

Example:

```
/* The two instructions between DO and END will both */
/* be executed if A has the value 3.                  */
if a=3 then do
    a=a+2
    say 'Smile!'
end
```

Here, only the first *instructionlist* is used. This forms the *body* of the group.

The instructions in the *instructionlists* may be any assignment, method call, or keyword instruction, including any of the more complex constructions such as loop, if, select, and the do instruction itself.

19.0.1 Label phrase

If label is used to specify a *name* for the group, then a leave which specifies that name may be used to leave the group, and the end that ends the group may optionally specify the name of the group for additional checking.

Example:

```
do label sticky
  x=ask
  if x='quit' then leave sticky
  say 'x was' x
end sticky
```

19.0.2 Protect phrase

If `protect` is given it must be followed by a *term* that evaluates to a value that is not just a type and is not of a primitive type; while the `do` construct is being executed, the value (object) is protected - that is, all the instructions in the `do` construct have exclusive access to the object.

Both `label` and `protect` may be specified, in any order, if required.

19.0.3 Exceptions in do groups

Exceptions that are raised by the instructions within a `do` group may be caught using one or more `catch` clauses that name the *exception* that they will catch. When an exception is caught, the exception object that holds the details of the exception may optionally be assigned to a variable, *vare*.

Similarly, a `finally` clause may be used to introduce instructions that will always be executed at the end of the group, even if an exception is raised (whether caught or not).

The *Exceptions* section (see page 149) has details and examples of `catch` and `finally`.

19.0.4 Binary

A group of one or more statements in a `do binary` group will follow the semantics of binary statements in binary classes or methods; the scope is limited to the `do binary` group. 3.04

Exit instruction

```
exit [expression];
```

`exit` is used to unconditionally leave a program, and optionally return a result to the caller. The entire program is terminated immediately.

If an *expression* is given, it is evaluated and the result of the evaluation is then passed back to the caller in an implementation-dependent manner when the program terminates. Typically this value is expected to be a small whole number; most implementations will accept values in the range 0 through 250. If no expression is given, a default result (which depends on the implementation, and is typically zero) is passed back to the caller.

Example:

```
j=3  
exit j*4  
/* Would exit with the value '12' */
```

“Running off the end” of a program is equivalent to the instruction **return**; In the case where the program is simply a stand-alone application with no `class` or `method` instructions, this has the same effect as **exit**, in that it terminates the whole program and returns a default result.

If instruction

```
if expression[;]
  then[;] instruction
  [else[;] instruction]
```

The `if` construct is used to conditionally execute an instruction or group of instructions. It can also be used to select between two alternatives. The expression is evaluated and must result in either 0 or 1. If the result was 1 (true) then the instruction after the `then` is executed. If the result was 0 (false) and an `else` was given then the instruction after the `else` is executed. **Example:**

```
if answer='Yes' then say 'OK!'
  else say 'Why not?'
```

Remember that if the `else` clause is on the same line as the last clause of the `then` part, then you need a semicolon to terminate that clause. **Example:**

```
if answer='Yes' then say 'OK!'; else say 'Why not?'
```

The `else` binds to the nearest `then` at the same level. This means that any `if` that is used as the instruction following the `then` in an `if` construct that has an `else` clause, must itself have an `else` clause (which may be followed by the dummy instruction, `nop`). **Example:**

```
if answer='Yes' then if name='Fred' then say 'OK, Fred.'
  else say 'OK.'
  else say 'Why not?'
```

To include more than one instruction following `then` or `else`, use a grouping instruction (`do`, `loop`, or `select`). **Example:**

```
if answer='Yes' then do
  say 'Line one of two'
  say 'Line two of two'
end
```

In this instance, both `say` instructions are executed when the result of the `if` expression is 1.

Multiple expressions, separated by commas, can be given on the `if` clause, which then has the syntax:


```
if expression[, expression]... [;]
```

In this case, the expressions are evaluated in turn from left to right, and if the result of any evaluation is 1 then the test has succeeded and the instruction following the associated then clause is executed. If all the expressions evaluate to 0 and an else was given then the instruction after the else is executed.

Note that once an expression evaluation has resulted in 1, no further expressions in the clause are evaluated. So, for example, in:

```
-- assume 'name' is a string  
if name=null, name='' then say 'Empty'
```

then if name does not refer to an object it will compare equal to null and the say instruction will be executed without evaluating the second expression in the if clause.

Notes:

1. An *instruction* may be any assignment, method call, or keyword instruction, including any of the more complex constructions such as do, loop, select, and the if instruction itself. A null clause is not an instruction, however, so putting an extra semicolon after the then or else is not equivalent to putting a dummy instruction. The nop instruction is provided for this purpose.
2. The keyword then is treated specially, in that it need not start a clause. This allows the expression on the if clause to be terminated by the then, without a ";" being required - were this not so, people used to other computer languages would be inconvenienced. Hence the symbol then cannot be used as a variable name within the expression.⁵¹

Import instruction

```
import name;
```

where *name* is one or more non-numeric *symbols* separated by periods, with an optional trailing period.

The import instruction is used to simplify the use of classes from other packages. If a class is identified by an import instruction, it can then be referred to by its short name, as given on the class instruction (see page 62) , as well as by its fully qualified name.

There may be zero or more import instructions in a program. They must precede any class instruction (or any instruction that would start the default class).

In the following description, a *package name* names a package as described under the package instruction (see page 96). The import *name* must be one of:

- A qualified class name, which is a package name immediately followed by a period which is immediately followed by a short class name - in this case, the individual class identified is imported.
- A package name - in this case, all the classes in the specified package are imported. The name may have a trailing period.
- A partial package name (a package name with one or more parts omitted from the right, indicated by a trailing period after the parts that are present) - in this case, all classes in the package hierarchy below the specified point are imported.

Examples:

```
import java.lang.String  
import java.lang  
import java.
```

The first example above imports a single class (which could then be referred to simply as "**String**"). The second example imports all classes in the "**java.lang**" package. The third example imports all classes in all the packages whose name starts with "**java.**".

When a class is imported explicitly, for example, using

```
import java.awt.List
```

this indicates that the short name of the class (`List`, in this example) may be used to refer to the class unambiguously. That is, using this short name will not report an ambiguous reference warning (as it would without the `import` instruction, because a `java.util.List` class was added in Java 1.2).

It follows that:

- Two classes imported explicitly cannot have the same short name.
- No class in a program being compiled can have the same short name as a class that is imported explicitly.

because in either of these situations a use of the short name would be ambiguous.

Note also that an explicit import does not import the minor or dependent classes associated with a name; they each require their own explicit import (unless the entire package is imported).

In the reference implementation, the fundamental NetREXX and Java package hierarchies are automatically imported by default, as though the instructions:

```
import netrexx.lang.  
import java.lang.  
import java.io.  
import java.util.  
import java.net.  
import java.awt.  
import java.applet.  
import javax.swing
```

4.01 *had been executed before the program begins. For JDK9+ environments, NetREXX looks for packages and classes in JPMS' jrt:/ file system In addition, classes in the current (working) directory are imported if no package instruction is specified. If a package instruction is specified then all classes in that package are imported.*

Iterate instruction

```
iterate [name];
```

where *name* is a non-numeric *symbol*.

`iterate` alters the flow of control within a loop construct. It may only be used in the body (the first *instructionlist*) of the construct.

Execution of the instruction list stops, and control is passed directly back up to the loop clause just as though the last clause in the body of the construct had just been executed. The control variable (if any) is then stepped (iterated) and termination conditions tested as normal and the instruction list is executed again, unless the loop is terminated by the loop clause.

If no *name* is specified, then `iterate` will step the innermost active loop.

If a *name* is specified, then it must be the name of the label, or control variable if there is no label, of a currently active loop (which may be the innermost), and this is the loop that is iterated. Any active `do`, `loop`, or `select` constructs inside the loop selected for iteration are terminated (as though by a `leave` instruction).

Example:

```
loop i=1 to 4
  if i=2 then iterate i
  say i
end
/* Would display the numbers: 1, 3, 4 */
```

Notes:

1. A loop is active if it is currently being executed. If a method (even in the same class) is called during execution of a loop, then the loop becomes inactive until the method has returned. `iterate` cannot be used to step an inactive loop.
2. The *name* symbol, if specified, must exactly match the label (or the name of the control variable, if there is no label) in the loop clause in all respects except case.

Leave instruction

```
leave [name];
```

where *name* is a non-numeric *symbol*.

`leave` causes immediate exit from one or more `do`, `loop`, or `select` constructs. It may only be used in the body (the first *instructionlist*) of the construct.

Execution of the instruction list is terminated, and control is passed to the end clause of the construct, just as though the last clause in the body of the construct had just been executed or (if a loop) the termination condition had been met normally, except that on exit the control variable (if any) will contain the value it had when the `leave` instruction was executed.

If no *name* is specified, then `leave` must be within an active loop and will terminate the innermost active loop.

If a *name* is specified, then it must be the name of the label (or control variable for a loop with no label), of a currently active `do`, `loop`, or `select` construct (which may be the innermost). That construct (and any active constructs inside it) is then terminated. Control then passes to the clause following the end clause that matches the `do`, `loop`, or `select` clause identified by the *name*.

Example:

```
loop i=1 to 5
  say i
  if i=3 then leave
end i
/* Would display the numbers: 1, 2, 3 */
```

Notes:

1. If any construct being left includes a `finally` clause, the *instructionlist* following the `finally` will be executed before the construct is left.
2. A `do`, `loop`, or `select` construct is active if it is currently being executed. If a method (even in the same class) is called during execution of an active construct, then the construct becomes inactive until the method has returned. `leave` cannot be used to leave an inactive construct.
3. The *name* symbol, if specified, must exactly match the label (or the name of the control variable, for a loop with no label) in the `do`, `loop`, or `select` clause in all respects except case.

Loop instruction

```
loop [label name] [protect term] [repetitor] [conditional];  
    instructionlist  
    [catch [vare =] exception;  
        instructionlist]...  
    [finally [;]  
        instructionlist]  
end [name];
```

where *repetitor* is one of:

```
varc = expri [to exprt] [by exprb] [for exprf]  
varo over termo  
for exprr  
forever
```

and *conditional* is either of:

```
while exprw  
until expru
```

and *name* is a non-numeric *symbol*

and *instructionlist* is zero or more *instructions*

and *expri*, *exprt*, *exprb*, *exprf*, *exprr*, *exprw*, and *expru* are *expressions*.

The loop instruction is used to group instructions together and execute them repetitively. The loop may optionally be given a label, and may protect an object while the instructions in the loop are executed; exceptional conditions can be handled with `catch` and `finally`. `loop` is the most complicated of the NetREXX keyword instructions. It can be used as a simple indefinite loop, a predetermined repetitive loop, as a loop with a bounding condition that is recalculated on each iteration, or as a loop that steps over the contents of a collection of values.

25.1 Syntax notes:

- The `label` and `protect` phrases may be in any order. They must precede any *repetitor* or *conditional*.
- The first *instructionlist* is known as the *body* of the loop.
- The `to`, `by`, and `for` phrases in the first form of *repetitor* may be in any order, if used, and will be evaluated in the order they are written.
- Any instruction allowed in a method is allowed in an *instructionlist*, including assignments, method call instructions, and keyword instructions (including any of the more complex constructions such as `if`, `do`, `select`, or the `loop` instruction itself).
- If `for` or `forever` start the *repetitor* and are followed by an "=" character, they are taken as control variable names, not keywords (as for assignment instructions).
- The expressions *expri*, *expri*, *exprb*, or *exprf* will be ended by any of the keywords `to`, `by`, `for`, `while`, or `until` (unless the word is the name of a variable).
- The expressions *exprw* or *expru* will be ended by either of the keywords `while` or `until` (unless the word is the name of a variable).

25.2 Indefinite loops

If neither *repetitor* nor *conditional* are present, or the *repetitor* is the keyword `forever`, then the loop is an *indefinite loop*. It will be ended only when some instruction in the first *instructionlist* causes control to leave the loop.

Example:

```
/* This displays "Go caving!" at least once */
loop forever
  say 'Go caving!'
  if ask='' then leave
end
```

25.3 Bounded loops

If a *repetitor* (other than `forever`) or *conditional* is given, the first *instructionlist* forms a *bounded loop*, and the instruction list is executed according to any *repetitor phrase*, optionally modified by a *conditional phrase*.

Simple bounded loops When the *repetitor* starts with the keyword `for`, the expression *exprr* is evaluated immediately (with 0 added, to effect any rounding) to give a repetition count, which must be a whole number that is zero or positive. The loop is then executed that many times, unless it is terminated by some other condition.

Example:

```

/* This displays "Hello" five times */
loop for 5
  say 'Hello'
end

```

A *controlled loop* begins with an *assignment*, which can be identified by the "=" that follows the name of a control variable, *varc*. The control variable is assigned an initial value (the result of *expri*, formatted as though 0 had been added) before the first execution of the instruction list. The control variable is then stepped (by adding the result of *exprb*) before the second and subsequent times that the instruction list is executed.

The name of the control variable, *varc*, must be a non-numeric symbol that names an existing or new variable in the current method or a property in the current class (that is, it cannot be element of an array, the property of a superclass, or a more complex term). It is further restricted in that it must not already be used as the name of a control variable or label in a loop (or do or select construct) that encloses the new loop.

The instruction list in the body of the loop is executed repeatedly while the end condition (determined by the result of *exprt*) is not met. If *exprb* is positive or zero, then the loop will be terminated when *varc* is greater than the result of *exprt*. If negative, then the loop will be terminated when *varc* is less than the result of *exprt*. The expressions *exprt* and *exprb* must result in numbers. They are evaluated once only (with 0 added, to effect any rounding), in the order they appear in the instruction, and before the loop begins and before *expri* (which must also result in a number) is evaluated and the control variable is set to its initial value.

The default value for *exprb* is 1. If no *exprt* is given then the loop will execute indefinitely unless it is terminated by some other condition. **Example:**

```

loop i=3 to -2 by -1
  say i
end
/* Would display: 3, 2, 1, 0, -1, -2 */

```

Note that the numbers do not have to be whole numbers: **Example:**

```

x=0.3
loop y=x to x+4 by 0.7
  say y
end
/* Would display: 0.3, 1.0, 1.7, 2.4, 3.1, 3.8 */

```

The control variable may be altered within the loop, and this may affect the iteration of the loop. Altering the value of the control variable in this way is normally considered to be suspect programming practice, though it may be appropriate in certain circumstances. Note that the end condition is tested at the start of each iteration (and after the control variable is stepped, on the second and subsequent iterations). It is therefore possible for the body of the loop to be skipped entirely if the end condition is met immediately. The execution of a controlled loop may further be bounded by a for phrase. In this case, *expri* must be given and must evaluate to a non-negative whole number. This acts just like the repetition count in a

simple bounded loop, and sets a limit to the number of iterations around the loop if it is not terminated by some other condition.

exprf is evaluated along with the expressions *exprt* and *exprb*. That is, it is evaluated once only (with 0 added), when the loop instruction is first executed and before the control variable is given its initial value; the three expressions are evaluated in the order in which they appear. Like the *to* condition, the *for* count is checked at the start of each iteration, as shown in the programmer's (see page 80) model:ea.. **Example:**

```
loop y=0.3 to 4.3 by 0.7 for 3
  say y
end
/* Would display: 0.3, 1.0, 1.7 */
```

In a controlled loop, the symbol that describes the control variable may be specified on the end clause (unless a label is specified, see below). NetREXX will then check that this symbol exactly matches the *varc* of the control variable in the loop clause (in all respects except case). If the symbol does not match, then the program is in error - this enables the nesting of loops to be checked automatically. **Example:**

```
loop k=1 to 10
  ...
end k /* Checks this is the END for K loop */
```

Note: The values taken by the control variable may be affected by the numeric settings, since normal NetREXX arithmetic rules apply to the computation of stepping the control variable.

Controlled bounded loop over When the second token of the *repetitor* is the keyword *over*, the control variable, *varo*, is used to work through the sub-values in the collection of indexed strings identified by *termo*. In this case, the loop instruction takes a "snapshot" of the indexes that exist in the collection at the start of the loop, and then for each iteration of the loop the control variable is set to the next available index from the snapshot.

The number of iterations of the loop will be the number of indexes in the collection, unless the loop is terminated by some other condition. **Example:**

```
mycoll=' '
mycoll['Tom']=1
mycoll['Dick']=2
mycoll['Harry']=3
loop name over mycoll
  say mycoll[name]
end
/* might display: 3, 1, 2 */
```

Notes:

1. The order in which the values are returned is undefined; all that is known is that all indexes available when the loop started will be recorded and assigned to *varo* in turn as the loop iterates.
2. The same restrictions apply to *varo* as apply to *varc*, the control variable for controlled loops (see above).

3. Similarly, the symbol *varo* may be used as a name for the loop and be specified on the end clause (unless a label is specified, see below).

*In the reference implementation, the over form of repetitor may also be used to step through the contents of any object that is of a type that is a subclass of `java.util.Dictionary`, such as an object of type `java.util.Hashtable`. In this case, *termo* specifies the dictionary, and a snapshot (enumeration) of the keys to the Dictionary is taken at the start of the loop. Each iteration of the loop then assigns a new key to the control variable *varo* which must be (or will be given, if it is new) the type `java.lang.Object`.*

Conditional phrases Any of the forms of loop syntax can be followed by a *conditional* phrase which may cause termination of the loop.

If *while* is specified, *exprw* is evaluated, using the latest values of all variables in the expression, before the instruction list is executed on every iteration, and after the control variable (if any) is stepped. The expression must evaluate to either 0 or 1, and the instruction list will be repeatedly executed while the result is 1 (that is, the loop ends if the expression evaluates to 0).

Example:

```
loop i=1 to 10 by 2 while i<6
  say i
end
/* Would display: 1, 3, 5 */
```

If *until* is specified, *expru* is evaluated, using the latest values of all variables in the expression, on the second and subsequent iterations, and before the control variable (if any) is stepped.⁵² The expression must evaluate to either 0 or 1, and the instruction list will be repeatedly executed until the result is 1 (that is, the loop ends if the expression evaluates to 1). **Example:**

```
loop i=1 to 10 by 2 until i>6
  say i
end
/* Would display: 1, 3, 5, 7 */
```

Note that the execution of loops may also be modified by using the *iterate* or *leave* instructions.

25.4 Label phrase

The *label* phrase may be used to specify a *name* for the loop. The name can then optionally be used on

- a *leave* instruction, to specify the name of the loop to leave
- an *iterate* instruction, to specify the name of the loop to be iterated
- the end clause of the loop, to confirm the identity of the loop that is being ended, for additional checking.

Example:

⁵²Thus, it appears that the *until* condition is tested after the instruction list is executed on each iteration. However, it is the loop clause that carries out the evaluation.

```

loop label pooks i=1 to 10
  loop label hill while j<3
    ...
    if a=b then leave pooks
    ...
  end hill
end pooks

```

In this example, the leave instruction leaves both loops.

If a label is specified using the label keyword, it overrides any name derived from the control variable name (if any). That is, the variable name cannot be used to refer to the loop if a label is specified.

25.5 Protect phrase

The protect phrase may be used to specify a term, *term*, that evaluates to a value that is not just a type and is not of a primitive type; while the loop construct is being executed, the value (object) is protected - that is, all the instructions in the loop construct have exclusive access to the object. **Example:**

```

loop protect myobject while a<b
  ...
end

```

Both label and protect may be specified, in any order, if required.

25.6 Exceptions in loops

Exceptions that are raised by the instructions within a loop construct may be caught using one or more catch clauses that name the *exception* that they will catch. When an exception is caught, the exception object that holds the details of the exception may optionally be assigned to a variable, *var*.

Similarly, a finally clause may be used to introduce instructions that will always be executed when the loop ends, even if an exception is raised (whether caught or not).

The *Exceptions* section (see page 149) has details and examples of catch and finally.

25.7 Programmer's model - how a typical loop is executed

This model forms part of the definition of the loop instruction. For the following loop:

```

loop varc = expri to expri by exprb while exprw
  ...
  instruction list

```

```
...
end
```

NetREXX will execute the following:

```
$tempt=exprt+0 /* ($variables are internal and */
$tempb=exprb+0 /* are not accessible.) */
varc=expri+0
Transfer control to the point identified as $start:

$loop:
/* An UNTIL expression would be tested here, with: */
/* if expru then leave */
varc=varc + $tempb
$start:
if varc > $tempt then leave /* leave quits a loop */
/* A FOR count would be checked here */
if \ exprw then leave
...
instruction list
...
Transfer control to the point identified as $loop:
```

Notes:

1. This example is for *exprb* ≥ 0 . For a negative *exprb*, the test at the start point of the loop would use " $<$ " rather than " $>$ ".
2. The upwards transfer of control takes place at the end of the body of the loop, immediately before the end clause (or any catch or finally clause). The end clause is only reached when the loop is finally completed.

Method instruction

```
method name([[arglist]])
            [visibility] [modifier] [protect] [binary] [deprecated]
            [returns termr]
            [signals signallist];
```

where *arglist* is a list of one or more *assignments*, separated by commas

and *visibility* is one of:

inheritable
private
public
shared

and *modifier* is one of:

abstract
constant
final
native
static

and *signallist* is a list of one or more *terms*, separated by commas.

The method instruction is used to introduce a method within a class, as described in *Program structure* (see page 114), and define its attributes. The method must be given a *name*, which must be a non-numeric symbol. This is its *short name*.

If the short name of a method matches the short name of the class in which it appears, it is a *constructor method*. Constructor methods are used for constructing values (objects), and are described in detail in *Methods and Constructors* (see page 32).

The *body* of the method consists of all clauses following the method instruction (if any) until the next method or class instruction, or the end of the program.

The *visibility*, *modifier*, and *protect* keywords, and the *returns* and *signals* phrases, may appear in any order.

26.1 Arguments

The *arglist* on a method instruction, immediately following the method name, is optional and defines a list of the arguments for the method. An *argument* is a value that was provided by the caller when the method was invoked.

If there are no arguments, this may optionally be indicated by an "empty" pair of parentheses.

In the *arglist*, each argument has the syntax of an *assignment* (see page 49) , where the "=" and the following *expression* may be omitted. The name in the assignment provides the name for the argument (which must not be the same as the name of any property in the class). Each argument is also optionally assigned a type, or type and default value, following the usual rules of assignment. If there is no assignment, the argument is assigned the NetREXX string type, **REXX**.

If there is no assignment (that is, there is no "=") or the expression to the right of the "=" returns just a type, the argument is *required* (that is, it must always be specified by the caller when the method is invoked).

If an explicit value is given by the expression then the argument is *optional*; when the caller does not provide an argument in that position, then the expression is evaluated when the method is invoked and the result is provided to the method as the argument.

Optional arguments may be omitted "from the right" only. That is, arguments may not be omitted to the left of arguments that are not omitted. **Examples:**

```
method fred
method fred()
method fred(width, height)
method fred(width=int, height=int 10)
```

In these examples, the first two method instructions are equivalent, and take no arguments. The third example takes two arguments, which are both strings of type **REXX**. The final example takes two arguments, both of type **int**; the second argument is optional, and if not supplied will default to the value 10 (note that any valid expression could be used for the default value).

26.2 Visibility

Methods may be public, inheritable, private, or shared:

- A *public method* is visible to (that is, may be used by) all other classes to which the current class is visible.
- An *inheritable method* is visible to (that is, may be used by) all classes in the same package and also those classes that extend (that is, are subclasses of) the current class.
- A *private method* is visible only within the current class.

- A *shared method* is visible within the current package but is not visible outside the package. Shared methods cannot be inherited by classes outside the package.

By default (*i.e.*, if no visibility keyword is specified), methods are public.

26.3 Modifier

Most methods consist of instructions that follow the method instruction and implement the method; the method is associated with an object constructed by the class. These are called *standard methods*. The *modifier* keywords define that the method is not a standard method - it is special in some way. Only one of the following modifier keywords is allowed:

abstract An *abstract method* has the name of the method and the types (but not values) of its arguments defined, but no instructions to implement the method are provided (or permitted).

If a class contains any abstract methods, an object cannot be constructed from it, and so the class itself must be abstract; the abstract keyword must be present on the class instruction (see page 62) .

Within an interface class, the abstract keyword is optional on the methods of the class, as all methods must be abstract. No other *modifier* is allowed on the methods of an interface class.

constant A *constant method* is a static method that cannot be overridden by a method in a subclass of the current class. That is, it is both final and static (see below).

final A *final method* is considered to be complete; it cannot be overridden by a subclass of the current class. private methods are implicitly final.⁵³

native A *native method* is a method that is implemented by the environment, not by instructions in the current class. Such methods have no NetREXX instructions to implement the method (and none are permitted), and they cannot be overridden by a method in a subclass of the current class.

Native methods are used for accessing primitive operations provided by the underlying operating system or by implementation-dependent packages.

static A *static method* is a method that is not a constructor and is associated with the class, rather than with an object constructed by the class. It cannot use properties directly, except those that are also static (or constant).

Static methods may be invoked in the following ways:

1. Within the initialization expression of a static or constant property (such methods are invoked when the class is first loaded).
2. By qualifying the name of the method with the name of its class (qualified by the package name if necessary), for example, using "**Math.Sin(1.3)**" or "**java.lang.Math.Sin(1.3)**". Methods called in this way are in effect *functions*.

⁵³This modifier may allow compilers to improve the performance of methods that are final, but may also reduce the reusability of the class.

3. By implicitly qualifying the name by including the name of its class in the uses phrase in the `class` instruction for the current class. Static methods in classes listed in this way can be used directly, without qualification, for example, as "**Sin(1.3)**". They may be still be qualified, if preferred.

In the reference implementation, stand-alone applications are started by the `java` command invoking a static method called `main` which takes a single argument (of type `java.lang.String[]`) and returns no result.

26.4 Protect

The keyword `protect` indicates that the method protects the current object (or class, for a static method) while the instructions in the method are executed. That is, the instructions in the method have exclusive access to the object; if some other method (or construct) is executing in parallel with the invocation of the method and is protecting the same object then the method will not start execution until the object is no longer protected.

Note that if a method or construct protecting an object invokes a method (or starts a new construct) that protects the same object then execution continues normally. The inner method or construct is not prevented from executing, because it is not executing in parallel.

26.5 Binary

The keyword `binary` indicates that the method is a *binary method*.

In binary methods, literal strings and numeric symbols are assigned native string or binary (primitive) types, rather than NetRExx types, and native binary operations are used to implement operators where possible. When `binary` is not in effect (the default), terms in expressions are converted to NetRExx types before use by operators. The section *Binary values and operations* (see page 146) *operations:ea*. describes the implications of binary methods and classes in detail.

Notes:

1. Only the instructions inside the body of the method are affected by the `binary` keyword; any arguments and expressions on the method instruction itself are not affected (this ensures that a single rule applies to all the method signatures in a class).
2. All methods in a binary class are binary methods; the `binary` keyword on methods is provided for classes in which only the occasional method needs to be binary (perhaps for performance reasons). It is not an error to specify `binary` on a method in a binary class.

26.6 Deprecated

The keyword `deprecated` indicates that the method is *deprecated*, which implies that a better alternative is available and documented. A compiler can use this information to warn of out-of-date or other use that is not recommended.

Note that individual methods in interface classes cannot be deprecated; the whole class should be deprecated in this case.

26.7 Returns

The `returns` keyword is followed by a term, *term*, that must evaluate to a type. This type is used to define the type of values returned by `return` instructions within the method.

The `returns` phrase is only required if the method is to return values of a type that is not NetREXX strings (class `REXX`). If `returns` is specified, all `return` instructions (see page 102) within the method must specify an expression. **Example:**

```
method filer(path, name) returns File
    return File(path||name)
```

This method always returns a value which is a **File** object.

26.8 Signals

The `signals` keyword introduces a list of terms that evaluate to types that are Exceptions (see page 149). This list enumerates and documents the exceptions that are signalled within the method (or by a method which is called from the current method) but are not caught by a `catch` clause in a control construct.

Example:

```
method soup(cat) signals IOException, DivideByZero
```

It is considered good programming practice to use this list to document “unusual” exceptions signalled by a method. Implementations that support the concept of *checked exceptions* (see page 151) must report as an error any checked exception that is incorrectly included in the list (that is, if the exception is never signalled or would always be caught). Such implementations may also offer an option that enforces the listing of all or some checked exceptions.

26.9 Duplicate methods

Methods may not duplicate properties or other methods in the same class. Specifically:

- The short name of a method must not be the same as the name of any property in the same class.
- The number (zero or more) and types of the arguments of a method (or any subset permitted by omitting optional arguments) must not be the same as those of any other method of the same name in the class (also checking any subset permitted by omitting optional arguments).

Note that the second rule does allow multiple methods with the same name in a class, provided that the number of arguments differ or at least one argument differs in type.

Nop instruction

```
nop;
```

`nop` is a dummy instruction that has no effect. It can be useful as an explicit “do nothing” instruction following a `then` or `else` clause. **Example:**

```
select
  when a=b then nop           -- Do nothing
  when a>b then say 'A > B'
  otherwise    say 'A < B'
end
```

Note: Putting an extra semicolon instead of the `nop` would merely insert a null clause, which would just be ignored by NetREXX. The second `when` clause would then immediately follow the `then`, and hence would be reported as an error. `nop` is a true instruction, however, and is therefore a valid target for the `then` clause.

Numeric instruction

```
numeric digits [exprd];  
    form [scientific];  
           [engineering];
```

The `numeric` instruction is used to change the way in which arithmetic operations are carried out by a program. The effects of this instruction are described in detail in the section on *Numbers and Arithmetic* (see page 136) .

numeric digits controls the precision under which arithmetic operations will be evaluated (see page 138) .

If no expression *exprd* is given then the default value of 9 is used. Otherwise the result of the expression is rounded, if necessary, according to the current setting of `numeric digits` before it is used. The value used must be a positive whole number. There is normally no limit to the value for `numeric digits` (except the constraints imposed by the amount of storage and other resources available) but note that high precisions are likely to be expensive in processing time. It is recommended that the default value be used wherever possible.

Note that small values of `numeric digits` (for example, values less than 6) are generally only useful for very specialized applications. The setting of `numeric digits` affects all computations, so even the operation of loops may be affected by rounding if small values are used.

If an implementation does not support a requested value for `numeric digits` then the instruction will fail with an exception (which may, as usual, be caught with the catch clause of a control construct).

The current setting of `numeric digits` may be retrieved with the **digits** special word (see page 122) .

numeric form controls which form of exponential notation (see page 144) is to be used for the results of operations. This may be either *scientific* (in which case only one, non-zero, digit will appear before the decimal point), or *engineering* (in which case the power of ten will always be a multiple of three, and the part before the decimal point will be in the range 1 through 999). The default notation is *scientific*. The form is set directly by the sub-keywords `scientific` or `engineering`; if neither sub-keyword is given, `scientific` is assumed. The current setting of `numeric form` may be retrieved with the **form** special word (see page 122) .

If an implementation does not support a requested value for `numeric form` then the instruction will fail with an exception (which may, as usual, be caught with the `catch` clause of a control construct).

The `numeric` instruction may be used where needed as a dynamically executed instruction in a method.

It may also appear, more than once if necessary, before the first method in a class, in which case it forms the default setting for the initialization of subsequent properties in the class and for all methods in the class. In this case, any exception due to the `numeric` instruction is raised when the class is first loaded.

Further, one or more `numeric` instructions may be placed before the first `class` instruction in a program; they do not imply the start of a class. The `numeric` settings then apply to all classes in the program (except interface classes), as though the `numeric` instructions were placed immediately following the `class` instruction in each class (except that they will not be traced).

Options instruction

```
options wordlist;
```

where *wordlist* is one or more *symbols* separated by blanks.

The options instruction is used to pass special requests to the language processor (for example, an interpreter or compiler).

Individual words, known as *option words*, in the *wordlist* which are meaningful to the language processor will be obeyed (these might control optimizations, enforce standards, enable implementation-dependent features, *etc.*); those which are not recognized will be ignored (they are assumed to be instructions to a different language processor). Option words in the list that are known will be recognized independently of case.

There may be zero or more options instructions in a program. They apply to the whole program, and must come before the first class instruction (or any instruction that starts a class).

In the reference implementation, the known option words are:

annotations *All annotations are passed through to Java, this is the default. When option `-noannotations` is in effect, no annotations will be passed through, although they are in the program source. This implies that the `@SuppressWarnings("unchecked")` annotation, which is passed through every method, is also left out, which possibly yields (harmless) `javac` warnings (on the subject of generics).*

binary *All classes in this program will be binary (see page 64) classes. In binary classes, literals are assigned binary (primitive) or native string types, rather than NetREXX types, and native binary operations are used to implement operators where appropriate, as described in "Binary values and operations" (see page 146) . In classes that are not binary, terms in expressions are converted to the NetREXX string type, **REXX**, before use by operators.*

comments *Comments from the NetREXX source program will be passed through to the the Java output file (which may be saved with a `.java.keep` extension by using the `-keep` command option).*

*Line comments become Java line comments (introduced by `"/"/`). Block comments become Java block comments (delimited by `"/**"` and `"*/"`), with nested block comments having their delimiters changed to `"(-"` and `"-)"`.*

classpath *The `-classpath` operand allows dynamic specification of the classpath used* 3.02

by the NetRexxC compiler without having to depend on the CLASSPATH environment variable. There is no `-noclasspath` counterpart.

compact Requests that warnings and error messages be displayed in compact form. This format is more easily parsed than the default format, and is intended for use by editing environments.

Each error message is presented as a single line, prefixed with the error token identification enclosed in square brackets. The error token identification comprises three words, with one blank separating the words. The words are: the source file specification, the line number of the error token, the column in which it starts, and its length. For example (all on one line):

```
[D:  
test  
test.nrx 3 8 5] Error: The external name  
'class' is a Java reserved word, so would not be  
usable from Java programs
```

Any blanks in the file specification are replaced by a null (`'\0'`) character. Additional words could be added to the error token identification later.

console Requests that compiler messages be written to console (the default). Use `-noconsole` to prevent messages being written to the console.

This option only has an effect as a compiler option, and applies to all programs being compiled.

crossref Requests that cross-reference listings of variables be prepared, by class.

decimal Decimal arithmetic may be used in the program. If `nodecimal` is specified, the language processor will report operations that use (or, like normal string comparison, might use) decimal arithmetic as an error. This option is intended for performance-critical programs where the overhead of inadvertent use of decimal arithmetic is unacceptable.

diag Requests that diagnostic information (for experimental use only) be displayed. The `diag` option word may also have side-effects.

3.04 **ecj** Indicates to the translator a preference for using the ecj compiler, if available

explicit Requires that all local variables must be explicitly declared (by assigning them a type but no value) before assigning any value to them. This option is intended to permit the enforcement of "house styles" (but note that the NetREXX compiler always checks for variables which are referenced before their first assignment, and warns of variables which are set but not used).

format Requests that the translator output file (Java source code) be formatted for improved readability. Note that if this option is in effect, line numbers from the input file will not be preserved (so run-time errors and exception trace-backs may show incorrect line numbers).

java Requests that Java source code be produced by the translator. If `nojava` is specified, no Java source code will be produced; this can be used to save a little time when checking of a program is required without any compilation or Java code resulting.

3.04 **javac** Indicates to the translator a preference for using the javac compiler, if available

- keepasjava** Requests that Java source code is kept as `[programfile].java`. Implies `-replace`. This option only has an effect as a compiler option, and applies to all programs being compiled.
- logo** Requests that the language processor display an introductory logotype sequence (name and version of the compiler or interpreter, etc.).
- mod** Requests that the language processor produces a formatted representation of the source program - see *Programming Guide*.
- replace** Requests that replacement of the translator output (`.java`) file be allowed. The default, `noreplace`, prevents an existing `.java` file being accidentally overwritten.
- savelog** Requests that compiler messages be written to the file `NetREXXC.log` in the current directory. The messages are also displayed on the console, unless `-noconsole` is specified.
This option only has an effect as a compiler option, and applies to all programs being compiled.
- sourcedir** Requests that all `.class` files be placed in the same directory as the source file from which they are compiled. Other output files are already placed in that directory. Note that using this option will prevent the `-run` command option from working unless the source directory is the current directory.
- strictargs** Requires that method invocations always specify parentheses, even when no arguments are supplied. Also, if `strictargs` is in effect, method arguments are checked for usage - a warning is given if no reference to the argument is made in the method.
- strictassign** Requires that only exact type matches be allowed in assignments (this is stronger than Java requirements). This also applies to the matching of arguments in method calls.
- strictcase** Requires that local and external name comparisons for variables, properties, methods, classes, and special words match in case (that is, names must be identical to match).
- strictimport** Requires that all imported packages and classes be imported explicitly using `import` instructions. That is, if in effect, there will be no automatic imports (see page 71), except those related to the package instruction.
This option only has an effect as a compiler option, and applies to all programs being compiled.
- strictprops** Requires that all properties, including those local to the current class, be qualified in references. That is, if in effect, local properties cannot appear as simple names but must be qualified by `this.` (or equivalent) or the class name (for static properties).
- strictsignal** Requires that all checked exceptions (see page 151) signalled within a method but not caught by a `catch` clause be listed in the `signals` phrase of the method instruction.
- symbols** Symbol table information (names of local variables, etc.) will be included in any generated `.class` file. This option is provided to aid the production of classes that are easy to analyse with tools that can understand the symbol table information. The use of this option increases the size of `.class` files.
- trace, traceX** If given as `trace`, `trace1`, or `trace2`, then `trace` instructions are accepted. The trace output is directed according to the option word: `trace1` requests

that trace output is written to the standard output stream, `trace` or `trace2` imply that the output should be written to the standard error stream (the default). If `notrace` is given, then trace instructions are ignored. The latter can be useful to prevent tracing overheads while leaving trace instructions in a program.

utf8 If given, clauses following the `options` instruction are expected to be encoded using UTF-8, so all Unicode characters may be used in the source of the program. In UTF-8 encoding, Unicode characters less than `'\u0080'` are represented using one byte (whose most-significant bit is 0), characters in the range `'\u0080'` through `'\u07FF'` are encoded as two bytes, in the sequence of bits:

```
110xxxxx 10xxxxxx
```

where the eleven digits shown as `x` are the least significant eleven bits of the character, and characters in the range `'\u0800'` through `'\uFFFF'` are encoded as three bytes, in the sequence of bits:

```
1110xxxx 10xxxxxx 10xxxxxx
```

where the sixteen digits shown as `x` are the sixteen bits of the character.

If `noutf8` is given, following clauses are assumed to comprise only Unicode characters in the range `'\x00'` through `'\xFF'`, with the more significant byte of the encoding of each character being 0.

Note: this option only has an effect as a compiler option, and applies to all programs being compiled. If present on an `options` instruction, it is checked and must match the compiler option (this allows processing with or without `utf8` to be enforced).

verbose, verboseX Sets the "noisiness" of the language processor. The digit `X` may be any of the digits `0` through `5`; if omitted, a value of `3` is used. The options `noverbose` and `verbose0` both suppress all messages except errors and warnings.

warnexit0 Exit the translator with returncode 0 even if warnings are issued. This option only has an effect as a compiler option, and applies to all programs being compiled.

Prefixing any of the above with "no" turns the selected option off. Example:

```
options binary nocrossref nostrictassign strictargs
```

The default settings of the various options are:

```
address nobinary nocomments nocompact console crossref decimal nodiag noexplicit  
noformat java logo noreplace nosavelog nosourcedir nostrictargs  
nostrictassign nostrictcase nostrictimport nostrictprops nostrictsignal  
nosymbols trace2 noutf8 verbose3
```

When an option word is repeated (in the same `options` instruction or not), or conflicting option words are specified, then the last use determines the state of the option.

All option words may also be set as command line options when invoking the processor, by prefixing them with "-": Example:

```
java COM.ibm.netrexx.process.NetRexxC -format foo.nrx
```

In this case, any options may come before, after, or between file specifications.

With the exception of the `utf8` option (see above), options set with the `options` instruction override command-line settings, following the "last use" rule.

For more information, see the installation and user documentation for your implementation.

All options are listed in Appendix C on page 193.

Package instruction

```
package name;
```

where *name* is one or more non-numeric *symbols* separated by periods.

The package instruction is used to define the package to which the class or classes in the current program belong. Classes that belong to the same package have privileged access to other classes in the same package, in that each class is visible to all other classes in the same package, even if not declared public. Packages also conveniently group classes for use by the `import` instruction (see page 71).

The *name* must specify a *package name*, which is one or more non-numeric symbols, separated by periods, with no blanks.

There must be at most one package instruction in a program. It must precede any class instruction (or any instruction that would start the default class).

If a program contains no package instruction then its package is implementation-defined. Typically it is grouped with other programs in some implementation-defined logical collection, such as a directory in a file system. **Examples:**

```
package testpackage  
package com.ibm.venta
```

When a class is identified as belonging to a package, it has a *qualified class name*, which is its short name, as given on the `class` instruction (see page 62), prefixed with the package name and a period. For example, if the short name of a class is "**RxLanguage**" and the package name is "**com.ibm.venta**" then the qualified name of the class would be "**com.ibm.venta.RxLanguage**".

In the reference implementation, packages are kept in a hierarchy derived from the Java classpath, where the segments of a package name correspond to a path in the hierarchy. The hierarchy is typically the directories in a file system, or some equivalent (such as a "Zip" archive file), and so package names should be considered case-sensitive (as some Java implementations use case-sensitive file systems).

Parse instruction

```
parse term template;
```

where *template* is one or more non-numeric *symbols* separated by blanks and/or *patterns*, and a *pattern* is one of:

```
literalstring  
[indicator] number  
[indicator] (symbol)
```

and *indicator* is one of **+**, **-**, or **=**.

The parse instruction is used to assign characters (from a string) to one or more variables according to the rules and templates described in the section *Parsing templates* (see page 129).

The value of the *term* is expected to be a string; if it is not a string, it will be converted to a string.

Any variables used in the *template* are named by non-numeric *symbols* (that is, they cannot be an array reference or other term); they refer to a variable or property in the current class. Any values that are used in patterns during the parse are converted to strings before use.

Any variables set by the parse instruction must have a known string type, or are given the NetREXX string type, **REXX**, if they are new.

The term itself is not changed unless it is a variable which also appears in the template and whose value is changed by being in the template.

Example:

```
parse wordlist word1 wordlist
```

In this idiomatic example, the first word is removed from **wordlist** and is assigned to the variable **word1**, and the remainder is assigned back to **wordlist**.

Notes:

1. The special words *ask*, *source*, and *version*, as described in the section *Special names and methods* (see page 121), allow:

```
parse ask x      -- parses a line from input stream
```

```
parse asknoecho x -- parses a line from input stream without
echo
parse source x -- parses 'Java method filename'
parse version x -- parses 'NetRexx version date'
```

These special words may also be used within expressions.

2. Similarly, it is recommended that the initial (main) method in a stand-alone application place the command string passed to it in a variable called **arg**.⁵⁴

If this is done, the instruction:

```
parse arg template
```

will work, in a stand-alone application, in the same way as in REXX (even though **arg** is not a keyword in this case).⁵⁵

⁵⁴In the reference implementation, this is automatic if the *main* method is generated by the NetRexx language processor.

⁵⁵Note, though, that the command string may have been edited by the environment; certain characters may not be allowed, multiple blanks may have been reduced to single blanks, etc.

Properties instruction

```
properties [visibility] [modifier] [deprecated] [unused];
```

where *visibility* is one of:

```
inheritable  
private  
public  
shared
```

and *modifier* is one of:

```
constant  
static  
transient  
volatile
```

and there must be at least one *visibility* or *modifier* keyword.

The `properties` instruction is used to define the attributes of following *property* variables, and therefore must precede the first method instruction in a class. A `properties` instruction replaces any previous `properties` instruction (that is, the attributes specified on `properties` instructions are not cumulative).

The *visibility*, *modifier*, `deprecated`, and `unused` keywords may be in any order.

Note: An unqualified `properties` statement (one that has no *visibility* or *modifier* keyword), is not in error, but generates a variable `properties`, which is most probably not the intention of the programmer. *The reference implementation issues a warning but allows this practice.*

3.03

An example of the use of `properties` instructions may be found in the *Program Structure* section (see page 114) .

32.1 Visibility

Properties may be public, inheritable, private, or shared:⁵⁶

- A *public property* is visible to (that is, may be used by) all other classes to which the current class is visible.
- An *inheritable property* is visible to (that is, may be used by) all classes in the same package and also those classes that extend (that is, are subclasses of) the current class, and which qualify the property using an object of the subclass, or either this or super.
- A *private property* is visible only within the current class.
- A *shared property* is visible within the current package but is not visible outside the package. Shared properties cannot be inherited by classes outside the package.

By default, if no `properties` instruction is used, or *visibility* is not specified, properties are inheritable (but not public).⁵⁷

32.2 Modifier

Properties may also be constant, static, transient, or volatile:

- A *constant property* is associated with the class, rather than with an instance of the class (an object). It is initialized when the class is loaded and may not be changed thereafter.
- A *static property* is associated with the class, rather than with an instance of the class (an object). It is initialized when the class is loaded, and may be changed thereafter.
- A *transient property* is a property which should not be saved when an instance of the class is saved (made persistent).
- A *volatile property* may change asynchronously, outside the control of the class, even when no method in the class is being executed. If an implementation does not allow asynchronous modification of properties, it should ignore this keyword.

Constant and static properties exist from when the class is first loaded (used), even if no object is constructed by the class, and there will only be one copy of each property. Other properties are constructed and initialized only when an object is constructed by the class; each object then has its own copy of such properties.

By default, if no `properties` instruction is used, or *modifier* is not specified, properties are associated with an object constructed by the class.

⁵⁶An experimental option for *visibility*, `indirect`, is described in Appendix B (see page 125) .

⁵⁷The default, here, was chosen to encourage the "encapsulation" of data within classes.

32.3 Deprecated

The keyword `deprecated` indicates that any property introduced by this instruction is *deprecated*, which implies that a better alternative is available and documented. A compiler can use this information to warn of out-of-date or other use that is not recommended.

32.4 Unused

The keyword `unused` indicates that the private properties which follow are not referenced explicitly in the code for the class, and so a language processor should not warn that they exist but have not been used. If a *visibility* keyword is specified it must be `private`.

For example:

```
properties private constant unused
-- Serialization version
serialVersionUID=long 8245355804974198832
```

32.5 Properties in interface classes

In interface classes (see page 64) , properties must be both `public` and `constant`. In such classes, these attributes for properties are the default and the properties instruction must not be used.

Return instruction

```
return [expression];
```

`return` is used to return control (and possibly a result) from a NetREXX program or method to the point of its invocation.

The expression (if any) is evaluated, active control constructs are terminated (as though by a `leave` instruction), and the value of the expression is passed back to the caller.

The result passed back to the caller is a string of type **REXX**, unless a different type was specified using the `returns` keyword on the method instruction (see page 82) for the current method. In this case, the type of the value of the expression must match (or be convertible to, as by the rules for assignment) the type specified by the `returns` phrase.

Within a method, the use of expressions on `return` must be consistent. That is, either all `return` instructions must specify an expression, or none may. If a `returns` phrase is given on the method instruction for the current method then all `return` instructions must specify an expression.

Say instruction

```
say [expression];
```

`say` writes a string to the default output character stream. This typically causes it to be displayed (or spoken, or typed, *etc.*) to the user.

Example:

```
data=100
say data 'divided by 4 =>' data/4
/* would display: "100 divided by 4 => 25" */
```

The result of evaluating the *expression* is expected to be a string; if it is not a string, it will be converted to a string. This result string is written from the program via an implementation-defined output stream.

By default, the result string is treated as a “line” (an implementation-dependent mechanism for indicating line termination is effected after the string is written). If, however, the string ends in the NUL character (`'\-'` or `'\0'`) then that character is removed and line termination is not indicated.

The result string may be of any length. If no expression is specified, or the expression result is **null**, then an empty line is written (that is, as though the expression resulted in a null string).

Select instruction

```
select [label name] [protect term] [case expression];
    whenlist
    [otherwise[;] instructionlist]
    [catch [vare =] exception;
        instructionlist]...
    [finally[;]
        instructionlist]
end [name];
```

where *name* is a non-numeric *symbol*

and *whenlist* is one or more *whenconstructs*

and *whenconstruct* is:

```
when expression[, expression]... [;] then[;] instruction
```

and *instructionlist* is zero or more *instructions*.

`select` is used to conditionally execute one of several alternatives. The construct may optionally be given a label, and may protect an object while the instructions in the construct are executed; exceptional conditions can be handled with `catch` and `finally`, which follow the body of the construct.

Starting with the first `when` clause, each expression in the clause is evaluated in turn from left to right, and if the result of any evaluation is 1 (or equals the case expression, see below) then the test has succeeded and the instruction following the associated `then` (which may be a complex instruction such as `if`, `do`, `loop`, or `select`) is executed and control will then pass directly to the `end`.

If the result of all the expressions in a `when` clause is 0, control will pass to the next `when` clause.

Note that once an expression evaluation in a `when` clause has resulted in a successful test, no further expressions in the clause are evaluated.

If none of the `when` expressions result in 1, then control will pass to the instruction list (if any) following `otherwise`. In this situation, the absence of an

otherwise is a run-time error.⁵⁸ **Notes:**

1. An *instruction* may be any assignment, method call, or keyword instruction, including any of the more complex constructions such as `do`, `loop`, `if`, and the `select` instruction itself. A null clause is not an instruction, however, so putting an extra semicolon after the `then` is not equivalent to putting a dummy instruction (as it would be in C or PL/I). The `nop` instruction is provided for this purpose.
2. The keyword `then` is treated specially, in that it need not start a clause. This allows the expression on the `when` clause to be terminated by the `then`, without a `;` being required - were this not so, people used to other computer languages would be inconvenienced. Hence the symbol `then` cannot be used as a variable name within the expression.⁵⁹

35.1 Label phrase

If `label` is used to specify a *name* for the `select` group, then a `leave` instruction (see page 74) which specifies that name may be used to leave the group, and the end that ends the group may optionally specify the name of the group for additional checking. **Example:**

```
select label roman
  when a=b then say 'same'
  when a<b then say 'lo'
  otherwise
    say 'hi'
    if a=0 then leave roman
    say 'a non-0'
end roman
```

In this example, if the variable `a` has the value 0 and `b` is negative then just `"hi"` is displayed.

35.2 Protect phrase

If `protect` is given it must be followed by a *term* that evaluates to a value that is not just a type and is not of a primitive type; while the `select` construct is being executed, the value (object) is protected - that is, all the instructions in the `select` construct have exclusive access to the object.

Both `label` and `protect` may be specified, in any order, if required.

⁵⁸In the reference implementation, a `NoOtherwiseException` is raised.

⁵⁹Strictly speaking, `then` should only be recognized if not the name of a variable. In this special case, however, NetRexx language processors are permitted to treat `then` as reserved in the context of a `when` clause, to provide better performance and more useful error reporting.

35.3 Case phrase

If case is given it must follow any label or protect phrase, and must be followed by an *expression*.

When case is used, the expression following it is evaluated at the start of the select construct. The result of the expression is then compared, using the strict equality operator (`==`), to the result of evaluating the expression or expressions in each of the when clauses in turn until a match is found. As usual, if no match is found then control will pass to the instruction list (if any) following otherwise, and in this situation the absence of an otherwise is a run-time error. For example, in:

```
select case i+1
  when 1 then say 'one'
  when 1+1 then say 'two'
  when 3, 4, 5 then say 'many'
end
```

then if `i` had the value 1 then the message displayed would be "two".

The third when clause in the example demonstrates the use of the multiple expressions in a when clause in this context. Similar to a select without case, each expression is evaluated in turn from left to right and is then compared to the result of the case expression. As soon as one matches that result, execution of the when clause stops (any further expressions are not evaluated) and the instruction following the associated then clause is executed.

Notes:

1. When case is used, the result of evaluating the expression following each when no longer has to be 0 or 1. Instead, it must be possible to compare each result to the result of the case expression.
2. The case expression is evaluated only on entry to the select construct; it is not re-evaluated for each when clause.
3. An exception raised during evaluation of the case expression will be caught by a suitable catch clause in the construct, if one is present. Similarly, evaluation of the case expression is protected by the protect phrase, if one is present.
4. *In the reference implementation, a select case construct will be translated into a Java switch construct provided that it meets the following criteria:*
 - *The type of the case expression is byte, char, int, or short.*
 - *The value of all the expressions on the when clauses are primitive constants (that is, they consist of only constants of primitive types and operators valid for them and so may be evaluated at compile time).*
 - *No two expressions on the when clauses evaluate to the same value.*
 - *It is not subject to tracing.*

Under these conditions the semantics of the switch construct match those defined for select. The example shown above would be translated to a switch construct if `i` had type `int` and options `binary` were in effect.

35.4 Exceptions in select constructs

Exceptions that are raised by the instructions within the body of the group, or during evaluation of the case expression, may be caught using one or more catch clauses that name the *exception* that they will catch. When an exception is caught, the exception object that holds the details of the exception may optionally be assigned to a variable, *vare*.

Similarly, a *finally* clause may be used to introduce instructions that will always be executed at the end of the select group, even if an exception is raised (whether caught or not).

The *Exceptions* section (see page 149) has details and examples of catch and finally.

Signal instruction

```
signal term;
```

The `signal` instruction causes an “abnormal” change in the flow of control, by raising an *exception*.

The exception *term* may be a term that constructs or evaluates to an exception object, or it may be expressed as the name of an exception type (in which case the default constructor, with no arguments, for that type is used to construct an exception object). The exception object then represents the exception and is available, if required, when the exception is handled.

The handling of exceptions is detailed in the *Exceptions section* (see page 149). In summary, when an exception is signalled, all active pending `do` groups, `loop` loops, `if` constructs, and `select` constructs may be ended. For each one in turn, from the innermost:

1. No further clauses within the body of the construct will be executed (in this respect, `signal` acts like a `leave` for the construct).
2. The *instructionlist* following the first `catch` clause that matches the exception, if any, is executed.
3. The *instructionlist* following the `finally` clause for the construct, if any, is executed.

If a `catch` matched the exception the exception is deemed handled, and execution resumes as though the construct ended normally (unless a new exception was signalled in the `catch` or `finally` instruction lists, in which case it is processed). Otherwise, any enclosing construct is ended in the same manner. If there is no enclosing construct, then the current method is ended and the exception is signalled in the caller.

Examples:

```
signal RxErrorTrace  
signal DivideException('Divide by zero')
```

In the reference implementation, the term must either

- *evaluate to an object that is assignable to the type **Throwable** (for example, a subclass of **Exception** or **RuntimeException**).*
- *be a type that is a subclass of **Throwable**, in which case the default constructor (with no arguments) for the given type is used to construct the exception object.*

Trace instruction

```
trace traceoption;
```

where *traceoption* is one of:

```
tracesetting  
var [varlist]
```

where *tracesetting* is one of:

```
all  
methods  
off  
results
```

and *varlist* is one or more variable *names*
(optionally prefixed with a + or -)

The trace instruction is used to control the tracing of the execution of NetRexx methods, and is primarily used for debugging. It may change either the general trace *setting* or may select or deselect the tracing of individual variables.

Within methods, the trace instruction changes the trace setting or variables tracing when it is executed, and affects the tracing of all clauses in the method which are then executed (until changed by a later trace instruction).

One or more trace instructions may appear before the first method in a class, one of which may set the initial trace setting for all methods in the class (the default is off) and others may set up variables tracing that applies to all the methods in the class. These act as though the trace instructions were placed immediately following the method instruction in each method (except that they will not be traced).

Similarly, one or more trace instructions may be placed before the first class instruction in a program; they do not imply the start of a class. One of these may set the initial trace setting and others may set up variables tracing for all classes in the program (except interface classes) and act as though the trace instructions were placed immediately following the class instruction in each class.

37.1 Tracing clauses

The trace *setting* controls the tracing of clauses in a program, and may be one of the following:

all All clauses (except null clauses without commentary) which are in methods and which are executed after the trace instruction will be traced. If trace `all` is placed before the first method in the current class, the method instructions in the class, together with the values of the arguments passed to each method, will be traced when the method is invoked (that is, trace `all` implies trace `methods`).

methods All method clauses in the class will be traced when the method they introduce is invoked, together with the values of the arguments passed to each method; no other clauses, or results, will be traced. The trace `methods` instruction must be placed before the first method in the current class (as otherwise it would have no effect).

off Turns tracing off; no following clauses, variables, or results will be traced.

results All clauses (except null clauses without commentary) which are in methods and which are executed after the trace instruction will be traced, as though trace `all` had been requested. In addition, the results of all *expression* evaluations and any results assigned to a variable by an assignment, loop, or parse instruction are also traced.

If trace `results` is placed before the first method in the current class, the method instructions in the class will be traced when the method is invoked, together with the values of the arguments passed to each method.

Notes:

1. Tracing of clauses shows the data from the source of the program, starting at the first character of the first token of the clause and including any commentary from that point until the end of the clause.
2. When a loop is being traced, the loop clause itself will be traced on every iteration of the loop, as indicated by the programmer's model (see page 80) ; the end clause is only traced once, when the loop completes normally.
3. With trace `results`, an expression is not traced if it is immediately used for an assignment (in an assignment instruction, or when the control variable is initialized in a loop instruction). The assignment will trace the result of the expression.

37.2 Tracing variables

The `var` option adds names to a list of monitored variables; it can also remove names from the list. If the name of a variable in the current class or method is in the list, then trace `results` is turned on for any assignment, loop, or parse clause that assigns a new value to the named variable.

Variable names are specified by listing them after the `var` keyword. Each name may be optionally prefixed by a `+` or a `-` sign. A `+` sign indicates that the variable is to be added to the list of monitored variables (the default), and a `-` sign indicates that the variable is to be removed from the list. Blanks may be added before and after variable names and signs to separate the tokens and to improve readability. For example:

```
trace var a b c
-- now variables a, b, and c will be traced
trace var -b -c d
-- now variables a and d will be traced
```

Notes:

1. Names in the list following the `var` keyword are simple symbols that name variables in the current class or current method. The variables may be properties, method arguments, or local variables, and may be of any type, including arrays. The names are not case-sensitive; any variables whose names match, independent of case, will be monitored.
2. No variable name can appear more than once in the list on one `trace var` instruction. However, it is not an error to add the name of a variable which does not exist or is not then assigned a value. Similarly, it is not an error to remove a name which is not currently being monitored.
3. One or more `trace var` instructions (along with one other `trace` instruction) are allowed before the first method in a class. They all modify an initial list of monitored variables which is then used for all methods in the class. Similarly, `trace var` instructions are allowed before the first class in a program, in which case they apply to all classes (except interface classes).
4. Other `trace` instructions do not affect the list of monitored variables. The `trace off` instruction may be used to turn off tracing completely; in this case `trace var` (with or without any variable names) will then turn the tracing of variables back on, using the current (or modified) variable list.
5. For a `parse` instruction, only monitored variables have their assignments traced (unless `trace results` is already in effect).

37.3 The format of trace output

Trace output is either clauses from the program being traced, or results (such as the results from expressions).

The first clause or result traced on any line will be preceded by its line number in the program; this is right-justified in a space which allows for the largest line number in the program, plus one blank. Following clauses or results from the same line are preceded by white space of the same width; however, any change of line number causes the line number to be included.

Clauses that are traced will be displayed with the formatting (indentation) and layout used in the original source stream for the program, starting with the first character of the first token of the clause.

TABLE 8: Trace identifier tags

**	identifies the first line of the source of a single clause, <i>i.e.</i> , the data actually in the program.
-	identifies a continuation line from the source of a single clause. Continuations may be due to the use of a continuation character (see page 20) or to the use of a block comment (see page 17) which spans more than one line.
>a>	Identifies a value assigned to a method argument of the current method. The name of the argument is included in the trace.
>p>	Identifies a value assigned to a property. The name of the property is included in the trace if the property is in the current class.
>v>	Identifies a value assigned to a local variable in the current method. The name of the variable is included in the trace.
»>	Identifies the result of an expression evaluation that is not used for an assignment (for example, an argument expression in a method call).
+++	Reserved for error messages that are not supplied by the environment underlying the implementation.

Results (if requested) are converted to a string for tracing if necessary, are not indented, and have a double quote prefixed and suffixed so that leading and trailing blanks are apparent; if, however, the result being traced is **null** (see page 122) then the string "[**null**]" is shown (without quotes). For results with an associated name (the values assigned to local variables, method arguments, or properties in the current class), the name of the result precedes the data, separated by a single blank.

For clarity, implementations may replace "control codes" in the encoding of results (for example, EBCDIC values less than '\x40', or Unicode values less than '\x20') by a question mark ("?"). All lines displayed during tracing have a three character tag to identify the type of data being traced. This tag follows the line number (or the space for a line number), and is separated from the line number by a single blank. The traced clause or result follows the tag, after another blank. The identifier tags are listed in table 8.

If a trace line is produced in a different context (program or thread) from the preceding trace line (if any) then a *trace context* line is shown. This shows the name of the program that produced the trace line, and also the name of the thread (and thread group) of the context.

The thread group name is not shown if it is `main`, and in this case the thread name is then also suppressed if its name is `main`.

Examples: If the following instructions, starting on line 53 of a 120-line program, were executed:

```
trace all
if i=1 then say 'Hello'
```

```
    else say 'i<>1'  
say -  
  'A continued line'
```

the trace output (if *i* were 1) would be:

```
54 ** if i=1  
    **      then  
    **      say 'Hello'  
56 ** say -  
57 *- 'A continued line'
```

Similarly, for the 3-line program:

```
trace results  
number=1/7  
parse number before '.' after
```

the trace output would be:

```
2 ** number=1/7  
  >v> number "0.142857143"  
3 ** parse number before '.' after  
  >v> before "0"  
  >v> after "142857143"
```

Notes:

1. Trace output is written to an implementation-defined output stream (typically the "standard error" output stream, which lets it be redirected to a destination separate from the default destination for output which is used by the `say` instruction).
2. In some implementations, the use of `trace` instructions may substantially increase the size of classes and the execution time of methods affected by tracing.⁶⁰
3. With some implementations it may be possible to switch tracing on externally, without requiring modification to the program.

⁶⁰In the reference implementation, options `notrace` may be used to disable all `trace` instructions and hence ensure that tracing overhead is not accidentally incurred.

Program structure

A NetREXX *program* is a collection of clauses (see page 16) derived from a single implementation-defined source stream (such as a file). When a program is processed by a language processor⁶¹ it defines one or more classes. Classes are usually introduced by the `class` instruction (see page 62), but if the first is a standard class, intended to be run as a stand-alone application, then the `class` instruction can be omitted. In this case, NetREXX defines an implied class and initialization method that will be used.

The implied class and method permits the writing of “low boilerplate” programs, with a minimum of syntax. The simplest, documented, NetREXX program that has an effect might therefore be:

Example:

```
/* This is a very simple NetRexx program */  
say 'Hello World!'
```

In more detail, a NetREXX program consists of:

1. An optional *prolog* (package, import, and options instructions). Only one package instruction is permitted per program.
2. One or more class definitions, each introduced by a `class` instruction.

A *class definition* comprises:

1. The `class` instruction which introduces the class (which may be inferred, see below).
2. Zero or more property variable assignments, along with optional `properties` instructions that can alter their attributes, and optional `numeric` and `trace` instructions. Property variable assignments take the form of an *assignment* (see page 49), with an optional “=” and expression, which may:
 - just name a property (by omitting the “=” and expression of the assignment), in which case it refers to a string of type **REXX**
 - assign a type to the property (when the expression evaluates to just a type)
 - assign a type and initial value to the property (when the expression returns a value).
3. Zero or more method definitions, each introduced by a `method` instruction (which may be inferred if the `class` instruction is inferred, see below).

A *method definition* comprises:

⁶¹Such as a compiler or interpreter.

- Any NetREXX instructions, except the class, method, and properties instructions and those allowed in the prolog (the package, import, and options instructions).

Example:

```

/* A program with two classes */
import java.applet.  -- for example

class testclass extends Applet
  properties public
    state          -- property of type 'Rexx'
    i=int          -- property of type 'int'
  properties constant
    j=int 3        -- property initialized to '3'

  method start
    say 'I started'
    state='start'

  method stop
    say 'I stopped'
    state='stop'

class anotherclass
  method testing
    loop i=1 to 10
      say '1, 2, 3, 4...'
      if i=7 then return
    end
    return

  method anotherstest
    say '1, 2, 3, 4'

```

This example shows a prolog (with just an import instruction) followed by two classes. The first class includes two public properties, one constant property, and two methods. The second class includes no properties, but also has two methods.

Note that a return instruction implies no static scoping; the content of a method is ended by a method (or class) instruction, or by the end of the source stream. The return instruction at the end of the **testing** method is, therefore, unnecessary.

38.1 Program defaults

The following defaults are provided for NetREXX programs:

1. If, while parsing prolog instructions, some instruction that is not valid for the prolog and is not a class instruction is encountered, then a default class instruction (with an implementation-provided short name, typically derived from the name of the source stream) is inserted. If the instruction was not a method instruction, then a default method instruction (with a name

and attributes appropriate for the environment, such as **main**) is also inserted.

In this latter case, it is assumed that execution of the program will begin by invocation of the default method. In other words, a "stand-alone" application can be written without explicitly providing the class and method instructions for the first method to be executed. An example of such a program is given in Appendix A (see page 177) .

*In the reference implementation, the **main** method in a stand-alone application is passed the words forming the command string as an array of strings of type **java.lang.String** (one word to each element of the array). When the NetREXX reference implementation provides the **main** method instruction by default, it also constructs a NetREXX string of type **REXX** from this array of words, with a blank added between words, and assigns the string to the variable **arg**.*

The command string may also have been edited by the underlying operating system environment; certain characters may not be allowed, multiple blanks or whitespace may have been reduced to single blanks, etc.

2. If a method ends and the last instruction at the outer level of the method scope is not return then a return instruction is added if it could be reached. In this case, if a value is expected to be returned by the method (due to other return instructions returning values, or there being a returns keyword on the method instruction), an error is reported.

Language processors may provide options to prevent, or warn of, these defaults being applied, as desired.

Minor and Dependent classes

A *minor class* in NetREXX is a class whose name is qualified by the name of another class, called its *parent*, and a *dependent class* is a minor class that has a link to its parent class that allows a child object simplified access to its parent object and its properties.

39.1 Minor classes

A *minor class* in NetREXX is a class whose name is qualified by the name of another class, called its *parent*. This qualification is indicated by the form of the name of the class: the short name of the minor class is prefixed by the name of its parent class (separated by a period). For example, if the parent is called `Foo` then the full name of a minor class `Bar` would be written `Foo.Bar`. The short name, `Bar`, is used for the name of any constructor method for the class; outside the class it can only be used to identify the class in the context of the parent class (or from children of the minor class, see below).

The names of minor classes may be used in exactly the same way as other class names (types) in programs. For example, a property might be declared and initialized thus:

```
abar=Foo.Bar null    -- this has type Foo.Bar
```

or, if the class has a constructor, perhaps:

```
abar=Foo.Bar( )      -- constructs a Foo.Bar object
```

Minor classes must be in the same program (and hence in the same package) as their parent. They are introduced by a `class` instruction that specifies their full name, for example:

```
class Foo.Bar extends SomeClass
```

Minor classes must immediately follow their parent class.⁶²

Minor classes may have a parent which is itself a minor class, to any depth; the name and the positioning rules are extended as necessary. For example, the following classes might exist in a program:

```
class Foo
  class Foo.Bar
    class Foo.Bar.Nod
```

⁶²This allows compilers that generate Java source code to preserve line numbering.


```
class Foo.Bar.Pod
class Foo.Car
```

As before, the children of `Foo.Bar` immediately follow their parent. The list of children of `Foo` can be continued after the children of `Foo.Bar` have all been specified.

Note that the short name (last part of the name) of a minor class may not be the same as the short name of any of its parents (a class `Foo.Bar.Foo` or a class `Foo.Bar.Bar` would be in error, for example). This allows minor classes to refer to their parent classes by their short name without ambiguity.

39.1.1 Constructing objects in minor classes

A parent class can construct an object of a child class in the usual manner, by simply specifying its constructor (identified by its short name, full name, or qualified name). For example, a method in the `Foo.Bar` class above could construct an object of type `Foo.Bar.Nod` using:

```
anod=Nod( )
```

(assuming the `Foo.Bar.Nod` class has a constructor that takes no arguments).

Similarly, minor classes can refer to the types and constructors of any of its parents by simply using their short names. Hence, the `Foo.Bar.Nod` class could construct objects of its parents' types thus:

```
abar=Bar( )
afoo=Foo( )
```

(again assuming the parent classes have constructors that take no arguments).

Classes other than the parent or an immediate child must use the full name (if necessary, qualified by the package name) to refer to a minor class or its constructor.

39.2 Dependent classes

As described in the last section, minor classes provide an enhanced packaging (naming) mechanism for classes, allowing classes to be structured within packages. A stronger link between a child class and its parent is indicated by the modifier keyword `dependent` on the child class, which indicates that the child is a *dependent class*. For example:

```
class Foo.Dep dependent extends SomeClass
method Dep -- this is the constructor
```

An object constructed from a dependent class (a *dependent object*) is linked to the context of an object of its parent type (its *parent object*). The linkage thus provided allows the child object simplified access to the parent object and its properties.

In the example, an object of type `Foo.Dep` can only be constructed in the context of a parent object, which must be of type `Foo`.

39.2.1 Constructing dependent objects

A parent class can construct a dependent object in the same way as when constructing objects of other child types; that is, by simply specifying its constructor. In this case, however, the current object (`this`) becomes the parent object of the newly constructed object. For example, a method in the `Foo` class above could construct a dependent object of type `Foo.Dep` using:

```
adep=Dep( )
```

(assuming the `Dep` class has a constructor that takes no arguments).

In general, for a class to construct an object from a dependent class, it must have a reference to an object of the parent class (which will become the parent of the new object), and the constructor must be called (by its short name) in the context of that parent object. For example:

```
parentObject=Foo( )  
adep=parentObject.Dep( )
```

(In the same way, the first example could have been written:

```
adep=this.Dep( )
```

within the parent class the `this.` is implied.)

In order to subclass a dependent class, the constructor of the dependent class must be invoked by the subclass constructor in a similar manner. In this case, a qualified call to the usual special constructor `super` is used, for example:

```
class ASub extends Foo.Dep  
  method ASub(afoo=Foo)  
    afoo.super( )
```

The qualifier (`afoo` in the example) must be either the name of an argument to the constructor, or the special word `parent` (if the classes share a common parent class), or the short name of a parent class followed by `.this` (see below). The call to `super` must be the first instruction in the method, as usual, and it must be present (it will not be generated automatically by the compiler).

39.2.2 Access to parent objects and their properties

Dependent classes have simplified access to their parent objects and their properties. In particular:

- The special word `parent` may be used to refer to the parent object of the current object. It may appear alone in a term, or at the start of a compound term. It can only be used in non-static contexts in a dependent class.
- In general, any of the objects in the chain of parents of a dependent object may be referred to by qualifying the special word `this` with the short

name of the parent class. For example, extending the previous example, if the class `Foo.Dep.Ent` was a dependent class it could contain references to `Foo.this` (the parent of its parent) or `Dep.this` (the latter being the same as specifying parent). If preferred, the full name or the fully qualified name of the parent class may be used instead of the short name.

Like `parent`, this construct can only be used at the start of a term in non-static contexts in a dependent class.

- As usual, properties external to the current class must always be qualified in some way (for example, the prefix `parent.` can be used in a term such as `parent.aprop`).

39.3 Restrictions

Minor classes may have any of the attributes (`public`, `interface`, *etc.*) of other classes, and behave in every way like other classes, with the following restrictions:

- If a class is a static class (that is, it contains only static or constant properties and methods) then any children cannot be dependent classes (because no object of the parent class can be constructed). Similarly, interface classes and abstract classes cannot have dependent classes.
- Dependent classes may not be interfaces.
- Dependent classes may not contain static or constant properties (or methods).⁶³ These must be placed in a parent which is not a dependent class.
- Minor classes may be `public` only if their parent is also `public`. (Note that this is the only case where more than one `public` class is permitted in a program.) In general: a minor class cannot be more visible than its parent.

⁶³This restriction allows compilation for the Java platform.

Special names and methods

For convenience, NetREXX provides some *special names* for naming commonly-used concepts within terms. These are only recognized if there is no variable of the same name previously seen in the current scope, as described in the section on *Terms* (see page 26) . This allows the set of special words to be expanded in the future, if necessary, without invalidating existing variables. Therefore, these names are not reserved; they may be used as variable names instead, if desired. There are also two "special methods" that are used when constructing objects.

40.1 Special names

The following special names are allowed in NetREXX programs, and are recognized independently of case.⁶⁴ With the exception of **length** and **class**, these may only be used alone as a term or at the start of a compound term.

ask Returns a string of type **REXX**, read as a line from the implementation-defined default input stream (often the user's "console").

Example:

```
if ask='yes' then say 'OK'
```

ask can only appear alone, or at the start of a compound term.⁶⁵

asknoecho Returns a string of type **REXX**, read as a line from the implementation-defined default input stream (often the user's "console"), without an echo to the screen of the typed character(s).

Example:

```
if asknoecho='yes' then say 'OK'
```

asknoecho can only appear alone, or at the start of a compound term.

⁶⁶

class The object of type **Class** that describes a specific type. This word is only recognized as the second part of a compound term, where the evaluation of the first part of the term resulted in a type or qualified type.

Example:

```
obj=String.class
say obj.isInterface /* would say '0' */
```

⁶⁴Unless options *strictcase* is in effect.

⁶⁵In the reference implementation, **ask** is simply a shorthand for `REXXIO.Ask()`.

⁶⁶In the reference implementation, **asknoecho** is shorthand for `System.console().ReadPassword()`.

digits The current setting of numeric digits (see page 89) , returned as a string of type **REXX**. This will be one or more Arabic numerals, with no leading blanks, zeros, or sign, and no trailing blanks or exponent. **digits** can only appear alone, or at the start of a compound term.

form The current setting of numeric form (see page 89) , returned as a string of type **REXX**. This will have either the value "**scientific**" or the value "**engineering**". **form** can only appear alone, or at the start of a compound term.

length The length of an array (see page 54) , returned as an implementation-dependent binary type or string. This word is only recognized as the last part of a compound term, where the evaluation of the rest of the term resulted in an array of dimension 1.

Example:

```
foo=char[7]
say foo.length      /* would say '7' */
```

Note that you can get the length of a NetREXX string with the same syntax.
⁶⁷

In that case, however, a **length()** method is being invoked.

null The *empty reference*. This is a special value that represents "no value" and may be assigned to variables (or returned from methods) except those whose type is both primitive and undimensioned. It may also be used in a comparison for equality (or inequality) with values of suitable type, and may be given a type. **Examples:**

```
blob=int[3]    -- 'blob' refers to array of 3 ints
blob=null     -- 'blob' is still of type int[],
               -- but refers to no real object
mob=Mark null -- 'mob' is type 'Mark'
```

The **null** value may be considered to represent the state of being uninitialized. It can only appear as simple symbol, not as a part of a compound term.

RC The special variable named *RC* holds the returncode of the last executed command sent to the *address* environment. When the environment was never addressed, *RC* holds the value 'RC'.

source Returns a string of type **REXX** identifying the source of the current class. The string consists of the following words, with a single blank between the words and no trailing or leading blanks:

1. the name of the underlying environment (*e.g.*, **Java**)
2. either **method** (if the term is being used within a method) or **class** (if the term is being used within a property assignment, before the first method in a class)
3. an implementation-dependent representation of the name of the source stream for the class (*e.g.*, **Fred.nrx**).

source can only appear alone, or at the start of a compound term.

⁶⁷Unless *options strictargs* is in effect.

sourceline The line number of the first token of the current clause in the NetREXX program, returned as a string of type **Rexx**. This will be one or more Arabic numerals, with no leading blanks, zeros, or sign, and no trailing blanks or exponent. **sourceline** can only appear alone, or at the start of a compound term.

super Returns a reference to the current object, with a type that is the type of the class that the current object's class extends. This means that a search for methods or properties which **super** qualifies will start from the superclass rather than in the current class. This is used for invoking a method or property (in the superclass or one of its superclasses) that has been overridden in the current class. **Example:**

```
method printit(x)
  say 'it'           -- modification
  super.printit(x)  -- now the usual processing
```

If a property being referenced is in fact defined by a superclass of the current class, then the prefix "**super.**" is perhaps the clearest way to indicate that name refers to a property of a superclass rather than to a local variable. (You could also qualify it by the name of the superclass.) **super** can only appear alone, or at the start of a compound term.

this Returns a reference to the current object. When a method is invoked, for example in:

```
word=R\textsc{exx} "hello"  -- 'word' refers to "hello"
say word.substr(3)         -- invokes substr on "hello"
```

then the method **substr** in the class **Rexx** is invoked, with argument '3', and with the properties of the value (object) "**hello**" available to it. These properties may be accessed simply by name, or (more explicitly) by prefixing the name with "**this.**". Using "**this.**" can make a method more readable, especially when several objects of the same type are being manipulated in the method. **this** can only appear alone, or at the start of a compound term.

trace The current trace (see page 109) setting, returned as a NetREXX string. This will be one of the words:

```
off var methods all results
```

(var is returned when clause tracing is off but variable tracing has then been turned on using a trace var instruction.) **trace** can only appear alone, or at the start of a compound term.

version Returns a string of type **Rexx** identifying the version of the NetREXX language in effect when the current class was processed. The string consists of the following words, with a single blank between the words and no trailing or leading blanks:

1. A word describing the language. The first seven letters will be the characters **NetREXX**, and the remainder may be used to identify a particular implementation or language processor. This word may not include any periods.
2. The language level description, which must be a number with no sign or exponential part. For example, "**4.02-GA**" is the language level of this definition.

3. Three words describing the language processor release date in the same format as the default for the REXX "**date()**" function.⁶⁸ For example, "**22 May 2009**".
- version** can only appear alone, or at the start of a compound term.

40.2 Special methods

Constructors (methods used for constructing objects) in NetREXX must invoke a constructor of their superclass before making any modifications to the current object (or invoke another constructor in the current class).

This is simplified and made explicit by the provision of the special method names **super** and **this**, which refer to constructors of the superclass and current class respectively. These special methods are only recognized when used as the first, method call, instruction in a constructor, as described in *Methods and constructors* (see page 32) . Their names will be recognized independently of case.⁶⁹

In addition, NetREXX provides special constructor methods for the primitive types that allow binary construction of primitives. These are described in *Binary values and arithmetic* (see page 148).

⁶⁸As defined in **American National Standard for Information Technology - Programming Language REXX, X3.274-1996**, American National Standards Institute, New York, 1996. See also `Date()` on page 187.

⁶⁹Unless options *strictcase* is in effect.

JavaBean Support

This chapter describes the *indirect properties* feature.

The intention of this feature is to make it easier to write a certain kind of class known as a *JavaBean*. Almost all JavaBeans will have *properties*, which are data items that a user of a JavaBean is expected to be able to customize (for example, the text on a pushbutton). The names and types of the properties of a JavaBean are inferred from "*design patterns*" (in this context, conventions for naming methods) or from `PropertyDescriptor` objects associated with the JavaBean.

The JavaBean properties do not necessarily correspond to instance variables in the class - although very often they do. The JavaBean specification does not guarantee that JavaBean properties that can be set can also be inspected, nor does it describe how ambiguities of naming and method signatures are to be handled.

The NetREXXC compiler allows a more rigorous treatment of JavaBean properties, by allowing an optional attribute of properties in a class that declares them to be *indirect properties*. Indirect properties are properties of a known type that are private to the class, but which are expected to be publicly accessible indirectly, though certain conventional method calls.

Declaring properties to be indirect offers the following advantages:

- For many simple cases, the access methods for the properties can be generated automatically; there is no need to explicitly code them in the source file for the class. This is especially helpful for Indexed Properties (where four methods are needed, in general).
- Where access methods are explicitly provided in the class, they can be checked for correct form, signature and accessibility. This detects errors at compile time that otherwise would only be determined by testing.
- Similarly, attention can be drawn to the presence of methods that may be intended to be an access method for an indirect property, but will not be recognized as such by builders.

The next section describes the use of indirect properties in more detail.

41.1 Indirect properties

The `properties` instruction (see page 99) is used to define the attributes of following *property* variables. The *visibility* of properties may include a new alter-

native: indirect. Properties with this form of visibility are known as *indirect properties*. These are properties of a known type that are private to the class, but which are expected to be publicly accessible indirectly, though certain conventional method calls.

For example, consider the simple program:

```
class Sandwich extends Canvas implements Serializable
    properties indirect
        slices=Color.gray
        filling=Color.red

    method Sandwich
        resize(100,30)

    method paint(g=Graphics)
        g.setColor(slices)
        g.fillRect(0, 0, size.width, size.height)
        g.setColor(filling)
        g.fillRect(12, 12, size.width-12, size.height-12)
```

This declares the **Sandwich** class as having two indirect properties, called **slices** and **filling**, both being of type **java.awt.Color**.

In the example, no access methods are provided for the properties, so the compiler will add them. By implementation-dependent convention, the names are prefixed with verbs such as **get** and **set**, *etc.*, and have the first character of their name uppercased to form the method names. Hence, in this Java-based example, the following four methods are added:

```
method getSlices returns java.awt.Color
    return slices
method getFilling returns java.awt.Color
    return filling
method setSlices($1=java.awt.Color)
    slices=$1
method setFilling($2=java.awt.Color)
    filling=$2
```

(where **\$1** and **\$2** are "hidden" names used for accessing the method arguments).

Note that the indirect attribute for a property is an alternative to the public, private, and inheritable, and shared attributes. Like private properties, indirect properties can only be accessed directly by name from within the class in which they occur; other classes can only access them using the access methods (or other methods that may use, or have a side-effect on, the properties).

Indirect properties may be constant (implying that only a get method is generated or allowed, though the private property may be changed by methods within the class) or transient (see page 100) . They may not be static or volatile.

In detail, the rules used for generating automatic methods for a property whose name is **xxxx** are as follows:

1. A method called **getXxxx** which returns the value of the property is gener-

- ated. The returned value will have the same type as `xxxx`.
2. If the type of `xxxx` is **boolean** then the generated method will be called `isXxxx` instead of `getXxxx`.
 3. If the property is not constant then a method for setting the property will also be generated. This will be called `setXxxx`, and take a single argument of the same type as `xxxx`. This assigns the argument to the property and returns no value.

If the property has an array type (for example, `char[]`), then it must only have a single dimension. Two further methods may then be generated, according to the rules:

1. A method called `getXxxx` which takes a single **int** as an argument and which returns an item from the property array is generated. The returned value will have the same type as `xxxx`, without the `[]`. The integer argument is used to index into the array.
2. As before, if the result type of the method would be **boolean** then the name of the method will be `isXxxx` instead of `getXxxx`.
3. If the property is not constant then a method for setting an item in the property array will also be generated. This will be called `setXxxx`, and take two arguments: the first is an **int** that is used to select the item to be changed, and the second is an undimensioned argument of the same type as `xxxx`. It assigns the second argument to the item in the property array indexed by the first argument, and returns no value.

For example, for an indirect property declared thus:

```
properties indirect
    fred=foo.Bar[]
```

the four methods generated would be:

```
method getFred returns foo.Bar[]; return fred
method getFred($1=int) returns foo.Bar; return fred[$1]
method setFred($2=foo.Bar[]); fred=$2
method setFred($3=int, $4=foo.Bar); fred[$3]=$4
```

Note that in all cases a method will only be generated if it would not exactly match a method explicitly coded in the current class.

41.1.1 Explicit provision of access methods

Often, for example when an indirect property has an on-screen representation, it is desirable to redraw the property when the property is changed (and in more complicated cases, there may be interactions between properties). These and other actions will require extra processing which will not be carried out by automatically generated methods. To add this processing the access methods will have to be coded explicitly. In the "Sandwich" example, we only need to supply the **set** methods, perhaps by adding the following to the example class above:

```

method setSlices(col=Color)
    slices=col      -- update the property
    this.repaint    -- redraw the component

method setFilling(col=Color)
    filling=col
    this.repaint

```

If we add these two methods, they will no longer be added automatically (the two **get** methods will continue to be provided automatically, however). Further, since the names match possible access methods for properties that are declared to be indirect, the compiler will check the method declaration: the method signatures and return type (if any) must be correct, for example. Also, since the names of access methods are case-sensitive (in a Java environment), you will be warned if a method appears to be intended to be an access method but the case of one or more letters is wrong.

Specifically, the checks carried out are as follows:

1. For methods whose names exactly match a potential access method for an indirect property (that is, start with **is**, **get**, or **set**, which is then followed by the name of an indirect property with the first character of the name uppercased):
 - The argument list for (signature of) the method must match one of those that could possibly be automatically generated for the property.
 - The returns type (if any) must match the expected returns type for that method.
 - If the returns type is simply **boolean**, then the method name must start with **is**. Conversely, if the method name starts with **is** then the returns type must be just **boolean**.
 - If the property is constant then the name of the method cannot start with **set**.
 - A warning is given if the method is not `public` (the default).
2. For methods whose names match a potential access method, as above, except in case:
 - A warning is given that the method in question may be intended to be an indirect property access method, but will not be recognized as such by builders.

These checks detect a wide variety of errors at compile time, hence speeding the development of classes that use indirect properties.

Parsing templates

The parse instruction allows a selected string to be parsed (split up) and assigned to variables, under the control of a *template*.

The various mechanisms in the template allow a string to be split up by explicit matching of strings (called *patterns*), or by specifying numeric positions (*positional patterns* - for example, to extract data from particular columns of a line read from a character stream). Once split into parts, each segment of the string can then be assigned to variables as a whole or by words (delimited by blanks).

This section first gives some informal examples of how the parsing template can be used, and then defines the algorithms in detail.

42.1 Introduction to parsing

The simplest form of parsing template consists of a list of variable names. The string being parsed is split up into words (characters delimited by blanks), and each word from the string is assigned to a variable in sequence from left to right. The final variable is treated specially in that it will be assigned whatever is left of the original string and may therefore contain several words. For example, in the parse instruction:

```
parse 'This is a sentence.' v1 v2 v3
```

the term (in this case a literal string) following the instruction keyword is parsed, and then: the variable *v1* would be assigned the value **"This"**, *v2* would be assigned the value **"is"**, and *v3* would be assigned the value **"a sentence."**.

Leading blanks are removed from each word in the string before it is assigned to a variable, as is the blank that delimits the end of the word. Thus, variables set in this manner (*v1* and *v2* in the example) will never have leading or trailing blanks, though *v3* could have both leading and trailing blanks. Note that the variables assigned values in a template are always given a new value and so if there are fewer words in the string than variables in the template then the unused variables will be set to the null string. The second parsing mechanism uses a literal string in a template as a pattern, to split up the string. For example:

```
parse 'To be, or not to be?' w1 ',' w2
```

would cause the string to be scanned for the comma, and then split at that point; the variable *w1* would be set to **"To be"**, and *w2* is set to **" or not to be?"**. Note that the pattern itself (and **only** the pattern) is removed from the string. Each

section of the string is treated in just the same way as the whole string was in the previous example, and so either section could be split up into words. Thus, in:

```
parse 'To be, or not to be?' w1 ',' w2 w3 w4
```

w2 and *w3* would be assigned the values "or" and "not", and *w4* would be assigned the remainder: "to be?".

If the string in the last example did not contain a comma, then the pattern would effectively "match" the end of the string, so the variable to the left of the pattern would get the entire input string, and the variables to the right would be set to a null string. The pattern may be specified as a variable, by putting the variable name in parentheses. The following instructions therefore have the same effect as the last example:

```
c=', '  
parse 'To be, or not to be?' w1 (c) w2 w3 w4
```

The third parsing mechanism is the numeric positional pattern. This works in the same way as the string pattern except that it specifies a column number. So:

```
parse 'Flying pigs have wings' x1 5 x2
```

would split the string at the fifth column, so *x1* would be "Flyi" and *x2* would start at column 5 and so be "ng pigs have wings". More than one pattern is allowed, so for example:

```
parse 'Flying pigs have wings' x1 5 x2 10 x3
```

would split the string at columns 5 and 10, so *x2* would be "ng pi" and *x3* would be "gs have wings". The numbers can be relative to the last number used, so:

```
parse 'Flying pigs have wings' x1 5 x2 +5 x3
```

would have exactly the same effect as the last example; here the +5 may be thought of as specifying the length of the string to be assigned to *x2*. As with literal string patterns, the positional patterns can be specified as a variable by putting the name of a variable, in parentheses, in place of the number. An absolute column number should then be indicated by using an equals sign ("=") instead of a plus or minus sign. The last example could therefore be written:

```
start=5  
length=5  
data='Flying pigs have wings'  
parse data x1 =(start) x2 +(length) x3
```

String patterns and positional patterns can be mixed (in effect the beginning of a string pattern just specifies a variable column number) and some very powerful things can be done with templates. The next section describes in more detail how the various mechanisms interact.

42.2 Parsing definition

This section describes the rules that govern parsing. In its most general form, a template consists of alternating pattern specifications and variable names. Blanks may be added between patterns and variable names to separate the tokens and to improve readability. The patterns and variable names are used strictly in sequence from left to right, and are used once only. In practice, various simpler forms are used in which either variable names or patterns may be omitted; we can therefore have variable names without patterns in between, and patterns without intervening variable names. In general, the value assigned to a variable is that sequence of characters in the input string between the point that is matched by the pattern on its left and the point that is matched by the pattern on its right. If the first item in a template is a variable, then there is an implicit pattern on the left that matches the start of the string, and similarly if the last item in a template is a variable then there is an implicit pattern on the right that matches the end of the string. Hence the simplest template consists of a single variable name which in this case is assigned the entire input string. Setting a variable during parsing is identical in effect to setting a variable in an assignment. The constructs that may appear as patterns fall into two categories; patterns that act by searching for a matching string (literal patterns), and numeric patterns that specify an absolute or relative position in the string (positional patterns). Either of these can be specified explicitly in the template, or alternatively by a reference to a variable whose value is to be used as the pattern. For the following examples, assume that the following sample string is being parsed; note that all blanks are significant - there are two blanks after the first word "is" and also after the second comma:

```
'This is  the text which, I think,  is scanned.'
```

42.2.1 Parsing with literal patterns

Literal patterns cause scanning of the data string to find a sequence that matches the value of the literal. Literals are expressed as a quoted string. The null string matches the end of the data. The template:

```
w1 ', ' w2 ', ' w3
```

when parsing the sample string, results in:

```
w1 has the value "This is  the text which"  
w2 has the value " I think"  
w3 has the value "  is scanned."
```

Here the string is parsed using a template that asks that each of the variables receive a value corresponding to a portion of the original string between commas; the commas are given as quoted strings. Note that the patterns themselves are removed from the data being parsed. A different parse would result with the template:

```
w1 ', ' w2 ', ' w3 ', ' w4
```

which would result in:

```
w1 has the value "This is the text which"  
w2 has the value " I think"  
w3 has the value " is scanned."  
w4 has the value "" (null string)
```

This illustrates an important rule. When a match for a pattern cannot be found in the input string, it instead "matches" the end of the string. Thus, no match was found for the third ',' in the template, and so *w3* was assigned the rest of the string. *w4* was assigned a null string because the pattern on its left had already reached the end of the string. Note that **all** variables that appear in a template in this way are assigned a new value.

42.2.2 Parsing strings into words

If a variable is directly followed by one or more other variables, then the string selected by the patterns is assigned to the variables in the following manner. Each blank-delimited word in the string is assigned to each variable in turn, except for the last variable in the group (which is assigned the remainder of the string). The values of the variables which are assigned words will have neither leading nor trailing blanks. Thus the template:

```
w1 w2 w3 w4 ', '
```

would result in:

```
w1 has the value "This'  
w2 has the value "is"  
w3 has the value "the"  
w4 has the value "text which"
```

Note that the final variable (*w4* in this example) could have had both leading blanks and trailing blanks, since only the blank that delimits the previous word is removed from the data. Also observe that this example is not the same as specifying explicit blanks as patterns, as the template:

```
w1 ' ' w2 ' ' w3 ' ' w4 ', '
```

would in fact result in:

```
w1 has the value "This'  
w2 has the value "is"  
w3 has the value "" (null string)  
w4 has the value "the text which"
```

since the third pattern would match the third blank in the data. In general, when a variable is followed by another variable then parsing of the input into individual words is implied. The parsing process may be thought of as first splitting the original string up into other strings using the various kinds of patterns, and then assigning each of these new strings to (zero or more) variables.

42.2.3 Use of the period as a placeholder

A period (separated from any symbols by at least one blank) acts as a placeholder in a template. It has exactly the same effect as a variable name, except that no variable is set. It is especially useful as a "dummy variable" in a list of variables, or to collect (ignore) unwanted information at the end of a string. Thus the template:

```
. . . word4 .
```

would extract the fourth word ("**text**") from the sample string and place it in the variable *word4*. Blanks between successive periods in templates may be omitted, so the template:

```
... word4 .
```

would have the same result as the last template.

42.2.4 Parsing with positional patterns

Positional patterns may be used to cause the parsing to occur on the basis of position within the string, rather than on its contents. They take the form of whole numbers, optionally preceded by a plus, minus, or equals sign which indicate relative or absolute positioning. These may cause the matching operation to "back up" to an earlier position in the data string, which can only occur when positional patterns are used. **Absolute positional patterns:** A number in a template that is **not** preceded by a sign refers to a particular (absolute) character column in the input, with 1 referring to the first column. For example, the template:

```
s1 10 s2 20 s3
```

results in:

```
s1 has the value "This is "  
s2 has the value "the text w"  
s3 has the value "hich, I think, is scanned."
```

Here *s1* is assigned characters from the first through the ninth character, and *s2* receives input characters 10 through 19. As usual the final variable, *s3*, is assigned the remainder of the input.

An equals sign ("=") may be placed before the number to indicate explicitly that it is to be used as an absolute column position; the last template could have been written:

```
s1 =10 s2 =20 s3
```

A positional pattern that has no sign or is preceded by the equals sign is known as an *absolute positional pattern*. **Relative positional patterns:** A number in a template that is preceded by a plus or minus sign indicates movement relative to the character position at which the previous pattern match occurred. This is a *relative positional pattern*. If a plus or minus is specified, then the position used

for the next match is calculated by adding (or subtracting) the number given to the last matched position. The last matched position is the position of the first character of the last match, whether specified numerically or by a string.

For example, the instructions:

```
parse '123456789' 3 w1 +3 w2 3 w3
```

result in

w1 has the value "345"

w2 has the value "6789"

w3 has the value "3456789"

The `+3` in this case is equivalent to the absolute number `6` in the same position, and may also be considered to be specifying the length of the data string to be assigned to the variable `w1`. This example also illustrates the effects of a positional pattern that implies movement to a character position to the left of (or to) the point at which the last match occurred. The variable on the left is assigned characters through the end of the input, and the variable on the right is, as usual, assigned characters starting at the position dictated by the pattern. A useful effect of this is that multiple assignments can be made:

```
parse x 1 w1 1 w2 1 w3
```

This results in assigning the (entire) value of `x` to `w1`, `w2`, and `w3`. (The first `"1"` here could be omitted as it is effectively the same as the implicit starting pattern described at the beginning of this section.) If a positional pattern specifies a column that is greater than the length of the data, it is equivalent to specifying the end of the data (*i.e.*, no padding takes place). Similarly, if a pattern specifies a column to the left of the first column of the data, this is not an error but instead is taken to specify the first column of the data. Any pattern match sets the "last position" in a string to which a relative positional pattern can refer. The "last position" set by a literal pattern is the position at which the match occurred, that is, the position in the data of the *first* character in the pattern. The literal pattern in this case is **not** removed from the parsed data. Thus the template:

```
',' -1 x +1
```

will:

1. Find the first comma in the input (or the end of the string if there is no comma).
2. Back up one position.
3. Assign one character (the character immediately preceding the comma or end of string) to the variable `x`.

One possible application of this is looking for abbreviations in a string. Thus the instruction:

```
/* Ensure options have a leading blank and are
   in uppercase before parsing. */
parse (' 'opts).upper ' PR' +1 prword ' '
```

will set the variable *prword* to the first word in *opts* that starts with "PR" (in any case), or will set it to the null string if no such word exists. **Notes:**

1. The positional patterns **+0** and **-0** are valid, have the same effect, and may be used to include the whole of a previous literal (or variable) pattern within the data string to be parsed into any following variables.
2. As illustrated in the last example, patterns may follow each other in the template without intervening variable names. In this case each pattern is obeyed in turn from left to right, as usual.
3. There may be blanks between the sign in a positional pattern and the number, because NetREXX defines that blanks adjacent to special characters are removed.

42.2.5 Parsing with variable patterns

It is sometimes desirable to be able to specify a pattern by using the value of a variable instead of a fixed string or number. This may be achieved by placing the name of the variable to be used as the pattern in parentheses (blanks are not necessary either inside or outside the parentheses, but may be added if desired). This is called a *variable reference*; the value of the variable is converted to string before use, if necessary. If the parenthesis to the left of the variable name is not preceded by an equals, plus, or minus sign ("=", "+", or "-") the value of the variable is then used as though it were a literal (string) pattern. The variable may be one that has been set earlier in the parsing process, so for example:

```
input="L/look for/1 10"  
parse input verb 2 delim +1 string (delim) rest
```

will set:

```
verb to 'L'  
delim to '/'  
string to 'look for'  
rest to '1 10'
```

If the left parenthesis is preceded by an equals, plus, or minus sign then the value of the variable is used as an absolute or relative positional pattern (instead of as a literal string pattern). In this case the value of the variable must be a non-negative whole number, and (as before) it may have been set earlier in the parsing process.

Numbers and Arithmetic

NetREXX arithmetic attempts to carry out the usual operations (including addition, subtraction, multiplication, and division) in as "natural" a way as possible. What this really means is that the rules followed are those that are conventionally taught in schools and colleges. However, it was found that unfortunately the rules used vary considerably (indeed much more than generally appreciated) from person to person and from application to application and in ways that are not always predictable. The NetREXX arithmetic described here is therefore a compromise which (although not the simplest) should provide acceptable results in most applications.

43.1 Introduction

Numbers can be expressed in NetREXX very flexibly (leading and trailing blanks are permitted, exponential notation may be used) and follow conventional syntax. Some valid numbers are:

```

12          /* A whole number          */
'-76'       /* A signed whole number             */
12.76      /* Some decimal places              */
' + 0.003 ' /* Blanks around the sign, etc.    */
17.        /* Equal to 17                      */
'.5'       /* Equal to 0.5                     */
4E+9       /* Exponential notation             */
0.73e-7    /* Exponential notation             */

```

(Exponential notation means that the number includes a sign and a power of ten following an "E" that indicates how the decimal point will be shifted. Thus **4E+9** above is just a short way of writing **4000000000**, and **0.73e-7** is short for **0.000000073**.) The arithmetic operators include addition (indicated by a "+"), subtraction ("-"), multiplication ("*"), power ("**"), and division ("/"). There are also two further division operators: integer divide ("%") which divides and returns the integer part, and remainder ("//") which divides and returns the remainder. Prefix plus ("+") and prefix minus ("-") operators are also provided. When two numbers are combined by an operation, NetREXX uses a set of rules to define what the result will be (and how the result is to be represented as a character string). These rules are defined in the next section, but in summary:

- Results will be calculated with up to some maximum number of significant digits. That is, if a result required more than 9 digits it would normally be rounded to 9 digits. For instance, the division of 2 by 3 would result

in 0.666666667 (it would require an infinite number of digits for perfect accuracy).

You can change the default of 9 significant digits by using the `numeric digits` instruction. This lets you calculate using as many digits as you need - thousands, if necessary.

- Except for the division and power operators, trailing zeros are preserved (this is in contrast to most electronic calculators, which remove all trailing zeros in the decimal part of results). So, for example:

```
2.40 + 2  =>  4.40
2.40 - 2  =>  0.40
2.40 * 2  =>  4.80
2.40 / 2  =>  1.2
```

This preservation of trailing zeros is desirable for most calculations (and especially financial calculations). If necessary, trailing zeros may be easily removed with the `strip` method (see page 171), or by division by 1.

- A zero result is always expressed as the single digit '0'.
- Exponential form is used for a result depending on its value and the setting of `numeric digits` (the default is 9 digits). If the number of places needed before the decimal point exceeds this setting, or the absolute value of the number is less than **0.000001**, then the number will be expressed in exponential notation; thus

```
1e+6 * 1e+6
```

results in "**1E+12**" instead of "**1000000000000**", and

```
1 / 3E+10
```

results in "**3.33333333E-11**" instead of "**0.000000000333333333**".

- Any mixture of Arabic numerals (0-9) and Extra digits (see page 18) can be used for the digits in numbers used in calculations. The results are expressed using Arabic numerals.

43.2 Definition

This definition describes arithmetic for NetREXX strings (type **REXX**). The arithmetic operations are identical to those defined in the ANSI standard for REXX.

⁷⁰

43.2.1 Numbers

A *number* in NetREXX is a character string that includes one or more decimal digits, with an optional decimal point. The decimal point may be embedded in the digits, or may be prefixed or suffixed to them. The group of digits (and optional point) thus constructed may have leading or trailing blanks, and an optional sign ("**+**" or "**-**") which must come before any digits or decimal point. The sign may also have leading or trailing blanks. Thus:

⁷⁰American National Standard for Information Technology - Programming Language REXX, X3.274-1996, American National Standards Institute, New York, 1996.

```

sign    ::= + | -
digit   ::= 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9
digits  ::= digit [digit]...
numeric ::= digits . [digits]
          | [.] digits
number  ::= [blank]... [sign [blank]...]
          numeric [blank]...

```

where if the implementation supports extra digits (see page 18) these are also accepted as *digits*, providing that they represent values in the range zero through nine. In this case each extra digit is treated as though it were the corresponding character in the range 0-9. Note that a single period alone is not a valid number.

43.2.2 Precision

The maximum number of significant digits that can result from an arithmetic operation is controlled by the `digits` keyword on the `numeric` instruction (see page 89):

```
numeric digits [expression];
```

The expression is evaluated and must result in a positive whole number. This defines the precision (number of significant digits) to which arithmetic calculations will be carried out; results will be rounded to that precision, if necessary. If no expression is specified, then the default precision is used. The default precision is 9, that is, all implementations must support at least nine digits of precision. An implementation-dependent maximum (equal to or larger than 9) may apply: an attempt to exceed this will cause execution of the instruction to terminate with an exception. Thus if an algorithm is defined to use more than 9 digits then if the `numeric digits` instruction succeeds then the computation will proceed and produce identical results to any other implementation. Note that `numeric digits` may set values below the default of nine. Small values, however, should be used with care - the loss of precision and rounding thus requested will affect all NetREXX computations, including (for example) the computation of new values for the control variable in loops.

In the remainder of this section, the notation **digits** refers to the current setting of `numeric digits`. This setting may also be referred to in expressions in programs by using the **digits** special word (see page 121).

43.2.3 Arithmetic operators

NetREXX arithmetic is effected by the operators "+", "-", "*", "/", "%", "//", and "**" (add, subtract, multiply, divide, integer divide, remainder, and power) which all act upon two terms, together with the prefix operators "+" and "-" (plus and minus) which both act on a single term. The result of all these operations is a NetREXX string, of type **REXX**. This section describes the way in which

these operations are carried out. Before every arithmetic operation, the term or terms being operated upon have any extra digits converted to the corresponding Arabic numeral (the digits 0-9). They then have leading zeros removed (noting the position of any decimal point, and leaving just one zero if all the digits in the number are zeros) and are then truncated to **digits+1** significant digits ⁷¹ (if necessary) before being used in the computation. The operation is then carried out under up to double that precision, as described under the individual operations below. When the operation is completed, the result is rounded if necessary to the precision specified by the numeric `digits` instruction. Rounding is done in the "traditional" manner, in that the extra (guard) digit is inspected and values of 5 through 9 are rounded up, and values of 0 through 4 are rounded down. ⁷² A conventional zero is supplied preceding a decimal point if otherwise there would be no digit before it. Trailing zeros are retained for addition, subtraction, and multiplication, according to the rules given below, except that a result of zero is always expressed as the single character '0'. For division, insignificant trailing zeros are removed after rounding.

The **format** method (see page 166) is defined to allow a number to be represented in a particular format if the standard result provided by NetREXX does not meet requirements.

43.2.4 Arithmetic operation rules - basic operators

The basic operators (addition, subtraction, multiplication, and division) operate on numbers as follows:

Addition and subtraction If either number is zero then the other number, rounded to **digits** digits if necessary, is used as the result (with sign adjustment as appropriate). Otherwise, the two numbers are extended on the right and left as necessary up to a total maximum of **digits+1** digits. The number with smaller absolute value may therefore lose some or all of its digits on the right. ⁷³ The numbers are then added or subtracted as appropriate. For example:

```
xxxx.xxx + yy.yyyyy
```

becomes:

```
  xxxx.xxx00
+ 00yy.yyyyy
-----
  zzzz.zzzzz
```

`.sumadd` The result is then rounded to **digits** digits if necessary, taking into account any extra (carry) digit on the left after an addition, but otherwise counting from the position corresponding to the most significant digit of the terms being added or subtracted. Finally, any insignificant leading ze-

⁷¹ That is, to the precision set by numeric `digits`, plus one extra "guard" digit.

⁷² Even/odd rounding would require the ability to calculate to arbitrary precision (that is, to a precision not governed by the setting of numeric `digits`) at any time and is therefore not the mechanism defined for NetREXX.

⁷³ In the example, the number `yy.yyyyy` would have three digits truncated if **digits** were 5.

ros are removed. The *prefix operators* are evaluated using the same rules; the operations "+number" and "-number" are calculated as "0+number" and "0-number", respectively.

Multiplication The numbers are multiplied together ("long multiplication") resulting in a number which may be as long as the sum of the lengths of the two operands. For example:

xxx.xxx * yy.yyyyy

becomes:

zzzzz.zzzzzzzz

and the result is then rounded to **digits** digits if necessary, counting from the first significant digit of the result.

Division For the division:

yyy / xxxxx

the following steps are taken: first, the number "yyy" is extended with zeros on the right until it is larger than the number "xxxxx" (with note being taken of the change in the power of ten that this implies). Thus in this example, "yyy" might become "yyy00". Traditional long division then takes place, which can be written:

```

      zzzz
      .-----
xxxxx | yyy00

```

The length of the result ("zzzz") is such that the rightmost "z" will be at least as far right as the rightmost digit of the (extended) "y" number in the example. During the division, the "y" number will be extended further as necessary, and the "z" number (which will not include any leading zeros) may increase up to **digits+1** digits, at which point the division stops and the result is rounded. Following completion of the division (and rounding if necessary), insignificant trailing zeros are removed.

Examples:

```

/* With 'numeric digits 5' */
12+7.00    == 19.00
1.3-1.07   == 0.23
1.3-2.07   == -0.77
1.20*3     == 3.60
7*3        == 21
0.9*0.8    == 0.72
1/3        == 0.33333
2/3        == 0.66667
5/2        == 2.5
1/10       == 0.1
12/12      == 1
8.0/2      == 4

```

Note: With all the basic operators, the position of the decimal point in the terms

being operated upon is arbitrary. The operations may be carried out as integer operations with the exponent being calculated and applied afterwards. Therefore the significant digits of a result are not in any way dependent on the position of the decimal point in either of the terms involved in the operation.

43.2.5 Arithmetic operation rules - additional operators

The operation rules for the power ("**"), integer division ("%"), and remainder ("//") operators are as follows:

Power The "**" (power) operator raises a number (on the left of the operator) to a power (on the right of the operator). The term on the right is rounded to **digits** digits (if necessary), and must, after any rounding, be a whole number, which may be positive, negative, or zero. If negative, the absolute value of the power is used, and then the result is inverted (divided into 1). For calculating the power, the number is effectively multiplied by itself for the number of times expressed by the power, and finally trailing zeros are removed (as though the result were divided by one). In practice (see note below for the reasons), the power is calculated by the process of left-to-right binary reduction. For "x**n": "n" is converted to binary, and a temporary accumulator is set to 1. If "n" has the value 0 then the initial calculation is complete. Otherwise each bit (starting at the first non-zero bit) is inspected from left to right. If the current bit is 1 then the accumulator is multiplied by "x". If all bits have now been inspected then the initial calculation is complete, otherwise the accumulator is squared by multiplication and the next bit is inspected. When the initial calculation is complete, the temporary result is divided into 1 if the power was negative.

The multiplications and division are done under the normal arithmetic operation rules, detailed earlier in this section, using a precision of **digits+length+1** digits. Here, **length** is the length in decimal digits of the integer part of the whole number "n" (*i.e.*, excluding any sign, decimal part, decimal point, or insignificant leading zeros, as though the operation `n%1` had been carried out and any sign removed). Finally, the result is rounded to **digits** digits, if necessary, and insignificant trailing zeros are removed.

Integer division The "%" (integer divide) operator divides two numbers and returns the integer part of the result. The result returned is defined to be that which would result from repeatedly subtracting the divisor from the dividend while the dividend is larger than the divisor. During this subtraction, the absolute values of both the dividend and the divisor are used: the sign of the final result is the same as that which would result if normal division were used. The result returned will have no fractional part (that is, no decimal point or zeros following it). If the result cannot be expressed exactly within **digits** digits, the operation is in error and will fail - that is, the result cannot have more digits than the current setting of numeric digits. For example, `1000000000%3` requires ten digits to express the result exactly (`3333333333`) and would therefore fail if **digits** were 9 or smaller.

Remainder The `"/"` (remainder) operator will return the remainder from integer division, and is defined as being the residue of the dividend after the operation of calculating integer division as just described. The sign of the remainder, if non-zero, is the same as that of the original dividend. This operation will fail under the same conditions as integer division (that is, if integer division on the same two terms would fail, the remainder cannot be calculated).

Examples:

```
/* Again with 'numeric digits 5' */
2**3      == 8
2**-3     == 0.125
1.7**8    == 69.758
2\%3      == 0
2.1//3    == 2.1
10\%3     == 3
10//3     == 1
-10//3    == -1
10.2//1   == 0.2
10//0.3   == 0.1
3.6//1.3  == 1.0
```

Notes:

1. A particular algorithm for calculating powers is described, since it is efficient (though not optimal) and considerably reduces the number of actual multiplications performed. It therefore gives better performance than the simpler definition of repeated multiplication. Since results could possibly differ from those of repeated multiplication, the algorithm must be defined here so that different implementations will give identical results for the same operation on the same values. Other algorithms for this (and other) operations may always be used, so long as they give identical results to those described here.
2. The integer divide and remainder operators are defined so that they may be calculated as a by-product of the standard division operation (described above). The division process is ended as soon as the integer result is available; the residue of the dividend is the remainder.

43.2.6 Numeric comparisons

Any of the comparative operators (see page 43) may be used for comparing numeric strings. However, the strict comparisons (for example, `"=="` and `">>"`) are not numeric comparative operators and should not normally be used for comparing numbers, since they compare from left to right and leading and trailing blanks (and leading zeros) are significant for these operators. Numeric comparison, using the normal comparative operators, is effected by subtracting the two numbers (calculating the difference) and then comparing the result with `'0'` - that is, the operation:

A ? B

where "?" is any normal comparative operator, is identical to:

```
(A - B) ? '0'
```

It is therefore the *difference* between two numbers, when subtracted under NetREXX subtraction rules, that determines their equality.

43.2.7 Exponential notation

The definition of numbers above (see page 137) describes "pure" numbers, in the sense that the character strings that describe numbers can be very long. **Examples:**

```
say 10000000000 * 10000000000
/* would display: 10000000000000000000 */
```

```
say 0.000000000001 * 0.000000000001
/* would display: 0.000000000000000000000001 */
```

For both large and small numbers some form of exponential notation is useful, both to make such long numbers more readable and to make evaluation possible in extreme cases. In addition, exponential notation is used whenever the "pure" form would give misleading information. For example:

```
numeric digits 5
say 54321*54321
```

would display "2950800000" if long form were to be used. This is misleading, as it appears that the result is an exact multiple of 100000, and so NetREXX would express the result in exponential notation, in this case "2.9508E+9". The definition of *number* (see above) is therefore extended by replacing the description of **numeric** by the following:

```
mantissa ::= digits . [digits]
           | [.] digits
numeric  ::= mantissa [E sign digits]
```

In other words, the numeric part of a number may be followed by an "E" (indicating an exponential part), a sign, and an integer following the sign that represents a power of ten that is to be applied. The "E" may be in uppercase or lowercase. Note that no blanks are permitted within this part of a number, but the integer may have leading zeros. **Examples:**

```
12E+11 = 1200000000000
12E-5  = 0.00012
12e+4  = 120000
```

All valid numbers may be used as data for arithmetic. The results of calculations will be returned in exponential form depending on the setting of numeric digits. If the number of places needed before the decimal point exceeds **digits**, or if the absolute value of the result is less than **0.000001**, then exponential form will be used. The exponential form generated by NetREXX always has a sign following the "E". If the exponent is 0 then the exponential part is omitted - that is,

an exponential part of "E+0" will never be generated. If the default format for a number is not satisfactory for a particular application, then the **format** method may be used to control its format. Using this, numbers may be explicitly converted to exponential form or even forced to be returned in "pure" form. Different exponential notations may be selected with the `numeric form` instruction (see page 89). This instruction allows the selection of either scientific or engineering notation. *Scientific notation* adjusts the power of ten so there is a single non-zero digit to the left of the decimal point. *Engineering notation* causes powers of ten to be expressed as a multiple of three - the integer part may therefore range from 1 through 999. **Examples:**

```
numeric form scientific
say 123.45 * 1e11
/* would display: 1.2345E+13 */

numeric form engineering
say 123.45 * 1e11
/* would display: 12.345E+12 */
```

The default exponential notation is scientific.

43.2.8 Whole numbers

Within the set of numbers understood by NetREXX it is useful to distinguish the subset defined as *whole numbers*.

A *whole number* in NetREXX is a number that has a decimal part which is all zeros (or that has no decimal part).

43.2.9 Numbers used directly by NetREXX

As discussed above, the result of any arithmetic operation is rounded (if necessary) according to the setting of `numeric digits`. Similarly, when a number (which has not necessarily been involved in an arithmetic operation) is used directly by NetREXX then the same rounding is also applied, just as though the operation of adding the number to 0 had been carried out. After this operation, the integer part of the number must have no more digits than the current setting of `numeric digits`.

In the following cases, the number used must be a whole number and an implementation restriction on the largest number that can be used may apply:

- positional patterns, including variable positional patterns, in parsing templates (see page 129)
- the power value (right hand operand) of the power operator (see page 141).
- the values of *expr* and *exprf* (following the `for` keyword) in the loop instruction (see page 75)
- the value of *exprd* (following the `digits` keyword) in the `numeric` instruction (see page 89).

Implementation minimum: A minimum length of 9 digits must be supported for these uses of whole numbers by a NetREXX language processor.

43.2.10 Implementation independence

The NetREXX arithmetic rules are defined in detail, so that when a given program is run the results of all computations are sufficiently defined that the same answer will result for all correct implementations. Differences due to the underlying machine architecture will not affect computations. This contrasts with most other programming languages, and with binary arithmetic (see page 146) in NetREXX, where the result obtained may depend on the implementation because the precision and algorithms used by the language processor are defined by the implementation rather than by the language.

43.2.11 Exceptions and errors

The following exceptions and errors may be signalled during arithmetic:

- **Divide exception** This exception will be signalled if division by zero was attempted, or if the integer result of an integer divide or remainder operation had too many digits.
- **Overflow/Underflow exception** This exception will be signalled if the exponential part of a result (from an operation that is not an attempt to divide by zero) would exceed the range that can be handled by the language processor, when the result is formatted according to the current settings of numeric digits and numeric form. The language defines a minimum capability for the exponential part, namely exponents whose absolute value is at least as large as the largest number that can be expressed as an exact integer in default precision. Thus, since the default precision is nine, implementations must support exponents in the range **-999999999** through **999999999**.
- **Insufficient storage** Storage is needed for calculations and intermediate results, and on occasion an arithmetic operation may fail due to lack of storage. This is considered an operating environment error as usual, rather than an arithmetical exception.

*In the reference implementation, the exceptions and error types used for these three cases are **DivideException**, **ExponentOverflowException**, and **OutOfMemoryError**, respectively.*

Binary values and operations

By default, arithmetic and string operations in NetREXX are carried out using the NetREXX string class, **REXX**, which offers the robust set of operators described in *Expressions and operators* (see page 41).

NetREXX implementations, however, may also provide *primitive* datatypes, as described in *Types and Classes* (see page 23). These primitive types are used for compact storage of numbers and for fast binary arithmetic, features which are built-in to the hardware of most computers.

To make use of binary arithmetic, a class is declared to be a *binary class* (see page 64) by using the `binary` keyword on the `class` instruction. In such a class, literal strings and numeric symbols are assigned native string or primitive types, rather than NetREXX types, where appropriate, and native binary operations are used to implement operators where possible, as detailed below. Implementations may also provide a keyword on the `options` (see page 91) instruction that indicates that all classes in a program are binary classes.⁷⁴

Alternatively, individual methods within a class may be declared to be a *binary method* (see page 85) by using the `binary` keyword on the `method` instruction.

3.04 Alternatively, individual `do` groups within a method may be declared to be a *do binary group* (see page 66) by using the `binary` keyword on the `do` instruction.

Binary classes and methods should be used with care. Although binary arithmetic can have a considerable performance advantage over arithmetic that is not implemented in hardware, it can give incorrect or unexpected results. In particular, whole numbers (integers) are often held in fixed-sized data areas (of 8, 16, 32, or 64 bits), and overflowing the data area during a calculation can result in a positive number becoming negative and vice versa. Similarly, binary numbers that are not whole numbers (floating-point numbers) cannot exactly represent common numbers in the decimal system (**0.1**, **0.2**, *etc.*), and hence can give unexpected results.

44.1 Operations in binary classes and methods

In a binary class or method, the following (and only the following) rules differ from the usual rules:

Dyadic operations in expressions If the operands of a dyadic operator both

⁷⁴In the reference implementation, `options binary` is used.

have primitive numeric types⁷⁵ then binary operations are carried out. The type of the result is implementation defined, and is typically the type of the more precise of the two operands, or of some minimum precision.⁷⁶ Arithmetic operations follow the usual rules of binary arithmetic, as defined for the underlying environment of the implementation.

Note that NetREXX provides both divide and integer divide operators; in a binary class or method, the divide operator ("`/`") converts its operands to floating-point types and returns a floating-point result, whereas the integer divide operator ("`%`") converts its operands to integer types and returns an integer result. The remainder operator must accept both integer and floating-point types.

Logical operations (*and*, *or*, and *exclusive or*) apply to all the bits of the operands, and are not permitted on floating-point types.

Prefix operations in expressions If the operand of a prefix operator has a primitive numeric type, then the type of the result is the type of the operand, subject to the same minimum as dyadic operations. Prefix plus and minus follow the rules of dyadic operators (because they are defined as being zero plus or minus the operand) with the additional rule that if acting on a literal number (a constant in the program) then the result is also considered to be a literal constant. Logical not (prefix "`\`") does not apply to all the bits of its operand; instead, it changes a **0** to **1** and vice versa.

Assignments In assignments where the value being assigned is the result of an expression which comprises a string or number literal constant, the type of the result is defined as follows:

1. Strings are given the native string type, even for a single-character literal.⁷⁷
2. Numbers are given the smallest possible primitive numeric type that will contain the literal without loss of information (or minimal loss of information for numbers with decimal or exponential parts). If this is smaller than the implementation-defined minimum precision used for the result of adding the literal to **0**, then the type of that minimum precision is used.

If the constant is an integer, and no primitive integer binary type has sufficient precision to hold the number without loss of information, then the number is treated as a literal string (that is, as though it were enclosed in quotes). NetREXX arithmetic would then be used if it were involved in an arithmetic operation.

These rules can apply in assignment instructions, the initial assignment to the control variable in the loop instruction, or the assignment of a default value to the argument of a method; the result type may define the type of the variable (if new, or a method argument).

Control variables in loops In the loop instruction, if the control variable has a primitive integer type, and the increment (by value) has a primitive integer

⁷⁵In the reference implementation, *boolean* is considered to be a numeric type (having the values **0** or **1**) but *char* is not. Characters, and strings or arrays of characters, always use the rules defined for NetREXX strings.

⁷⁶In the reference implementation, the minimum precision is 32 bits, so an *int* is returned for results that would otherwise be *byte* or *short*. If both operands are *boolean*, however, and the operation is a logical operation, then the type of the result is *boolean*.

⁷⁷In the reference implementation, this type is *java.lang.String*.

type, then binary arithmetic will be used for stepping the control variable, following the rules for binary arithmetic in expressions described above. Similarly, if the control variable has a primitive integer type, and the end (to) value has a primitive integer type, then binary arithmetic will be used for the comparison that tests for loop termination.

Numeric instruction The `numeric` instruction does not affect binary operations. It has the usual effects on operations carried out using NetREXX arithmetic.

Note: At all times (whether in binary classes, binary methods, or anywhere else) implementations may use primitive types and operations, and techniques such as late binding of types, as an optimization providing that the results obtained are identical to those defined in this language definition.

44.2 Binary constructors

NetREXX provides special constructors for implementation-defined primitive types that allow bit-wise construction of primitives. These *binary constructors* are especially useful for manipulating the binary encodings of individual characters.

The binary constructors follow the same syntax as other constructors, with the name being that of a primitive type. All binary constructors take one argument, which must have a primitive type.

The bits of the value of the argument are extended or truncated on the left to the same length as the bits required for the type of the constructor (following the usual binary rules of sign extension if the argument type is a signed numeric type), and a value with the type of the constructor is then constructed directly from those bits and returned.

Example: This example illustrates types from the reference implementation, with 32-bit signed integers of type `int` and 16-bit Unicode characters of type `char`.

```
i=int 77    -- i is now the integer 77
c=char(i)  -- c is now the character 'M'
j=int(c)    -- j is now the integer 77
```

Note that the conversion

```
j=int c
```

would have failed, as "M" is not a number.

Exceptions

Exceptional conditions, including errors, in NetRexx are handled by a mechanism called *Exceptions*. When an exceptional condition occurs, a *signal* takes place which may optionally be *caught* by an enclosing control construct, as detailed below.

An exception can be signalled by:

1. the program's environment, when some processing error occurs (such as running out of memory, or a problem discovered when reading or writing a file)
2. a method called by a NetRexx program (if, for example, it is passed incorrect arguments)
3. the signal instruction (see page 108) .

In all cases, the signal is handled in exactly the same way. First, execution of the current clause ceases; no further operations within the clause will be carried out.⁷⁸ Next, an object that represents the exception is constructed. The type of the exception object is implementation-dependent, as described for the signal instruction (see page 108) , and defines the type of the exception. The object constructed usually contains information about the Exception (such as a descriptive string).

Once the object has been constructed, all active `do` groups, `loop` loops, `if` constructs, and `select` constructs in the active method are "unwound", starting with the innermost, until the exception is caught by a control construct that specifies a suitable catch clause (see below) for handling the exception.

This unwinding takes place as follows:

1. No further clauses within the body of the construct will be executed (in this respect, the signal acts like a `leave` for the construct).
2. If a catch clause specifies a type to which the exception object can be assigned (that is, it matches or is a superclass of the type of exception object), then the *instructionlist* following that clause is executed, and the exception is considered to be handled (no further control constructs will be unwound). If more than one catch clause specifies a suitable type, the first is used.
3. The *instructionlist* following the `finally` clause for the construct, if any, is executed.

⁷⁸This is the only case in which an expression will not be wholly evaluated, for example.

4. The end clause is executed, hence completing execution of the construct. (The only effect of this is that it is seen when tracing.)
5. If the exception was handled, then execution resumes as though the construct completed normally. If it was not handled, then the process is repeated for any enclosing constructs.

If the exception is not caught by any of the control constructs enclosing the original point of the exception signal, then the current active method is terminated, without returning any data, and the exception is then signalled at the point where the method was invoked (that is, in the caller).

The process of unwinding control constructs and terminating the method is then repeated in each calling method until the exception is caught or the initial program invocation method (the main method) is terminated, in which case the program ends and the environment receives the signal (it would usually then display diagnostic information).

45.1 Syntax and example

The constructs that may be used to handle (catch) an exception are `do` groups, `loop` loops, and `select` constructs. Specifically, as shown in the syntax diagrams (*q.v.*), where the end clause can appear in these constructs, zero or more catch clauses can be used to define exception handlers, followed by zero or one finally clauses that describe "clean-up" code for the construct. The whole construct continues to be ended by an end clause.

The syntax of a catch clause is shown in the syntax diagrams. It always specifies an *exception* type, which may be qualified. It may optionally specify a symbol, *vare*, which is followed by an equals sign. This indicates that when the exception is caught then the object representing the exception will be assigned to the variable *vare*. If `new`, the type of the variable will be *exception*. Here is an example of a program that handles some of the exceptions signalled by methods in the **REXX** class; the `trace results` instruction is included to show the flow of execution:

```
trace results
do                                -- could be LOOP i=1 to 10, etc.
  say 1/arg
catch DivideException
  say 'Divide exception'
catch ex=NumberFormatException
  /* 'ex' is assigned the exception object */
  say 'Bad number for division:' ex.getMessage
finally
  say 'Done!'
end
```

In this example, if the argument passed to the program (and hence placed in **arg**) is a valid number, then its inverse is displayed. If the argument is 0, then "Divide exception" would be displayed. If the argument were an invalid number, the message describing the bad number would be displayed. For any other

exception (such as an **ExponentOverflowException**), the program would end and the environment would normally report the exception.

In **all** cases, the message "**Done!**" would be displayed; this would be true even if the body of the `do` construct executed a `return`, `leave`, or `iterate` instruction. Only an `exit` instruction (see page 68) would cause immediate termination of the construct (and the program).

Note: The `finally` keyword, like otherwise in the `select` construct, implies a semicolon after it, so the last `say` instruction in the example could have appeared on the same line as the `finally` without an intervening semicolon.

45.2 Exceptions after catch and finally clauses

If an exception is signalled in the *instructionlist* following a `catch` or `finally` clause, then the current exception is considered handled, the *instructionlist* is terminated, and the new exception is signalled. It will not be caught by `catch` clauses in the current construct. If it occurs in the *instructionlist* following a `catch` clause, then any `finally` instructions will be executed, as usual.

Similarly, executing a `return` or `exit` instruction within either of the *instructionlists* completes the handling of any pending signal.

45.3 Checked exceptions

NetREXX implementations may define certain exceptions as *checked exceptions*. These are exceptions that the implementation considers it useful to check; the checked exceptions that each method may signal are recorded. Within `do` groups, `loop` loops, and `select` constructs, for example, it is then possible to report if a `catch` clause tries to catch a checked exception that is not signalled within the body of the construct.

Checked exceptions that are signalled within a method (by a `signal` instruction or a method invocation) but not caught by a `catch` clause in the method are automatically added to the `signals` list for a method. Implementations that support checked exceptions are encouraged to provide options that list the uncaught checked exceptions for methods or enforce the explicit inclusion of some or all checked exceptions in the `signals` list on the method instruction.

In the reference implementation, all exceptions are checked except those that are subclasses of `java.lang.RuntimeException` or `java.lang.Error`. These latter are considered so ubiquitous that almost all methods would signal them.

Expressions assigned as the initial values of properties must not invoke methods that may signal checked exceptions.

The `strictsignal` option on the `options` instruction may be used to enforce the inclusion of all uncaught checked exceptions in methods' `signals` lists; this may be used to assure that any uncaught checked exceptions are intentional.

Thread Pool Support

Support for thread pooling is built into the NetRexx language.

rtp=RexxThreadPool(size,number) *size* is the number of parallel threads desired - default is the current number of available processors *number* is the number of the threadpool if using multiple pools - default is pool number 0

rtp.start(runtask) *runtask* needs to be a class that implements the *Runnable* interface

rtp.start(maintask,mainparm) *maintask* is a NetRexx Java class with a standard "main" method *mainparm* is a string parm to pass to the main method at startup

rtp.start("taskname",mainparm) *taskname* is a string identifying a Java class with a standard main method *mainparm* is a string parm to pass to the main method at startup⁷⁹

rtp.waituntildone Blocks until all threads in the pool are finished

rtp.waitforallpools Blocks until all threads in all task pools are complete

46.1 Examples

```
RexxThreadPool(3,1).start(Test(7)).start(Test(8)).start("TestMain", "9")
    .start("enviroscan")
RexxThreadPool(9).start(Test(1)).start(Test(2)).start(TestMain("3"), "3")
    .start(enviroscan.class)
RexxThreadPool().start(Test(4)).start(Test(5)).start(TestMain("6"), "6")
    .waituntildone
RexxThreadPool().start(Test(4)).start(Test(5)).start(TestMain("6"), "6")
    .waitforallpools
```

⁷⁹the start method returns the RexxThreadPool instance it is called on so that multiple calls can be stacked. Due to reflection use when starting "main" methods that call format cannot be interpreted - runnables interpret ok

Structured Lists Interface

A Structured List class is used to create objects which can contain structured or string encoded lists with associated methods for easy access. A Structured List class is a subclass of the native NetRexx "Rexx" class data type which implements the StructuredList interface as described below and can be used in many cases as a normal NetRexx data item. Assuming a class named "AStructuredList" is available to support structured lists:

- A StructuredList class should have a constructor that sets the encoded list value

This creates an unexamined (unstructured) list object. The object is an object of type Rexx with additional methods for processing lists. You must pass a Ruleset to the list object (see the "buildlist" method below) in order to obtain a structured version of the list object which can be accessed with list methods. Note that due to the immutability rule for base Rexx object values, all methods which alter the content of a structured list return a new structured list object. This rule also means that if a list is altered by a Java List view, the new structured list object must be obtained from the "currentlist" property of the Java List view.

Structured List formats known to NetRexx include

- WORDLIST
- DSV
- CSV
- TSV
- XML
- JSON
- PYTHON

47.1 Essential List Processing Methods

This section lists the methods provided by the *StructuredList* class.

buildlist(ruleset) Returns a structured list form of the encoded list contained in the object's string value after examination with the provided ruleset.

join(anotherStructuredList) Returns a structured list containing the elements of the structured list with the elements from another structured list appended to it.

islist Returns 1 if the item is a structured list, otherwise returns 0. If a structured list is empty, index "elements" will have a 0 value.

elementcount Returns the count of elements in a structured list.

getelement(index) Returns the element at the specified location in a structured list.

getJavaList Returns a Java List interface view of the structured list.

index(element,start) Returns the numeric index of the first occurrence of the given element in the list with index equal to or higher than start or 0 if not found.

insertelement(index, element) Adds an element to a list at the specified location. Returns modified list.

replaceanelement(index, element) Replaces an element of a list at the specified location. Returns modified list.

deleteanelement(index) Removes an element from a list at the specified location. Returns modified list.

sublist(fromIndex,toIndex) Creates a list which is a subset of the list containing the elements starting at the from index and ending at the to index.

47.2 Convenience Methods

append(element) Adds an element to the end of a list. Returns modified list.

head Returns the first element in a list.

tail Returns a list containing all elements except the first in a list.

count(value) Returns a count of how many elements have the provided value.

minval Returns the lowest value in a list. This is a runtime error if not all list elements are numeric.

maxval Returns the highest value in a list. This is a runtime error if not all list elements are numeric.

sum Returns the sum of all list elements. This is a runtime error if not all list elements are numeric.

reverselist Returns a structured list with the order of the elements reversed.

flatlist Returns a structured list with only the list elements (all metadata is removed). Nested sublists are also flattened.

47.3 NetRexx Structured List Format

A structured list is an ordered sequence of items stored in a NetRexx native data object (class `Rexx`) along with "meta data" describing the list. Each element of the list can be accessed with a whole number ranging from 1 to n where n is the number of elements in the list. The string encoded (aka "serialized") form of the string is the base value of the `Rexx` object. If an element of a list is itself a list (ie a nested list), then elements of the sublist can also be accessed by whole number indexes so that for example the 3rd element of the nested list which is the second element of the first list could be found at `myList[2,3]`

where `myList` is the object containing the structured list. Note that although a structured list is a Rexx object, changing the list directly rather than through the `StructuredList` interface methods will cause loss of metadata and unpredictable results for further list access.

47.3.1 Accessing structured lists with NetRexx facilities

Assuming `myList` is a Rexx object containing a structured list, then:

`myList["elements"]`

Provides a count of the number of elements in a structured list. This is meaningless if the item is not actually in the structured list format (ie if `islist` returns 0).

`myList[n,m]`

A list element can be accessed by numeric index using NetRexx indexed variable syntax. If the element is also a list, sub-elements can be accessed using the multiple index syntax.

Additional indexed values

`"rules"` The rules used to structure this list (the rules are also a structured list).

The following indexes must be prefixed with `"#"` to use if they are whole numbers according to the NetRexx datatype BIF - this avoids conflict with list index numbers.

`elementname` If an element has a name, using the name as an index will return the associated element value.

`element` The element string will return the index of the first location of that element in the list. **`elementvalue`** If an element has a name the named value string will return the index of the first location of that value in the list.

47.4 Structured List Ruleset Description

A list ruleset specifies a set of delimiters and options and can be provided in one of four ways:

1. A null ruleset (or the string `"default"`) signifies a default ruleset. Default rules are start/end delimiters `"(" and ")"`, a separator `" "` (a blank), an escape character `"` (double quote), and the option `"escape is quoted string mode"`.
2. A string such as `"CSV"`, which is a well known list format name, selects a built-in ruleset. For example, a list in CSV format could be decoded like this: `inputlist=inputstring.getlist("CSV")`. Formats known to a Structured List include `"WORDLIST"`, `"DSV"`, `"CSV"`, `"TSV"`, `"XML"`, `"JSON"`, `"PYTHON"`.
3. A string which provides a human readable custom set of list rules that is itself a decomposable list according to the default ruleset. Rulesets contain

two sections: delimiters and options which are specified as separate sublists as in the following example:

```
'delimiters(start("<") end(">") separator(",") meta("/") escape("\\")
nameseparator("=") )
options(separators-must-follow-sublists "adjacent separators reduce to one")'
```

4. A ruleset string that is itself an encoded list according to a known ruleset can simply be prepared before use like in this example:

```
inputlist=inputstring.buildlist(rulesetstring.buildlist("default"))
```

47.4.1 Delimiters

Any type of delimiter can be specified but the following (along with the ruleset options) define the structure of a list. Other delimiters can be provided and will be recognized and recorded as list elements when they occur which means they are "defacto" separators for the list elements.

Start A delimiter which signals the start of a sublist. Example: start("")

End A delimiter which signals the end of a sublist. Example: End("")

Escape An escape character used to include delimiters in the list data elements.
Example: escape("|")

Separator A delimiter used to separate list elements. Example: Separator(",")

NameSeparator A delimiter used to associate element names with element values. Example: Namesep("=")

47.4.2 Options

The following options are recognized (A "0" in front of the option indicates a default value and a "1" indicates an option which overrides the default value. The additional descriptions in parenthesis are not part of the option. Options can be specified as quoted text or with dashes substituted for spaces. Options can be abbreviated as long as they are unique. Options are not case-sensitive.

option 1 0 = separators follow sublists
1 = sublists are separators

option 2 0 = adjacent separators reduce to one
1 = produce empty elements for adjacent separators

option 3 0 = translate escape sequences
1 = do not translate escape sequences

option 4 0 = whitespace is translated to blank (TAB,FF,LF,CR,VT)
1 = treat whitespace as text

option 5 0 = escape is quoted string mode (ie "text, delimiters or double escape like t
1 = single character escape (ie \x)

- option 6** 0 = attribute names are implicit (ie fun(x,y))
1 = explicit attribute names (ie with delimiter as in fun=(x,y) or fun:(x,y))
- option 7** 0 = delimiters are implicit (do not record structural delimiters)
1 = tokens are delimiters (save delimiters as separate elements)
- option 8** 0 = keep leading and trailing whitespace
1 = strip leading and trailing whitespace

Methods for NetREXX strings

This section describes the set of methods defined for the NetREXX string class, **REXX**. These are called *built-in methods*, and include character manipulation, word manipulation, conversion, and arithmetic methods.

Implementations will also provide other methods for the **REXX** class (for example, to implement the NetREXX operators or to provide constructors with primitive arguments), but these are not part of the NetREXX language.⁸⁰

48.1 General notes on the built-in methods:

1. All methods work on a NetREXX string of type **REXX**; this is referred to by the name *string* in the descriptions of the methods. For example, if the **word** method were invoked using the term:

```
"Three word phrase".word(2)
```

then in the description of **word** the name *string* refers to the string **"Three word phrase"**, and the name *n* refers to the string **"2"**.

2. All method arguments are of type **REXX** and all methods return a string of type **REXX**; if a number is returned, it will be formatted as though 0 had been added with no rounding.
3. The first parenthesis in a method call must immediately follow the name of the method, with no space in between.
4. The parentheses in a method call can be omitted if no arguments are required and the method call is part of a *compound term* (see page 27).⁸¹
5. A position in a string is the number of a character in the string, where the first character is at position 1, *etc.*
6. Where arguments are optional, commas may only be included between arguments that are present (that is, trailing commas in argument lists are not permitted).
7. A *pad* argument, if specified, must be exactly one character long.
8. If a method has a sub-option selected by the first character of a string, that character may be in upper or lowercase.
9. Conversion between character encodings and decimal or hexadecimal is dependent on the machine representation (encoding) of characters and hence will return appropriately different results for Unicode, ASCII, EBCDIC, and other implementations.

⁸⁰Details of the methods provided in the reference implementation are included in Appendix C (see page 180).

⁸¹Unless an implementation-provided option to disallow parenthesis omission is in force.

48.2 The built-in methods

abbrev(*info* [*length*]) returns 1 if *info* is equal to the leading characters of *string* and *info* is not less than the minimum length, *length*; 0 is returned if either of these conditions is not met. *length* must be a non-negative whole number; the default is the length of *info*. **Examples:**

```
'Print'.abbrev('Pri')    == 1
'PRINT'.abbrev('Pri')    == 0
'PRINT'.abbrev('PRI',4)  == 0
'PRINT'.abbrev('PRY')    == 0
'PRINT'.abbrev('')       == 1
'PRINT'.abbrev('',1)     == 0
```

Note: A null string will always match if a length of 0 (or the default) is used. This allows a default keyword to be selected automatically if desired.

Example:

```
say 'Enter option: '; option=ask
select /* keyword1 is to be the default */
  when 'keyword1'.abbrev(option) then ...
  when 'keyword2'.abbrev(option) then ...
  ...
otherwise ...
end
```

abs() returns the absolute value of *string*, which must be a number. Any sign is removed from the number, and it is then formatted by adding zero with a digits setting that is either nine or, if greater, the number of digits in the mantissa of the number (excluding leading insignificant zeros). Scientific notation is used, if necessary.

Examples:

```
'12.3'.abs                == 12.3
'-0.307'.abs              == 0.307
'123.45E+16'.abs          == 1.2345E+18
'- 1234567.7654321'.abs  == 1234567.7654321
```

- 3.02 **b2d**([*n*]) Binary to decimal. Converts *string*, a string of at least one binary (0 and/or 1) digits, to an equivalent string of decimal characters (a number), without rounding. The returned string will use digits, and will not include any blanks. If the number of binary digits in the string is not a multiple of four, then up to three '0' digits will be added on the left before conversion to make a total that is a multiple of four. If *string* is the null string, 0 is returned. If *n* is not specified, *string* is taken to be an unsigned number.

Examples:

```
'01110'.b2d == 14
'10000001'.b2d == 129
'111110000001'.b2d == 3969
'111111110000001'.b2d == 65409
'1100011011110000'.b2d == 50928
```

If *n* is specified, *string* is taken as a signed number expressed in *n* binary characters. If the most significant (left-most) bit is zero then the number

is positive; otherwise it is a negative number in twos-complement form. In both cases it is converted to a NetRexx number which may, therefore, be negative. If *n* is 0, 0 is always returned.

If necessary, *string* is padded on the left with '0' characters (note, not "signextended"), or truncated on the left, to length *n* characters; (that is, as though `string.right(n, '0')` had been executed.)

Examples:

```
'10000001'.b2d(8) == -127
'10000001'.b2d(16) == 129
'1111000010000001'.b2d(16) == -3967
'1111000010000001'.b2d(12) == 129
'1111000010000001'.b2d(8) == -127
'1111000010000001'.b2d(4) == 1
'0000000000110001'.b2d(0) == 0
```

b2x() Binary to hexadecimal. Converts *string*, a string of at least one binary (0 and/or 1) digits, to an equivalent string of hexadecimal characters. The returned string will use uppercase Roman letters for the values A-F, and will not include any blanks. If the number of binary digits in the string is not a multiple of four, then up to three '0' digits will be added on the left before conversion to make a total that is a multiple of four.

Examples:

```
'11000011'.b2x == 'C3'
'10111'.b2x == '17'
'0101'.b2x == '5'
'101'.b2x == '5'
'111110000'.b2x == '1F0'
```

center(length [,pad]) or

centre(length [,pad]) returns a string of length *length* with *string* centered in it, with *pad* characters added as necessary to make up the required length. *length* must be a non-negative whole number. The default *pad* character is blank. If the string is longer than *length*, it will be truncated at both ends to fit. If an odd number of characters are truncated or added, the right hand end loses or gains one more character than the left hand end.

Examples:

```
'ABC'.centre(7) == ' ABC '
'ABC'.center(8, '-') == '--ABC--'
'The blue sky'.centre(8) == 'e blue s'
'The blue sky'.center(7) == 'e blue '
```

Note: This method may be called either **centre** or **center**, which avoids difficulties due to the difference between the British and American spellings.

changestr(needle, new) returns a copy of *string* in which each occurrence of the *needle* string is replaced by the *new* string. Each unique (non-overlapping) occurrence of the *needle* string is changed, searching from left to right and starting from the first (leftmost) position in *string*. Only the original *string* is searched for the *needle*, and each character in *string* can only be included in one match of the *needle*.

If the *needle* is the null string, the result is a copy of *string*, unchanged.

Examples:

```
'elephant'.changestr('e','X')    == 'XlXphant'
'elephant'.changestr('ph','X')   == 'eleXant'
'elephant'.changestr('ph','hph') == 'elehphant'
'elephant'.changestr('e','')     == 'lphant'
'elephant'.changestr('','!!!')   == 'elephant'
```

The **countstr** method (see page 162) can be used to count the number of changes that could be made to a string in this fashion.

compare(target [,pad]) returns 0 if *string* and *target* are the same. If they are not, the returned number is positive and is the position of the first character that is not the same in both strings. If one string is shorter than the other, one or more *pad* characters are added on the right to make it the same length for the comparison. The default *pad* character is a blank.

Examples:

```
'abc'.compare('abc')    == 0
'abc'.compare('ak')    == 2
'ab '.compare('ab')    == 0
'ab '.compare('ab',' ') == 0
'ab '.compare('ab','x') == 3
'ab-- '.compare('ab','-') == 5
```

copies(n) returns *n* directly concatenated copies of *string*. *n* must be positive or 0; if 0, the null string is returned.

Examples:

```
'abc'.copies(3) == 'abcabcabc'
'abc'.copies(0) == ''
''.copies(2)    == ''
```

copyindexed(sub) copies the collection of indexed sub-values (see page 53) of *sub* into the collection associated with *string*, and returns the modified *string*. The resulting collection is the union of the two collections (that is, it contains the indexes and their values from both collections). If a given index exists in both collections then the sub-value of *string* for that index is replaced by the sub-value from *sub*.

The non-indexed value of *string* is not affected.

Example: Following the instructions:

```
foo='def'
foo['a']=1
foo['b']=2
bar='ghi'
bar['b']='B'
bar['c']='C'
merged=foo.copyIndexed(bar)
```

then:

```
merged['a'] == '1'
merged['b'] == 'B'
merged['c'] == 'C'
merged['d'] == 'def'
```

countstr(needle) returns the count of non-overlapping occurrences of the *needle* string in *string*, searching from left to right and starting from the first (leftmost) position in *string*.

If the *needle* is the null string, **0** is returned.

Examples:

```
'elephant'.countstr('e') == '2'  
'elephant'.countstr('ph') == '1'  
'elephant'.countstr('') == '0'
```

The **changestr** method (see page 161) can be used to change occurrences of *needle* to some other string.

c2d() Coded character to decimal. Converts the encoding of the character in *string* (which must be exactly one character) to its decimal representation. The returned string will be a non-negative number that represents the encoding of the character and will not include any sign, blanks, insignificant leading zeros, or decimal part.

Examples:

```
'M'.c2d == '77' -- ASCII or Unicode  
'7'.c2d == '247' -- EBCDIC  
'\textbackslash{r}'.c2d == '13' -- ASCII or Unicode  
'\textbackslash{0}'.c2d == '0'
```

The **c2x** method (see page 163) can be used to convert the encoding of a character to a hexadecimal representation.

c2x() Coded character to hexadecimal. Converts the encoding of the character in *string* (which must be exactly one character) to its hexadecimal representation (unpacks). The returned string will use uppercase Roman letters for the values A-F, and will not include any blanks. Insignificant leading zeros are removed.

Examples:

```
'M'.c2x == '4D' -- ASCII or Unicode  
'7'.c2x == 'F7' -- EBCDIC  
'\textbackslash{r}'.c2x == 'D' -- ASCII or Unicode  
'\textbackslash{0}'.c2x == '0'
```

The **c2d** method (see page 163) can be used to convert the encoding of a character to a decimal number.

datatype(option) returns 1 if *string* matches the description requested with the *option*, or 0 otherwise. If *string* is the null string, 0 is always returned.

Only the first character of *option* is significant, and it may be in either uppercase or lowercase. The following *option* characters are recognized:

A (Alphanumeric); returns 1 if *string* only contains characters from the ranges "a-z", "A-Z", and "0-9".

B (Binary); returns 1 if *string* only contains the characters "0" and/or "1".

D (Digits); returns 1 if *string* only contains characters from the range "0-9".

L (Lowercase); returns 1 if *string* only contains characters from the range "a-z".

M (Mixed case); returns 1 if *string* only contains characters from the ranges "a-z" and "A-Z".

N (Number); returns 1 if *string* is a syntactically valid NetREXX number that could be added to '0' without error,

- S** (Symbol); returns 1 if *string* only contains characters that are valid in non-numeric symbols (the alphanumeric characters and underscore), and does not start with a digit. Note that both uppercase and lowercase letters are permitted.
- U** (Uppercase); returns 1 if *string* only contains characters from the range "A-Z".
- W** (Whole Number); returns 1 if *string* is a syntactically valid NetREXX number that can be added to '0' without error, and whose decimal part after that addition, with no rounding, is zero.
- X** (heXadecimal); returns 1 if *string* only contains characters from the ranges "a-f", "A-F", and "0-9".

Examples:

```
'101'.datatype('B')    == 1
'12.3'.datatype('D')   == 0
'12.3'.datatype('N')   == 1
'12.3'.datatype('W')   == 0
'LaArca'.datatype('M') == 1
''.datatype('M')       == 0
'Llanes'.datatype('L') == 0
'3 d'.datatype('s')    == 0
'BCd3'.datatype('X')   == 1
'BCgd3'.datatype('X')  == 0
```

Note: The **datatype** method tests the meaning of the characters in a string, independent of the encoding of those characters. Extra letters and Extra digits cause **datatype** to return 0 except for the number tests ("N" and "W"), which treat extra digits whose value is in the range 0-9 as though they were the corresponding Arabic numeral.

date() see page 187

delstr(n [,length]) returns a copy of *string* with the sub-string of *string* that begins at the *n*th character, and is of length *length* characters, deleted. If *length* is not specified, or is greater than the number of characters from *n* to the end of the string, the rest of the string is deleted (including the *n*th character). *length* must be a non-negative whole number, and *n* must be a positive whole number. If *n* is greater than the length of *string*, the string is returned unchanged.

Examples:

```
'abcd'.delstr(3)      == 'ab'
'abcde'.delstr(3,2)  == 'abe'
'abcde'.delstr(6)    == 'abcde'
```

delword(n [,length])

returns a copy of *string* with the sub-string of *string* that starts at the *n*th word, and is of length *length* blank-delimited words, deleted. If *length* is not specified, or is greater than number of remaining words in the string, it defaults to be the remaining words in the string (including the *n*th word). *length* must be a non-negative whole number, and *n* must be a positive whole number. If *n* is greater than the number of words in *string*, the string is returned unchanged. The string deleted includes any blanks following

the final word involved, but none of the blanks preceding the first word involved.

Examples:

```
'Now is the time'.delword(2,2) == 'Now time'  
'Now is the time '.delword(3)  == 'Now is '  
'Now time'.delword(5)         == 'Now time'
```

d2b([n]) Decimal to binary. Returns a string of binary characters of length as needed or of length *n*, which is the binary representation of the decimal number. The returned string will use 0 and 1 characters for binary values. *string* must be a whole number, and must be non-negative unless *n* is specified, or an error will result. If *n* is not specified, the length of the result returned is such that there are no leading 0 characters, unless *string* was equal to 0 (in which case '0' is returned).

If *n* is specified it is the length of the final result in characters; that is, after conversion the input string will be sign-extended to the required length (negative numbers are converted assuming twos-complement form). If the number is too big to fit into *n* characters, it will be truncated on the left. *n* must be a nonnegative whole number.

Examples:

```
'0'.d2b == 0  
'9'.d2b == 1001  
'19'.d2b == 10011  
'129'.d2b == 10000001  
'129'.d2b(1) == 1  
'129'.d2b(8) == 10000001  
'127'.d2b(12) == 000001111111  
'129'.d2b(16) == 0000000010000001  
'257'.d2b(8) == 00000001  
'-127'.d2b(8) == 10000001  
'-127'.d2b(16) == 1111111110000001  
'12'.d2b(0) ==
```

d2c() Decimal to coded character. Converts the *string* (a NetREXX *number*) to a single character, where the number is used as the encoding of the character. *string* must be a non-negative whole number. An error results if the encoding described does not produce a valid character for the implementation (for example, if it has more significant bits than the implementation's encoding for characters).

Examples:

```
'77'.d2c == 'M' -- ASCII or Unicode  
'+77'.d2c == 'M' -- ASCII or Unicode  
'247'.d2c == '7' -- EBCDIC  
'0'.d2c == '\textbackslash 0'
```

d2x([n])

Decimal to hexadecimal. Returns a string of hexadecimal characters of length as needed or of length *n*, which is the hexadecimal (unpacked) representation of the decimal number. The returned string will use uppercase Roman letters for the values A-F, and will not include any blanks. *string* must be a whole number, and must be non-negative unless *n* is specified,

or an error will result. If n is not specified, the length of the result returned is such that there are no leading 0 characters, unless *string* was equal to 0 (in which case '0' is returned).

If n is specified it is the length of the final result in characters; that is, after conversion the input string will be sign-extended to the required length (negative numbers are converted assuming twos-complement form). If the number is too big to fit into n characters, it will be truncated on the left. n must be a non-negative whole number.

Examples:

```
'9'.d2x      == '9'  
'129'.d2x    == '81'  
'129'.d2x(1) == '1'  
'129'.d2x(2) == '81'  
'127'.d2x(3) == '07F'  
'129'.d2x(4) == '0081'  
'257'.d2x(2) == '01'  
'-127'.d2x(2) == '81'  
'-127'.d2x(4) == 'FF81'  
'12'.d2x(0)  == ''
```

exists(index) returns 1 if *index* names a sub-value (see page 53) of *string* that has explicitly been assigned a value, or 0 otherwise.

Example: Following the instructions:

```
vowel=0  
vowel['a']=1  
vowel['b']=1  
vowel['b']=null -- drops previous assignment
```

then:

```
vowel.exists('a') == '1'  
vowel.exists('b') == '0'  
vowel.exists('c') == '0'
```

format([before [,after]])

formats (lays out) *string*, which must be a number.

The number, *string*, is first formatted by adding zero with a digits setting that is either nine or, if greater, the number of digits in the mantissa of the number (excluding leading insignificant zeros). If no arguments are given, the result is precisely that of this operation.

The arguments *before* and *after* may be specified to control the number of characters to be used for the integer part and decimal part of the result respectively. If either of these is omitted (with no arguments specified to its right), or is **null**, the number of characters used will be as many as are needed for that part.

before must be a positive number; if it is larger than is needed to contain the integer part, that part is padded on the left with blanks to the requested length. If *before* is not large enough to contain the integer part of the number (including the sign, for negative numbers), an error results.

after must be a non-negative number; if it is not the same size as the decimal part of the number, the number will be rounded (or extended with

zeros) to fit. Specifying 0 for *after* will cause the number to be rounded to an integer (that is, it will have no decimal part or decimal point).

Examples:

```
' - 12.73'.format          == '-12.73'  
'0.000'.format            == '0'  
'3'.format(4)             == '  3'  
'1.73'.format(4,0)        == '  2'  
'1.73'.format(4,3)        == '  1.730'  
'-.76'.format(4,1)        == ' -0.8'  
'3.03'.format(4)          == '  3.03'  
' - 12.73'.format(null,4) == '-12.7300'
```

Further arguments may be passed to the `format` method to control the use of exponential notation. The full syntax of the method is then:

`format([before[,after[,expplaces[,exdigits[,exform]]]])` The first two arguments are as already described. The other three (*expplaces*, *exdigits*, and *exform*) control the exponent part of the result. The default for any of the arguments may be selected by omitting them (if there are no arguments to be specified to their right) or by using the value **null**.

expplaces must be a positive number; it sets the number of places (digits after the sign of the exponent) to be used for any exponent part, the default being to use as many as are needed. If *expplaces* is specified and is not large enough to contain the exponent, an error results. If *expplaces* is specified and the exponent will be 0, then *expplaces*+2 blanks are supplied for the exponent part of the result.

exdigits sets the trigger point for use of exponential notation. If, after the first formatting, the number of places needed before the decimal point exceeds *exdigits*, or if the absolute value of the result is less than 0.000001, then exponential form will be used, provided that *exdigits* was specified. When *exdigits* is not specified, exponential notation will never be used. The current setting of numeric digits may be used for *exdigits* by specifying the special word **digits** (see page 122). If 0 is specified for *exdigits*, exponential notation is always used unless the exponent would be 0.

exform sets the form for exponential notation (if needed). *exform* may be either '**Scientific**' (the default) or '**Engineering**'. Only the first character of *exform* is significant and it may be in uppercase or in lowercase. The current setting of numeric form may be used by specifying the special word **form** (see page 122). If engineering form is in effect, up to three digits (plus sign) may be needed for the integer part of the result (*before*).

Examples:

```
'12345.73'.format(null,null,2,2) == '1.234573E+04'  
'12345.73'.format(null,3,null,0) == '1.235E+4'  
'1.234573'.format(null,3,null,0) == '1.235'  
'123.45'.format(null,3,2,0)      == '1.235E+02'  
'1234.5'.format(null,3,2,0,'e') == '1.235E+03'  
'1.2345'.format(null,3,2,0)      == '1.235'  
'12345.73'.format(null,null,3,6) == '12345.73'  
'12345e+5'.format(null,3)        == '1234500000.000'
```

Implementation minimum: If exponents are supported in an implementation, then they must be supported for exponents whose absolute value

is at least as large as the largest number that can be expressed as an exact integer in default precision, *i.e.*, 999999999. Therefore, values for *explaces* of up to 9 should also be supported.

insert(new [,n [,length [,pad]]]) inserts the string *new*, padded or truncated to length *length*, into a copy of the target *string* after the *n*th character; the string with any inserts is returned. *length* and *n* must be a non-negative whole numbers. If *n* is greater than the length of the target string, padding is added before the *new* string also. The default value for *n* is 0, which means insert before the beginning of the string. The default value for *length* is the length of *new*. The default *pad* character is a blank.

Examples:

```
'abc'.insert('123')           == '123abc'
'abcdef'.insert(' ',3)        == 'abc def'
'abc'.insert('123',5,6)       == 'abc 123 '
'abc'.insert('123',5,6,'+')   == 'abc++123+++'
'abc'.insert('123',0,5,'-')  == '123--abc'
```

lastpos(needle [,start]) returns the position of the last occurrence of the string *needle* in *string* (the "haystack"), searching from right to left. If the string *needle* is not found, or is the null string, 0 is returned. By default the search starts at the last character of *string* and scans backwards. This may be overridden by specifying *start*, the point at which to start the backwards scan. *start* must be a positive whole number, and defaults to the value *string.length* if larger than that value or if not specified (with a minimum default value of one).

Examples:

```
'abc def ghi'.lastpos(' ')   == 8
'abc def ghi'.lastpos(' ',7) == 4
'abcdefghi'.lastpos(' ')     == 0
'abcdefghi'.lastpos('cd')    == 3
''.lastpos('?')              == 0
```

left(length [,pad]) returns a string of length *length* containing the left-most *length* characters of *string*. The string is padded with *pad* characters (or truncated) on the right as needed. The default *pad* character is a blank. *length* must be a non-negative whole number. This method is exactly equivalent to *string.substr(1, length [, pad])*.

Examples:

```
'abc d'.left(8)           == 'abc d   '
'abc d'.left(8,'.')       == 'abc d...'
'abc defg'.left(6)        == 'abc de'
```

length() returns the number of characters in *string*.

Examples:

```
'abcdefgh'.length == 8
''.length         == 0
```

lower([n [,length]])

returns a copy of *string* with any uppercase characters in the sub-string of *string* that begins at the *n*th character, and is of length *length* characters, replaced by their lowercase equivalent.

n must be a positive whole number, and defaults to 1 (the first character in *string*). If *n* is greater than the length of *string*, the string is returned unchanged.

length must be a non-negative whole number. If *length* is not specified, or is greater than the number of characters from *n* to the end of the string, the rest of the string (including the *n*th character) is assumed.

Examples:

```
'SumA'.lower      == 'suma'  
'SumA'.lower(2)   == 'Suma'  
'SuMB'.lower(1,1) == 'suMB'  
'SUMB'.lower(2,2) == 'SumB'  
''.lower         == ''
```

max(number) returns the larger of *string* and *number*, which must both be numbers. If they compare equal (that is, when subtracted, the result is 0), then *string* is selected for the result.

The comparison is effected using a numerical comparison with a *digits* setting that is either nine or, if greater, the larger of the number of digits in the mantissas of the two numbers (excluding leading insignificant zeros). The selected result is formatted by adding zero to the selected number with a *digits* setting that is either nine or, if greater, the number of digits in the mantissa of the number (excluding leading insignificant zeros). Scientific notation is used, if necessary.

Examples:

```
0.max(1)          ==1  
'-1'.max(1)       ==1  
'+1'.max(-1)      ==1  
'1.0'.max(1.00)   == '1.0'  
'1.00'.max(1.0)   == '1.00'  
'1234567000000'.max(1234567E+5) == '1234567000000'  
'1234567E+5'.max('1234567000000') == '1.234567E+11'
```

min(number) returns the smaller of *string* and *number*, which must both be numbers. If they compare equal (that is, when subtracted, the result is 0), then *string* is selected for the result.

The comparison is effected using a numerical comparison with a *digits* setting that is either nine or, if greater, the larger of the number of digits in the mantissas of the two numbers (excluding leading insignificant zeros). The selected result is formatted by adding zero to the selected number with a *digits* setting that is either nine or, if greater, the number of digits in the mantissa of the number (excluding leading insignificant zeros). Scientific notation is used, if necessary.

Examples:

```
0.min(1)          ==0  
'-1'.min(1)       == '-1'  
'+1'.min(-1)      == '-1'  
'1.0'.min(1.00)   == '1.0'  
'1.00'.min(1.0)   == '1.00'  
'1234567000000'.min(1234567E+5) == '1234567000000'  
'1234567E+5'.min('1234567000000') == '1.234567E+11'
```

overlay(new [,n [,length [,pad]]]) overlays the string *new*, padded or truncated to length *length*, onto a copy of the target *string* starting at the *n*th character; the string with any overlays is returned. Overlays may extend beyond the end of the original *string*. If *length* is specified it must be a non-negative whole number. If *n* is greater than the length of the target string, padding is added before the *new* string also. The default *pad* character is a blank, and the default value for *n* is 1. *n* must be greater than 0. The default value for *length* is the length of *new*.

Examples:

```
'abcdef'.overlay(' ',3)      == 'ab def'  
'abcdef'.overlay('.',3,2)   == 'ab. ef'  
'abcd'.overlay('qq')       == 'qqcd'  
'abcd'.overlay('qq',4)     == 'abcqq'  
'abc'.overlay('123',5,6,'+') == 'abc+123+++'
```

pos(needle [,start]) returns the position of the string *needle*, in *string* (the "haystack"), searching from left to right. If the string *needle* is not found, or is the null string, 0 is returned. By default the search starts at the first character of *string* (that is, *start* has the value 1). This may be overridden by specifying *start* (which must be a positive whole number), the point at which to start the search; if *start* is greater than the length of *string* then 0 is returned. **Examples:**

```
'Saturday'.pos('day')      == 6  
'abc def ghi'.pos('x')    == 0  
'abc def ghi'.pos(' ')    == 4  
'abc def ghi'.pos(' ',5)  == 8
```

reverse() returns a copy of *string*, swapped end for end.

Examples:

```
'ABC.'.reverse           == '.cBA'  
'XYZ '.reverse           == ' ZYX'  
'Tranquility'.reverse   == 'ytiliuqnaT'
```

right(length [,pad]) returns a string of length *length* containing the right-most *length* characters of *string* - that is, padded with *pad* characters (or truncated) on the left as needed. The default *pad* character is a blank. *length* must be a non-negative whole number.

Examples:

```
'abc d'.right(8) == ' abc d'  
'abc def'.right(5) == 'c def'  
'12'.right(5,'0') == '00012'
```

sequence(final) returns a string of all characters, in ascending order of encoding, between and including the character in *string* and the character in *final*. *string* and *final* must be single characters; if *string* is greater than *final*, an error is reported.

Examples:

```
'a'.sequence('f')      == 'abcdef'  
'\x0'.sequence('\x03') == '\x00\x01\x02\x03'  
'\ufffe'.sequence('\uffff') == '\ufffe\uffff'
```

sign() returns a number that indicates the sign of *string*, which must be a number. *string* is first formatted, just as though the operation "**string+0**" had been carried out with sufficient digits to avoid rounding. If the number then starts with '-' then '-1' is returned; if it is '0' then '0' is returned; and otherwise '1' is returned.

Examples:

```
'12.3'.sign    == 1
'0.0'.sign    == 0
'-0.307'.sign == -1
```

soundex() returns the normalized soundex value of the string. This implementation is for the English language. 3.08

Examples:

```
'EULER'.soundex() == 'E460'
```

space([n [,pad]]) returns a copy of *string* with the blank-delimited words in *string* formatted with *n* (and only *n*) *pad* characters between each word. *n* must be a non-negative whole number. If *n* is 0, all blanks are removed. Leading and trailing blanks are always removed. The default for *n* is 1, and the default *pad* character is a blank.

Examples:

```
'abc def '.space    == 'abc def '
' abc def '.space(3) == 'abc  def '
'abc def '.space(1)  == 'abc def '
'abc def '.space(0)  == 'abcdef '
'abc def '.space(2, '+') == 'abc++def '
```

strip([option [,char]]) returns a copy of *string* with Leading, Trailing, or Both leading and trailing characters removed, when the first character of *option* is L, T, or B respectively (these may be given in either uppercase or lowercase). The default is B. The second argument, *char*, specifies the character to be removed, with the default being a blank. If given, *char* must be exactly one character long.

Examples:

```
' ab c '.strip      == 'ab c '
' ab c '.strip('L') == 'ab c '
' ab c '.strip('t') == ' ab c '
'12.70000'.strip('t',0) == '12.7'
'0012.700'.strip('b',0) == '12.7'
```

substr(n [,length [,pad]]) returns the sub-string of *string* that begins at the *n*th character, and is of length *length*, padded with *pad* characters if necessary. *n* must be a positive whole number, and *length* must be a non-negative whole number. If *n* is greater than *string.length*, then only *pad* characters can be returned. If *length* is omitted it defaults to be the rest of the string (or 0 if *n* is greater than the length of the string). The default *pad* character is a blank.

Examples:

```
'abc'.substr(2)      == 'bc '
'abc'.substr(2,4)    == 'bc  '
'abc'.substr(5,4)    == '    '
```

```
'abc'.substr(2,6,'.') == 'bc....'  
'abc'.substr(5,6,'.') == '.....'
```

Note: In some situations the positional (numeric) patterns of parsing templates are more convenient for selecting sub-strings, especially if more than one sub-string is to be extracted from a string.

subword(*n* [,*length*]) returns the sub-string of *string* that starts at the *n*th word, and is up to *length* blank-delimited words long. *n* must be a positive whole number; if greater than the number of words in the string then the null string is returned. *length* must be a non-negative whole number. If *length* is omitted it defaults to be the remaining words in the string. The returned string will never have leading or trailing blanks, but will include all blanks between the selected words.

Examples:

```
'Now is the time'.subword(2,2) == 'is the'  
'Now is the time'.subword(3)  == 'the time'  
'Now is the time'.subword(5)  == ''
```

time() see page 189

translate(*tableo*, *tablei* [,*pad*]) returns a copy of *string* with each character in *string* either unchanged or translated to another character.

The **translate** method acts by searching the input translate table, *tablei*, for each character in *string*. If the character is found in *tablei* (the first, leftmost, occurrence being used if there are duplicates) then the corresponding character in the same position in the output translate table, *tableo*, is used in the result string; otherwise the original character found in *string* is used. The result string is always the same length as *string*.

The translate tables may be of any length, including the null string. The output table, *tableo*, is padded with *pad* or truncated on the right as necessary to be the same length as *tablei*. The default *pad* is a blank.

Examples:

```
'abc'.translate('\&', 'b')          == 'a\&\&c'  
'abcdef'.translate('12', 'ec')     == 'ab2d1f'  
'abcdef'.translate('12', 'abcd', '.') == '12..ef'  
'4123'.translate('abcd', '1234')   == 'dabc'  
'4123'.translate('hods', '1234')   == 'shod'
```

Note: The last two examples show how the **translate** method may be used to move around the characters in a string. In these examples, any 4-character string could be specified as the first argument and its last character would be moved to the beginning of the string. Similarly, the term:

```
'gh.ef.abcd'.translate(19970827, 'abcdefgh')
```

(which returns "27.08.1997") shows how a string (in this case perhaps a date) might be re-formatted and merged with other characters using the **translate** method.

trunc([*n*]) returns the integer part of *string*, which must be a number, with *n* decimal places (digits after the decimal point). *n* must be a non-negative whole number, and defaults to zero.

The number *string* is formatted by adding zero with a digits setting that is either nine or, if greater, the number of digits in the mantissa of the number (excluding leading insignificant zeros). It is then truncated to *n* decimal places (or trailing zeros are added if needed to make up the specified length). If *n* is 0 (the default) then an integer with no decimal point is returned. The result will never be in exponential form.

Examples:

```
'12.3'.trunc          == 12
'127.09782'.trunc(3) == 127.097
'127.1'.trunc(3)     == 127.100
'127'.trunc(2)       == 127.00
'0'.trunc(2)         == 0.00
```

upper([*n* [,*length*]]) returns a copy of *string* with any lowercase characters in the sub-string of *string* that begins at the *n*th character, and is of length *length* characters, replaced by their uppercase equivalent.

n must be a positive whole number, and defaults to 1 (the first character in *string*). If *n* is greater than the length of *string*, the string is returned unchanged.

length must be a non-negative whole number. If *length* is not specified, or is greater than the number of characters from *n* to the end of the string, the rest of the string (including the *n*th character) is assumed.

Examples:

```
'Fou-Baa'.upper      == 'FOU-BAA'
'Mad Sheep'.upper    == 'MAD SHEEP'
'Mad sheep'.upper(5) == 'Mad SHEEP'
'Mad sheep'.upper(5,1) == 'Mad Sheep'
'Mad sheep'.upper(5,4) == 'Mad SHEEP'
'tinganon'.upper(1,1) == 'Tinganon'
''.upper             == ''
```

verify(reference [,option [,start]]) verifies that *string* is composed only of characters from *reference*, by returning the position of the first character in *string* that is not also in *reference*. If all the characters were found in *reference*, 0 is returned. The *option* may be either **'Nomatch'** (the default) or **'Match'**. Only the first character of *option* is significant and it may be in uppercase or in lowercase. If **'Match'** is specified, the position of the first character in *string* that **is** in *reference* is returned, or 0 is returned if none of the characters were found. The default for *start* is 1 (that is, the search starts at the first character of *string*). This can be overridden by giving a different *start* point, which must be positive. If *string* is the null string, the method returns 0, regardless of the value of the *option*. Similarly if *start* is greater than *string.length*, 0 is returned. If *reference* is the null string, then the returned value is the same as the value used for *start*, unless **'Match'** is specified as the *option*, in which case 0 is returned.

Examples:

```
'123'.verify('1234567890') == 0
'1Z3'.verify('1234567890') == 2
'AB4T'.verify('1234567890','M') == 3
'1P3Q4'.verify('1234567890','N',3) == 4
```



```
'ABCDE'.verify('', 'n', 3) == 3
'AB3CD5'.verify('1234567890', 'm', 4) == 6
```

word(*n*) returns the *n*th blank-delimited word in *string*. *n* must be positive. If there are fewer than *n* words in *string*, the null string is returned. This method is exactly equivalent to *string*.**subword**(*n*,1).

Examples:

```
'Now is the time'.word(3) == 'the'
'Now is the time'.word(5) == ''
```

wordindex(*n*) returns the character position of the *n*th blank-delimited word in *string*. *n* must be positive. If there are fewer than *n* words in the string, 0 is returned.

Examples:

```
'Now is the time'.wordindex(3) == 8
'Now is the time'.wordindex(6) == 0
```

wordlength(*n*) returns the length of the *n*th blank-delimited word in *string*. *n* must be positive. If there are fewer than *n* words in the string, 0 is returned.

Examples:

```
'Now is the time'.wordlength(2) == 2
'Now comes the time'.wordlength(2) == 5
'Now is the time'.wordlength(6) == 0
```

wordpos(*phrase* [*start*]) searches *string* for the first occurrence of the sequence of blank-delimited words *phrase*, and returns the word number of the first word of *phrase* in *string*. Multiple blanks between words in either *phrase* or *string* are treated as a single blank for the comparison, but otherwise the words must match exactly. Similarly, leading or trailing blanks on either string are ignored. If *phrase* is not found, or contains no words, 0 is returned. By default the search starts at the first word in *string*. This may be overridden by specifying *start* (which must be positive), the word at which to start the search.

Examples:

```
'now is the time'.wordpos('the') == 3
'now is the time'.wordpos('The') == 0
'now is the time'.wordpos('is the') == 2
'now is the time'.wordpos('is the') == 2
'now is the time'.wordpos('is time') == 0
'To be or not to be'.wordpos('be') == 2
'To be or not to be'.wordpos('be', 3) == 6
```

words() returns the number of blank-delimited words in *string*.

Examples:

```
'Now is the time'.words == 4
''.words == 0
''.words == 0
```

x2b() Hexadecimal to binary. Converts *string* (a string of at least one hexadecimal characters) to an equivalent string of binary digits. Hexadecimal characters may be any decimal digit character (0-9) or any of the first six alphabetic characters (a-f), in either lowercase or uppercase. *string* may be

of any length; each hexadecimal character will be converted to a string of four binary digits. The returned string will have a length that is a multiple of four, and will not include any blanks.

Examples:

```
'C3'.x2b == '11000011'  
'7'.x2b  == '0111'  
'1C1'.x2b == '000111000001'
```

x2c() Hexadecimal to coded character. Converts the *string* (a string of hexadecimal characters) to a single character (packs). Hexadecimal characters may be any decimal digit character (0-9) or any of the first six alphabetic characters (a-f), in either lowercase or uppercase.

string must contain at least one hexadecimal character; insignificant leading zeros are removed, and the string is then padded with leading zeros if necessary to make a sufficient number of hexadecimal digits to describe a character encoding for the implementation.

An error results if the encoding described does not produce a valid character for the implementation (for example, if it has more significant bits than the implementation's encoding for characters). **Examples:**

```
'004D'.x2c == 'M' -- ASCII or Unicode  
'4d'.x2c  == 'M' -- ASCII or Unicode  
'A2'.x2c  == 's' -- EBCDIC  
'0'.x2c   == '\textbackslash 0'
```

The **d2c** method (see page 165) can be used to convert a NetREXX number to the encoding of a character.

x2d([n]) Hexadecimal to decimal. Converts the *string* (a string of hexadecimal characters) to a decimal number, without rounding. If *string* is the null string, 0 is returned.

If *n* is not specified, *string* is taken to be an unsigned number.

Examples:

```
'0E'.x2d == 14  
'81'.x2d == 129  
'F81'.x2d == 3969  
'FF81'.x2d == 65409  
'c6f0'.x2d == 50928
```

If *n* is specified, *string* is taken as a signed number expressed in *n* hexadecimal characters. If the most significant (left-most) bit is zero then the number is positive; otherwise it is a negative number in twos-complement form. In both cases it is converted to a NetREXX number which may, therefore, be negative. If *n* is 0, 0 is always returned.

If necessary, *string* is padded on the left with '0' characters (note, not "sign-extended"), or truncated on the left, to length *n* characters; (that is, as though *string*.right(*n*, '0') had been executed.)

Examples:

```
'81'.x2d(2) == -127  
'81'.x2d(4) == 129  
'F081'.x2d(4) == -3967  
'F081'.x2d(3) == 129
```

```
'F081'.x2d(2) == -127  
'F081'.x2d(1) == 1  
'0031'.x2d(0) == 0
```

The **c2d** method (see page 163) can be used to convert a character to a decimal representation of its encoding.

Appendix A - A Sample NetREXX Program

This appendix includes a short program, called **qtime**, which is an example of a "real" NetREXX program. The programs included elsewhere in this book have been contrived to illustrate specific points. By contrast, **qtime** is a simple but useful tool that genuinely improves the human factors of computer systems. People frequently wish to know the time of day, and this program presents this information in a natural way.

The style used for this example is the same as that used throughout the book, with all symbols except those describing classes being written in lower case. Other NetREXX programming styles are possible, of course; NetREXX syntax is designed to permit a wide variety of styles with a minimum of punctuation.

The **qtime** program is a modification of one of the first REXX programs ever written (much of the program is identical). The main changes are:

- Indexed variables (brackets notation) are used instead of REXX stems.
- The **word** method from the **REXX** class is used instead of the **word** REXX built-in function.
- The Java **Date** class is used to determine the current time.

qtime.nrx - Query Time

```

/*-----*/
/* QTIME. This program displays the time in real English. */
/* If "?" is given as the first argument word then the */
/* program displays a description of itself. */
/*-----*/

/*----- First process any argument words -----*/
parse arg parm . /* get the first argument word */
select
  when parm='?' then tell /* say what we do */
  when parm='' then nop /* OK (no first argument) */
  otherwise
    say 'The only valid argument to QTIME is "?". The word'
    say 'that you supplied ("'parm'") has been ignored.'
    tell /* usually helpful to describe the program */
end

/*----- Now start processing in earnest -----*/
/* Nearness phrases - using associative array lookup */
near='' /* default */
near[0]='' /* exact */
near[1]=' just gone'; near[2]=' just after' /* after */
near[3]=' nearly'; near[4]=' almost' /* before */

```

```

/* Extract the hours, minutes, and seconds from the time. */
/* Use the Java Date class as Rexx.Date not yet implemented */
parse Date() . . . now . /* time is fourth word */
parse now hour':'min':'sec

if sec>29 then min=min+1 /* round up minutes */
mod=min//5 /* where we are in 5 minute bracket */
out="It's"near[mod] /* start building the result */
if min>32 then hour=hour+1 /* we are TO the hour... */
min=min+2 /* shift minutes to straddle a 5-minute point */

/* Now special-case the result for Noon and Midnight hours */
if hour//12=0 & min//60<=4 then do
  if hour=12 then say out 'Noon.'
  else say out 'Midnight.'
  return /* we are finished here */
end

min=min-(min//5) /* find nearest 5 mins */
if hour>12
then hour=hour-12 /* get rid of 24-hour clock */
else
  if hour=0 then hour=12 /* .. and allow for midnight */

/* Determine the phrase to use for each 5-minute segment */
select
  when min=0 then nop /* add "o'clock" later */
  when min=60 then min=0 /* ditto */
  when min= 5 then out=out 'five past'
  when min=10 then out=out 'ten past'
  when min=15 then out=out 'a quarter past'
  when min=20 then out=out 'twenty past'
  when min=25 then out=out 'twenty-five past'
  when min=30 then out=out 'half past'
  when min=35 then out=out 'twenty-five to'
  when min=40 then out=out 'twenty to'
  when min=45 then out=out 'a quarter to'
  when min=50 then out=out 'ten to'
  when min=55 then out=out 'five to'
end

numbers='one two three four five six'- /* (continuation) */
'seven eight nine ten eleven twelve '
out=out numbers.word(hour) /* add the hour number */
if min=0 then out=out "o'clock" /* .. and o'clock if exact */

say out'.' /* display the final result */

/*-----*/
/* Subroutine that describes the purpose of the program */
/*-----*/
method tell static
say 'QTIME will display the current time in real English.'
say 'Call without any arguments to display the time, or with'
say ' "?" to display this information.'
say 'British English idioms are used in this program.'
say /* blank line - we are about to continue and show time */

```

return

/* Mike Cowlshaw, December 1979 - January 1985
/* NetRexx version March 1996

*/
*/

Appendix B - The `netrexx.lang` Package

This appendix documents the `netrexx.lang` package, which includes the classes used for creating string objects of type `REXX` along with several classes that are often used while running NetREXX programs.

This appendix describes the public methods and properties of these classes, as implemented by the reference implementation. It does not include those "built-in" Methods for NetREXX (see page 159) strings in the `REXX` class that form part of the NetREXX language, or those classes and methods that are internal "helper" components (which, for example, are used as repositories for rarely-executed code).

The classes in the `netrexx.lang` package are:

- The Exception classes (see page 180)
- `REXX` (see page 181)
- `REXXIO` (helper class, see page 186)
- `REXXNode` (helper class, for indexed strings)
- `REXXOperators` interface (see page 191)
- `REXXParse` (helper class, for parse)
- `REXXSet` (see page 191)
- `REXXTrace` (helper class, for trace)
- `REXXUtil` (helper class, for the `REXX` class)
- `REXXWords` (helper class, for the `REXX` class)

50.1 Exception classes

The classes provided for exceptions in the `netrexx.lang` package are all subclasses of `java.lang.RuntimeException` and all have the same content. Each has two constructors: one taking no argument and the other taking a string of type `java.lang.String`, which is used for additional detail describing the exception.

The Exceptions are signalled as follows.

BadArgumentException signalled when an argument to a method is incorrect.

BadColumnException signalled when a column number in a parsing template is not valid (for example, not a number).

BadNumericException signalled when a numeric `digits` instruction tries to set a value that is not a whole number, or is not positive, or is more than nine digits.

DivideException signalled when an error occurs during a division. This may be due to an attempt to divide by zero, or when the intermediate result of an integer divide or remainder operation is not valid.

ExponentOverflowException signalled when the exponent resulting from an operation would require more than nine digits.

NoOtherwiseException signalled when a `select` construct does not supply an `otherwise` clause and none of expressions on the `when` clauses resulted in `'1'`.

NotCharacterException signalled when a conversion from a string to a single character was attempted but the string was not exactly one character long.

NotLogicException signalled when a conversion from a string to a boolean was attempted but the string was neither the string `'0'` nor the string `'1'`.

Other exceptions, from the `java.lang` package, may also be signalled, for example **NumberFormatException** or **NullPointerException**.

50.2 The REXX class

The class `netrexx.lang.REXX` implements the NetREXX string class, and includes the "built-in" Methods for NetREXX strings (see page 159).

Described here are the platform-dependent methods as provided in the reference implementation: constructors (see page 181) for the class, the methods for arithmetic operations (see page 182) , and miscellaneous methods (see page 185) intended for general use.

The class `netrexx.lang.REXX` is serializable.

50.3 REXX constructors

These constructors all create a string of type `netrexx.lang.REXX`.

`REXX(arg=boolean)`

Constructs a string which will have the value `'1'` if `arg` is 1 (*true*) or the value `'0'` if `arg` is 0 (*false*).

`REXX(arg=byte)`

Constructs a string which is the decimal representation of the 8-bit signed binary integer `arg`. The string will contain only decimal digits, prefixed with a leading minus sign (hyphen) if `arg` is negative. A leading zero will be present only if `arg` is zero.

`REXX(arg=char)`

Constructs a string of length 1 whose first and only character is a copy of `arg`.

`REXX(arg=char[])`

Constructs a string by copying the characters of the character array `arg` in sequence. The length of the string is the number of elements in the character array (that is, `arg.length`).

REXX(arg=int)

Constructs a string which is the decimal representation of the 32-bit signed binary integer *arg*. The string will contain only decimal digits, prefixed with a leading minus sign (hyphen) if *arg* is negative. A leading zero will be present only if *arg* is zero.

REXX(arg=double)

Constructs a string which is the decimal representation of the 64-bit signed binary floating point number *arg*. (*The precise format of the result may change and will be defined later.*)

REXX(arg=float)

Constructs a string which is the decimal representation of the 32-bit signed binary floating point number *arg*. (*The precise format of the result may change and will be defined later.*)

REXX(arg=long)

Constructs a string which is the decimal representation of the 64-bit signed binary integer *arg*. The string will contain only decimal digits, prefixed with a leading minus sign (hyphen) if *arg* is negative. A leading zero will be present only if *arg* is zero.

REXX(arg=REXX)

Constructs a string which is copy of *arg*, which is of type **netrex.lang.REXX**. *arg* must not be **null**. Any sub-values (see page 53) are ignored (that is, they are not present in the object returned by the constructor).

REXX(arg=short)

Constructs a string which is the decimal representation of the 16-bit signed binary integer *arg*. The string will contain only decimal digits, prefixed with a leading minus sign (hyphen) if *arg* is negative. A leading zero will be present only if *arg* is zero.

REXX(arg=String)

Constructs a NetREXX string by copying the characters of *arg*, which is of type **java.lang.String**, in sequence. The length of the string is same as the length of *arg* (that is, **arg.length()**). *arg* must not be **null**.

REXX(arg=String[])

Constructs a NetREXX string by concatenating the elements of the **java.lang.String** array *arg* together in sequence with a blank between each pair of elements. This may be used for converting the argument word array passed to the **main** method of a Java application into a single string. If the number of elements of *arg* is zero then an empty string (of length 0) is returned. Otherwise, the length of the string is the sum of the lengths of the elements of *arg*, plus the number of elements of *arg*, less one. *arg* must not be **null**.

50.4 REXX arithmetic methods

These methods implement the NetREXX arithmetic operators, as described in the section on *Numbers and (see page 136) arithmetic*. Each corresponds to and implements a method in the **REXXOperators** interface class (see page 191) .

Each of the methods here takes a `RexxSet` (see page 191) object as an argument. This argument provides the numeric settings for the operation; if `null` is provided for the argument then the default settings are used (`digits=9`, `form=scientific`).

For monadic operators, only the `RexxSet` argument is present; the operation acts upon the current object. For dyadic operators, the `RexxSet` argument and a `Rexx` argument are present; the operation acts with the current object being the left-hand operand and the second argument being the right-hand operand. For example, under default numeric settings, the expression:

```
award+extra
```

(where *award* and *extra* are references to objects of type `Rexx`) could be written as:

```
award.OpAdd(null, extra)
```

which would return the result of adding *award* and *extra* under the default numeric settings.

`OpAdd(set=RexxSet, rhs=Rexx)`

Implements the NetRexx `+` (Add) operator, and returns the result as a string of type `Rexx`.

`OpAnd(set=RexxSet, rhs=Rexx)`

Implements the NetRexx `&` (And) operator, and returns a result (0 or 1) of type `boolean`.

`OpCc(set=RexxSet, rhs=Rexx)`

Implements the NetRexx `||` or *abuttal* (Concatenate without blank) operator, and returns the result as a string of type `Rexx`.

`OpCcblank(set=RexxSet, rhs=Rexx)`

Implements the NetRexx *blank* (Concatenate with blank) operator, and returns the result as a string of type `Rexx`.

`OpDiv(set=RexxSet, rhs=Rexx)`

Implements the NetRexx `/` (Divide) operator, and returns the result as a string of type `Rexx`.

`OpDivI(set=RexxSet, rhs=Rexx)`

Implements the NetRexx `%` (Integer divide) operator, and returns the result as a string of type `Rexx`.

`OpEq(set=RexxSet, rhs=Rexx)`

Implements the NetRexx `=` (Equal) operator, and returns a result (0 or 1) of type `boolean`.

`OpEqS(set=RexxSet, rhs=Rexx)`

Implements the NetRexx `==` (Strictly equal) operator, and returns a result (0 or 1) of type `boolean`.

`OpGt(set=RexxSet, rhs=Rexx)`

Implements the NetRexx `>` (Greater than) operator, and returns a result (0 or 1) of type `boolean`.

- OpGtEq(set=REXXSet, rhs=REXX)
 Implements the NetREXX \geq (Greater than or equal) operator, and returns a result (0 or 1) of type **boolean**.
- OpGtEqS(set=REXXSet, rhs=REXX)
 Implements the NetREXX \gg (Strictly greater than or equal) operator, and returns a result (0 or 1) of type **boolean**.
- OpGtS(set=REXXSet, rhs=REXX)
 Implements the NetREXX \gg (Strictly greater than) operator, and returns a result (0 or 1) of type **boolean**.
- OpLt(set=REXXSet, rhs=REXX)
 Implements the NetREXX $<$ (Less than) operator, and returns a result (0 or 1) of type **boolean**.
- OpLtEq(set=REXXSet, rhs=REXX)
 Implements the NetREXX \leq (Less than or equal) operator, and returns a result (0 or 1) of type **boolean**.
- OpLtEqS(set=REXXSet, rhs=REXX)
 Implements the NetREXX \ll (Strictly less than or equal) operator, and returns a result (0 or 1) of type **boolean**.
- OpLtS(set=REXXSet, rhs=REXX)
 Implements the NetREXX \ll (Strictly less than) operator, and returns a result (0 or 1) of type **boolean**.
- OpMinus(set=REXXSet)
 Implements the NetREXX **Prefix -** (Minus) operator , and returns the result as a string of type **REXX**.
- OpMult(set=REXXSet, rhs=REXX)
 Implements the NetREXX ***** (Multiply) operator , and returns the result as a string of type **REXX**.
- OpNot(set=REXXSet)
 Implements the NetREXX **Prefix ** (Not) operator, and returns a result (0 or 1) of type **boolean**.
- OpNotEq(set=REXXSet, rhs=REXX)
 Implements the NetREXX \neq (Not equal) operator, and returns a result (0 or 1) of type **boolean**.
- OpNotEqS(set=REXXSet, rhs=REXX)
 Implements the NetREXX \neq (Strictly not equal) operator, and returns a result (0 or 1) of type **boolean**.
- OpOr(set=REXXSet, rhs=REXX)
 Implements the NetREXX **|** (Inclusive or) operator, and returns a result (0 or 1) of type **boolean**.
- OpPlus(set=REXXSet)
 Implements the NetREXX **Prefix +** (Plus) operator , and returns the result as a string of type **REXX**.
- OpPow(set=REXXSet, rhs=REXX)
 Implements the NetREXX ****** (Power) operator , and returns the result as a string of type **REXX**.

OpRem(set=R_{EXX}Set, rhs=R_{EXX})

Implements the NetR_{EXX} // (Remainder) operator , and returns the result as a string of type **R_{EXX}**.

OpSub(set=R_{EXX}Set, rhs=R_{EXX})

Implements the NetR_{EXX} - (Subtract) operator, and returns the result as a string of type **R_{EXX}**.

OpXor(set=R_{EXX}Set, rhs=R_{EXX})

Implements the NetR_{EXX} && (Exclusive or) operator, and returns a result (0 or 1) of type **boolean**.

50.5 REXX miscellaneous methods

These methods provide standard Java methods for the class, together with various conversions.

charAt(offset=int)

Returns the character from the string at *offset* (that is, if *offset* is 0 then the first character is returned, *etc.*). The character is returned as type **char**.

If *offset* is negative, or is greater than or equal to the length of the string, an exception is signalled.

equals(item=Object)

Compares the string with the value of *item*, using a strict character-by-character comparison, and returns a result of type **boolean**.

If *item* is **null** or is not an instance of one of the types **R_{EXX}**, **java.lang.String**, or **char[]**, then 0 is returned. Otherwise, *item* is converted to type **R_{EXX}** and the OpEqS (see page 183) method (or equivalent) is used to compare the current string with the converted string, and its result is returned.

hashCode()

Returns a hashcode of type **int** for the string. This hashcode is suitable for use by the **java.util.Hashtable** class.

toboolean()

Converts the string to type **boolean**. If the string is neither "0" nor "1" then a **NotLogicException** (see page 181) is signalled.

tobyte()

Converts the string to type **byte**. If the string is not a number, has a non-zero decimal part, or is out of the possible range for a **byte** (8-bit signed integer) result then a **NumberFormatException** is signalled.

toByteArray()

Converts the string to type **byte[]**, observing the standard codepage for the platform. 3.04

tochar()

Converts the string to type **char**. If the string is not exactly one character in length then a **NotCharacterException** (see page 181) is signalled.

toCharArray()

Converts the string to type **char[]**. A character array object of the same length as the string is created, and the characters of the string are copied to the array in sequence. The character array is then returned.

`toDouble()`

Converts the string to type **double**. If the string is not a number, or is out of the possible range for a **double** (64-bit signed floating point) result then a **NumberFormatException** is signalled.

`toFloat()`

Converts the string to type **float**. If the string is not a number, or is out of the possible range for a **float** (32-bit signed floating point) result then a **NumberFormatException** is signalled.

`toint()`

Converts the string to type **int**. If the string is not a number, has a non-zero decimal part, or is out of the possible range for an **int** (32-bit signed integer) result then a **NumberFormatException** is signalled.

`tolong()`

Converts the string to type **long**. If the string is not a number, has a non-zero decimal part, or is out of the possible range for a **long** (64-bit signed integer) result then a **NumberFormatException** is signalled. [%hide

`toREXX(arg=char[])` static

Takes *arg*, an array of characters, and returns a copy of it as a string of type **netrexx.lang.REXX**. If the argument is **null**, then **null** is returned (not a null string). This is a static method (a function).

`toREXX(arg=String)` static

Takes *arg*, a **java.lang.String**, and returns a copy of it as a string of type **netrexx.lang.REXX**. If the argument is **null**, then **null** is returned (not a null string). This is a static method (a function).

`toshort()`

Converts the string to type **short**. If the string is not a number, has a non-zero decimal part, or is out of the possible range for a **short** (16-bit signed) result then a **NumberFormatException** is signalled.

`toString()`

Converts the string to type **java.lang.String**. A String object of the same length as the string is created, and the characters of the string are copied to the new string in sequence. The String is then returned.

50.6 The REXXIO class

The `RexxIO` class implements a number of helper methods, for example `RexxIO.say`, a call to which is generated when a program contains a `say` statement.

Ask() static returns **REXX** get a line of text from the console

AskOne() static returns **REXX** get one character from the console (still requires a return)

Say(obj=Object) static returns boolean put a line out to the console If the line ends in the NUL character ('-' or '0') then no line termination is provided (and the NUL is deleted). If the write succeeds 0 is returned, otherwise 1 is returned.

Say(str=String) static returns boolean put a line out to the console

Say(line=REXX) static returns boolean put a line out to the console

Say(c=char) static returns boolean put a line out to the console

Say(n=long) static returns boolean put a line out to the console

Say(f=float) static returns boolean put a line out to the console

Say(d=double) static returns boolean put a line out to the console

Say(b=boolean) static returns boolean put a line out to the console

Say(aline=char[]) static returns boolean
put a line out to the console

setOutputStream(out=OutputStream) static change the outputstream for say 3.07
to use

pushOutputStream(out=OutputStream) static push an outputstream on the 3.07
decque, for say to use

popOutputStream() static remove an outputstream from the decque, will not 3.07
be used anymore

File(nm) returns REXXIO define a file to the REXXIO instance 3.07

forEachline(c=LineHandler) define a callback that calls an instance of the 3.07
LineHandler interface

forEachline(c=LineHandler,numLines) define a callback that calls an instance 3.07
of the LineHandler interface that is only called a number of times
as specified in numLines

50.7 The REXXDate class

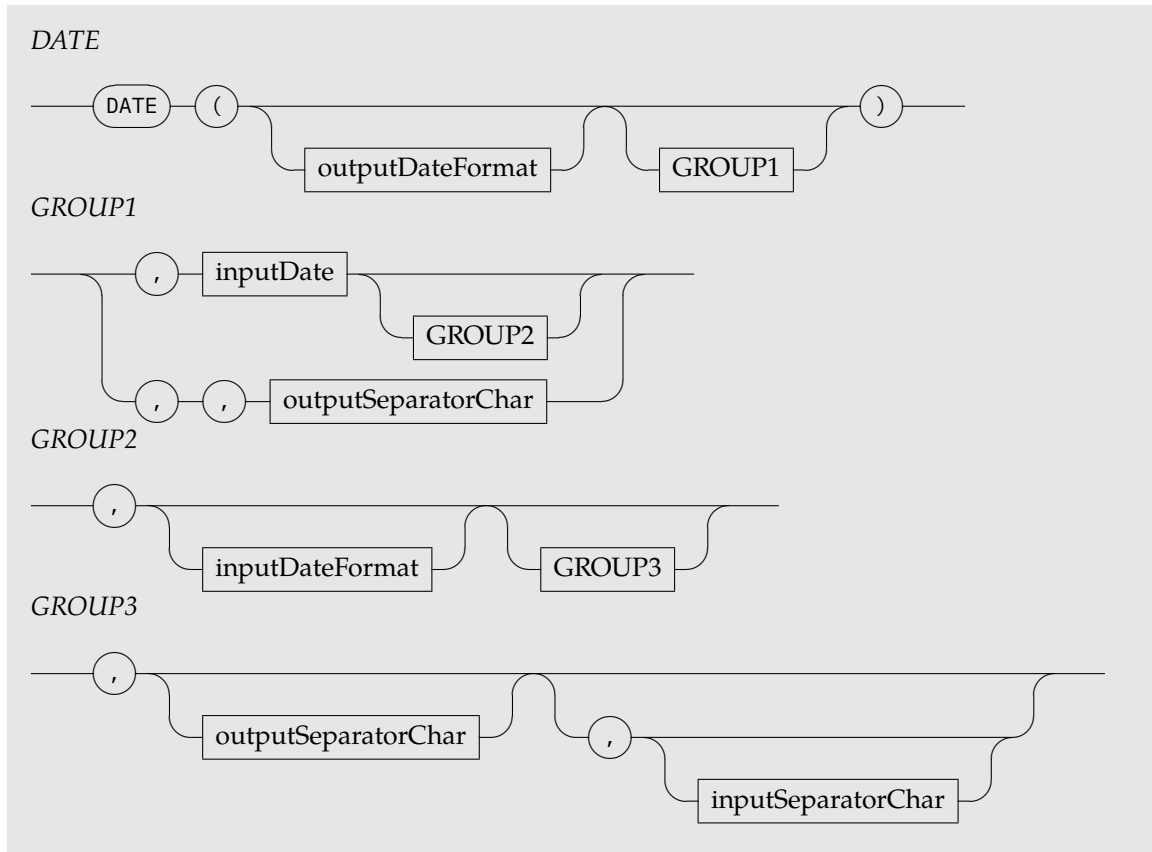
The REXXDate class inherits from REXXTime which implements a superset of the Classic REXX Date() and Time() functions⁸², including some of the options that were available in REXX/VM but left out of the REXX ANSI/ISO/INCITS standard.

The date() function can be called standalone when using a uses REXXDate option on the class statement. You can use the following options to obtain specific date formats. (Only the capitalized letter is needed; all characters following it are ignored.)

50.7.1 Options

Base the number of complete days (that is, not including the current day) since and including the base date, 1 January 0001, in the format: dddddd (no leading zeros or blanks). The expression DATE('B')//7 returns a number in the range 0–6 that corresponds to the current day of the week, where 0

⁸²At the 4.02 level, including the input and conversion functions



is Monday and 6 is Sunday. Thus, this function can be used to determine the day of the week independent of the national language in which you are working. Note: The base date of 1 January 0001 is determined by extending the current Gregorian calendar backward (365 days each year, with an extra day every year that is divisible by 4 except century years that are not divisible by 400). It does not take into account any errors in the calendar system that created the Gregorian calendar originally.

Century the number of days, including the current day, since and including January 1 of the last year that is a multiple of 100 in the form: ddddd (no leading zeros). Example: A call to DATE('C') on March 13 1992 returns 33675, the number of days from 1 January 1900 to 13 March 1992. Similarly, a call to DATE('C') on 2 January 2000 returns 2, the number of days from 1 January 2000 to 2 January 2000. Note: When the Century option is used for input, the output may change, depending on the current century. For example, if DATE('S','1','C') was entered on any day between 1 January 1900 and 31 December 1999, the result would be 19000101. However, if DATE('S','1','C') was entered on any day between 1 January 2000 and 31 December 2099, the result would be 20000101. It is important to understand the above, and code accordingly.

Days the number of days, including the current day, so far in the current year in the format: ddd (no leading zeros or blanks).

Julian date in the format: yyyyddd (yy and ddd must have leading zeros).

European date in the format: dd/mm/yy (dd, mm, and yy must have leading

zeros).

Month full name of the current month. Only valid for OutputDateFormat.

Normal date in the format: dd mon yyyy. This is the default (dd cannot have any leading zeros or blanks; yyyy must have leading zeros but cannot have any leading blanks). If Normal is specified for input_date_format, the input_date must have the month (mon) specified in English (for example, Jan, Feb, Mar, and so on).

Ordered date in the format: yy/mm/dd (suitable for sorting, and so forth; yy, mm, and dd must have leading zeros).

Standard date in the format: yyyyymmdd (suitable for sorting, and so forth; yyyy, mm, and dd must have leading zeros).

Usa date in the format: mm/dd/yy (mm, dd, and yy must have leading zeros).

Weekday the name for the day of the week.

50.7.2 Examples


```
class TestDate uses RexxDate

method main(args=String[]) static
  say date('b', '10 Mar 1962') -- 716308
  say date('w', '10 Mar 1962', 'n') -- Saturday
  say date('w', '716308', 'b') -- Saturday
  say date('s', '716308', 'b') -- 19620310
  say date('s', '716308', 'b', '/') -- 1962/03/10
  say date('s', '716308', 'b', '-') -- 1962-03-10
  say date('w', 7688, 'c') -- Sunday
  say date('c', '1 Feb 2021') -- 7703
  say date('j', '18 Jan 2021') -- 2021018
  say date('j', '10 Mar 1962') -- 1962069
```

50.8 The REXXTime class

The REXXTime class implements the Classic REXX Time() function.

TIME



The time() function can be called standalone when using a uses REXXTime option on the class statement. You can use the following options to obtain specific time formats. (Only the capitalized letter is needed; all characters following it are ignored.)

50.8.1 Options

Civil returns the time in Civil format: hh:mmxx. The hours may take the values 1 through 12, and the minutes the values 00 through 59. The minutes are followed immediately by the letters am or pm. This distinguishes times in the morning (12 midnight through 11:59 a.m.—appearing as 12:00am through 11:59am) from noon and afternoon (12 noon through 11:59 p.m.—appearing as 12:00pm through 11:59pm). The hour has no leading zero. The minute field shows the current minute (rather than the nearest minute) for consistency with other TIME results.

Elapsed returns ssssssss.uuuuuu, the number of seconds.microseconds since the elapsed-time clock (described later) was started or reset. The number has no leading zeros or blanks, and the setting of NUMERIC DIGITS does not affect the number. The fractional part always has six digits.

Hours returns up to two characters giving the number of hours since midnight in the format: hh (no leading zeros or blanks, except for a result of 0).

Long returns time in the format: hh:mm:ss.uuuuuu (uuuuuu is the fraction of seconds, in microseconds). The first eight characters of the result follow the same rules as for the Normal form, and the fractional part is always six digits.

Minutes returns up to four characters giving the number of minutes since midnight in the format: mmmm (no leading zeros or blanks, except for a result of 0).

Normal returns the time in the default format hh:mm:ss, as described previously. The hours can have the values 00 through 23, and minutes and seconds, 00 through 59. All these are always two digits. Any fractions of seconds are ignored (times are never rounded up). This is the default.

Reset returns ssssssss.uuuuuu, the number of seconds.microseconds since the elapsed-time clock (described later) was started or reset and also resets the elapsed-time clock to zero. The number has no leading zeros or blanks, and the setting of NUMERIC DIGITS does not affect the number. The fractional part always has six digits.

Seconds returns up to five characters giving the number of seconds since midnight in the format: sssss (no leading zeros or blanks, except for a result of 0).

50.8.2 Examples

```
class TestTime uses RexxTime

method main(args=String[]) static
  say time()           -- 22:16:33
  say time('C')       -- 10:16pm
  say time('E')       -- 0.000000
  say time('R')       -- 0
  say time('H')       -- 22
  say time('L')       -- 22:16:33.836725
  say time('M')       -- 1336
```

```

say time('N')    -- 22:16:33
say time('O')    -- 22:16:33
say time('R')    -- 0.001204
say time('E')    -- 0.000271
say time('S')    -- 80193

```

50.9 The REXXOperators interface class

The **REXXOperators** interface class defines the signatures of the methods that implement the NetREXX (and REXX) operators. These methods are described in the section REXX arithmetic methods (see page 182)

In the future this interface may be used to allow the overloading of operators for objects of types other than **REXX**. The current NetREXX language definition does not permit operator overloading.

50.10 The REXXSet class

The **REXXSet** class is used to provide the numeric settings for the methods described in the section REXX (see page 182) arithmetic methods. When provided, a REXXSet Object supplies the numeric settings for the operation; when **null** is provided then the default settings are used (`digits=9`, `form=SCIENTIFIC`).

50.10.1 Public properties

These properties supply the numeric settings and certain values they may take. After construction, the `digits` and `form` values should only be changed by using the `setDigits` and `setForm` methods.

DEFAULT_DIGITS

A constant of type **int** that describes the default number of digits for a numeric operation (9).

DEFAULT_FORM

A constant of type **byte** that describes the default exponential format for a numeric operation (`SCIENTIFIC`).

digits

A value of type **int** that describes the numeric digits to be used for a numeric operation. The REXX arithmetic (see page 182) methods:ea. use this value to determine the significance of results. `digits` must always be greater than zero.

ENGINEERING

A constant of type **byte** that signifies that engineering exponential formatting should be used for a numeric operation.

form

A value of type **byte** that describes the exponential format to be used for a numeric operation. The REXX arithmetic (see page 182) methods:ea. use

this value to determine the formatting of results that require an exponent. form must be either ENGINEERING or SCIENTIFIC.

SCIENTIFIC

A constant of type **byte** that signifies that scientific exponential formatting should be used for a numeric operation.

50.10.2 Constructors

These constructors are used to set the initial values of a REXXSet object.

REXXSet()

Constructs a REXXSet object which has default digits and form properties.

REXXSet(newdigits=int)

Constructs a REXXSet object which has its digits property set to *newdigits* and its form property set to DEFAULT_DIGITS.

REXXSet(newdigits=int, newform=byte)

Constructs a REXXSet object which has its digits property set to *newdigits* and its form property set to *newform*.

REXXSet(arg=REXXSet)

Constructs a REXXSet object which is copy of *arg*, which is of type **netrex.lang.REXXSet**. *arg* must not be **null**.

50.10.3 Methods

The REXXSet class has the following additional methods:

formword()

Returns a string of type **netrex.lang.Rexx** that describes the form property. This will either be the string '**engineering**' or the string '**scientific**', corresponding to the form value ENGINEERING or SCIENTIFIC, respectively.

setDigits(newdigits=REXX)

Sets the digits value for the **RexxSet** object, from *newdigits*, after rounding and checking as defined for the numeric instruction; *newdigits* must be a positive whole number with no more than nine digits. No value is returned.

setForm(newformword=REXX)

Sets the form value for the **RexxSet** object, from *newformword*. This must equal either the string '**engineering**' or the string '**scientific**', corresponding to the form value ENGINEERING or SCIENTIFIC, respectively. No value is returned.

Appendix C - Translator Options

There are a number of options for the translator, some of which can be specified on the translator command line, and others also in the program source on the **option** statement. In the following table, c stands for *commandline only*, and b stands for *both source and commandline*. On the commandline, options are prefixed with a *dash* ("-"), while in programsource they are not - there they are preceded by the option statement.

TABLE 9: Options

Option	Meaning	Place
address	address instruction is allowed (default)	c
annotations	annotations are allowed (default)	c
arg words	interpret; remaining words are arguments	c
binary	classes are binary classes	b
classpath	specify a classpath	c
compile	compile (default; -nocompile implies -keep)	c
comments	copy comments across to generated .java	b
compact	display error messages in compact form	b
console	display messages on console (default)	c
crossref	generate cross-reference listing	b
decimal	allow implicit decimal arithmetic	b
diag	show diagnostic messages	b
ecj	prefer the ecj compiler	c
exec	interpret with no argument words	c
explicit	local variables must be explicitly declared	b
format	format output file (pretty-print)	b
java	generate Java source code for this program	b
javac	prefer the javac compiler	c
keep	keep any completed .java file (as xxx.java.keep)	c
keepasjava	keep any completed .java file (as xxx.java)	c
logo	display logo (banner) after starting	b
prompt	prompt for new request after processing	c
savelog	save messages in NetRexxC.log	c
replace	replace .java file even if it exists	b
sourcedir	force output files to source directory	b
strictargs	empty argument lists must be specified as ()	b

Continued on next page

Table 9 – continued from previous page

strictassign	assignment must be cost-free	b
strictcase	names must match in case	b
strictimport	all imports must be explicit	b
strictmethods	superclass methods are not compared to local methods for best match	b
strictprops	even local properties must be qualified	b
strictsignal	signals list must be explicit	b
symbols	include symbols table in generated .class files	b
time	display timings	c
trace[n]	trace stream [1 or 2], or 0 for NOTRACE	b
utf8	source file is in UTF8 encoding	b
verbose[n]	verbosity of progress reports [0-5]	b
warnexit0	exit with a zero returncode on warnings	c

Options valid for the options statement and on the commandline

These are the options that can be used on the **options** statement:

address This option defaults to address. When noaddress is specified, the use of the address instruction in a program will be flagged as an error.

binary All classes in this program will be binary classes. In binary classes, literals are assigned binary (primitive) or native string types, rather than NetREXX types, and native binary operations are used to implement operators where appropriate, as described in “Binary values and operations”. In classes that are not binary, terms in expressions are converted to the NetREXX string type, REXX, before use by operators.

comments Comments from the NetREXX source program will be passed through to the Java output file (which may be saved with a .java.keep or .java extension by using the -keep and -keepasjava command options, respectively).

compact Requests that warnings and error messages be displayed in compact form. This format is more easily parsed than the default format, and is intended for use by editing environments. Each error message is presented as a single line, prefixed with the error token identification enclosed in square brackets. The error token identification comprises three words, with one blank separating the words. The words are: the source file specification, the line number of the error token, the column in which it starts, and its length. For example (all on one line):

```
[D:\test\test.nrx 3 8 5] Error: The external name
'class' is a Java reserved word, so would not be
usable from Java programs
```

Any blanks in the file specification are replaced by a null ('\0') character. Additional words could be added to the error token identification later.

crossref Requests that cross-reference listings of variables be prepared, by class.

- decimal** Decimal arithmetic may be used in the program. If `nodecimal` is specified, the language processor will report operations that use (or, like normal string comparison, might use) decimal arithmetic as an error. This option is intended for performance-critical programs where the overhead of inadvertent use of decimal arithmetic is unacceptable.
- diag** Requests that diagnostic information (for experimental use only) be displayed. The `diag` option word may also have side-effects.
- explicit** Requires that all local variables must be explicitly declared (by assigning them a type but no value) before assigning any value to them. This option is intended to permit the enforcement of “house styles” (but note that the NetRexx compiler always checks for variables which are referenced before their first assignment, and warns of variables which are set but not used).
- format** Requests that the translator output file (Java source code) be formatted for improved readability. Note that if this option is in effect, line numbers from the input file will not be preserved (so run-time errors and exception trace-backs may show incorrect line numbers).
- java** Requests that Java source code be produced by the translator. If `nojava` is specified, no Java source code will be produced; this can be used to save a little time when checking of a program is required without any compilation or Java code resulting.
- logo** Requests that the language processor display an introductory logotype sequence (name and version of the compiler or interpreter, etc.).
- sourcedir** Requests that all `.class` files be placed in the same directory as the source file from which they are compiled. Other output files are already placed in that directory. Note that using this option will prevent the `-run` command option from working unless the source directory is the current directory.
- strictargs** Requires that method invocations always specify parentheses, even when no arguments are supplied. Also, if `strictargs` is in effect, method arguments are checked for usage – a warning is given if no reference to the argument is made in the method.
- strictassign** Requires that only exact type matches be allowed in assignments (this is stronger than Java requirements). This also applies to the matching of arguments in method calls.
- strictcase** Requires that local and external name comparisons for variables, properties, methods, classes, and special words match in case (that is, names must be identical to match).
- strictimport** Requires that all imported packages and classes be imported explicitly using `import` instructions. That is, if in effect, there will be no automatic imports, except those related to the `package` instruction.
- strictmethods** Superclass methods are not compared to local methods for best match.
- strictprops** Requires that all properties, including those local to the current class, be qualified in references. That is, if in effect, local properties cannot appear as simple names but must be qualified by `this`. (or equivalent) or the class name (for static properties).

strictsignal Requires that all checked exceptions signalled within a method but not caught by a catch clause be listed in the signals phrase of the method instruction.

symbols Symbol table information (names of local variables, etc.) will be included in any generated .class file. This option is provided to aid the production of classes that are easy to analyse with tools that can understand the symbol table information. The use of this option increases the size of .class files.

trace, traceX If given as **-trace**, **-trace1**, or **-trace2**, then trace instructions are accepted. The trace output is directed according to the option word: **-trace1** requests that trace output is written to the standard output stream, **-trace** or **-trace2** imply that the output should be written to the standard error stream (the default).

utf8 If given, clauses following the options instruction are expected to be encoded using UTF-8, so all Unicode characters may be used in the source of the program. In UTF-8 encoding, Unicode characters less than '\u0080' are represented using one byte (whose most-significant bit is 0), characters in the range '\u0080' through '\u07FF' are encoded as two bytes, in the sequence of bits:

110xxxxx 10xxxxxx

where the eleven digits shown as x are the least significant eleven bits of the character, and characters in the range '\u0800' through '\uFFFF' are encoded as three bytes, in the sequence of bits:

1110xxxx 10xxxxxx 10xxxxxx

where the sixteen digits shown as x are the sixteen bits of the character. If **noutf8** is given, following clauses are assumed to comprise only Unicode characters in the range '\x00' through '\xFF', with the more significant byte of the encoding of each character being 0. Note: this option only has an effect as a compiler option, and applies to all programs being compiled. If present on an options instruction, it is checked and must match the compiler option (this allows processing with or without utf8 to be enforced).

verbose, verboseX Sets the “noisiness” of the language processor. The digit X may be any of the digits 0 through 5; if omitted, a value of 3 is used. The options **-noverbose** and **verbose0** both suppress all messages except errors and warnings

Options valid on the commandline

The translator also implements some additional option words, which control compilation features. These cannot be used on the **options** instruction⁸³, and are:

⁸³Although at the moment, there will be no indication of this

- arg** The **-arg** *words* option is used when interpreting programs, it indicates that after the **-arg** statement, commandline arguments for the interpreted program follow
- classpath** The **-classpath** option allows dynamic specification of the classpath used by the NetREXXC compiler without having to depend on the CLASSPATH environment variable. (since: NetREXX 3.02) .
- exec** The **-exec** *words* option is used when interpreting programs. With this option, no commandline arguments are possible.
- ecj** prefer the ecj compiler when available
- keep** keep the intermediate *.java* file for each program. It is kept in the same directory as the NetREXX source file as *xxx.java.keep*, where *xxx* is the source file name. The file will also be kept automatically if the *javac* compilation fails for any reason.
- javac** prefer the javac compiler when available
- keepasjava** keep the intermediate *.java* file for each program. It is kept in the same directory as the NetREXX source file as *xxx.java*, where *xxx* is the source file name. Implies **-replace**. Note: use this option carefully in mixed-source projects where you might have *.java* source files around.
- nocompile** do not compile (just translate). Use this option when you want to use a different Java compiler. The *.java* file for each program is kept in the same directory as the NetREXX source file, as the file *xxx.java.keep* (where *xxx* is the source file name).
- noconsole** do not display compiler messages on the console (command display screen). This is usually used with the *savelog* option.
- savelog** write compiler messages to the file *NetREXXC.log*, in the current directory. This is often used with the *noconsole* option.
- time** display translation, *javac* or *ecj* compile, and total times (for the sum of all programs processed).
- run** run the resulting Java class as a stand-alone application, provided that the compilation had no errors.
- warnexit0** Exit the translator with returncode 0 even if warnings are issued. Useful with build tools that would otherwise exit a build.

Index

- * multiplication operator,, 42, 140
- *-* tracing flag,, 112
- *=* tracing flag,, 112
- + plus sign,addition operator, 42, 139
- + plus sign,in parsing template, 133
- ++ invalid sequence,, 20
- +++ tracing flag,, 112
- ,, 16, 17, 23, 32, 36, 38, 42, 44, 49, 62, 65, 69, 73-75, 81, 82, 84, 103, 104, 107-109, 117, 118, 120, 131, 135, 139, 145, 163, 172, 177
- continuation character,, 20
- minus sign,in parsing template, 133
- minus sign,subtraction operator, 42, 139
- . (period),as placeholder in parsing, 133
- . (period),in numbers, 138
- . (period),in terms, 26
- = equals sign,assignment indicator, 49
- = equals sign,equal operator, 43
- = equals sign,in LOOP instruction, 75
- = equals sign,in parsing template, 133
- >> tracing flag,, 112
- >a> tracing flag,, 112
- >p> tracing flag,, 112
- >v> tracing flag,, 112
- \$ dollar sign,in symbols, 18
- & and operator,, 44
- && exclusive or operator,, 44
- _ underscore,in symbols, 18
- NetRexx,language definition, 12
- \backslash,escape character, 18
- \backslash,not operator, 44
- \= not equal operator,, 43
- \\invalid sequence,, 20
- Rexx, 35, 52
 - address, 60
 - arg, 55, 98, 150, 177
 - binary, 58
 - case, 106
 - catch, 150
 - class, 35, 37, 58, 115, 117-119, 121, 126, 153, 189, 190
 - constant, 101, 115
 - dependent, 118
 - digits, 143
 - do, 66, 67, 69, 150, 178
 - else, 51, 69, 112, 178
 - end, 66, 67, 69, 73, 74, 76-80, 88, 105, 106, 115, 150, 160, 177, 178
 - engineering, 144
 - exit, 68
 - extends, 35, 37, 115, 117-119, 126
 - finally, 150
 - for, 77, 78
 - forever, 76
 - form, 144
 - if, 51, 66, 67, 69, 70, 73, 74, 76, 80, 105, 112, 113, 115, 121, 178
 - implements, 126
 - import, 71, 72, 115
 - indirect, 126
 - iterate, 73
 - label, 67, 80, 105
 - leave, 67, 74, 76, 80, 105
 - loop, 57, 73, 74, 76-80, 115
 - method, 35, 37, 51, 58, 83, 86, 115, 118, 119, 123, 126-128, 178, 189, 190
 - nop, 88, 177, 178
 - numeric, 143, 144
 - options, 58
 - otherwise, 88, 105, 160, 177
 - over, 78
 - package, 96
 - parse, 97, 98, 113, 129, 130, 134, 135, 177, 178
 - private, 37, 58, 101
 - properties, 37, 58, 101, 115, 126
 - protect, 80
 - public, 115
 - return, 35, 58, 86, 115, 126, 127, 178
 - returns, 58, 86, 126, 127
 - say, 6, 31, 51, 53, 54, 56-61, 66, 67, 69, 70, 73, 74, 76-79, 88, 103, 105, 106, 112-115, 121-123, 143, 144, 150, 160, 177, 178, 189-191
 - scientific, 144
 - select, 88, 105, 106, 160, 177, 178
 - signal, 108
 - signals, 86
 - static, 58, 178, 189, 190
 - super, 37, 52, 119, 123
 - then, 51, 66, 67, 69, 70, 73, 74, 76, 80, 88, 105, 106, 112, 113, 115, 121, 160,

177, 178
 this, 35, 37, 52, 119, 128
 to, 73, 74, 77-80, 115
 trace, 111-113, 150
 until, 79
 upper, 134, 173
 when, 88, 105, 106, 160, 177, 178
 while, 57, 79, 80

ABBREV method,, 160
 Abbreviations,testing with ABBREV method,
 160
 ABS method,, 160
 Absolute,column specification in parsing,
 133
 Absolute,positional pattern, 133
 Absolute,value, finding using ABS method,
 160
 Abstract classes,, 63
 Abstract methods,, 63, 84
 ABSTRACT,on CLASS instruction, 63
 ABSTRACT,on METHOD instruction, 84
 Abuttal concatenation operator,, 42, 44
 Active constructs,, 73, 74
 Adapter classes,, 63
 ADAPTER,on CLASS instruction, 63
 Addition,definition, 139
 Addition;.pi ,Subtraction;.pi
 /Multiplication;.pi /Division, 42
 address option, 194
 Address,instruction, 60
 Algebraic precedence,, 46
 ALL,TRACE setting, 110
 Alphabets,checking with DATATYPE, 163
 Alphanumerics,checking with DATATYPE, 163
 AND,logical operator, 44
 Annotate,instruction, 58
 ANNOTATIONS, in OPTIONS instruction, 91
 ANSI standard,arithmetic definition, 137
 Arbitrary precision arithmetic,, 136
 arg words option, 196
 Arguments,of methods, 32
 Arguments,on METHOD instruction, 83
 Arguments,optional, 83
 Arguments,passing to methods, 32
 Arguments,provided by caller, 83
 Arguments,required, 83
 Arithmetic,comparisons, 142
 Arithmetic,errors, 145
 Arithmetic,exceptions, 145
 Arithmetic,implementation independence,
 145
 Arithmetic,NUMERIC settings, 89
 Arithmetic,operation rules, 139
 Arithmetic,operators, 42, 136, 138
 Arithmetic,overflow, 145
 Arithmetic,precision, 138
 Arithmetic,underflow, 145
 Array initializer,in terms, 26, 55
 Arrays,, 54
 Arrays,constructors, 54
 Arrays,in terms, 31
 Arrays,initializing, 55
 Arrays,references, 54
 ASCII,coded character set, 14
 Ask method,, 186
 ASK special word,, 121
 ASKNOECHO special word,, 121
 AskOne method,, 186
 Assignment,, 48, 49
 Assignment,binary, 147
 Assignment,instruction, 48, 49
 Assignment,of literals, 147
 Assignment,property initialization, 114

B2D method, 160
 B2X method, 161
 Backslash character,escape sequence, 18
 Backslash character,in strings, 18
 Backslash character,not operator, 44
 BadArgumentException,, 180
 BadColumnException,, 180
 BadNumericException,, 180
 Binary blocks,, 146
 Binary classes,, 64, 146
 Binary classes,assignment, 147
 Binary classes,binary methods, 85
 Binary classes,control variables, 147
 Binary classes,LOOP instruction, 147
 Binary classes,NUMERIC instruction, 148
 Binary constructors,, 148
 Binary literals,, 147
 Binary methods,, 85, 146
 Binary methods,assignment, 147
 Binary methods,control variables, 147
 Binary methods,LOOP instruction, 147
 Binary methods,NUMERIC instruction, 148
 Binary numbers,, 39, 146
 Binary numeric symbol,, 19, 22
 Binary operations,dyadic, 146
 Binary operations,monadic, 147
 Binary operations,prefix, 147
 binary option, 194
 Binary,arithmetic, 146
 Binary,checking with DATATYPE, 163
 Binary,conversion to decimal, 160
 Binary,conversion to hexadecimal, 161
 Binary,from decimal, 165
 BINARY,in OPTIONS instruction, 91
 BINARY,on CLASS instruction, 64
 BINARY,on METHOD instruction, 85
 Binary,operations, 146
 Binary,values, 146
 Bits,binary operators, 44
 Bits,checking with DATATYPE, 163
 Blank,, 16
 Blank,adjacent to operator character, 20
 Blank,adjacent to special character, 20
 Blank,as concatenation operator, 42
 Blank,as type conversion operator, 44

Blank,operator, 42, 44
 Blank,removal with SPACE method, 171
 Blank,removal with STRIP method, 171
 Block comments,, 17
 Body,of a loop, 76
 Body,of classes, 62
 Body,of group, 66
 Body,of methods, 82
 Body,of select, 104
 Boolean operations,, 44
 boolean type, value of,, 39
 Bottom of program, reaching during
 execution,, 68
 Bounded loop,, 76
 Bounded loop,controlled, 77
 Bounded loop,over values, 78
 Bounded loop,simple, 76
 Brackets,in array initializers, 26, 55
 Brackets,in array references, 54
 Brackets,in indexed references, 26
 Brackets,in indexed strings, 53
 Brackets,in terms, 26
 BY phrase of LOOP instruction,, 75

 C2D method,, 163
 C2X method,, 163
 Carriage return character,escape sequence,
 18
 Case,of names, 21
 CASE,on SELECT instruction, 106
 Casting,to type, 44
 CATCH,on DO instruction, 67
 CATCH,on LOOP instruction, 80
 CATCH,on SELECT instruction, 107
 CATCH,use of, 150
 Caught exceptions,, 149
 CENTER method,, 161
 CENTRE method,, 161
 CHANGESTR method,, 161
 Changing strings,using CHANGESTR, 161
 Changing strings,using TRANSLATE, 172
 char,as a string, 39
 Character sets,, 14
 Character,, 14
 Character,appearance, 14
 Character,conversion to decimal, 163
 Character,conversion to hexadecimal, 163
 Character,converting to binary, 148
 Character,encodings, 14, 148
 Character,from a number, 165, 175
 Character,from decimal, 165
 Character,from hexadecimal, 175
 Character,glyphs, 14
 Character,removal with STRIP method, 171
 charAt method,, 185
 Checked exceptions,, 151
 Class,, 23
 Class,body of, 62
 Class,definition, 114
 Class,filename of, 122
 Class,instances of, 36
 Class,name of, 62
 Class,names, case of, 21
 Class,package of, 96
 Class,qualified name of, 96
 Class,short name of, 62
 CLASS,special word, 121
 Class,starting, 62
 Classes,abstract, 63
 Classes,adapter, 63
 Classes,and subclasses, 65
 Classes,and superclasses, 65
 Classes,binary, 64
 Classes,dependent, 25, 118
 Classes,final, 64
 Classes,interface, 64
 Classes,minor, 24, 117
 Classes,parent, 24, 117
 Classes,private, 63
 Classes,public, 63
 Classes,shared, 63
 Classes,standard, 63
 classpath option, 197
 classpath option,, 91
 Clauses,, 16
 Clauses,continuation of, 20
 Clauses,null, 48
 Coded character set,ASCII, 14
 Coded character set,EBCDIC, 14
 Coded character set,Unicode, 14
 Coded character,, 14
 Coded character,conversion to decimal, 163
 Coded character,conversion to hexadecimal,
 163
 Coded character,from decimal, 165
 Coded character,from hexadecimal, 175
 Collating sequence, using SEQUENCE,, 170
 Column specification in parsing,, 133
 Comma,in array references, 54
 Comma,in indexed strings, 53
 Comma,in method calls, 32
 Command line options,, 94
 comments option, 194
 COMMENTS option,, 91
 Comments,, 16
 Comments,block, 17
 Comments,line, 16
 Comments,nesting, 17
 Comments,starting a program with, 17
 compact option, 194
 COMPACT option,, 92
 Comparative operators,, 43
 COMPARE method,, 162
 Comparison,of numbers, 43, 142
 Comparison,of strings and numbers, 43
 Comparison,of strings/using COMPARE, 162
 Compiler options,, 91
 Compound terms,, 27
 Concatenation,of strings, 42
 Concatenation,of types, 44

- Conditional loops,, 75
- Conditional phrase,, 76, 79
- CONSOLE option,, 92
- Console, writing to with SAY,, 103
- Constant methods,, 84
- CONSTANT,on METHOD instruction, 84
- CONSTANT,on PROPERTIES instruction, 100
- Constants,, 100
- Constants,used by classes, 65
- Constants,using properties, 100
- Constructor,Rexx(boolean), 181
- Constructor,Rexx(byte), 181
- Constructor,Rexx(char), 181
- Constructor,Rexx(char[]), 181
- Constructor,Rexx(double), 182
- Constructor,Rexx(float), 182
- Constructor,Rexx(int), 182
- Constructor,Rexx(long), 182
- Constructor,Rexx(Rexx), 182
- Constructor,Rexx(short), 182
- Constructor,Rexx(String), 182
- Constructor,Rexx(String[]), 182
- Constructor,RexxSet(), 192
- Constructor,RexxSet(int), 192
- Constructor,RexxSet(int,byte), 192
- Constructor,RexxSet(RexxSet), 192
- Constructors,, 36, 82
- Constructors,array, 54
- Constructors,binary, 148
- Constructors,default, 36
- Constructors,in minor classes, 117
- Constructors,method, 82
- Constructors,of dependent objects, 119
- Constructors,of minor classes, 118
- Constructors,qualified, 119
- Constructors,special, 124
- Constructs,active, 74
- Continuation,character, 20
- Continuation,of clauses, 20
- Control variable,, 77, 78
- Controlled loops,, 77
- Conversion,automatic, 38
- Conversion,binary constructors, 148
- Conversion,binary to decimal, 160
- Conversion,binary to hexadecimal, 161
- Conversion,character to decimal, 163
- Conversion,character to hexadecimal, 163
- Conversion,coded character to decimal, 163
- Conversion,coded character to hexadecimal, 163
- Conversion,cost of, 40
- Conversion,decimal to binary, 165
- Conversion,decimal to character, 165
- Conversion,decimal to hexadecimal, 165
- Conversion,explicit, 40
- Conversion,formatting numbers, 166
- Conversion,hexadecimal to binary, 174
- Conversion,hexadecimal to character, 175
- Conversion,hexadecimal to decimal, 175
- Conversion,of characters, 148
- Conversion,of types, 38
- Conversion,of well-known types, 38
- COPIES method,, 162
- COPYINDEXED method,, 162
- Copying a string using COPIES,, 162
- Copying indexed variables,, 162
- Counting,strings, using COUNTSTR, 162
- Counting,words, using WORDS, 174
- COUNTSTR method,, 162
- crossref option, 194
- CROSSREF option,, 92
- D2B method,, 165
- D2C method,, 165
- D2X method,, 165
- Data,conversions, 38
- Data,length of, 41, 168
- Data,terms, 26, 41
- Data,type checking, 41
- Data,types, 23
- DATATYPE method,, 163
- Datatypes,, 23, 38, 41
- decimal option, 194
- DECIMAL option,, 92
- Decimal,conversion to binary, 165
- Decimal,conversion to character, 165
- Decimal,conversion to hexadecimal, 165
- Declarations,of variables, 50
- DEFAULT_DIGITS property,, 191
- DEFAULT_FORM property,, 191
- Deleting,part of a string, 164
- Deleting,words from a string, 164
- Delimiters,for comments, 16
- Delimiters,for strings, 17
- DELSTR method,, 164
- DELWORD method,, 164
- Dependent classes,, 25
- Dependent classes,restrictions, 120
- Dependent object,, 118
- Dependent object,constructing, 119
- DEPENDENT,on CLASS instruction, 118
- DEPRECATED,on CLASS instruction, 64
- DEPRECATED,on METHOD instruction, 86
- DEPRECATED,on PROPERTIES instruction, 101
- diag option, 195
- DIAG option,, 92
- Diagrams, of syntax,, 13
- digits property,, 191
- Digits,checking with DATATYPE, 163
- DIGITS,effect on whole numbers, 144
- Digits,in numbers, 137
- DIGITS,on NUMERIC instruction, 89, 138
- DIGITS,rounding when numbers used, 144
- DIGITS,special word, 122
- Dimension,of arrays, 24
- Dimension,of types, 24
- Dimensioned types,, 24
- DivideException,, 181
- Division,definition, 140
- Division,integer, 136

DO group,, 66
 DO group,naming of, 66
 DO instruction,LABEL, 66
 Dollar sign,in symbols, 18
 Double-quote,escape sequence, 18
 Double-quote,string delimiter, 17
 Dummy instruction, NOP,, 88
 Duplicate methods,, 86
 Dyadic operators,, 41

E-notation,, 45, 144
 E-notation,definition, 143
 E-notation,in symbols, 19
 EBCDIC,coded character set, 14
 ECJ option, 92
 Empty reference, null,, 122
 Encodings, of characters,, 14
 Encodings,binary, 148
 Encodings,of characters, 14
 END clause,specifying control variable, 78
 End condition of a LOOP loop,, 77
 End-of-file character,, 16
 Engineering notation,, 89, 144
 ENGINEERING property,, 191
 ENGINEERING value for NUMERIC FORM,, 89
 EOF character,, 16
 Equality,of objects, 43
 Equality,testing of, 43
 equals method,, 185
 Errors during arithmetic,, 145
 Escape sequences in strings,, 18
 Euro character,, 18
 Euro character,in symbols, 18
 Evaluation,of expressions, 41
 Evaluation,of terms, 28
 Example>Hello World, 114
 Example,of constructors, 37
 Example,of exception handling, 150
 Example,of two classes, 115
 Example,program, 177
 Exception,BadArgumentException, 180
 Exception,BadColumnException, 180
 Exception,BadNumericException, 180
 Exception,DivideException, 181
 Exception,ExponentOverflowException, 181
 Exception,NoOtherwiseException, 181
 Exception,NotCharacterException, 181
 Exception,NotLogicException, 181
 Exception,NullPointerException, 181
 Exception,NumberFormatException, 181
 Exceptions,, 149
 Exceptions,after CATCH clause, 151
 Exceptions,after FINALLY clause, 151
 Exceptions,checked, 151
 Exceptions,during arithmetic, 145
 Exceptions,during conversions, 39
 Exceptions,listed on METHOD instruction,
 86
 Exceptions,raising, 108
 Exceptions,signalling, 108

Exceptions,throwing, 108
 Exclusive OR,logical operator, 44
 exec option, 197
 EXISTS method,, 166
 EXIT instruction,, 68
 explicit option, 195
 EXPLICIT option,, 92
 Exponential notation,, 45, 89, 136, 144
 Exponential notation,definition, 143
 Exponential notation,in symbols, 19
 Exponentiation,definition, 141
 ExponentOverflowException,, 181
 Expressions,evaluation, 41
 Expressions,examples, 46
 Expressions,results of, 41
 EXTENDS,on CLASS instruction, 65
 Extra digits,in numbers, 138
 Extra digits,in numeric symbols, 18, 19
 Extra digits,in symbols, 18
 Extra letters, in symbols,, 18
 Extracting,a sub-string, 171
 Extracting,words from a string, 172

False value,, 44
 File method,, 187
 Final classes,, 64
 Final methods,, 84
 FINAL,on CLASS instruction, 63
 FINAL,on METHOD instruction, 84
 FINALLY,on DO instruction, 67
 FINALLY,on LOOP instruction, 80
 FINALLY,on SELECT instruction, 107
 FINALLY,reached by LEAVE, 74
 FINALLY,use of, 150
 Finding a mismatch using COMPARE,, 162
 Finding a string in another string,, 168,
 170
 Fixed size, of arrays,, 54
 flag, address, 194
 flag, binary, 194
 flag, nocompile, 197
 flag, noconsole, 197
 flag, run, 197
 flag, savelog, 197
 flag, time, 197
 flag,arg words, 196
 flag,classpath, 197
 flag,comments, 194
 flag,compact, 194
 flag,crossref, 194
 flag,decimal, 194
 flag,diag, 195
 flag,exec, 197
 flag,explicit, 195
 flag,format, 195
 flag,java, 195
 flag,keep, 197
 flag,keepasjava, 197
 flag,logo, 195
 flag,sourcedir, 195

- flag,strictargs, 195
- flag,strictassign, 195
- flag,strictcase, 195
- flag,strictimport, 195
- flag,strictmethods, 195
- flag,strictprops, 195
- flag,strictsignal, 195
- flag,symbols, 196
- flag,trace, traceX, 196
- flag,utf8, 196
- flag,verbose, verboseX, 196
- flag,warnexit0, 197
- Floating-point numbers, binary,, 146
- Flow control,abnormal, with SIGNAL, 108
- Flow control,with DO construct, 66
- Flow control,with IF construct, 69
- Flow control,with LOOP construct, 75
- Flow control,with SELECT construct, 104
- FOR,phrase of LOOP instruction, 75
- FOR,repititor on LOOP instruction, 75
- forEachline method,, 187
- FOREVER,loops, 76
- FOREVER,repititor on LOOP instruction, 75
- Form feed character,, 16
- form property,, 191
- FORM,option of NUMERIC instruction, 89, 144
- FORM,special word, 122
- format option, 195
- FORMAT,method, 166
- FORMAT,option, 92
- Formatting,numbers for display, 166
- Formatting,numbers with TRUNC, 172
- Formatting,of output during tracing, 111
- Formatting,text centering, 161
- Formatting,text left justification, 168
- Formatting,text right justification, 170
- Formatting,text spacing, 171
- formword() method,, 192
- Full name,of classes, 117
- Fully-qualified name, of classes,, 96
- Functions,numeric arguments of, 144
- Functions,return from, 102
- Functions,used by classes, 65

- Glyphs,, 14
- Group, DO,, 66
- Guard digit in arithmetic,, 139

- hashCode method,, 185
- Hexadecimal numeric symbol,, 19, 21
- Hexadecimal,checking with DATATYPE, 163
- Hexadecimal,conversion to binary, 174
- Hexadecimal,conversion to character, 175
- Hexadecimal,conversion to decimal, 175
- Hexadecimal,digits in escapes, 18
- Hexadecimal,escape sequence, 18
- Hyphen,as continuation character, 20

- IF instruction,, 69
- IMPLEMENTS,on CLASS instruction, 65

- Implied semicolons,, 20
- IMPORT instruction,, 71
- Imports,automatic, 72
- Imports,explicit, 71
- Indefinite loops,, 75, 76
- Indentation during tracing,, 111
- Index strings,for sub-values, 53
- Index strings,testing for, 166
- Indexed references,arrays, 54
- Indexed references,in terms, 26
- Indexed references,indexed strings, 53
- Indexed strings,, 53
- Indexed strings,copying, 162
- Indexed strings,merging, 162
- Indexed strings,testing for, 166
- Indirect properties,, 125
- INDIRECT,on PROPERTIES instruction, 125
- Inequality, testing of,, 43
- Infinite loops,, 75
- INHERITABLE,on METHOD instruction, 83
- INHERITABLE,on PROPERTIES instruction, 100
- Initializing arrays,, 55
- INSERT method,, 168
- Inserting a string into another,, 168
- Instance, of a class,, 36
- Instructions,, 57
- Instructions,Address, 60
- Instructions,Annotate, 58
- Instructions,assignment, 48, 49
- Instructions,CLASS, 62
- Instructions,DO, 66
- Instructions,EXIT, 68
- Instructions,IF, 69
- Instructions,IMPORT, 71
- Instructions,ITERATE, 73
- Instructions,keyword, 48, 57
- Instructions,LEAVE, 74
- Instructions,LOOP, 75
- Instructions,METHOD, 82, 85
- Instructions,method call, 48
- Instructions,NOP, 88
- Instructions,NUMERIC, 89
- Instructions,OPTIONS, 91
- Instructions,PACKAGE, 96
- Instructions,PARSE, 97
- Instructions,PROPERTIES, 99, 125
- Instructions,RETURN, 102
- Instructions,SAY, 103
- Instructions,SELECT, 104
- Instructions,SIGNAL, 108
- Instructions,TRACE, 109
- Integer division,, 136
- Integer division,definition, 141
- Integers, binary,, 146
- Interface classes,, 64
- Interface classes,properties in, 101
- INTERFACE,on CLASS instruction, 63
- Interfaces,implemented by classes, 65
- Internal functions,return from, 102
- Interpreter options,, 91

ITERATE instruction,, 73
ITERATE instruction,use of variable on, 73

java option, 195
JAVA option,, 92
Java,in reference implementation, 12
JavaBean properties,, 125
JAVAC option, 92
javac option, 197

keep option, 197
KEEPASJAVA option, 92
keepasjava option, 197
Keyword instructions,, 48, 57
Keywords,, 48
Keywords,mixed case, 57

LABEL,on DO instruction, 66
LABEL,on LOOP instruction, 79
LABEL,on SELECT instruction, 105
Language processor options,, 91
LASTPOS method,, 168
Leading blanks,removal with STRIP method,
171
Leading zeros,adding with the RIGHT
method, 170
Leading zeros,removal with STRIP method,
171
LEAVE instruction,, 74
LEAVE instruction,use of variable on, 74
LEFT method,, 168
LENGTH,method, 168
Length,of arrays, 31
Length,of comments, 17
LENGTH,special word, 31, 122
Letters,checking with DATATYPE, 163
Line comments,, 16
Line ends, effect of,, 20
Line feed character,escape sequence, 18
Line numbers, in tracing,, 111
Line, displaying,, 103
lists,, 154
Literal patterns,, 131
Literal strings,, 17
Literal strings,in terms, 26
Literals, binary,, 147
Local variables,, 51
Locating,a string in another string, 168,
170
Locating,a word or phrase in a string, 174
Logical operations,, 44
logo option, 195
LOGO option,, 93
Loops,active, 73, 74
Loops,execution model, 80
Loops,in binary classes and methods, 147
Loops,label, 79
Loops,modification of, 73
Loops,naming of, 79
Loops,repertitive, 75, 76
Loops,termination of, 74

LOWER method,, 168
Lowercase,checking with DATATYPE, 163
Lowercase,names, 21
Lowercasing strings,, 168

Mantissa of exponential numbers,, 143
Matching methods,, 34
Mathematical method,ABS, 160
Mathematical method,DATATYPE options, 163
Mathematical method,FORMAT, 166
Mathematical method,MAX, 169
Mathematical method,MIN, 169
Mathematical method,SIGN, 170
MAX method,, 169
Merging indexed variables,, 162
Method call instructions,, 33, 48
METHOD instruction,, 85
Method,, 23
Method, built-in,ABBREV, 160
Method, built-in,ABS, 160
Method, built-in,B2D, 160
Method, built-in,B2X, 161
Method, built-in,C2D, 163
Method, built-in,C2X, 163
Method, built-in,CENTER, 161
Method, built-in,CENTRE, 161
Method, built-in,CHANGESTR, 161
Method, built-in,COMPARE, 162
Method, built-in,COPIES, 162
Method, built-in,COPYINDEXED, 162
Method, built-in,COUNTSTR, 162
Method, built-in,D2B, 165
Method, built-in,D2C, 165
Method, built-in,D2X, 165
Method, built-in,DATATYPE, 163
Method, built-in,DELSTR, 164
Method, built-in,DELWORD, 164
Method, built-in,EXISTS, 166
Method, built-in,FORMAT, 166
Method, built-in,INSERT, 168
Method, built-in,LASTPOS, 168
Method, built-in,LEFT, 168
Method, built-in,LENGTH, 168
Method, built-in,LOWER, 168
Method, built-in,MAX, 169
Method, built-in,MIN, 169
Method, built-in,OVERLAY, 169
Method, built-in,POS, 170
Method, built-in,REVERSE, 170
Method, built-in,RIGHT, 170
Method, built-in,SEQUENCE, 170
Method, built-in,SIGN, 170
Method, built-in,SOUNDEX, 171
Method, built-in,SPACE, 171
Method, built-in,STRIP, 171
Method, built-in,SUBSTR, 171
Method, built-in,SUBWORD, 172
Method, built-in,TRANSLATE, 172
Method, built-in,TRUNC, 172
Method, built-in,UPPER, 173

Method, built-in, VERIFY, 173
 Method, built-in, WORD, 174
 Method, built-in, WORDINDEX, 174
 Method, built-in, WORDLENGTH, 174
 Method, built-in, WORDPOS, 174
 Method, built-in, WORDS, 174
 Method, built-in, X2B, 174
 Method, built-in, X2C, 175
 Method, built-in, X2D, 175
 Method, argument variables, 51
 Method, Ask, 186
 Method, AskOne, 186
 Method, body of, 82
 Method, calls in terms, 26
 Method, charAt, 185
 Method, definition, 114
 Method, equals, 185
 Method, File, 187
 Method, forEachline, 187
 Method, formword(), 192
 Method, hashCode, 185
 Method, names, case of, 21
 Method, NotEqS, 184
 Method, OpAdd, 183
 Method, OpAnd, 183
 Method, OpCc, 183
 Method, OpCcbblank, 183
 Method, OpDiv, 183
 Method, OpDivI, 183
 Method, OpEq, 183
 Method, OpEqS, 183
 Method, OpGt, 183
 Method, OpGtEq, 184
 Method, OpGtEqS, 184
 Method, OpGtS, 184
 Method, OpLt, 184
 Method, OpLtEq, 184
 Method, OpLtEqS, 184
 Method, OpLtS, 184
 Method, OpMinus, 184
 Method, OpMult, 184
 Method, OpNot, 184
 Method, OpNotEq, 184
 Method, OpOr, 184
 Method, OpPlus, 184
 Method, OpPow, 184
 Method, OpRem, 185
 Method, OpSub, 185
 Method, OpXor, 185
 Method, popOutputStream, 187
 Method, pushOutputStream, 187
 Method, Say, 187
 Method, setDigits(Rexx), 192
 Method, setForm(Rexx), 192
 Method, short name of, 82
 Method, starting, 82
 Method, toboolean, 185
 Method, tobyte, 185
 Method, tobytearray, 185
 Method, tochar, 185
 Method, todouble, 186
 Method, tofloat, 186
 Method, toint, 186
 Method, tolong, 186
 Method, toRexx, 186
 Method, toshort, 186
 Method, toString, 186
 Methods,, 32
 Methods, abstract, 63, 84
 Methods, arguments of, 83
 Methods, binary, 85
 Methods, constant, 84
 Methods, constructor, 36, 82
 Methods, duplicate, 86
 Methods, final, 84
 Methods, inheritable, 83
 Methods, invocation of, 32
 Methods, native, 84
 Methods, overloading, 86
 Methods, overriding, 35
 Methods, private, 83
 Methods, protected, 85
 Methods, public, 83
 Methods, resolution of, 33
 Methods, return values, 86
 Methods, searching for, 34
 Methods, shared, 83
 Methods, special, 124
 Methods, standard, 84
 Methods, static, 84
 METHODS, TRACE setting, 110
 MIN method,, 169
 Minor classes,, 24
 Minor classes, constructing, 118
 Minor classes, naming of, 117
 Minor classes, nesting of, 117
 Minor classes, restrictions, 120
 Mixed case, checking with DATATYPE, 163
 Mixed case, names, 21
 MOD option,, 93
 Model, of loop execution, 80
 Monadic (prefix) operators,, 41
 Moving characters, with TRANSLATE method,,
 172
 Multiplication, definition, 140
 Names, special, class, 121
 Names, special, sourceline, 123
 Names, case of, 21
 Names, of variables, 49
 Names, on ITERATE instructions, 73
 Names, on LEAVE instructions, 74
 Names, special/ask, 121
 Names, special/asknoecho, 121
 Names, special/digits, 122
 Names, special/form, 122
 Names, special/length, 122
 Names, special/null, 122
 Names, special/RC, 122
 Names, special/source, 122

Names,special/super, 123
 Names,special/this, 123
 Names,special/trace, 123
 Names,special/version, 123
 Native methods,, 84
 NATIVE,on METHOD instruction, 84
 Negation,of logical values, 44
 Negation,of numbers, 42
 Nesting of comments,, 17
 netrexx.lang package,, 180
 netrexx.lang,Exceptions, 180
 netrexx.lang,Rexx arithmetic methods, 182
 netrexx.lang,Rexx class, 181
 netrexx.lang,Rexx constructors, 181
 netrexx.lang,Rexx miscellaneous methods,
 185
 netrexx.lang,RexxDate class, 187
 netrexx.lang,RexxIO class, 186
 netrexx.lang,RexxOperators class, 191
 netrexx.lang,RexxSet class, 191
 netrexx.lang,RexxSet constructors, 192
 netrexx.lang,RexxSet methods, 192
 netrexx.lang,RexxSet properties, 191
 netrexx.lang,RexxTime class, 189
 Newline character,escape sequence, 18
 NOANNOTATIONS option,, 91
 NOBINARY option,, 91
 NOCOMMENTS option,, 91
 NOCOMPACT option,, 92
 nocompile option, 197
 noconsole option, 197
 NOCONSOLE option,, 92
 NOCROSSREF option,, 92
 NODECIMAL option,, 92
 NODIAG option,, 92
 NOEXPLICIT option,, 92
 NOFORMAT option,, 92
 NOJAVA option,, 92
 NOLOGO option,, 93
 NOMOD option,, 93
 NoOtherwiseException,, 181
 NOP instruction,, 88
 NOREPLACE option,, 93
 Normal comparative operators,, 43
 Normalizing a string by its sound,SOUNDEX,
 171
 NOSAVELOG option,, 93
 NOSOURCEDIR option,, 93
 NOSTRICTARGS option,, 93
 NOSTRICTASSIGN option,, 93
 NOSTRICTCASE option,, 93
 NOSTRICTIMPORT option,, 93
 NOSTRICTPROPS option,, 93
 NOSTRICTSIGNAL option,, 93
 NOSYMBOLS option,, 93
 NOT operator,, 44
 Notation,engineering, 89, 144
 Notation,scientific, 89, 144
 Notations,in text, 13
 Notations,syntax, 13
 NotCharacterException,, 181
 NotEqS method,, 184
 NotLogicException,, 181
 NOTRACE option,, 93
 NOUTF8 option,, 94
 NOVERBOSE option,, 94
 Null character,escape sequence, 18
 Null clauses,, 48
 Null instruction, NOP,, 88
 NULL special word,, 122
 Null strings,, 17
 NullPointerException,, 181
 NumberFormatException,, 181
 Numbers,, 45, 136
 Numbers,arithmetic on, 42, 136, 138
 Numbers,as symbols, 19
 Numbers,checking with DATATYPE, 163
 Numbers,comparison of, 43, 142
 Numbers,conversion to character, 165, 175
 Numbers,conversion to hexadecimal, 165
 Numbers,definition, 137, 143
 Numbers,examples of, 45
 Numbers,formatting for display, 166
 Numbers,in LOOP instruction, 75
 Numbers,rounding, 166
 Numbers,truncating, 172
 Numbers,use of by NetRexx, 144
 Numeric symbols,, 19, 26
 Numeric symbols,binary, 22
 Numeric symbols,hexadecimal, 21
 NUMERIC,DIGITS, 138
 NUMERIC,FORM, 144
 NUMERIC,in binary classes and methods, 148
 NUMERIC,instruction, 89
 Numeric,part of a number, 137, 143
 Objects,comparing, 43
 Objects,constructing, 36
 Objects,equality, 43
 OFF,TRACE setting, 110
 OpAdd method,, 183
 OpAnd method,, 183
 OpCc method,, 183
 OpCcblank method,, 183
 OpDiv method,, 183
 OpDivI method,, 183
 OpEq method,, 183
 OpEqS method,, 183
 Operators,, 41
 Operators,arithmetic, 42, 136, 138
 Operators,blank, 42, 44
 Operators,characters used for, 20
 Operators,comparative, 43, 142
 Operators,composition of, 41
 Operators,concatenation, 42
 Operators,logical, 44
 Operators,precedence (priorities) of, 46
 Operators,type, 44
 OpGt method,, 183
 OpGtEq method,, 184

- OpGtEqS method,, 184
- OpGtS method,, 184
- OpLt method,, 184
- OpLtEq method,, 184
- OpLtEqS method,, 184
- OpLtS method,, 184
- OpMinus method,, 184
- OpMult method,, 184
- OpNot method,, 184
- OpNotEq method,, 184
- OpOr method,, 184
- OpPlus method,, 184
- OpPow method,, 184
- OpRem method,, 185
- OpSub method,, 185
- Option words,, 91
- option, address, 194
- option, binary, 194
- option, nocompile, 197
- option, noconsole, 197
- option, run, 197
- option, savelog, 197
- option, time, 197
- option,arg words, 196
- option,classpath, 197
- option,comments, 194
- option,compact, 194
- option,crossref, 194
- option,decimal, 194
- option,diag, 195
- option,exec, 197
- option,explicit, 195
- option,format, 195
- option,java, 195
- option,keep, 197
- option,keepasjava, 197
- option,logo, 195
- option,sourcedir, 195
- option,strictargs, 195
- option,strictassign, 195
- option,strictcase, 195
- option,strictimport, 195
- option,strictmethods, 195
- option,strictprops, 195
- option,strictsignal, 195
- option,symbols, 196
- option,trace, traceX, 196
- option,utf8, 196
- option,verbose, verboseX, 196
- option,warnexit0, 197
- Optional arguments,, 83
- OPTIONS,instruction, 91
- Options,on command line, 94
- OpXor method,, 185
- OR,logical exclusive, 44
- OR,logical inclusive, 44
- Over loops,, 78
- OVER repetitor on LOOP instruction,, 75
- Overflow, arithmetic,, 145
- OVERLAY method,, 169
- Overlaying a string onto another,, 169
- Overloaded methods,, 86
- Overriding methods,, 35
- PACKAGE instruction,, 96
- Package,, 23, 96
- Package,name of, 71, 96
- Package,netrexx.lang, 180
- Packing a string,with B2D, 160
- Packing a string,with B2X, 161
- Packing a string,with X2C, 175
- Parent class,, 117
- Parent object,, 118
- Parent,of dependent object, 119
- PARENT,special word, 119
- Parentheses,adjacent to blanks, 20
- Parentheses,in expressions, 41, 46
- Parentheses,in method calls, 26, 32
- Parentheses,in parsing templates, 135
- Parentheses,in terms, 26
- Parentheses,omitting from method calls, 26, 27
- PARSE,instruction, 97
- PARSE,parsing rules, 129
- Parsing templates,in PARSE instruction, 97
- Parsing,absolute columns, 133
- Parsing,definition, 130
- Parsing,general rules, 129, 131
- Parsing,introduction, 129
- Parsing,literal patterns, 131
- Parsing,positional patterns, 133
- Parsing,selecting words, 132
- Parsing,variable patterns, 135
- Period,as placeholder in parsing, 133
- Period,in numbers, 138
- Period,in terms, 26
- popOutputStream method,, 187
- POS position method,, 170
- Positional patterns,, 133
- Power operator,, 42
- Power operator,definition, 141
- Powers of ten in numbers,, 45, 143
- Precedence of operators,, 46
- Precision,arbitrary, 136
- Precision,of arithmetic, 138
- Prefix operators,, 41
- Prefix operators,+, 42
- Prefix operators,+/with types, 45
- Prefix operators,-, 42
- Prefix operators,-/with types, 45
- Prefix operators,\, 44
- Prefix operators,\/with types, 45
- Prefix operators,arithmetic, 139
- Primitive types,, 23, 146
- Primitive types,conversions, 38
- Priorities of operators,, 46
- PRIVATE,on CLASS instruction, 63
- PRIVATE,on METHOD instruction, 83
- PRIVATE,on PROPERTIES instruction, 100
- Program,filename of, 122

Program,prolog, 114
 Program,structure, 114
 Programmer's model of LOOP,, 80
 Programs,, 114
 Programs,examples, 177
 Programs,structure, 114
 Prolog, of a program,, 114
 PROPERTIES instruction,, 99, 125
 Properties,, 23, 50, 99
 Properties,case of names, 21
 Properties,constant, 100
 Properties,deprecated, 101
 Properties,for JavaBeans, 125
 Properties,in dependent classes, 120
 Properties,in interface classes, 101
 Properties,in minor classes, 120
 Properties,indirect, 125
 Properties,inheritable, 100
 Properties,initialization, 114
 Properties,modifiers, 100
 Properties,naming, 99
 Properties,private, 100
 Properties,public, 100
 Properties,shared, 100
 Properties,static, 100
 Properties,transient, 100
 Properties,unused, 101
 Properties,visibility, 100
 Properties,volatile, 100
 Property,DEFAULT_DIGITS, 191
 Property,DEFAULT_FORM, 191
 Property,digits, 191
 Property,ENGINEERING, 191
 Property,form, 191
 Property,SCIENTIFIC, 192
 PROTECT,on DO instruction, 67
 PROTECT,on LOOP instruction, 80
 PROTECT,on METHOD instruction, 85
 PROTECT,on SELECT instruction, 105
 Protected methods,, 85
 PUBLIC,on CLASS instruction, 63
 PUBLIC,on METHOD instruction, 83
 PUBLIC,on PROPERTIES instruction, 100
 Pure numbers,, 143
 pushOutputStream method,, 187

 qtime example program,, 177
 Qualified name, of classes,, 96
 Qualified types,, 23
 Quotes in strings,, 17

 REXX(boolean) constructor,, 181
 REXX(byte) constructor,, 181
 REXX(char) constructor,, 181
 REXX(char[]) constructor,, 181
 REXX(double) constructor,, 182
 REXX(float) constructor,, 182
 REXX(int) constructor,, 182
 REXX(long) constructor,, 182
 REXX(REXX) constructor,, 182
 REXX(short) constructor,, 182
 REXX(String) constructor,, 182
 REXX(String[]) constructor,, 182
 REXX,arithmetic, 136
 REXX,class/NetREXX strings, 23
 REXX,class/conversions, 39
 REXX,class/methods of, 159
 REXX,class/use by PARSE, 97
 REXXSet() constructor,, 192
 REXXSet(int) constructor,, 192
 REXXSet(int,byte) constructor,, 192
 REXXSet(REXXSet) constructor,, 192
 Raising exceptions,, 108
 RC special word,, 122
 Re-ordering characters,with TRANSLATE
 method, 172
 Real numbers, binary,, 146
 Reference implementation,, 12
 References,in terms, 26
 References,null, 122
 References,to arrays, 54
 References,to current object, 123
 References,to indexed strings, 53
 References,to methods, 32
 Relative column specification in parsing,,
 133
 Relative positional pattern,, 133
 Remainder operator,, 136
 Remainder operator,definition, 141
 Remainder operator;.pi ,Integer
 division;.pi /Exponentiation, 42
 Repeating a string with COPIES,, 162
 Repetitive loops,, 76
 Repetitor phrase,, 76
 REPLACE option,, 93
 Replacing strings,using CHANGESTR, 161
 Replacing strings,using TRANSLATE, 172
 Required arguments,, 83
 Resolution of methods,, 33
 Results,of methods, 86
 Results,returned by RETURN, 102
 Results,size of, 41
 RESULTS,TRACE setting, 110
 Return character,escape sequence, 18
 Return code, setting on exit,, 68
 RETURN instruction,, 102
 Return string, setting on exit,, 68
 Return Types, 36
 RETURNS,on METHOD instruction, 86
 REVERSE method,, 170
 RIGHT method,, 170
 Rounding,, 136
 Rounding,definition, 139
 Rounding,when numbers used, 144
 run option, 197
 Running off the end of a program,, 68

 saveLog option, 197
 SAVELOG option,, 93
 Say method,, 187

SAY,instruction, 103
 Scientific notation,, 89, 144
 SCIENTIFIC property,, 192
 SCIENTIFIC value for NUMERIC FORM,, 89
 Search order,for methods, 33
 Search order,for term evaluation, 28
 Searching a string for a word or phrase,,
 170, 174
 Select,label, 105
 Select,naming of, 105
 Semicolons,, 16
 Semicolons,can be omitted, 13
 Semicolons,implied, 20
 SEQUENCE method,, 170
 setDigits(Rexx) method,, 192
 setForm(Rexx) method,, 192
 setOutputStream method,, 187
 SHARED,on CLASS instruction, 63
 SHARED,on METHOD instruction, 83
 SHARED,on PROPERTIES instruction, 100
 Short name,of classes, 62, 117
 Short name,of methods, 82
 SIGN method,, 170
 SIGNAL instruction,, 108
 Signals,, 149
 SIGNALS,on METHOD instruction, 86
 Significant of exponential numbers,, 143
 Significant digits, in arithmetic,, 138
 Signs in parsing templates,, 133
 Simple D0 group,, 66
 Simple number,, 19
 Simple repetitor phrase,, 76
 Simple terms,, 26
 Single-quote,escape sequence, 18
 Single-quote,string delimiter, 17
 SOUNDEX method,, 171
 SOURCE special word,, 122
 sourcedir option, 195
 SOURCEDIR option,, 93
 SOURCELINE,special word, 123
 SPACE method,, 171
 Special characters,, 20
 Special characters,used for operators, 20
 Special methods,, 124
 Special methods,super, 119, 124
 Special methods,this, 124
 Special words,, 121
 Special words,ask, 121
 Special words,asknoecho, 121
 Special words,class, 121
 Special words,digits, 122
 Special words,form, 122
 Special words,length, 122
 Special words,null, 122
 Special words,parent, 119
 Special words,RC, 122
 Special words,source, 122
 Special words,sourceline, 123
 Special words,super, 123
 Special words,this, 119, 123
 Special words,trace, 123
 Special words,version, 123
 Square brackets,in array initializers, 26,
 55
 Square brackets,in indexed references, 26
 Standard classes,, 63
 Standard methods,, 84
 Static methods,, 84
 Static methods,used by classes, 65
 Static variable typing,, 51
 STATIC,on METHOD instruction, 84
 STATIC,on PROPERTIES instruction, 100
 stderr, used by TRACE,, 113
 stdin, reading with ASK,, 121
 stdin, reading with ASKNOECHO,, 121
 stdout, writing to with SAY,, 103
 Strict comparative operators,, 43
 strictargs option, 195
 STRICTARGS option,, 93
 strictassign option, 195
 STRICTASSIGN option,, 93
 strictcase option, 195
 STRICTCASE option,, 93
 strictimport option, 195
 STRICTIMPORT option,, 93
 strictmethods option, 195
 strictprops option, 195
 STRICTPROPS option,, 93
 strictsignal option, 195
 STRICTSIGNAL option,, 93
 Strings,, 17
 Strings,as literal constants, 17
 Strings,comparison of, 43
 Strings,concatenation of, 42
 Strings,escapes in, 18
 Strings,in terms, 26
 Strings,indexed, 53
 Strings,length of, 168
 Strings,lowercasing, 168
 Strings,moving with TRANSLATE method, 172
 Strings,null, 17
 Strings,quotes in, 17
 Strings,sub-values of, 53
 Strings,types of, 39
 Strings,uppercasing, 173
 Strings,verifying contents of, 173
 STRIP method,, 171
 Structured Lists, 154
 Stub, of term,, 27
 Sub-expressions, in terms,, 26
 Sub-keywords,, 57
 Sub-string, extracting,, 171
 Sub-values, of strings,, 53
 Subclass of a class,, 65
 Subroutines,calling, 33
 Subroutines,passing back values from, 102
 Subroutines,return from, 102
 Substitution,in expressions, 41
 SUBSTR method,, 171
 Subtraction,definition, 139

SUBWORD method,, 172
SUPER,special method, 119, 124
SUPER,special word, 123
Superclass of a class,, 65
Symbol characters,checking with DATATYPE,
163
symbols option, 196
SYMBOLS option,, 93
Symbols,, 18
Symbols,assigning values to, 49
Symbols,case of, 21
Symbols,in terms, 26
Symbols,numeric, 19, 26
Symbols,use of, 49
Symbols,valid names, 18
Syntax diagrams,notation for, 13
Syntax notation,, 13
System-dependent options,, 91

Tab character,, 16
Tab character,escape sequence, 18
Tabulation character,, 16
Templates, parsing,general rules, 129
Templates, parsing,in PARSE instruction,
97
Ten, powers of,, 143
Terminal, writing to with SAY,, 103
Terms,, 26, 41
Terms,compound, 27
Terms,evaluation of, 28
Terms,in assignments, 52
Terms,on left of =, 52
Terms,parsing of, 97
Terms,simple, 26
Terms,stub of, 27
Testing for indexed variables,, 166
THEN,following IF clause, 69
THEN,following WHEN clause, 104
THIS,special method, 124
THIS,special word, 119, 123
Thread,tracing, 112
Threads, 153
time option, 197
TO phrase of LOOP instruction,, 75
toboolean AskOne,, 186
toboolean method,, 185
tobyte method,, 185
tobytearray method,, 185
tochar method,, 185
todouble method,, 186
tofloat method,, 186
toint method,, 186
Tokens,, 17
tolong method,, 186
toRexx method,, 186
toshort method,, 186
toString method,, 186
Trace setting,, 110
Trace setting,altering with TRACE
instruction, 109

trace, traceX option, 196
Trace,context, 112
TRACE,instruction, 109
TRACE,option, 93
TRACE,special word, 123
Tracing,clauses, 110
Tracing,data identifiers, 112
Tracing,execution of programs, 109
Tracing,line numbers, 111
Tracing,variables, 110
Trailing blanks,removal with STRIP method,
171
Trailing zeros,, 139
TRANSIENT,on PROPERTIES instruction, 100
TRANSLATE method,, 172
Translation,with TRANSLATE method, 172
Trapping of exceptions,, 108
True value,, 44
TRUNC method,, 172
Truncating numbers,, 172
Types,, 23
Types,checking instances of, 44
Types,checking with DATATYPE, 163
Types,concatenation of, 44
Types,conversions, 38
Types,declaring, 50
Types,dimensioned, 24
Types,of terms, 41
Types,of values, 41
Types,operations on, 44
Types,primitive, 23, 146
Types,qualified, 23
Types,simplification, 38

Underflow, arithmetic,, 145
Underscore,in symbols, 18
Unicode,coded character set, 14
Unicode,escape sequence, 18
Unicode,UTF-8 encoding, 94
Unpacking a string,with C2X, 163
Unpacking a string,with X2B, 174
UNTIL phrase of LOOP instruction,, 75
UNUSED,on PROPERTIES instruction, 101
UPPER method,, 173
Uppercase,checking with DATATYPE, 163
Uppercase,names, 21
Uppercasing strings,, 173
USES,on CLASS instruction, 65
UTF-8 encoding,, 94
utf8 option, 196
UTF8 option,, 94

Variable reference,in parsing template,
135
Variables,, 49
Variables,controlling loops, 77
Variables,in parsing patterns, 135
Variables,indexed, 53
Variables,local, 51
Variables,method arguments, 51

- Variables,names of, 49
- Variables,parsing of, 97
- Variables,properties, 50
- Variables,scope of, 51
- Variables,setting new value, 49
- Variables,static typing of, 51
- Variables,subscripts, 53
- Variables,type of, 49
- Variables,valid names, 49
- Variables,visibility, 51
- VERBOSE option,, 94
- verbose, verboseX option, 196
- VERBOSEn option,, 94
- VERIFY method,, 173
- VERSION special word,, 123
- Visibility,of classes, 63
- Visibility,of methods, 83
- Visibility,of properties, 100
- VOLATILE,on PROPERTIES instruction, 100

- WARNEXIT0 option, 94
- warnexit0 option, 197
- Well-known conversions,, 38
- WHILE phrase of LOOP instruction,, 75
- White space,, 16
- Whole numbers,, 46
- Whole numbers,checking with DATATYPE, 163
- Whole numbers,definition, 144
- WORD method,, 174
- WORDINDEX method,, 174
- WORDLENGTH method,, 174
- WORDPOS method,, 174
- WORDS method,, 174
- Words, special,class, 121
- Words, special,sourceline, 123
- Words,counting, using WORDS, 174
- Words,deleting from a string, 164
- Words,extracting from a string, 172, 174
- Words,finding in a string, 174
- Words,finding length of, 174
- Words,in parsing, 132
- Words,locating in a string, 174
- Words,special/ask, 121
- Words,special/digits, 122
- Words,special/form, 122
- Words,special/length, 122
- Words,special/null, 122
- Words,special/RC, 122
- Words,special/source, 122
- Words,special/super, 123
- Words,special/this, 123
- Words,special/trace, 123
- Words,special/version, 123

- X2B method,, 174
- X2C method,, 175
- X2D method,, 175
- XOR, logical operator,, 44

- Zero character,escape sequence, 18
- Zeros,adding on the left, 170
- Zeros,padding, 170
- Zeros,removal with STRIP method, 171

ISBN 978-90-819090-1-3

