

NetRexx Programming Guide

RexxLA

Version 3.02 of March 23, 2013

**THE REXX LANGUAGE ASSOCIATION
NetRexx Programming Series
ISBN 978-90-819090-0-6**

Publication Data

©Copyright The Rexx Language Association, 2011-2013

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 978-90-819090-0-6

ISBN 978-90-819090-0-6



9 789081 909006 >

Contents

The NetRexx Programming Series	i
Typographical conventions	iii
Introduction	v
1 Meet the Rexx Family	1
1.1 Once upon a Virtual Machine	1
1.2 Once upon another Virtual Machine	1
1.3 Features of NetRexx	2
2 Learning to program	3
2.1 Console Based Programs	3
2.2 Comments in programs	5
2.3 Strings	5
2.4 Clauses	6
2.5 When does a Clause End?	6
2.6 Long Lines	6
2.7 Loops	7
2.8 Special Variables	9
3 NetRexx as a Scripting Language	11
4 NetRexx as an Interpreted Language	13
5 NetRexx as a Compiled Language	15
5.1 Compiling from another program	15
5.2 Compiling from memory strings	16
6 Calling non-JVM programs	17
7 Using NetRexx classes from Java	21
8 Classes	23
8.1 Classes	23
8.2 Dependent Classes	24
8.3 Properties	24
8.4 Methods	24
8.5 Inheritance	24
8.6 Overriding Methods	24
8.7 Overriding Properties	24

9	Using Packages	25
9.1	The package statement	25
9.2	Translator performance consequences	25
9.3	Some NetRexx package history	25
9.4	CLASSPATH	26
10	Programming Patterns	27
10.1	Events	27
10.2	Recursive Parse	27
10.3	Observer	27
11	Incorporating Class Libraries	29
11.1	The Collection Classes	29
12	Input and Output	31
12.1	The File Class	31
12.2	Streams	31
12.3	Line mode I/O	31
12.4	Byte Oriented I/O	31
12.5	Data Oriented I/O	31
12.6	Object Oriented I/O using Serialization	31
12.7	The NIO Approach	31
13	Algorithms in NetRexx	33
13.1	Factorial	33
13.2	Fibonacci	34
14	Using Parse	37
15	Using Trace	39
16	Concurrency	41
16.1	Threads	41
17	User Interfaces	43
17.1	AWT	43
17.2	Web Applets using AWT	43
17.3	Swing	47
17.4	Web Frameworks	47
18	Network Programming	49
18.1	Using Uniform Resource Locators (URL)	49
18.2	TCP/IP Socket I/O	49
18.3	RMI: Remote Method Interface	49
19	Database Connectivity with JDBC	51
20	WebSphere MQ	55
21	MQTT	61
21.1	Pub/Sub with MQ Telemetry	61

22	Component Based Programming: Beans	65
23	Using the NetRexxA API	67
23.1	The NetRexxA constructor	68
23.2	The parse method	68
23.3	The getClassObject method	69
23.4	The exiting method	69
23.5	Interpreting programs contained in memory strings	69
24	Interfacing to Open Object Rexx	73
24.1	BSF4ooRexx	73
25	NetRexx Tools	75
25.1	Editor support	75
25.2	Java to Nrx (java2nrx)	76
26	Using Eclipse for NetRexx Development	77
26.1	Downloading Eclipse	77
26.2	Setting up the workspace	77
26.3	Shellshock	78
26.4	Installing SVN	78
26.5	Downloading the NetRexx project from the SVN repository	78
26.6	Setting up the builds	79
26.7	Using the NetRexx version of the NetRexx Ant task	79
26.8	Setting up the Eclipse NetRexx Editor Plugin (Optional)	80
27	Platform dependent issues	81
27.1	Mobile Platforms	81
27.2	IBM Mainframe: Using NetRexx programs in z/OS batch	82
28	Translator inner workings	83
28.1	Method resolution	84
	Index	85

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 the formal definition for the language, documenting its syntax and semantics, and prescribing minimal functionality for language implementors. It is the definitive answer to any question on the language, and as such, is subject to approval of the NetRexx Architecture Review Board on any release of the language (including its NRL).
NJPipes Reference	The Data Flow oriented companion to NetRexx, with its CMS Pipes compatible syntax, is documented in this manual. It discusses installing and running Pipes for NetRexx, and has ample examples of defining your own stages in NetRexx.

Typographical conventions

In general, the following conventions have been observed in the NetRexx publications:


- Body text is in this font
- Examples of language statements are in a **bold** type
- Variables or strings as mentioned in source code, or things that appear on the console, are in a `typewriter` type
- Items that are introduced, or emphasized, are in an *italic* type
- Included program fragments are listed in this fashion:

Listing 1: Example Listing

```
1 — salute the reader
2 say 'hello reader'
```

- Syntax diagrams take the form of so-called *Railroad Diagrams* to convey structure, mandatory and optional items

Properties



Introduction

The *Programming Guide* is the book that has the broadest scope of the publications in the *NetRexx Programming Series*. Where the *Language Reference* and the *Quick Beginnings* need to be limited to a formal description and definition of the NetRexx language for the former, and a Quick Tour and Installation instructions for the latter, this book has no such limitations. It teaches programming, discusses computer language history and comparative linguistics, and shows many examples on how to make NetRexx work with diverse technologies as TCP/IP, Relational Database Management Systems, Messaging and Queuing (MQ[™]) systems, J2EE Containers as JBOSS[™] and IBM WebSphere Application Server[™], discusses various rich- and thin client Graphical User Interface Options, and discusses ways to use NetRexx on various operating platforms. For many people, the best way to learn is from examples instead of from specifications. For this reason this book is rich in example code, all of which is part of the NetRexx distribution, and tested and maintained. This has had its effect on the volume of this book, which means that unlike the other publications in the series, it is probably not a good idea to print it out in its entirety; its size will relegate it to being used electronically.

Terminology

The *NetRexx Language Reference (NRL)* is the source of the definitive truth about the language. In this *Programming Guide*, terminology is sometimes used more loosely than required for the more formal approach of the NRL. For example, there is a fine line distinguishing *statement*, *instruction* and *clause*, where the latter is a more Rexx-like concept that is not often mentioned in relation to other languages (if they are not COBOL or SQL). While we try not to be confusing, *clause* and *statement* will be interchangeably used, as are *instruction* and *keyword instruction*.

Acknowledgements

As this book is a compendium of decades of Rexx and NetRexx knowledge, it stands upon the shoulders of many of its predecessors, many of which are not available in print anymore in their original form, or will never be upgraded or actualized; we are indebted to many anonymous (because unacknowledged in the original publications) authors of IBM product documentation, and many others that we do know, and will thank in the following. If anyone knows of a name not mentioned here that should be, please be in touch.

A big IOU goes out to Alan Sampson, who singlehandedly contributed more than one hundred NetRexx programming examples. The Redbook authors (Peter Heuchert, Frederik Haesbrouck, Norio Furukawa, Ueli Wahli, Kris Buelens, Bengt Heijnesson, Dave Jones and Salvador Torres) have provided some important documents that has shown, in an early stage, how almost everything on the JVM is better and easier done in NetRexx. Kermit Kiser also provided examples and did maintenance on the translator. Bill Finlason provided the Eclipse instructions. If anyone feels their copyright is violated, please do let us know, so we can take out offending passages or paraphrase them beyond recognition. As the usage of all material in this publication is quoted for educational use, and consists of short fragments, a fair use clause will apply in most jurisdictions.

Meet the Rexx Family

1.1 Once upon a Virtual Machine

On the 22nd of March 1979, to be precise, Mike Cowlshaw of IBM had a vision of an easier to use command processor for VM, and wrote down a specification over the following days. VMTM (now called z/VM) is the original Virtual Machine operating system, stemming from an era in which time sharing was acknowledged to be the wave of the future and when systems as CTSS (on the IBM 704) and TSS (on the IBM 360 Family of computers) were early timesharing systems, that offered the user an illusion of having a large machine for their exclusive use, but fell short of virtualising the entire hardware. The CP/CMS system changed this; CP virtualised the hardware completely and CMS was the OS running on CP. CMS knew a succession of command interpreters, called EXEC, EXEC2 and RexxTM (originally REX - until it was found out, by the IBM legal department, that a product of another vendor had a similar name) - the EXEC roots are the explanation why some people refer to a NetRexx program as an “exec”. As a prime example of a *backronym*, Rexx stands for “Restructured Extended Executor”. It can be defended that Rexx came to be as a reaction on EXEC2, but it must be noted that both command interpreters shipped around the same time. From 1988 on Rexx was available on MVS/TSO and other systems, like DOS, Amiga and various Unix systems. Rexx was branded the official SAA procedures language and was implemented on all IBM’s Operating Systems; most people got to know Rexx on OS/2. In the late eighties the Object-Oriented successor of Rexx, Object Rexx, was designed by Simon Nash and his colleagues in the IBM Winchester laboratory. Rexx was thereafter known as Classic Rexx. Several open source versions of Classic Rexx were made over the years, of which Regina is a good example.

1.2 Once upon another Virtual Machine

In 1995 Mike Cowlshaw ported JavaTM to OS/2TM and soon after started with an experiment to run Rexx on the JVMTM. With Rexx generally considered the first of the general purpose scripting languages, NetRexxTM is the first alternative language for the JVM. The 0.50 release, from April 1996, contained the NetRexx runtime classes and a translator written in Rexx but tokenized and turned into an OS/2 executable. The 1.00 release came available in January 1997 and contained a translator bootstrapped to NetRexx. The Rexx string type that can also handle unlimited precision numerics is called Rexx in Java and NetRexx. Where Classic Rexx was positioned as a system *glue* lan-

guage and application macro language, NetRexx is seen as the one language that does it all, delivering system level programs or large applications.

Release 2.00 became available in August 2000 and was a major upgrade, in which interpreted execution was added. Until that release, NetRexx only knew *ahead of time* compilation (AOT).

Mike Cowlshaw left IBM in March 2010. IBM announced the transfer of NetRexx source code to the Rexx Language Association (RexxLA) on June 8, 2011, 14 years after the v1.0 release.

On June 8th, 2011, IBM released the NetRexx source code to RexxLA under the ICU open source license. RexxLA shortly after released this as NetRexx 3.00 and has followed with updates.

1.3 Features of NetRexx

Ease of use The NetRexx language is easy to read and write because many instructions are meaningful English words. Unlike some lower level programming languages that use abbreviations, NetRexx instructions are common words, such as **say**, **ask**, **if...then...else**, **do...end**, and **exit**.

Free format There are few rules about NetRexx format. You need not start an instruction in a particular column, you can also skip spaces in a line or skip entire lines, you can have an instruction span many lines or have multiple instructions on one line, variables do not need to be pre-defined, and you can type instructions in upper, lower, or mixed case.

Convenient built-in functions NetRexx supplies built-in functions that perform various processing, searching, and comparison operations for both text and numbers. Other built-in functions provide formatting capabilities and arithmetic calculations.

Easy to debug When a NetRexx exec contains an error, messages with meaningful explanations are displayed on the screen. In addition, the **trace** instruction provides a powerful debugging tool.

Interpreted The NetRexx language is an interpreted language. When a NetRexx exec runs, the language processor directly interprets each language statement, or translates the program in JVM bytecode.

Extensive parsing capabilities NetRexx includes extensive parsing capabilities for character manipulation. This parsing capability allows you to set up a pattern to separate characters, numbers, and mixed input.

Seamless use of JVM Class Libraries NetRexx can use any class, and class library for the JVM (written in Java or other JVM languages) in a seamless manner, that is, without the need for extra declarations or definitions in the source code.

Learning to program

2.1 Console Based Programs

One way that a computer can communicate with a user is to ask questions and then compute results based on the answers typed in. In other words, the user has a conversation with the computer. You can easily write a list of NetRexx instructions that will conduct a conversation. We call such a list of instructions a program. The following listing shows a sample NetRexx program. The sample program asks the user to give his name, and then responds to him by name. For instance, if the user types in the name Joe, the reply Hello Joe is displayed. Or else, if the user does not type anything in, the reply Hello stranger is displayed. First, we shall discuss how it works; then you can try it out for yourself.

Listing 2.1: Hello Stranger

```
1 /* A conversation */
2 say "Hello! What's your name?"
3 who=ask
4 if who = '' then say "Hello stranger"
5 else say "Hello" who
```

Briefly, the various pieces of the sample program are:

/* ... */ A comment explaining what the program is about. Where Rexx programs on several platforms must start with a comment, this is not a hard requirement for NetRexx anymore. Still, it is a good idea to start every program with a comment that explains what it does.

say An instruction to display Hello! What's your name? on the screen.

ask An instruction to read the response entered from the keyboard and put it into the computer's memory.

who The name given to the place in memory where the user's response is put.

if An instruction that asks a question.

who = " A test to determine if who is empty.

then A direction to execute the instruction that follows, if the tested condition is true.

say An instruction to display Hello stranger on the screen.

else An alternative direction to execute the instruction that follows, if the tested condition is not true. Note that in NetRexx, else needs to be on a separate line.

say An instruction to display Hello, followed by whatever is in who on the screen.

The text of your program should be stored on a disk that you have access to with the help of an *editor* program. On Windows, notepad or (notepad++), jEdit, X2 or SlickEdit

are suitable candidates. On Unix based systems, including MacOSX, vim or emacs are plausible editors. If you are on z/VM or z/OS, XEDIT or ISPF/PDF are a given. More about editing NetRexx code in chapter 25.1, *Editor Support*, on page 75.

When the text of the program is stored in a file, let's say we called it `hello.nrx`, and you installed NetRexx as indicated in the *NetRexx QuickStart Guide*, we can run it with

```
nrc -exec hello
```

and this will yield the result:

```
NetRexx portable processor, version NetRexx after3.01, build 1-20120406-1326
Copyright (c) RexxLA, 2011. All rights reserved.
Parts Copyright (c) IBM Corporation, 1995,2008.
Program hello.nrx
===== Exec: hello =====
Hello! What's your name?
```

If you do not want to see the version and copyright message every time, which would be understandable, then start the program with:

```
nrc -exec -nologo hello
```

This is what happened when Fred tried it.

```
Program hello.nrx
===== Exec: hello =====
Hello! What's your name?
Fred
Hello Fred
```

The **ask** instruction paused, waiting for a reply. Fred typed Fred on the command line and, when he pressed the ENTER key, the **ask** instruction put the word Fred into the place in the computer's memory called "who". The **if** instruction asked, is "who" equal to nothing:

```
who = ''
```

meaning, is the value of "who" (in this case, Fred) equal to nothing:

```
"Fred = ''
```

This was not true; so, the instruction after `then` was not executed; but the instruction after `else`, was.

But when Mike tried it, this happened:

```
Program hello.nrx
===== Exec: hello =====
Hello! What's your name?

Hello stranger
Processing of 'hello.nrx' complete
```

Mike did not understand that he had to type in his name. Perhaps the program should have made it clearer to him. Anyhow, he just pressed ENTER. The **ask** instruction put "" (nothing) into the place in the computer's memory called "who". The **if** instruction asked, is:

```
who = ''
```

meaning, is the value of “who” equal to nothing:

```
'' = ''
```

In this case, it was true. So, the instruction after **then** was executed; but the instruction after **else** was not.

2.2 Comments in programs

When you write a program, remember that you will almost certainly want to read it over later (before improving it, for example). Other readers of your program also need to know what the program is for, what kind of input it can handle, what kind of output it produces, and so on. You may also want to write remarks about individual instructions themselves. All these things, words that are to be read by humans but are not to be interpreted, are called comments. To indicate which things are comments, use:

```
/* to mark the start of a comment
*/ to mark the end of a comment.
```

The `/*` causes the translator to stop compiling and interpreting; this starts again only after a `*/` is found, which may be a few words or several lines later. For example,

```
/* This is a comment. */
say text /* This is on the same line as the instruction */
/* Comments may occupy more
than one line. */
```

NetRexx also has line mode comments - those turn a line at a time into a comment. They are composed of two dashes (hyphens, in listings sometimes fused to a typographical *em dash* - remember that in reality they are two *n dashes*).

```
-- this is a line comment
```

2.3 Strings

When the translator sees a quote (either `”` or `’`) it stops interpreting or compiling and just goes along looking for the matching quote. The string of characters inside the quotes is used just as it is. Examples of strings are:

```
’Hello’
”Final result: ”
```

If you want to use a quotation mark within a string you should use quotation marks of the other kind to delimit the whole string.

```
”Don’t panic”
’He said, ”Bother”’
```

There is another way. Within a string, a pair of quotes (of the same kind as was used to delimit the string) is interpreted as one of that kind.

```
’Don’’t panic’ (same as ”Don’t panic” )
”He said, ””Bother””” (same as ’He said, ”Bother”’)
```

2.4 Clauses

Your NetRexx program consists of a number of *clauses*. A clause can be:

1. A *keyword instruction* that tells the interpreter to do something; for example,

```
say "the word"
```

In this case, the interpreter will display the word on the user's screen.

2. An *assignment*; for example,

```
Message = 'Take care!'
```

3. A *null* clause, such as a completely blank line, or

```
;
```

4. A *method call instruction* which invokes a *method* from a *class*

```
'hiawatha'.left(2)
```

2.5 When does a Clause End?

It is sometimes useful to be able to write more than one clause on a line, or to extend a clause over many lines. The rules are:

- Usually, each clause occupies one line.
- If you want to put more than one clause on a line you must use a semicolon (;) to separate the clauses.
- If you want a clause to span more than one line you must put a dash (hyphen) at the end of the line to indicate that the clause continues on the next line. If a line does not end in a dash, a semicolon is implied.

What will you see on the screen when this exec is run?

Listing 2.2: RAH Exec

```
1 /* Example: there are six clauses in this program */ say "Everybody cheer!"
2 say "2"; say "4" ; say "6" ; say "8" ; say "Who do we" —
3 "appreciate?"
```

2.6 Long Lines

Since the days of the punch card images are over the lines in program sources have become longer and longer, and with NetRexx being a free format language, there is no real technical reason to limit line length. Still, for readability and for ease access to words within a line, it is often indicated to keep lines relatively short and tidy. For this reason, the *continuation character* '-' can be used. This also makes it possible to split long literal strings over lines.

Listing 2.3: Long lines

```
1 say 'good' —
2 'night'
```

This example will concatenate 'good' and 'night' with a space inbetween. When you want to avoid that, use the '||' concatenation operator.

Listing 2.4: Long lines with string concatenation without space

```
1 say 'good' -  
2 || 'night'
```

2.7 Loops

We can go on and write clause after clause in a program source files, but some repetitive actions in which only a small change occurs, are better handled by the **loop** statement, reminding of an anecdote that Andy Hertzfield tells¹ Imagine an assignment to neatly print out a table of exchange rates for dollars and euros for reference in a shop. We could of course make the following program:

Listing 2.5: Without a loop

```
1 say 1 'euro equals' 1 * 2.34 'dollars'  
2 say 2 'euro equals' 2 * 2.34 'dollars'  
3 say 3 'euro equals' 3 * 2.34 'dollars'  
4 say 4 'euro equals' 4 * 2.34 'dollars'  
5 say 5 'euro equals' 5 * 2.34 'dollars'  
6 say 6 'euro equals' 6 * 2.34 'dollars'  
7 say 7 'euro equals' 7 * 2.34 'dollars'  
8 say 8 'euro equals' 8 * 2.34 'dollars'  
9 say 9 'euro equals' 9 * 2.34 'dollars'  
10 say 10 'euro equals' 10 * 2.34 'dollars'
```

This is valid, but imagine the alarming thought that the list is deemed a success and you are tasked with making a new one, but now with values up to 100. That will be a lot of typing.

The way to do this is using the **loop**² statement.

Listing 2.6: With a loop

```
1 loop i=1 to 100  
2 say i 'euro equals' i * 2.34 'dollars'  
3 end
```

Now the *loop index variable* `i` varies from 1 to 100, and the statements between `loop` and `end` are repeated, giving the same list, but now from 1 to 100 dollars.

We can do more with the **loop** statement, it is extremely flexible. The following diagram is a (simplified, because here we left out the *catch* and *finally* options) rundown of the ways we can loop in a program.

A few examples of what we can do with this:

- Looping forever - better put, without deciding beforehand how many times

¹<http://www.folklore.org>: Bob's background looked to be a lot stronger in hardware than software, so we were somewhat skeptical about his software expertise, but he claimed to be equally adept at both. His latest project was a rebellious, skunk-works type effort to make a low cost version of the Star called "Cub" that used an ordinary Intel microprocessor (the 8086), which was heresy to the PARC orthodoxy, who felt that you needed custom, bit-slice processors to get sufficient performance for a Star-type machine. Bob had written much of the software for Cub himself.

"I've got lots of software experience", he declared, "in fact I've personally written over 350,000 lines of code."


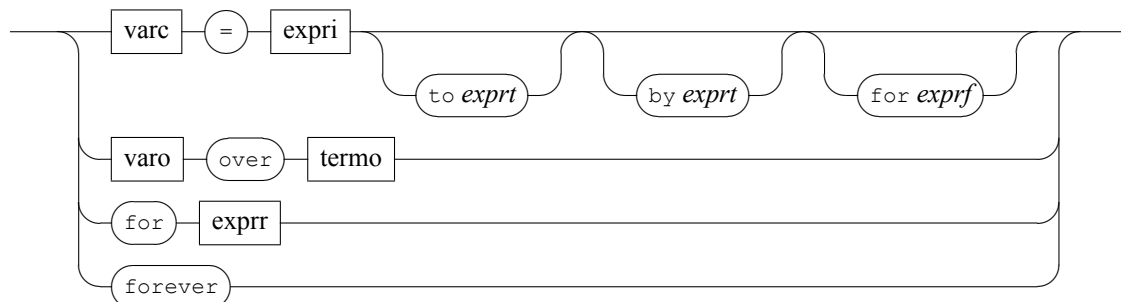

I thought that was pretty impressive, although I wondered how it was calculated. I couldn't begin to honestly estimate how much code I have written, since there are too many different ways to construe things.

That evening, I went out to dinner with my friend Rich Williams, who started at Apple around the same time that I did. Rich had a great sense of humor. I told him about the interview that I did in the afternoon, and how Bob Belleville claimed to have written over 350,000 lines of code.

"Well, I bet he did", said Rich, "but then he discovered loops!"

²Note that Classic Rexx uses **do** for this purpose. In recent Open Object Rexx versions **loop** can also be used.

FIGURE 1: Loop

loop*repetitor**conditional*

Listing 2.7: Loop Forever

```

1 loop forever
2   say 'another bonbon?'
3   x = ask
4   if x = 'enough already' then leave
5 end

```

The `leave` statement breaks the program out of the loop. This seems futile, but in the chapter about I/O we will see how useful this is when reading files, of which we generally do not know in advance how many lines we will read in the loop.

- Looping for a fixed number of times without needing a loop index variable

Listing 2.8: Loop for a fixed number of times without loop index variable

```

1 loop for 10
2   in.read() /* skip 10 lines from the input file */
3 end

```

- Looping back into the value of the loop index variable

Listing 2.9: Loop Forever

```

1 loop i = 100 to 90 by -2
2   say i
3 end

```

This yields the following output:

```

===== Exec: test =====
100
98
96
94
92
90
Processing of 'test.nrx' complete

```

2.8 Special Variables

We have seen that a *variable* is a place where some data, be it character data or numerical data, can be held. There are some special variables, as shown in the following program.

Listing 2.10: NetRexx Special Variables

```

1  /* NetRexx */
2  options replace format comments java crossref savelog symbols binary
3
4  class RCSpecialVariables
5
6  method RCSpecialVariables()
7    x = super.toString
8    y = this.toString
9    say '<super>'x'</super>'
10   say '<this>'y'</this>'
11   say '<class>'RCSpecialVariables.class'</class>'
12   say '<digits>'digits'</digits>'
13   say '<form>'form'</form>'
14   say '<[1, 2, 3].length>'
15   say [1, 2, 3].length
16   say '</[1, 2, 3].length>'
17   say '<null>'
18   say null
19   say '</null>'
20   say '<source>'source'</source>'
21   say '<sourceline>'sourceline'</sourceline>'
22   say '<trace>'trace'</trace>'
23   say '<version>'version'</version>'
24
25   say 'Type an answer:'
26   say '<ask>'ask'</ask>'
27
28   return
29
30 method main(args = String[]) public static
31
32   RCSpecialVariables()
33
34   return

```

this The special variables **this** and **super** refer to the current instance of the class and its superclass - what this means will be explained in detail in the chapter **Classes** on page 23, as is the case with the **class** variable.

digits The special variable **digits** shows the current setting for the number of decimal digits - the current setting of **numeric digits**. The related variable **form** returns the current setting of **numeric form** which is either `scientific` or `engineering`.

null The special variable **null** denotes the *empty reference*. It is there when a variable has no value.

source The **source** and **sourceline** variables are a good way to show the sourcefile and sourceline of a program, for example in an error message.

trace The **trace** variable returns the current trace setting, which can be one of the words `off var methods all results`.

version The **version** variable returns the version of the NetRexx translator that was in use at the time the clause we processed; in case of interpreted execution (see chapter 4 on 13, it returns the level of the current translator in use.

The result of executing this exec is as follows:

```
===== Exec: RCSpecialVariables =====
<super>RCSpecialVariables@4e99353f</super>
<this>RCSpecialVariables@4e99353f</this>
<class>class RCSpecialVariables</class>
<digits>9</digits>
<form>scientific</form>
<[1, 2, 3].length>
3
</[1, 2, 3].length>
<null>

</null>
<source>Java method RCSpecialVariables.nrx</source>
<sourceline>21</sourceline>
<trace>off</trace>
<version>NetRexx 3.02 27 Oct 2011</version>
Type an answer:
hello fifi
<ask>hello fifi</ask>
```

It might be useful to note here that these special variables are not fixed in the sense of that they are not *Reserved Variables*. NetRexx does not have reserved variables and any of these special variables can be used as an ordinary variable. However, when it is used as an ordinary variable, there is no way to retrieve the special behavior.

NetRexx as a Scripting Language

The term *scripting* is used here in the sense of using the programming language for quickly composed programs that interact with some application or environment to perform a number of simple tasks.

You can use NetRexx as a simple scripting language without having knowledge of, or using any of the features that is needed in a Java program that runs on the JVM - like defining a class name, and having a `main` method that is static and expects an array of `String` as its input.

Scripts can be written very fast. There is no overhead, such as defining a class, constructors and methods, and the programs contain only the necessary instructions. In this sense, a NetRexx script looks like an oo-version of a classic script, as the ceremonial aspects of defining class and method can be skipped. These will be automatically generated in the Java language source that is being generated for a script.

The scripting feature can be used for test purposes. It is an easy and convenient way of entering some statements and testing them. The scripting feature can also be used for the start sequence of a NetRexx application.

Scripts can be interpreted or compiled - there is no rule that a script needs to be interpreted. In both cases, interpreted or compiled, the NetRexx translator adds the necessary overhead to enable the JVM to execute the resulting program.

The scripting facility and its automatic generation of a class statement can lead to one surprising message when there is an error in the first part of the program: *class x already implied* when the automatically generated class statement (using the program file name) somehow clashes with the specified name that contains the error. When not using scripting mode, this error message nearly always indicates an error that occurred before the first class statement.

NetRexx as an Interpreted Language

In the JVM environment, compilation and interpretation are concepts that are not as straightforward as in other environments; JVM code is interpreted on several levels. When we are referring to *interpreted* NetRexx code, we indicate that there is no intermediate Java compilation step involved. A JVM .class file is always interpreted by the JVM runtime; the NetRexx translator is able to execute programs without generating either .java or .class files.

This enables a very quick edit-debug-run cycle, especially when combined with the command line feature that keeps the translator classes resident (the -prompt option), or one of the IDE plugins for NetRexx.

For NetRexx to deliver this functionality, the translator has been designed to have an analogous interpret facility for every code generation part.³

³This is the right way to explain this feature, because historically, the compiler was first (1996) and the interpretation facility was added later (in 2000).

NetRexx as a Compiled Language

5.1 Compiling from another program

The translator may be called from a NetRexx or Java program directly, by invoking the *main* method in the *org.netrexx.process.NetRexxC* class described as follows:

Listing 5.1: Invoking NetRexxC.main

```
1 method main(arg=Rexx, log=PrintWriter null) static returns int
```

The *Rexx* string passed to the method can be any combination of program names and options (except *-run*), as described above. Program names may optionally be enclosed in double-quote characters (and must be if the name includes any blanks in its specification).

A sample NetRexx program that invokes the NetRexx compiler to compile a program called *test* is:

Listing 5.2: Compiletest

```
1 /* compiletest.nrx */
2 s='test -keep -verbose4 -utf8'
3 say org.netrexx.process.NetRexxC.main(s)
```

Alternatively, the compiler may be called using the method:

Listing 5.3: Calling with Array argument

```
1 method main2(arg=String[], log=PrintWriter null) static returns int
```

in which case each element of the *arg* array must contain either a name or an option (except *-run*, as before). In this case, names must *not* be enclosed in double-quote characters, and may contain blanks.

For both methods, the returned *int* value will be one of the return values described above, and the second argument to the method is an optional *PrintWriter* stream. If the *PrintWriter* stream is provided, translator messages will be written to that stream (in addition to displaying them on the console, unless *-noconsole* is specified). It is the responsibility of the caller to create the stream (autoflush is recommended) and to close it after calling the compiler. The *-savelog* compiler option is ignored if a *PrintWriter* is provided (the *-savelog* option normally creates a *PrintWriter* for the file *NetRexxC.log*).

Note: NetRexxC is thread-safe (the only static properties are constants), but it is not known whether *javac* is thread-safe. Hence the invocation of multiple instances of NetRexxC on different threads should probably specify *-nocompile*, for safety.

5.2 Compiling from memory strings

Programs may also be compiled from memory strings by passing an array of strings containing programs to the translator using these methods:

Listing 5.4: From Memory

```
1 method main(arg=Rexx, programarray=String[], log=PrintWriter null) static returns int
2 method main2(arg=String[], programarray=String[], log=PrintWriter null) static returns
  int
```

Any programs passed as strings must be named in the arg parameter before any programs contained in files are named. For convenience when compiling a single program, the program can be passed directly to the compiler as a String with this method:

Listing 5.5: With String argument

```
1 method main(arg=Rexx, programstring=String, logfile=PrintWriter null) constant returns
  int
```

Here is an example of compiling a NetRexx program from a string in memory:

Listing 5.6: Example of compiling from String

```
1 import org.netrexx.process.NetRexxC
2 program = "say 'hello there via NetRexxC'"
3 NetRexxC.main("myprogram",program)
```

Calling non-JVM programs

Although NetRexx currently misses the `Address` facility that Classic Rexx and Object Rexx do have, it is easy to call non-JVM programs from a NetRexx program - not as easy as calling a JVM class of course, but if the following recipe is observed, it will show not to be a major problem. The following example is reusable for many cases.

Listing 6.1: Calling Non-JVM Programs

```

1  /* script\NonJava.nrx
2
3      This program starts an UNZIP program, redirect its output,
4      parses the output and shows the files stored in the zipfile */
5
6  parse arg unzip zipfile .
7
8  — check the arguments — show usage comments
9  if zipfile = '' then do
10     say 'Usage: Process unzipcommand zipfile'
11     exit 2
12 end
13
14 do
15     say "Files stored in" zipfile
16     say "____".left(39,"_") "____".left(39,"_")
17     child = Runtime.getRuntime().exec(unzip ' -v' zipfile) — program start
18
19     — read input from child process
20     in = BufferedReader(InputStreamReader(child.getInputStream()))
21     line = in.readLine
22
23     start = 0 — listing of files are not available yet
24     count = 0
25     loop while line \= null
26         parse line sep program
27         if sep = '_____' then start = \start
28         else
29             if start then do
30                 count = count + 1
31                 if count // 2 > 0 then say program.word(program.words).left(39) '\-'
32                 else say program.word(program.words)
33             end
34             line = in.readLine()
35         end
36
37     — wait for exit of child process and check return code
38     child.waitFor()
39     if child.exitValue() \= 0 then
40         say 'UNZIP return code' child.exitValue()
41
42     catch IOException
43         say 'Sorry cannot find' unzip
44     catch e2=InterruptedException
45         e2.printStackTrace()
46 end

```


Just firing off a program is no big deal, but this example (in script style) shows how easy it is to access the in- and output handles for the environment that executes the program, which enables you to capture the output the non-jvm program produces and do useful things with it.⁴ Line 17 starts the external command using the JVM `Runtime` class in a process called `child`. In line 20 we create a `BufferedReader` from the `child` processes' output. This is called an `InputStream` but it might as well have been called an `OutputStream`- everything regarding I/O is relative - but fortunately the designers of the JVM took care of deciding this for you. In lines 25-35 we loop through the results and show the files stored in the zipfile. Note that we **do** (line 14) have to **catch** (line 42) the `IOException` that ensues if the runtime cannot find the `unzip` program, maybe because it is not on the path or was not delivered with your operating system.

Starting from JVM 1.5 releases, there is a new way to accomplish the same goal, in a cleaner manner and with the added bonus of being able to redirect streams, and use environment variables. In this regard, the environment variable has made an important comeback from having its calls deprecated, to easy to use support in the *ProcessBuilder* class.

Listing 6.2: Use of ProcessBuilder

```

1  /**
2  * Class OSProcess implements ways to execute and get output from an OS Process
3  */
4  class OSProcess
5  properties indirect
6  pid = Process
7  returncode
8  commandList = ArrayList()
9  outList = ArrayList()
10
11  properties private
12  listeners = HashSet()
13  /**
14  * Default constructor
15  */
16  method OSProcess()
17  return
18
19  /**
20  * Method run starts an OS process from a command line in an ArrayList
21  * @param command is a List that has the command to be executed as elements
22  * @return List containing the output of the command
23  */
24  method outtrap(command_=ArrayList) returns ArrayList
25  this.commandList = command_
26  do
27  pb = ProcessBuilder(command_)
28  pb.redirectErrorStream(1)
29  this.pid = pb.start()
30  in = BufferedReader(InputStreamReader(this.pid.getInputStream()))
31  line = in.readLine()
32  loop while line <> null
33  this.outList.add(line)
34  line = in.readLine()
35  end
36  pid.waitFor()
37  returncode = pid.exitValue()
38  return this.outList
39  catch iox=IOException
40  say iox.getMessage()
41  return ArrayList()
42  catch InterruptedException
43  say "interrupted"
44  return ArrayList()

```

⁴This is akin to what one would do with *queue* on z/VM CMS and *outtrap* on z/OS TSO in Classic Rexx.

```

45     end — do
46
47     /**
48     * Method exec starts an OS process from a command line in an ArrayList
49     * @param then fires off outputEvent events to every registered listener
50     * @return void
51     */
52     method exec(command_=ArrayList)
53         this.commandList = command_
54         do
55             pb = ProcessBuilder(command_)
56             pb.redirectErrorStream(1)
57             this.pid = pb.start()
58             in = BufferedReader(InputStreamReader(this.pid.getInputStream()))
59             line = in.readLine()
60             loop while line <> null
61                 line = in.readLine()
62                 i = this.listeners.iterator()
63                 loop while i.hasNext()
64                     op = OutputEventListener i.next()
65                     op.outputReceived(OutputLineEvent(this,line,this.pid))
66                 end
67             end
68             pid.waitFor()
69             returncode = pid.exitValue()
70         catch iox=IOException
71             say iox.getMessage()
72         catch InterruptedException
73             say "interrupted"
74         end — do
75
76
77     /**
78     * Method addOutputEventListener supports registering an event listener
79     * @param listener_ is a OutputEventListener
80     */
81     method addOutputEventListener(listener_=OutputEventListener)
82         this.listeners.add(listener_)
83
84     /**
85     * Method removeOutputEventListener supports de-registering an event listener
86     * @param listener_ is a OutputEventListener
87     */
88     method removeOutputEventListener(listener_=OutputEventListener)
89         this.listeners.remove(listener_)

```

In the above sample, we are using two different ways to obtain the output from a process started by the JVM from our own program. The method *outtrap* waits until the invoked process is finished and returns all output lines in an `ArrayList`. Its name is not entirely coincidental with the similar TSO *outtrap* function.

Sometimes we cannot wait until the child process is finished, for example when it is a long running process and we need to capture the output on a line-by-line basis to see what is happening - in case of the example, this was done to capture the output as part of a testsuite for a multithreaded file transfer application, which has a server resident process that is not supposed to end, because one of its tasks is to poll a directory for incoming files with a specific pattern in the file names. This is implemented using an Event based pattern (as explained in 10.1 on page 27).

Listing 6.3: Output Line Event

```

1 import java.util.EventObject
2 /**
3  * Class OutputLineEvent embodies the OutputLineEvent
4  */
5 class OutputLineEvent extends EventObject
6
7 properties indirect
8 pid = Process

```

```

9   line
10  /**
11   * Default constructor
12   */
13  method OutputLineEvent(ob=Object,line_, pid_=Process)
14      super(ob)
15      this.line = line_
16      this.pid = pid_
17      return

```

Listing 6.4: Output Event Listener

```

1  import java.util.EventListener
2  /**
3   * Interface OutputEventListener specifies the one mandatory method for this interface
4   */
5  class OutputEventListener interface implements EventListener
6
7  method outputReceived(ob=OutputLineEvent)

```

The call would look something like this:

Listing 6.5: Example of calling the OSProcess class - registering an eventhandler

```

1  os = OSProcess()
2  os.addOutputEventListener(this)
3  os.exec(command)

```

The class must extend `OutputEvenListener`, and implement this method:

Listing 6.6: Example of implementing the listener method

```

1  method outputReceived(ob=OutputLineEvent)
2      this.counter = this.counter+1
3      say this.counter ob.getPid() ob.getLine()

```

Using NetRexx classes from Java

If you are a Java programmer, using a NetRexx class from Java is just as easy as using a Java class from NetRexx. NetRexx compiles to Java classes that can be used by Java programs. You should import the `netrexx.lang` package to be able to use the short class name for the Rexx (NetRexx string and numerics) class.

A NetRexx method without a `returns` keyword can return nothing, which is the void type in Java, or a Rexx string. NetRexx is case independent⁵; Java is case dependent. NetRexx generates the Java code with the case used in the class and method instructions. For example, if you named your class `Spider` in the NetRexx source file, the resulting Java class file is `Spider.class`. The public class name in your source program must match the NetRexx source file name. For example, if your source file is `SPIDER.NRX`, and your class is `Spider`, NetRexx generates a warning and changes the class name to `SPIDER` to match the file name. A Java program using the class name `Spider` would not find the generated class, because its name is `SPIDER.class` - if the compile succeeded, which is not guaranteed in case of casing mismatches. If you have problems, compile your NetRexx program with the **options -keepasjava -format**. You then can look at the generated java file for the correct spelling style and method parameters.

⁵With the default of `options nostrictcase` in effect.

Classes

Somewhere in the nineties Object Orientation became one of the mainstream ways to organize computer programs, and support for this was added to programming languages. C became C++ with a preprocessor that generates Cfront that is not entirely unlike the NetRexx translator produces Java. Java in itself is syntax-wise a cleaned up version of C++, but in essence an entirely different language. Its inventor and architect, James Gosling, has stated on various occasions that he was planning a fully different syntax for what finally became Java - but that Sun management more or less forced him to use a C++ derived syntax, because C++ compilers was what SUN did well at the time. With Brendan Eich having to base JavaScript on naming and syntax on Java, the circle that brought the world terse, curly braces based notations, is complete.

For an audience of Rexx programmers, the usual OO presentation goes into the advantages of the paradigm. Today, that is not really necessary, and OO is a given; it slightly deviates from earlier notation as result of trying to put data and procedure into *Objects*, but it is no great deal, and this NetRexx Programmer's Guide does not need a special section on the benefits of the OO paradigm. It is assumed that with a few examples everyone should be able to *get* it; some old programmers might resist but there is really no use in fighting the mainstream. Consequently, this section discusses the way to do this in NetRexx; the way NetRexx does it is for a very large part formed by the way the JVM dictates it, adapted to Rexx notational style and conventions.

8.1 Classes

Classes represent a blueprint, 'cookie cutter' approach in creating objects that do useful things. A class is defined in a file by the same name (exceptions here for dependent classes). So a class called `Cookie` is defined in a file called `Cookie.nrx`. Its *real*, which means its most specific name, including its package specification, is not given by the file name but by the combination of the `class=file` + the name given on the `package` statement. This enables one to put classes in different packages without having to change the file names.

8.2 Dependent Classes

8.3 Properties

8.4 Methods

8.5 Inheritance

8.6 Overriding Methods

8.7 Overriding Properties

Using Packages

Any non-toy, non-trivial program needs to be in a package. Only examples in programming books (present company included) have programs without package statements. The reason for this is that there is a fairly large chance that you will give something a name that is already used by someone else for something else. Things are not their names⁶, and the same names are given to wildly dissimilar things. The *package* construct is the JVM's approach to introducing *namespaces* into the total set of programs that programmers make. Different people will probably write some method that is called `listDifferences` sometime. With all my software in a package called `com.frob.nitz` and yours in a package called `com.frob.otzim`, there is no danger of our programs calling the wrong class and listing the wrong differences.

It is imperative to understand this chapter before continuing - it is a mechanical nuts-and-bolts issue but an essential one at that.

9.1 The package statement

The final words about the NetRexx **package** statement is in the NetRexx Language Reference, but the final statement about the package *mechanism* is in the JVM documentation.

9.2 Translator performance consequences

Because the NetRexx translator has to scan all packages that it can see (meaning a recursive scan of the directories below its own level in the directory tree, and on its classpath, it is often advisable (and certainly if `.` (a dot, representing the current directory) is part of the classpath) to do development in a subdirectory, instead of, for example, the top level home directory. If a large number of packages and classes are visible to the translator, compile times will be negatively impacted.

9.3 Some NetRexx package history

All IBM versions of NetRexx had the translator in a package called

`COM.ibm.netrexx.process`

⁶Willard Van Orman Quine, *Word and Object*, MIT Press, 1960, ISBN 0-262-67001-1

The official, SUN ordained convention for package names was, to prepend the reversed domain name of the vendor to the package name, while uppercasing the top level domain. NetRexx, being one of the first programs to make use of packages, followed this convention, that was quickly dropped by SUN afterwards, probably because someone experienced what trouble it could cause with version management software that adapted to case-*sensitive* and case-*insensitive* file systems. For NetRexx, which had started out keenly observing the rules, this insight came late, and it is a sober fact that as a result some needlessly profane language was uttered on occasion by some in some projects that suffered the consequences of this. With the first RexxLA release of NetRexx in 2011, the package name was changed to `org.netrexx`, while the runtime package name was kept as `netrexx.lang`, because some major other languages also follow this convention.

9.4 CLASSPATH

Most implementations of Java use an environment variable called CLASSPATH to indicate a search path for Java classes. The Java Virtual Machine and the NetRexx translator rely on the CLASSPATH value to find directories, zip files, and jar files which may contain Java classes. The procedure for setting the CLASSPATH environment variable depends on your operating system (and there may be more than one way).

- For Linux and Unix (BASH, Korn, or Bourne shell), use:

```
CLASSPATH=<newdir>:\$CLASSPATH
export CLASSPATH
```

- Changes for re-boot or opening of a new window should be placed in your `/etc/profile`, `.login`, or `.profile` file, as appropriate.
- For Linux and Unix (C shell), use:

```
setenv CLASSPATH <newdir>:\$CLASSPATH
```

Changes for re-boot or opening of a new window should be placed in your `.cshrc` file. If you are unsure of how to do this, check the documentation you have for installing the Java toolkit.

- For Windows operating systems, it is best to set the system wide environment, which is accessible using the Control Panel (a search for “environment” offsets the many attempts to relocate the exact dialog in successive Windows Control Panel versions somewhat).

Programming Patterns

Much has been made of patterns as aggregations of higher level embodiments of programming solutions. It has been observed that of a number of the C++ oriented patterns in Design Patterns⁷, a fair number owes its existence to complications in the C++ language and it not readily reproducible in a Java Patterns or Ruby Patterns book. The same goes for NetRexx- in this chapter we would like to present a number of Java patterns usable in NetRexx, and a number of patterns that are unique to NetRexx.

10.1 Events

10.2 Recursive Parse

This is a pattern unique to Rexx, by virtue of Rexx having the Parse statement. It also works in NetRexx.

10.3 Observer

The observer pattern can also be referred to as *Callback*, and the Java Event class delivers support for it. It is very usable if some result needs to be available for a set of callers, where the set is 0 to many. It works as follows: (see a simple implementation in section 6.4 on page 20) An object, maintains a list of its dependents, called observers, and notifies them automatically of any state changes, usually by calling one of their methods. It is mainly used to implement distributed event handling systems. The Observer pattern is also a key part in the familiar Model View Controller (MVC) architectural pattern. In the JVM, this object needs to implement the methods of the Listener interface; this interface specifies the `addListener` and `removeListener` methods; it keeps a collection in which references to the added listener objects are maintained. The listening is done to subclass Java Event classes. The event specifies the method to be called when 'firing off' and event. This means that this method is called on every listener.

One of the larger benefits: it decouples the observer from the subject. The subject doesn't need to know anything special about its observers. Instead, the subject simply allows observers to subscribe. When the subject generates an event, it simply passes it to each of its observers. Another benefit is that event consuming classes don't have to wait until a process is finished, and can consume events as they come in. The OSProcess

⁷Gamma, Helm, Johnson, Vlissides, Design Patterns: Elements of Reusable Object-Oriented Software, Addison-Wesley Professional; 1994

class on page 20) uses an event approach to consume output lines from a subprocess - in the version that puts the output in an ArrayList needs to wait for the subprocess to end, but the event driven version can monitor a long running process and analyze output lines whenever they are received.

Incorporating Class Libraries

11.1 The Collection Classes

Input and Output

12.1 The File Class

12.2 Streams

12.3 Line mode I/O

12.3.1 Line mode I/O using BufferedReader and PrintWriter

12.3.2 Line mode I/O using BufferedReader and BufferedWriter

12.4 Byte Oriented I/O

12.5 Data Oriented I/O

12.6 Object Oriented I/O using Serialization

12.7 The NIO Approach

Algorithms in NetRexx

13.1 Factorial

A *factorial* is the product of an integer and all the integers below it; the mathematical symbol used is ! (the exclamation mark). For example 4! is equal to 24 (because $4*3*2*1=24$). The following program illustrates a recursive (a method calling itself) and an iterative approach to calculating factorials.

Listing 13.1: Factorial

```

1  /* NetRexx */
2
3  options replace format comments java crossref savelog symbols nobinary
4
5  numeric digits 64 — switch to exponential format when numbers become larger than 64
   digits
6
7  say 'Input a number: \-'
8  say
9  do
10   n_ = long ask — Gets the number, must be an integer
11
12   say n_ '! =' factorial(n_) '(using iteration)'
13   say n_ '! =' factorial(n_, 'r') '(using recursion)'
14
15   catch ex = Exception
16     ex.printStackTrace
17 end
18
19 return
20
21 method factorial(n_ = long, fmethod = 'I') public static returns Rexx signals
   IllegalArgumentException
22
23 if n_ < 0 then —
24   signal IllegalArgumentException('Sorry, but' n_ 'is not a positive integer')
25
26 select
27   when fmethod.upper = 'R' then —
28     fact = factorialRecursive(n_)
29   otherwise —
30     fact = factorialIterative(n_)
31 end
32
33 return fact
34
35 method factorialIterative(n_ = long) private static returns Rexx
36
37 fact = 1
38 loop i_ = 1 to n_
39   fact = fact * i_
40 end i_
41
42 return fact
43
44 method factorialRecursive(n_ = long) private static returns Rexx

```

```

45
46   if n_ > 1 then -
47       fact = n_ * factorialRecursive(n_ - 1)
48   else -
49       fact = 1
50
51   return fact

```

Executing this program yields the following result:

```
===== Exec: RCFactorial =====
```

Input a number:

42

42! = 1405006117752879898543142606244511569936384000000000 (using iteration)

42! = 1405006117752879898543142606244511569936384000000000 (using recursion)

As you can see, fortunately, both approaches come to the same conclusion about the results. In the above program, both approaches are a bit intermingled; for more clarity about how to use recursion, have a look at this:

Listing 13.2: Factorial Recursive

```

1 class Factorial
2   numeric digits 64
3
4   method main(args=String[]) static
5       say factorial_(42)
6
7   method factorial_(number) static
8       if number = 0 then return 1
9       else return number * factorial_(number-1)

```

In this program we can clearly see that the `factorial_` method, that takes an argument `number` (which is of type `Rexx` if we do not specify it to be another type), calls itself in the method body. This means that at runtime, another copy of it is run, with as argument `number` that the first invocation returns (the result of `42*41`), and so on.

In general, a recursive algorithm is considered more elegant, while an iterative approach has a better runtime performance. Some language environments are optimized for recursion, which means that their processors can spot a recursive algorithm and optimize it by not making many useless copies of the code. Some day in the near future the JVM will be such an environment. Also, for some problems, for example the processing of tree structures, using a recursive algorithm seems much more natural, while an iterative algorithm seems complicated or forced.

13.2 Fibonacci

Listing 13.3: Fibonacci

```

1  /* NetRexx */
2  options replace format comments java crossref savelog symbols
3
4  numeric digits 210000                                /*prepare for some big ones.    */
5  parse arg x y .                                       /*allow a single number or range.*/
6  if x == '' then do                                    /*no input? Then assume -30-->+30*/
7      x = -30
8      y = -x
9  end
10
11 if y == '' then y = x                                  /*if only one number, show fib(n)*/
12 loop k = x to y                                       /*process each Fibonacci request.*/
13     q = fib(k)

```

```

14 w = q.length                                     /*if wider than 25 bytes, tell it*/
15 say 'Fibonacci' k="q
16 if w > 25 then say 'Fibonacci' k "has a length of" w
17 end k
18 exit
19
20 /*-----FIB subroutine (non-recursive)-----*/
21 method fib(arg) private static
22   parse arg n
23   na = n.abs
24
25   if na < 2 then return na                       /*handle special cases.      */
26   a = 0
27   b = 1
28
29   loop j = 2 to na
30     s = a + b
31     a = b
32     b = s
33   end j
34
35   if n > 0 | na // 2 == 1 then return s /*if positive or odd negative... */
36   else return -s /*return a negative Fib number. */

```

14

Using Parse

15

Using Trace

Concurrency

16.1 Threads

Threads are a built-in multitasking feature of the JVM. Where earlier JVM implementations sometime ran on so-called *Green Threads*, which is a library that implements thread support for OS'ses that do not have this facility (an early version of Java was called *GreenTalk* for this reason), modern versions all use native OS thread support.

A new thread is created when we create an instance of the Thread class. We cannot tell a thread which method to run, because threads are not references to methods. Instead we use the Runnable interface to create an object that contains the run method:

Every thread begins its concurrent life by executing the run method. The run method does not have any parameters, does not return a value, and is not allowed to signal any exceptions. Any class that implements the Runnable interface can serve as a target of a new thread. An object of a class that implements the Runnable interface is used as a parameter for the thread constructor.

Threads can be given a name that is visible when listing the threads in your system. It is good practice to name every thread, because if something goes wrong you can see which threads are still running. Additionally, threads are grouped by thread groups. If you do not supply a thread group, the new thread is added to the thread group of the currently executing thread. The threads of a group and their subgroups can be destroyed, stopped, resumed, or suspended by using the ThreadGroup object.

The next two samples are used in the following programs that illustrate thread usage.

Listing 16.1: Thread sample 1

```

1  /* thread/ThrdTst1.nrx */
2
3  h1 = Hello1('This is thread 1')
4  h2 = Hello1('This is thread 2')
5
6  Thread(h1, 'Thread Test Thread 1').start()
7  Thread(h2, 'Thread Test Thread 2').start()
8
9  class Hello1 implements Runnable
10     Properties inheritable
11     message = String
12
13     method Hello1( s = String)
14         message = s
15
16     method run()
17         loop for 50
18             say message
19         end

```

Listing 16.2: Thread sample 2

```
1  /* thread/ThrdTst2.nrx */
2
3  h1 = Hello2('This is thread 1')
4  h2 = Hello2('This is thread 2')
5
6  h1.start()
7  h2.start()
8
9  class Hello2 extends Thread
10   Properties inheritable
11     message = String
12
13   method Hello2( s = String)
14     super('Thread Test - Message' s)
15     message = s
16
17   method run()
18     loop for 50
19       say message
20     do
21       sleep(10)
22     catch InterruptedException
23   end
24 end
```

The second class, `Hello2`, does not *implement* the `Runnable` interface, but subclasses it, so it inherits its methods. This is a valid approach, and it is up to the developer to choose an implementation and worry about the semantics of an inherited thread interface. A newly created thread remains idle until the `start` method is invoked. The thread then wakes up and executes the `run` method of its target object. The `start` method can be called only once. The thread continues running until the `run` method completes or the `stop` method of the thread is called.

User Interfaces

17.1 AWT

17.2 Web Applets using AWT

Web applets can be written one of two styles:

- Lean and mean, where binary arithmetic is used, and only core Java classes (such as *java.lang.String*) are used. This is recommended for optimizing webpages which may be accessed by people using a slow internet connection. Several examples using this style are included in the NetRexx package like the two listed below.

Listing 17.1: Nervous Texxt

```

1  /* NervousText applet in NetRexx: Test of text animation.
2     Algorithms, names, etc. are directly from the Java version by
3     Daniel Wyszynski and kwalrath, 1995
4  */
5  options binary
6
7  class NervousTexxt extends Applet implements Runnable
8
9      separated = char[]
10     s = String
11     killme = Thread
12     threadSuspended = boolean 0
13
14     method init
15         resize(300,50)
16         setFont(Font("TimesRoman", Font.BOLD, 30))
17         s = getParameter("text")
18         if s = null then s = "NetRexx"
19
20         separated = char[s.length]
21         s.getChars(0, s.length, separated,0)
22
23     method start
24         if killme \= null then return
25         killme = Thread(this)
26         killme.start
27
28     method stop
29         killme = null
30
31     method run
32         loop while killme \= null
33             Thread.sleep(100)
34             this.repaint
35         catch InterruptedException
36         end
37         killme = null
38
39     method paint(g=Graphics)

```

```

40     loop i=0 to s.length-1
41         x_coord = int Math.random*10+15*i
42         y_coord = int Math.random*10+36
43         g.drawChars(separated, i, 1, x_coord, y_coord)
44     end
45
46     method mouseDown(evt=Event, x=int, y=int) returns boolean
47         if threadSuspended then killme.resume
48             else killme.suspend
49     threadSuspended = \threadSuspended
50     return 1

```

Listing 17.2: ArchText

```

1  /* ArchText applet: multi-coloured text on a white background */
2  /* Mike Cowlishaw April 1996, December 1996 */
3  options binary
4
5  class ArchText extends Applet implements Runnable
6
7      text ="NetRexx" /* default text */
8      tick =0 /* display counter */
9      timer =Thread null /* timer thread */
10     shadow=Image /* shadow image */
11     draw =Graphics /* where we can draw */
12
13     method init
14         s=getParameter("text") /* get any provided text */
15         if s\=null then text=s
16         shadow=createImage(getSize.width, getSize.height) /* image */
17         draw=shadow.getGraphics
18         draw.setColor(Color.white) /* background */
19         draw.fillRect(0, 0, getSize.width, getSize.height) /* .. */
20         draw.setFont(Font("TimesRoman", Font.BOLD, 30)) /* font */
21
22     method start
23         if timer=null then timer=Thread(this) /* new thread */
24         timer.setPriority(Thread.MAX_PRIORITY) /* time matters */
25         timer.start /* start the thread */
26
27     method stop
28         if timer=null then return /* have no thread */
29         timer.stop /* else stop it */
30         timer=null /* .. and discard */
31
32     method run /* this runs as thread */
33         loop while timer\=null
34             tick=tick+1 /* next update */
35             hue=((tick+133)//191)/191
36             draw.setColor(Color.getHSBColor(hue, 1, 0.7))
37             draw.drawString(text, 0, 30)
38             this.repaint /* .. and redraw */
39             Thread.sleep(119) /* wait awhile */
40         catch InterruptedException
41         end
42         timer=null /* discard */
43
44     method update(g=Graphics) /* override Applet's update */
45         paint(g) /* method to avoid flicker */
46
47     method paint(g=Graphics)
48         g.drawImage(shadow, 0, 0, null)

```

- Full-function, where decimal arithmetic is used, and advantage is taken of the full power of the NetRexx runtime *Rexx* class.

An example using this style is the below *WordClock.nrx*.

Listing 17.3: WordClock

```

1  /* WordClock — an applet that shows the time in English. */

```

```

2  /*                                                                 */
3  /*   Parameters:                                                  */
4  /*                                                                 */
5  /*     face — the font face to use                                */
6  /*     size — the font size to use                                */
7  /*                                                                 */
8  /* -----                                                        */
9  /* Based on the ancient QTIME.REXX, and typical Java applets. */
10
11 class WordClock extends Applet implements Runnable
12
13     timer=Thread null          /* the timer thread */
14     offsetx; offsety          /* text position */
15     now                        /* current time */
16
17 method init
18     /* Get parameters from the <applet> markup */
19     face=getParameter("face") /* font face */
20     if face=null then face="TimesRoman"
21     size=getParameter("size")
22     if size=null then size="20" /* font size */
23
24     setFont(Font(face, Font.BOLD, size))
25     resize(size*20, size*2)      /* set window size */
26     offsetx=size/2              /* and where text will start */
27     offsety=size*3/2            /* note Y is from top */
28     parse Date() . . . now .    /* initial time is fourth word */
29
30 method start
31     if timer=null then timer=Thread(this) /* new thread */
32     timer.setPriority(Thread.MAX_PRIORITY) /* time matters */
33     timer.start                       /* start the thread */
34
35 method stop
36     if timer\=null then do          /* have thread */
37         timer.stop                  /* .. so stop it */
38         timer=null                  /* .. and discard */
39     end
40
41 method run
42     /* Use the Java Date class to get the time */
43     loop while timer\=null
44         parse Date() . . . now .    /* time is fourth word */
45         this.repaint                /* redisplay */
46         parse now ':' ':' secs      /* where in minute */
47         wait=30-secs                /* calculate delay in seconds */
48         if wait<=0 then wait=wait+60
49         /* say 'secs, wait:' secs wait */
50         Thread.sleep(1000*wait)     /* wait for milliseconds */
51     catch InterruptedException
52         say 'Interrupted...'
53     end
54     timer=null                      /* done */
55
56 method paint(g=Graphics)
57     g.drawString(wordtime(now), offsetx, offsety) /* show it */
58
59 /* WORDTIME — a cut-down version of QTIME.REXX
60     Arg1 is the time string (hh:mm:ss)
61     Returns the time in english, as a REXX string
62     */
63 method wordtime(arg) static returns REXX
64     /* Extract the hours, minutes, and seconds from the time. */
65     parse arg hour ':' 'min' ':' 'sec
66     if sec>29 then min=min+1        /* round up minutes */
67
68     /* Nearness phrases — this time using an array */
69     near=REXX[5]                    /* five items */
70     near[0]=' '                     /* exact */
71     near[1]=' just gone'; near[2]=' just after' /* after */
72     near[3]=' nearly'; near[4]=' almost' /* before */
73
74     mod=min//5                      /* where we are in 5 minute bracket */
75     out="It's"near[mod]             /* start building the result */

```

```

76  if min>32 then hour=hour+1          /* we are TO the hour... */
77  min=min+2      /* shift minutes to straddle a 5-minute point */
78
79  /* Now special-case the result for Noon and Midnight hours */
80  if hour//12=0 & min//60<=4 then do
81      if hour=12 then return out 'Noon.'
82      return 'Midnight.'
83  end
84
85  min=min-(min//5)          /* find nearest 5 mins */
86  if hour>12
87  then hour=hour-12        /* get rid of 24-hour clock */
88  else
89      if hour=0 then hour=12    /* .. and allow for midnight */
90
91  /* Determine the phrase to use for each 5-minute segment */
92  select
93      when min=0 then nop          /* add "o'clock" later */
94      when min=60 then min=0      /* ditto */
95      when min= 5 then out=out 'five past'
96      when min=10 then out=out 'ten past'
97      when min=15 then out=out 'a quarter past'
98      when min=20 then out=out 'twenty past'
99      when min=25 then out=out 'twenty-five past'
100     when min=30 then out=out 'half past'
101     when min=35 then out=out 'twenty-five to'
102     when min=40 then out=out 'twenty to'
103     when min=45 then out=out 'a quarter to'
104     when min=50 then out=out 'ten to'
105     when min=55 then out=out 'five to'
106  end
107
108  numbers='one two three four five six'— /* continuation */
109  'seven eight nine ten eleven twelve '
110  out=out numbers.word(hour)          /* add the hour number */
111  if min=0 then out=out "o'clock" /* .. and o'clock if exact */
112
113  return out'.'          /* return the final result */
114
115  /* Mike Cowlishaw, December 1979 — January 1985. */
116  /* NetRexx version March 1996; applet April 1996. */

```

If you write applets which use the NetRexx runtime (or any other Java classes that might not be on the client browser), the rest of this section may help in setting up your Web server.

A good way of setting up an HTTP (Web) server for this is to keep all your applets in one subdirectory. You can then make the NetRexx runtime classes (that is, the classes in the package known to the Java Virtual Machine as *netrexx.lang*) available to all the applets by unzipping NetRexxR.jar into a subdirectory *netrexx/lang* below your applets directory.

For example, if the root of your server data tree is

D:\mydata

you might put your applets into

D:\mydata\applets

and then the NetRexx classes (unzipped from NetRexxR.jar) should be in the directory

D:\mydata\applets\netrexx\lang

The same principle is applied if you have any other non-core Java packages that you want to make available to your applets: the classes in a package called *iris.sort.quicksorts* would go in a subdirectory below *applets* called *iris/sort/quicksorts*, for example.

Note that since Java 1.1 or later it is possible to use the classes direct from the NetRexxR.jar file.

17.3 Swing

Swing is the most commonly used name for the second attempt from the SUN engineers to provide a graphical user interface library for the JVM. With AWT also acknowledged by SUN to be a quick attempt that was made just before release of the first Java package, it became clear that it was rather taxing on system resources without compensation by a pretty look. A case in point is the event mechanism, that indiscriminately sends around mouse and keyboard events even when nobody is listening to them. The architecture for Swing prescribes registering for events before they are produced, and tries to have the drawing done by the Java graphics engine instead of leaning heavily on the operating system's native GUI functionality. The user interface widgets that are produced by Java are called 'light' and their looks can be changed by applying different skins, called '*look-and-feel*' (LAF) libraries.

In the first months of its existence Swing gathered quite a bad reputation because it made the Java 1.2 releases that contained it very slow in starting up programs that used the library. Consequently, much was invested in performance studies by SUN engineers and these problems were solved. One of the things that came out is that dividing the libraries in a great many classes, done for performance reasons, worked counterproductive. All these problems were solved over the years, and developments in hardware and multithreading took care of the rest, and nowadays Swing is a valid way of producing a rich client user interface.

For esthetical reasons, it is best to research a bit in the third party look-and-feel libraries that can be obtained. Swing can be made to look beautiful, but it takes some care and the defaults are not helping.

17.3.1 Creating NetRexx Swing interfaces with NetBeans

17.4 Web Frameworks

17.4.1 JSF

Network Programming

18.1 Using Uniform Resource Locators (URL)

18.2 TCP/IP Socket I/O

18.3 RMI: Remote Method Interface

Database Connectivity with JDBC

For interfacing with Relational Database Management Systems (RDBMS) NetRexx uses the Java Data Base Connectivity (JDBC) model. This means that all important database systems, for which a JDBC driver has been made available, can be used from your NetRexx program. This is a large bonus when we compare this to the other open source scripting languages, that have been made go by with specific, nonstandard solutions and special drivers. In contrast, NetRexx programs can be made compatible with most database systems that use standard SQL, and, with some planning and care, can switch database implementations at will.

Listing 19.1: A JDBC Query example

```

1  /* jdbc\JdbcQry.nrx
2
3  This NetRexx program demonstrate DB2 query using the JDBC API.
4  Usage: Java JdbcQry [<DB-URL>] [<userprefix>] */
5
6  import java.sql.
7
8  parse arg url prefix          — process arguments
9  if url = '' then
10     url = 'jdbc:db2:sample'
11 else do                      — check for correct URL
12     parse url p1 ':' p2 ':' rest
13     if p1 \= 'jdbc' | p2 \= 'db2' | rest = '' then do
14         say 'Usage: java JdbcQry [<DB-URL>] [<userprefix>]'
15         exit 8
16     end
17 end
18 if prefix = '' then prefix = 'userid'
19
20 do                            — loading DB2 support
21     say 'Loading DB2 driver classes...'
22     Class.forName('COM.ibm.db2.jdbc.app.DB2Driver').newInstance()
23     — Class.forName('COM.ibm.db2.jdbc.net.DB2Driver').newInstance()
24 catch e1 = Exception
25     say 'The DB2 driver classes could not be found and loaded !'
26     say 'Exception ( ' e1 ' ) caught : \n' e1.getMessage()
27     exit 1
28 end                            — end : loading DB2 support
29
30 do                            — connecting to DB2 host
31     say 'Connecting to:' url
32     jdbcCon = Connection DriverManager.getConnection(url, 'userid', 'password')
33 catch e2 = SQLException
34     say 'SQLException(s) caught while connecting !'
35     loop while (e2 \= null)
36         say 'SQLState:' e2.getSQLState()
37         say 'Message: ' e2.getMessage()
38         say 'Vendor: ' e2.getErrorCode()
39         say
40         e2 = e2.getNextException()
41     end
42     exit 1

```

```

43 end                                     — end : connecting to DB2 host
44
45 do                                     — get list of departments with the managers
46 say 'Creating query...'
47 query = 'SELECT deptno, deptname, lastname, firstnme' —
48         'FROM' prefix'.DEPARTMENT dep,' prefix'.EMPLOYEE emp'-
49         'WHERE dep.mgrno=emp.empno ORDER BY dep.deptno'
50 stmt = Statement jdbcCon.createStatement()
51 say 'Executing query:'
52 loop i=0 to (query.length()-1)%75
53     say ' ' query.substr(i*75+1,75)
54 end
55 rs = ResultSet stmt.executeQuery(query)
56 say 'Results:'
57 loop row=0 while rs.next()
58     say rs.getString('deptno') rs.getString('deptname') —
59         'is directed by' rs.getString('lastname') rs.getString('firstnme')
60 end
61 rs.close()                             — close the ResultSet
62 stmt.close()                           — close the Statement
63 jdbcCon.close()                         — close the Connection
64 say 'Retrieved' row 'departments.'
65 catch e3 = SQLException
66     say 'SQLException(s) caught !'
67     loop while (e3 \= null)
68         say 'SQLState:' e3.getSQLState()
69         say 'Message:' e3.getMessage()
70         say 'Vendor:' e3.getErrorCode()
71     say
72     e3 = e3.getNextException()
73 end
74 end                                     — end: get list of departments

```

The first peculiarity of JDBC is the way the driver class is loaded. When most classes are 'pulled in' by the translator, a JDBC driver traditionally is loaded through the reflection API. This happens in line 22 with the `Class.forName` call. This implies that the library containing this class must be on the classpath.

In line 32 we connect to the database using a url and a userid/password combination. This is an easy way to do and test, but for most serious applications we do not want plaintext userids and passwords in the sourcecode, so most of the time we would store the connection info in a file that we store in encrypted form, or we use facilities of J2EE containers that can provide data sources that take care of this, while at the same time decoupling your application source from the infrastructure that it will run on.

In line 47 the query is composed by filling in variables in a Rexx string and making a `Statement` out of it, in line 50. In line 55, the `Statement` is executed, which yields a `ResultSet`. This has a *cursor* that moves forward with each `next` call. The `next` call returns *true* as long as there are rows from the resultset to return.

The `ResultSet` interface implements *getter* methods for all JDBC Types. In the above example, all returned results are of type `String`.

Listing 19.2: A JDBC Update example

```

1 /* jdbc\JdbcUpd.nrx
2
3     This NetRexx program demonstrate DB2 update using the JDBC API.
4     Usage: Java JdbcUpd [<DB-URL>] [<userprefix>] [U] */
5
6 import java.sql.
7
8 parse arg url prefix lowup             — process arguments
9 if url = '' then
10     url = 'jdbc:db2:sample'
11 else do                               — check for correct URL
12     parse url p1 ':' p2 ':' rest
13     if p1 \= 'jdbc' | p2 \= 'db2' | rest = '' then do

```

```

14         say 'Usage: java JdbcUpd [<DB-URL>] [<userprefix>] [U]'
15         exit 8
16     end
17 end
18 if prefix = '' then prefix = 'userid'
19 if lowup \= 'U' then lowup = 'L'
20
21 do                                     — loading DB2 support
22     say 'Loading DB2 driver classes...'
23     Class.forName('COM.ibm.db2.jdbc.app.DB2Driver').newInstance()
24     — Class.forName('COM.ibm.db2.jdbc.net.DB2Driver').newInstance()
25 catch e1 = Exception
26     say 'The DB2 driver classes could not be found and loaded !'
27     say 'Exception (' e1 ') caught : \n' e1.getMessage()
28     exit 1
29 end                                     — end : loading DB2 support
30
31 do                                     — connecting to DB2 host
32     say 'Connecting to:' url
33     jdbcCon = Connection DriverManager.getConnection(url, 'userid', 'password')
34 catch e2 = SQLException
35     say 'SQLException(s) caught while connecting !'
36     loop while (e2 \= null)
37         say 'SQLState:' e2.getSQLState()
38         say 'Message: ' e2.getMessage()
39         say 'Vendor: ' e2.getErrorCode()
40         say
41         e2 = e2.getNextException()
42     end
43     exit 1
44 end                                     — end : connecting to DB2 host
45
46 do                                     — retrieve employee, update firstname
47
48     say 'Preparing update...'           — prepare UPDATE
49     updateQ = 'UPDATE' prefix'.EMPLOYEE SET firstnme = ? WHERE empno = ?'
50     updateStmt = PreparedStatement jdbcCon.prepareStatement(updateQ)
51     say 'Creating query...'             — create SELECT
52     query = 'SELECT firstnme, lastname, empno FROM' prefix'.EMPLOYEE'
53     stmt = Statement jdbcCon.createStatement()
54     rs = ResultSet stmt.executeQuery(query) — execute select
55
56     loop row=0 while rs.next()           — loop employees
57         firstname = String rs.getString('firstnme')
58         if lowup = 'U' then firstname = firstname.toUpperCase()
59         else do
60             dChar = firstname.charAt(0)
61             firstname = dChar || firstname.substring(1).toLowerCase()
62         end
63         updateStmt.setString(1, firstname) — parms for update
64         updateStmt.setString(2, rs.getString('empno'))
65         say 'Updating' rs.getString('lastname') firstname ': \0'
66         say updateStmt.executeUpdate() 'row(s) updated' — execute update
67     end
68
69     rs.close()                           — close the ResultSet
70     stmt.close()                         — close the Statement
71     updateStmt.close()                  — close the PreparedStatement
72     jdbcCon.close()                     — close the Connection
73     say 'Updated' row 'employees.'
74 catch e3 = SQLException
75     say 'SQLException(s) caught !'
76     loop while (e3 \= null)
77         say 'SQLState:' e3.getSQLState()
78         say 'Message: ' e3.getMessage()
79         say 'Vendor: ' e3.getErrorCode()
80         say
81         e3 = e3.getNextException()
82     end
83 end                                     — end: employees

```

For database updates, we connect using the driver in the same way (line 23) and now prepare the statement used for the database update (line 50). In this example,

we loop through the cursor of a select statement and update the row in line 66. The `executeUpdate` method of `PreparedStatement` returns the number of updated rows as an indication of success.

From JDBC 2.0 on, cursors are updateable (and scrollable, so they can move back and forth), so we would not have to go through this effort - but it is a valid example of an update statement.

WebSphere MQ

WebSphere MQ (also and maybe better known as MQ Series) is IBM's messaging and queuing middleware, in use at a great many financial institutions and other companies. It has, from a programming point of view, two API's: JMS (Java Messaging Services), a generic messaging API for the Java world, and MQI, which is older and proprietary to IBM's product. The below examples show the MQI; other examples might show JMS applications.

This is the sample Java application for MQI, translated (and a lot shorter) to NetRexx.

Listing 20.1: MQ Sample

```

1 import com.ibm.mq.MQException
2 import com.ibm.mq.MQGetMessageOptions
3 import com.ibm.mq.MQMessage
4 import com.ibm.mq.MQPutMessageOptions
5 import com.ibm.mq.MQQueue
6 import com.ibm.mq.MQQueueManager
7 import com.ibm.mq.constants.MQConstants
8
9 class MQSample
10 properties private
11
12 qManager = "rjtestqm";
13 qName = "SYSTEM.DEFAULT.LOCAL.QUEUE"
14
15 method main(args=String[]) static binary
16     m = MQSample()
17     do
18         say "Connecting to queue manager: " m.qManager
19         qMgr = MQQueueManager(m.qManager)
20
21         openOptions = MQConstants.MQOO_INPUT_AS_Q_DEF | MQConstants.MQOO_OUTPUT
22
23         say "Accessing queue: " m.qName
24         queue = qMgr.accessQueue(m.qName, openOptions)
25
26         msg = MQMessage()
27         msg.writeUTF("Hello, World!")
28
29         pmo = MQPutMessageOptions()
30
31         say "Sending a message..."
32         queue.put(msg, pmo)
33
34         rcvMessage = MQMessage()
35
36         gmo = MQGetMessageOptions()
37
38         say "...and getting the message back again"
39         queue.get(rcvMessage, gmo)
40
41         msgText = rcvMessage.readUTF()
42         say "The message is: " msgText
43
44         say "Closing the queue"

```

```

45     queue.close()
46
47     say "Disconnecting from the Queue Manager"
48     qMgr.disconnect()
49     say "Done!"
50 catch ex=MQException
51     say "A WebSphere MQ Error occured : Completion Code " ex.completionCode "Reason
      Code " ex.reasonCode
52 catch ex2=java.io.IOException
53     say "An IOException occurred whilst writing to the message buffer: " ex2
54 end

```

This sample connects to the Queue Manager (called *rjtestqm*) in *bindings mode*, as opposed to *client mode*. Bindings mode is only a connection possibility for client programs that are running in the same OS image as the Queue Manager, on the server. Note that the application connects (line 19), accesses a queue (line 23), puts a message (line 32), gets it back (line 39) closes the queue (line 45) and disconnects (line 48) all without checking returncodes: the exceptionhandler takes care of this, and all irregularities will be reported from the catch MQException block starting at line 50).

The main method does in this case not follow the canonical form, but has 'binary' as an extra option. Option binary can be defined on the command line as an option to the translator, as a program option, as a class option and as a method option. Here the smallest scope is chosen. There is a good reason to make this method a binary method: accessing a queue in MQ Series requires some options that are set using a mask of binary flags - this works, in current NetRexx versions, only in binary mode, because the operators have other semantics in nobinary mode.

Listing 20.2: MQ Message Reader

```

1  import com.ibm.mq.
2
3  class MessageReader
4      properties private
5
6      qManager = "rjtestqm";
7      qName    = "TESTQUEUE1"
8
9      method main(args=String[]) static binary
10
11         m = MessageReader()
12         do
13             MQEnvironment.hostname = 'localhost'
14             MQEnvironment.port      = int 1414
15             MQEnvironment.channel   = 'CHANNEL1'
16
17             — exit assignment
18             exits = TimeoutChannelExit()
19             MQEnvironment.channelReceiveExit = exits
20             MQEnvironment.channelSendExit   = exits
21             MQEnvironment.channelSecurityExit = exits
22
23             say "Connecting to QM: " m.qManager
24             qMgr = MQQueueManager(m.qManager)
25
26             openOptions = MQConstants.MQOO_INPUT_AS_Q_DEF
27
28             say "Accessing Queue : " m.qName
29             queue = qMgr.accessQueue(m.qName, openOptions)
30
31             gmo = MQGetMessageOptions() — essential here is that we have MQGMO_WAIT;
32             otherwise we cannot timeout
33             gmo.Options = MQConstants.MQGMO_WAIT | MQConstants.MQGMO_FAIL_IF_QUIESCING |
34                         MQConstants.MQGMO_SYNCPOINT
35             gmo.WaitInterval = MQConstants.MQWI_UNLIMITED
36
37             loop forever
38                 rcvMessage = MQMessage()

```

```

37 queue.get(rcvMessage, gmo)
38 msgText = rcvMessage.readUTF()
39 say "Got a message; the message is: " msgText
40 say
41 end
42
43 catch ex=MQException
44   say "A WebSphere MQ Error occurred : Completion Code " ex.completionCode "Reason
      Code " ex.reasonCode
45   say "Closing the queue"
46   queue.close()
47   say "Disconnecting from the Queue Manager"
48   qMgr.disconnect()
49   say "Done!"
50 end

```

In contrast to the previous sample the MessageReader sample only has one import statement. This is always hotly debated in project teams, one school likes the succinctness of including only the top level import, and only goes deeper when there is ambiguity detected; another school spells out the all imports to the bitter end.

The MessageReader sample connects to another queue, called TESTQUEUE1 (specified in line 7) but here we connect in *client mode*, as indicated by lines 13-15 which specify an MQEnvironment. Other options are using an MQSERVER environment variable or a *Channel Definition Table*.

This program is also uncommon in that it uses MQConstants.MQGMO_WAIT as an option instead of being triggered as a process by a message on a trigger queue. Using this option means that the program waits (stays active, not really busy polling but depending on an OS event) until a new message arrives, which will be processed immediately.

In lines 18-21 a *Channel Exit* is specified. This exit is show in the following example.

Listing 20.3: MQ Java Channel Exit

```

1 import com.ibm.mq.
2 import java.nio.
3
4 class TimeoutChannelExit implements WMQSendExit, WMQReceiveExit, WMQSecurityExit
5
6   properties
7
8   tTask = WatchdogTimer
9   t = java.util.Timer
10  timeout = long
11  initialized = boolean
12
13  method TimeoutChannelExit()
14    say "TimeoutChannelExit Constructor Called"
15    t = java.util.Timer()
16    timeout = long 15000
17
18  method channelReceiveExit(channelExitParms=MQCXP, —
19    channelDefinition=MQCD, —
20    agentBuffer=ByteBuffer) returns ByteBuffer
21
22    do
23      this.tTask.cancel() — cancel the timer task whenever a message is read
24    catch NullPointerException — but catch the null pointer the first time
25    end
26    this.tTask = WatchdogTimer()
27    this.t.schedule(this.tTask, this.timeout)
28    return agentBuffer
29
30  method channelSecurityExit(channelExitParms=MQCXP, —
31    channelDefinition=MQCD, —
32    agentBuffer=ByteBuffer) returns ByteBuffer
33
34    return agentBuffer
35
36  method channelSendExit(channelExitParms=MQCXP, —
37    channelDefinition=MQCD, —

```



```

36         agentBuffer=ByteBuffer) returns ByteBuffer
37     return agentBuffer

```

Listing 20.4: WatchdogTimer

```

1 class WatchdogTimer extends TimerTask
2
3 method WatchdogTimer()
4 method run()
5     say 'WATCHDOG TIMER TIMEOUT: HPOpenView Alert Issued' Date()

```

MQ Series has traditional channel exits (programs that can look at the message contents before the application gets to it). In the MQI Java environment there is something akin to this functionality, but a Java channel exit for MQ Series has to be defined in the application, as shown in the previous example. The function of this particular exit is to implement a *Watchdog timer* - on a separate thread, as shown in the sample that follows the sample channel exit. The timer threatens here to have issues a HP OpenView alert, but that part has been left out.

This particular sample has been designed to do something that is normally a bit harder to do: signal the operations department when something does NOT happen - here the assumption is that there is a payment going over the queue at least once every 20 minutes - when that does not happen, an alert is issued. With every message that goes through, the timer thread is reset, and only when it is allowed to time out, action is undertaken.

Listing 20.5: Publish/Subscribe

```

1 import com.ibm.mq.
2
3 class MQPubSubSample
4
5     properties inheritable
6     queueManagerName = String
7     syncPoint        = Object()
8     props            = Hashtable
9     topicString      = String
10    topicObject       = String
11    subscribers       = Thread[]
12    subscriberCount   = int
13
14    properties volatile inheritable
15    readySubscribers = int 0 —must be defined volatile
16
17    method MQPubSubSample()
18        topicString      = null
19        topicObject      = System.getProperty("com.ibm.mq.pubSubSample.topicObject", "
20            TESTTOPIC")
21        queueManagerName = System.getProperty("com.ibm.mq.pubSubSample.queueManagerName", "
22            rjtestqm")
23        subscriberCount = Integer.getInteger("com.ibm.mq.pubSubSample.subscriberCount",
24            100).intValue()
25        this.props       = Hashtable()
26        this.props.put("hostname", "127.0.0.1")
27        this.props.put("port", Integer(1414))
28        this.props.put("channel", "SYSTEM.DEF.SVRCONN")
29
30    method main(agr:String[]) static binary
31        sample = MQPubSubSample()
32        sample.launchSubscribers()
33
34        /*
35         * wait until all the subscriber threads have finished the subscription
36         */
37        do protect sample.syncPoint
38            loop while sample.readySubscribers < sample.subscriberCount
39            do
31                sample.syncPoint.wait()

```

```

38     catch InterruptedException
39     end
40     end — loop while sample
41 end — do
42
43 sample.doPublish()
44
45 method launchSubscribers()
46   say "Launching the subscribers"
47   subscribers = Thread[subscriberCount]
48
49   threadNo = int 0
50   loop while threadNo < this.subscribers.length
51     this.subscribers[threadNo] = MQPubSubSample.Subscriber("Subscriber" threadNo)
52     this.subscribers[threadNo].start()
53     threadNo = threadNo + 1
54   end
55
56 method doPublish() signals IOException
57   say "method doPublish started"
58   destinationType = int CMQC.MQOT_TOPIC
59   do
60     queueManager = MQQueueManager(this.queueManagerName, this.props)
61     messageForPut = MQMessage()
62     say "***Publishing ***"
63     messageForPut.writeString("Hello world!")
64     queueManager.put(destinationType, topicObject, messageForPut)
65   catch e=MQException
66     say "Exception while publishing " e
67   end
68
69 class MQPubSubSample.Subscriber binary dependent extends Thread
70
71   properties private
72   myName = String
73   openOptionsForGet = int CMQC.MQSO_CREATE | CMQC.MQSO_FAIL_IF QUIESCING | CMQC.
74     MQSO_MANAGED | CMQC.MQSO_NON_DURABLE
75
76 method Subscriber(subscriberName=String)
77   super(subscriberName)
78   myName = subscriberName
79
80 method run()
81   do
82     say myName " — ***Subscribing***"
83     queueManager = MQQueueManager(parent.queueManagerName, parent.props)
84     destinationForGet = queueManager.accessTopic(parent.topicString, parent.
85       topicObject, CMQC.MQTOPIC_OPEN_AS.SUBSCRIPTION, openOptionsForGet)
86
87     do protect parent.syncpoint
88       parent.readySubscribers = parent.readySubscribers + 1
89       parent.syncPoint.notify()
90     end
91
92     mgmo = MQGetMessageOptions()
93     mgmo.options = CMQC.MQGMO_WAIT
94     mgmo.waitInterval = 30000
95     say myName " — ***Retrieving***"
96     messageForGet = MQMessage()
97
98     do protect getClass()
99       destinationForGet.get(messageForGet, mgmo)
100     end
101
102     messageDataFromGet = String messageForGet.readLine()
103     say myName " — Got [" messageDataFromGet "]"
104
105   catch e=Exception
106     say myName " " e
107     e.printStackTrace()
108   end
109   parent.readySubscribers = parent.readySubscribers - 1

```

This sample shows the publish-subscribe interfaces that at some time have been added to the product. This specific sample shows some Java thread complexity but is a good

example of doing publish/subscribe work in a multithreaded way, which is a natural fit for this type of work.

MQTT

21.1 Pub/Sub with MQ Telemetry

Publish/subscribe (pub/sub) is a model that lends itself very well to a number of one publisher, many subscriber type of applications; the tools to enter this technology have never been as available as they are now. Also, MQTT is a small protocol that needs to be taken seriously: Facebook has recently become one of the largest users.

Designed as a low-overhead on-the-wire protocol for brokers in the Internet-of-things age, MQTT is an exciting new development in the Messaging and Queueing realm. It is a good choice for any broker functionality, as the minimal message overhead is 2 bytes, but the maximum messages size, in one of the more popular open source brokers is a good 250MB, which give you a message size that is a lot higher than anything possible in the early years of MQ Series back in the nineties. It is now possible to do development with an entry level, entirely open source suite, and scale up to commercial, clustered and highly available implementations when needed, since the protocol has is supported by the base IBM WebSphere MQ product and is an added deliverable in WSMQ 7.5, after being available as an installable add-on for several years.

Here I will show how extremely straightforward it is to create a pub/sub application using this technology. These examples use NetRexx, the Eclipse PAHO Java client library and the open source Mosquitto broker; all these components are completely free and open source. I have installed Mosquitto on my MacBook using the brew system(fn), which makes it as much trouble as “sudo brew install mosquitto”. NetRexx is an excellent language for these examples, as it is compact and avoids the C-inspired ceremony of Java language syntax; if your project requires Java, you can just save the generated Java source (using the new `-keepasjava` option).

Mosquitto(fn) is written by Roger Light as an open source equivalent of IBM’s `rsmb` (real small message broker) example application, which is free but lacks source code. It is a small broker application that nevertheless runs production sized workloads. As MQTT, as opposed to the MQI or JMS API’s you use when developing a messaging application, is an on-the-wire protocol (commercial messaging systems tend to have their own, unpublished, on-the-wire protocols), we need an API to use it. This API consists of a set of calls that do the formatting of the messages to the requirements of the on-the-wire protocol for you. The messages themselves are just byte-arrays, which gives you the ultimate freedom in designing their content. It is not unusual for connected devices to encode their information in a few bits; on the other hand, there is no reason not to use extreme verbosity in messages; as long as you send the `.getBytes` that your String yields, MQTT will send it. When encoding information in a compact way, the

protocol design will really pay off, because the protocol overhead, in comparison with http and other chatty protocols, is very low. A limited set of quality of service options (qos) will indicate if you want send and pray, acknowledged delivery or acknowledged one-time-only delivery.

The API library that was chosen for these examples is that from the Eclipse PAHO project. This project, which is in its early stages, has C, Javascript and Java client libraries available. I chose the Java client because the JVM environment is where most of the organizations that I work for will use it. The PAHO Java client library is donated by IBM and written by Dave Locke; it is in active development. If you want to see how the protocol moves in packets over the network, I can recommend Wireshark, which does a good job of recognizing them (if you run on the standard port 1883) and showing you the message types (like ACK) and their bytes.

After having put the NetRexx(.jar) and paho client jars on your classpath, you are good to go. The first example here is the publisher – this is not a fragment, but the complete code. For production code we might add some more checks, as enterprise environments always are prone to suddenly run low on disk space and suffer missing authorizations, but it works as it stands. Do note that you do not have to define a message topic in advance – just think of one any use it, at least if you are in your own environment. With Mosquitto, there wasn't anything to define in advance, and the running Publisher (happily lifted from the Java example) in NetRexx was actually the first time I talked to Mosquitto on my MacBook.

Listing 21.1: MQTT Publish Sample

```
1 import java.sql.Timestamp
2 import org.eclipse.paho.client.mqttv3.
3
4 class Publish implements MqttCallback
5
6   method Publish()
7     conOpt = MqttConnectOptions()
8     conOpt.setCleanSession(0)
9     tmpDir = System.getProperty("java.io.tmpdir")
10    dataStore = MqttDefaultFilePersistence(tmpDir)
11    clientId = MqttClient.generateClientId()
12    topicName = "/world"
13    payload = "hello".toString().getBytes()
14    qos = 2
15
16    do
17      broker = "localhost"
18      port = "1883"
19      brokerUrl = "tcp://"broker":"port
20      client = MqttClient(brokerUrl,clientId, dataStore)
21      client.setCallback(this)
22    catch e=mqttException
23      say e.getMessage()
24      e.printStackTrace()
25    end — do
26
27    client.connect()
28    log("Connected to "brokerUrl" with client ID "client.getClientId())
29
30    — Get an instance of the topic
31    topic = client.getTopic(topicName)
32
33    message = MqttMessage(payload)
34    message.setQos(qos)
35
36    — Publish the message
37    time = Timestamp(System.currentTimeMillis()).toString()
38    log('Publishing at: 'time' to topic "'topicName'" with qos 'qos')
39    token = topic.publish(message)
```

```

40
41 — Wait until the message has been delivered to the server
42 token.waitForCompletion()
43
44 — Disconnect the client
45 client.disconnect()
46 log("Disconnected")
47
48
49 method log(line)
50   say line
51
52 method messageArrived(t=MqttTopic,m=MqttMessage)
53   log("Message Arrived: " t m)
54
55 method deliveryComplete(t=MqttDeliveryToken)
56   log("Delivery Complete: " t)
57
58 method connectionLost(t=Throwable)
59   log("Connection Lost:" t.getMessage())
60
61 method main(args=String[]) static
62   Publish()

```

Topics can have a hierarchical organization; this structure is put in by composing trees of topics, which are strings separated by '/'. In this way, it is easy to compose a /news/economics/today topic string that gives some structure to the publication. The classification is entirely up to the designer.

Messaging in its original form is an asynchronous technology, and for this reason the API offers a callback option, where the callback receives the results of your publish action in an asynchronous way. The broker assigns a message id which you receive back.

The second source fragment (and again, it is no fragment but the entire application program) shows the subscriber.

Listing 21.2: MQTT Subscribe Sample

```

1 import java.sql.Timestamp
2 import org.eclipse.paho.client.mqttv3.
3
4 class Subscribe implements MqttCallback
5
6 properties
7   client = MqttClient
8   conOpt = MqttConnectOptions()
9   tmpDir = System.getProperty("java.io.tmpdir")
10  clientId = MqttClient.generateClientId()
11  topicName = "/world"
12  qos = 2
13
14 method Subscribe()
15   do
16     connectAndSubscribe()
17   catch mx=MqttException
18     log(mx.getMessage())
19   end
20   — Block until Enter is pressed
21   log("Press <Enter> to exit");
22   do
23     System.in.read()
24   catch IOException
25   end
26
27   — Disconnect the client
28   client.disconnect()
29   log("Disconnected")
30
31 method connectAndSubscribe() signals MqttSecurityException,MqttException,
    MqttPersistenceException

```

```

32  conOpt.setCleanSession(1)
33  dataStore = MqttDefaultFilePersistence(tmpDir)
34  do
35    broker = "localhost"
36    port = "1883"
37    brokerUrl = "tcp://" + broker + ":" + port
38    client = MqttClient(brokerUrl, clientId, dataStore)
39    client.setCallback(this)
40  catch e=mqttException
41    say e.getMessage()
42    e.printStackTrace()
43  end — do
44
45  this.client.connect()
46  log("Connected to " + brokerUrl + " with client ID " + client.getClientId())
47
48  — Subscribe to the topic
49  log('Subscribing to topic "' + topicName + '" qos ' + qos)
50  this.client.subscribe(topicName, qos)
51
52  method log(line)
53    say line
54
55  method messageArrived(t=MqttTopic, m=MqttMessage)
56    log("Message Arrived: " + t + m)
57
58  method deliveryComplete(t=MqttDeliveryToken)
59    log("Delivery Complete: " + t)
60
61  method connectionLost(t=Throwable)
62    do
63      connectAndSubscribe()
64    catch mx=MqttException
65      log(mx.getMessage())
66    end
67
68  method main(args=String[]) static
69    Subscribe()

```

In the home setup, there is a Raspberry PI running the client while a server in the attic runs the Mosquitto broker. On the Raspberry, which runs Debian wheezy with the soft-float ABI that, at the moment of writing, is still necessary for the Oracle ARM Java implementation; everything done in NetRexx runs unchanged; I just move the classes to it using scp. The broker on the laptop takes care of the scenario in which I suddenly can do some development while not connected to the net, like when I have some moments to reflect on the code in the IKEA restaurant while my spouse runs the serious shopping business.

Security is outside of the scope of this introduction which shows you the sourcecode of a simple pub/sub application, but in Mosquitto the traffic can be secured using SSL certificates and userid/password combinations; also, the access to topics can be limited. In terms of availability, the Mosquitto configuration file offers an opportunity to send all messages for a defined set of topics to another connected broker, which might be in a different part of the world, or your home, to enable a redundant setup. While the broker does not offer the queue – transmission queue - channel setup with retrying channels that MQ does, the client API has some facilities to locally save the messages and retry if the communication was lost. Also, the last-will-and-testament facility is something that traditional MQ does not have.

Component Based Programming: Beans

JavaBeans is the name for the Java component model. It consists of two conventions, for the naming of *getter* and *setter* methods for properties, and the *event* mechanism for sending and receiving events. NetRexx adds support for the automatic generation of getter and setter methods, through the **properties indirect** option on the properties statement.

Using the NetRexxA API

As described elsewhere, the simplest way to use the NetRexx interpreter is to use the command interface (NetRexxC) with the *-exec* or *-arg* flags. There is also a more direct way to use the interpreter when calling it from another NetRexx (or Java) program, as described here. This way is called the *NetRexxA Application Programming Interface* (API).

The *NetRexxA* class is in the same package as the translator (that is, *org.netrexx.process*), and comprises a constructor and two methods. To interpret a NetRexx program (or, in general, call arbitrary methods on interpreted classes), the following steps are necessary:

1. Construct the interpreter object by invoking the constructor *NetRexxA()*. At this point, the environment's classpath is inspected and known compiled packages and extensions are identified.
2. Decide on the program(s) which are to be interpreted, and invoke the *NetRexxA parse* method to parse the programs. This parsing carries out syntax and other static checks on the programs specified, and prepares them for interpretation. A stub class is created and loaded for each class parsed, which allows access to the classes through the JVM reflection mechanisms.
3. At this point, the classes in the programs are ready for use. To invoke a method on one, or construct an instance of a class, or array, etc., the Java reflection API (in *java.lang* and *java.lang.reflect*) is used in the usual way, working on the *Class* objects created by the interpreter. To locate these *Class* objects, the API's *getClassObject* method must be used.

Once step 2 has been completed, any combination or repetition of using the classes is allowed. At any time (provided that all methods invoked in step 3 have returned) a new or edited set of source files can be parsed as described in step 2, and after that, the new set of class objects can be located and used. Note that operation is undefined if any attempt is made to use a class object that was located before the most recent call to the *parse* method.

Here's a simple example, a program that invokes the *main* method of the *hello.nrx* program's class:

Listing 23.1: Try the NetRexxA interface

```

1 options binary
2 import org.netrexx.process.NetRexxA
3
4 interpreter=NetRexxA()           — make interpreter
5
6 files=['hello.nrx']             — a file to interpret
7 flags=['nocrossref', 'verbose0'] — flags, for example

```

```

8 interpreter.parse(files, flags) — parse the file(s), using the flags
9
10 helloClass=interpreter.getClassObject(null, 'hello') — find the hello Class
11
12 — find the 'main' method; it takes an array of Strings as its argument
13 classes=[interpreter.getClassObject('java.lang', 'String', 1)]
14 mainMethod=helloClass.getMethod('main', classes)
15
16 — now invoke it, with a null instance (it is static) and an empty String array
17 values=[Object String[0]]
18
19 loop for 10 — let's call it ten times, for fun...
20   mainMethod.invoke(null, values)
21 end

```

Compiling and running (or interpreting!) this example program will illustrate some important points, especially if a **trace all** instruction is added near the top. First, the performance of the interpreter (or indeed the compiler) is dominated by JVM and other start-up costs; constructing the interpreter is expensive as the classpath has to be searched for duplicate classes, etc. Similarly, the first call to the parse method is slow because of the time taken to load, verify, and JIT-compile the classes that comprise the interpreter. After that point, however, only newly-referenced classes require loading, and execution will be very much faster.

The remainder of this section describes the constructor and the two methods of the NetRexxA class in more detail.

23.1 The NetRexxA constructor

Listing 23.2: Constructor

```
1 NetRexxA()
```

This constructor takes no arguments and builds an interpreter object. This process includes checking the classpath and other libraries known to the JVM and identifying classes and packages which are available.

23.2 The parse method

Listing 23.3: parse

```
1 parse(files=String[], flags=String[]) returns boolean
```

The parse method takes two arrays of Strings. The first array contains a list of one or more file specifications, one in each element of the array; these specify the files that are to be parsed and made ready for interpretation.

The second array is a list of zero or more option words; these may be any option words understood by the interpreter (but excluding those known only to the NetRexxC command interface, such as *time*).⁸ The parse method prefixes the *nojava* flag automatically, to prevent *.java* files being created inadvertently. In the example, *nocrossref* is supplied to stop a cross-reference file being written, and *verbose0* is added to prevent the logo and other progress displays appearing.

The *parse* method returns a boolean value; this will be 1 (true) if the parsing completed without errors, or 0 (false) otherwise. Normally a program using the API should test this

⁸Note that the option words are not prefixed with a -.

result an take appropriate action; it will not be possible to interpret a program or class whose parsing failed with an error.

23.3 The getClassObject method

Listing 23.4: getClassObject

```
1 getClassObject(package=String, name=String [,dimension=int]) returns Class
```

This method lets you obtain a Class object (an object of type *java.lang.Class*) representing a class (or array) known to the interpreter, including those newly parsed by a parse instruction.

The first argument, *package*, specifies the package name (for example, *com.ibm.math*).

For a class which is not in a package, *null* should be used (not the empty string, "").

The second argument, *name*, specifies the class name (for example, *BigDecimal*). For a minor (inner) class, this may have more than one part, separated by dots.

The third, optional, argument, specifies the number of dimensions of the requested class object. If greater than zero, the returned class object will describe an array with the specified number of dimensions. This argument defaults to the value 0.

An example of using the *dimension* argument is shown above where the *java.lang.String[]* array Class object is requested.

Once a Class object has been retrieved from the interpreter it may be used with the Java reflection API as usual. The Class objects returned are only valid until the parse method is next invoked.

23.4 The exiting method

Syntax:

Listing 23.5: exiting

```
1 exiting() returns boolean
```

If this method returns true, an interpreted program has invoked the NetRexx “exit” instruction to shut down the interpreter. If more programs need to be interpreted, a new instance of the interpreter will need to be created with the NetRexxA() constructor.

23.5 Interpreting programs contained in memory strings

RexxLA’s NetRexx 3.01 release adds some extensions to the NetRexxA API to support interpreting programs from memory strings. The first extension adds an optional array of strings containing programs to the standard parse API. It is mainly useful to IDE developers. It also serves as the basis to support two other extensions documented below.

Listing 23.6: parse program in memory buffer

```
1 method parse(filestrings=String[], programstrings=String[], flagstrings=String[],  
   logfile=PrintWriter null, outfile=PrintStream System.out) returns boolean
```

Parses a set of files, under specified flags:

- filestrings is a list of program names,
- programstrings is a list of program strings⁹,

⁹Note that program strings which are not named in the name list are ignored.

- flagstrings is a list of flags,
- logfile is a PrintWriter for parse output messages (optional),
- outfile is a PrintStream for console output messages (optional),

This method returns 1 if no error

The second extension is a new easy to use method to parse and interpret a program contained in a string.

Listing 23.7: parse program in string

```
1 method interpret(programname=String, programstring=String, argstring=String "",
    flagstring=String "", logfile=PrintWriter null, outfile=PrintStream System.out)
    returns boolean
```

A convenience method to interpret a single NetRexx program in a string:

- programname is the program name,
- programstring is the program string,
- argstring is the argument string (optional),
- flagstring is the translator flags string (optional),
- logfile is a PrintWriter for parse output messages (optional),
- outfile is a PrintStream for console output messages (optional),

This method returns 1 if no parse error. The default flag is -verbose0

Here is a simple example using the interpret method:

Listing 23.8: interpret from string

```
1 import org.netrexx.process.
2 netrexxapi=NetRexxA()
3 myprog="say 'argument string is' arg"
4 netrexxapi.interpret("myprog",myprog,"a passed argument")
```

The third extension is a slightly more complex “eval” method that allows a program to call a static method in a program string and receive an object back.

Listing 23.9: eval

```
1 method eval(programname=String, programstring=String, methodname=String, argstring=
    String "", flagstring=String "", logfile=PrintWriter null, outfile=PrintStream
    System.out) returns Object
```

A convenience method to interpret a method from a NetRexx program in a string and return an object:

- programname is the program name,
- programstring is the program string,
- methodname is the method name to call - the method must accept a String array like main methods,
- argstring is the argument string (optional),
- flagstring is the translator flags string (optional),
- logfile is a PrintWriter for parse output messages (optional),
- outfile is a PrintStream for console output messages (optional),

This method returns an object if no error. The default flag is -verbose0

Here is a simple example using the eval method:

Listing 23.10: eval example

```
1 import org.netrexx.process.  
2 netrexxapi=NetRexxA()  
3 termpgmstring='method term(sa=String[]) static returns rexx;i=Rexx(sa);return 1/i'  
4 say netrexxapi.eval("termpgm",termpgmstring,"term",99)
```

Interfacing to Open Object Rexx

24.1 BSF4ooRexx

NetRexx Tools

25.1 Editor support

This chapter lists editors that have plugin support for NetRexx, ranging from syntax coloring to full IDE support (specified), and Rexx friendly editors, that are extensible using Rexx as a macro language (which can be the first step to provide NetRexx editing support).

25.1.1 JVM - All Platforms

JEdit	Full support for NetRexx source code editing, to be found at http://www.jedit.org .
NetRexxDE	A revisions with additions of the NetRexx plugin for jEdit, moving to a full IDE for NetRexx. http://kenai.com/projects/netrexx-misc
Eclipse	Eclipse has a NetRexx plugin that provides a complete IDE environment for the development of NetRexx programs (in alpha release) by Bill Fenlason. The project is situated at SourceForge (http://eclipsenetrexx.sourceforge.net/). Chapter 26 on page 77 discusses the setup of Eclipse to build the translator itself; and has instructions for the setup of the NetRexx plugin.

25.1.2 Linux

Emacs	<code>netrexx-mode.el</code> (in the NetRexx package in the <code>tools</code> directory) runs on GNU Emacs, which is installed by default on most Linux developer distributions.
vim	<code>vi</code> with extensions

25.1.3 MS Windows

Emacs	<code>netrexx-mode.el</code> (in the NetRexx package in the <code>tools</code> directory) runs on GNU Emacs for Windows. http://www.gnu.org/software/emacs/windows/faq.html .
vim	<code>vi</code> with extensions

25.1.4 MacOSX

Aquamacs	A version of Emacs that is integrated with the MacOSX Aqua look and feel. (http://www.aquamacs.org). NetRexx mode is included in the NetRexx package in the <code>tools</code> directory.
Emacs	<code>netrexx-mode.el</code> (in the NetRexx package) runs on GNU Emacs for MacOSX. http://www.gnu.org/software/emacs .
Vim	Vi with extensions

25.2 Java to Nrx (java2nrx)

When working on a piece of Java code, or an example written in the language, sometimes it would be good if we could see the source in NetRexx to make it more readable. This is exactly what *java2nrx* by Marc Remes does. It has a Java 1.5 parser and an Abstract Syntax Tree that delivers a translation to NetRexx, to the extend of what is currently supported under NetRexx.

At the moment it is to be found at <http://kenai.org/NetRexx/contrib/java2nrx>

It is started by the `java2nrx.sh` script; for convenience, place `java2nrx.sh` and `java2nrx.jar` in the same directory. NetRexxC and java must be available on the path.

Usage: Alternatively:

FIGURE 2: Java2nrx 1





FIGURE 3: Java2nrx 2



-nrc runs NetRexxC compiler on output nrx file

-stdout prints NetRexx file on stdout

-run runs generated translated NetRexx output file

Using Eclipse for NetRexx Development

This is a guide for first time Eclipse users to set up a NetRexx development project. It is not a beginners guide to Eclipse, but is intended to explain how to download the NetRexx compiler source from SVN to be able to modify and build it using Eclipse¹⁰.

It is detailed and hopefully foolproof for someone who has never used Eclipse. It assumes a Windows user, but if you are a Linux or Mac user, you will no doubt understand what to do.

This guide is for Eclipse 4.2 (Juno), written August, 2012. New Eclipse releases occur every 4 months, so there may be differences depending on what the current version is.

26.1 Downloading Eclipse

There are many different preconfigured versions of Eclipse. As you become more experienced with it you may wish to use a different distribution, but the one specified here makes some things simple. It does contain some things that you may never use.

1. Make a new folder for the project. Name it appropriately (e.g. EclipseNetRexx)
2. Browse to eclipse.org, and click on “Download”.
3. Download the version named ECLIPSE IDE FOR JAVA DEVELOPERS for your your operating system.
4. The download is about 150 MB.
5. Unzip the downloaded file into your project folder.

26.2 Setting up the workspace

There are different strategies for managing Eclipse workspaces. Eclipse defaults to putting the workspace in your Windows documents folder - probably not what you want to do. The following is perhaps the most simple way.

1. Open the project folder. It will now contain a folder named eclipse.
2. Add a new folder named “workspace” in the project folder to go along with the eclipse folder.
3. Open the eclipse folder, and create a shortcut to eclipse.exe.
4. Move the shortcut to the desktop and rename it to something like “Eclipse NetRexx”.
5. Close the project folder, and double click the shortcut to start Eclipse.

¹⁰If you have questions or comments, feel free to contact Bill Fenlason at billfen@hvc.rr.com.

6. The “Select a workspace” dialog comes up - don’t use the default.
7. Browse to the workspace folder that you just created and select it.
8. Click (check) the “Use this as the default” box, and click OK.

26.3 Shellshock

If you have never used Eclipse, it can be a bit overwhelming. It is rather complicated, and has endless options, etc. In addition there are at least a thousand different plugins.

You will be greeted by a Welcome screen - you may find it interesting or boring. Exit from it via tback to the welcome screen from: Main Menu -> Help -> Welcome.

26.4 Installing SVN

This version of Eclipse comes with CVS and Git support built in, but the SVN support must be installed.

1. Click on Main Menu -> Help -> Eclipse MarketPlace.
2. Type SVN in the search box and hit Enter.
3. Locate Subversive - it will probably be the first entry - and click the Install button.
4. Click Next, I Accept the License and Finish. The SVN plugin will be downloaded.
5. Click Yes to restart Eclipse.
6. The SVN “Install connectors” dialog will start.
7. Select the SVN Kit 1.75.
8. Click Next, Accept the License, Finish, OK to unsigned content, and Yes to restart Eclipse.

26.5 Downloading the NetRexx project from the SVN repository

The SVN repository contains the NetRexx compiler/translator, documentation, examples, etc. These instructions assume you want only the compiler project.

1. The NetRexx SVN repository name is: <https://svn.kenai.com/svn/netrexx~netrexxc-repo>
2. Copy it (for pasting) from above, or get it from the kenai or netrexx.org site.
3. You do not need a period at the end.
4. Click on Main Menu -> File -> New -> Other -> SVN -> Project from SVN, then Next or double click.
5. Select Create a New Repository location, click Next
6. Paste (or type if you must) the repository name into the URL field and click Next
7. The Checkout from SVN - Select Resource dialog will come up. Click Browse
8. Double click on “netrexxc”, and then single click on “trunk” to select it. Click OK
9. Now click Finish in the checkout dialog to bring up the “Checkout As” dialog
10. Leave the selection at the default of “Checkout ... using the New ProjectWizard”, and Finish
11. The New Project dialog comes up - double click on Java and then Java Project (or use Next)

12. The New Java Project dialog comes up. Enter a project name, perhaps something like NetRexx301.
13. Click Finish, and the project is downloaded. It will show up in the Package Explorer on the left.

26.6 Setting up the builds

Ant support is built into Eclipse, but it must be configured to be able to access the bootstrap NetRexx compiler.

1. Double click on the build.xml file name in the package explorer. Note that its icon is an ant.
2. The build file will open in an editor window.
3. Right click in the window to bring up a context menu, and select Run As -> 2 Ant Build
4. Do NOT select 1 Ant Build.
5. The Ant configuration dialog comes up - it will show you all the targets, etc.
6. Click on the Classpath tab, and then click on User Entries.
7. Now click on Add External Jars to bring up the Jar Selection dialog.
8. Navigate to the lib folder in the project folder. Make sure you are not in the build folder.
9. Double click on NetRexxC.jar to select it.
10. Click on the Refresh tab, and check the Refresh resources on completion box.
11. Click Run to build the distribution. The messages will appear in the console listing below.
12. The java doc step may fail.
13. Close the build.xml file (X on the tab).

You can configure the ant build by using the configuration dialog in Run As -> 2 Ant Build. You may want to check “compile” and “jars” to run those steps. Use Apply to save the configuration.

There are two different builds. The second build.xml file is in the project -> tools -> ant-task folder. Open it up and repeat the above steps for that build.xml file. Each build file has its own ant configuration, and once set selecting Run As -> 1 Ant Build will run it. Or just hit F11.

26.7 Using the NetRexx version of the NetRexx Ant task

The above process uses the standard NetRexx Ant task, not the new one. To use the new one:

1. Main Menu -> Window -> Preferences -> Ant -> Runtime.
2. Open up and select Ant Home Entries. Then click on Add External Jars
3. Navigate to the lib folder in the project and select ant-netrexx.jar
4. The jar will appear at the bottom of the list.
5. Use the UP button to move it up (ahead) of the apache ant version, click OK

26.8 Setting up the Eclipse NetRexx Editor Plugin (Optional)

The NetRexx Editor plugin provides syntax coloring and error checking for nrx files, as well as one click compiling and translating.

1. Click on Main Menu -> Help -> Eclipse MarketPlace.
2. Type NetRexx in the search box and hit enter.
3. Click the Install button next to the Eclipse NetRexx package.
4. Click Next, Accept the License, Finish, OK to unsigned content, and Yes to restart Eclipse.
5. Click Main Menu -> Window -> Preferences -> NetRexx Editor to explore it

Platform dependent issues

27.1 Mobile Platforms

Android™ is a version of Linux and friendly to NetRexx programs. Indeed, with NetRexx performing so much better than the closest competition (jRuby, jython) on these devices, there might be a bright future for NetRexx in these environments.

However, there are some drawbacks, caused by the security architecture put in place. Free, unfettered programming like one can do on a desktop machine is a rare occurrence on these devices, and to get programs running on them requires some knowledge of the security architecture that has been put in place for mobile operating systems.

While Apple development still employs a closed model that allows programming only by buying a license with accompanying certificates, and vetting by the App Store employees, and an assumption you will program in Objective-C, Android allows programming but not as straightforward as we know it. To make simple command-line NetRexx programs, both device types need to be *rooted* to allow optimal access. Android allows the installation of applications without vetting by third parties, but dictates a programming model that incurs some overhead - which is a drawback for the occasional scripter.

27.1.1 Android

The security model of Android is based on *least needed privilege* and is implemented by assigning each application a different userid, so that applications on the same device (be it a phone or a tablet) cannot get to each others data. The consequence of this is that simple NetRexx programming and scripting

27.1.2 Apple IOS

Nonwithstanding the current policy of Apple to only allow Objective-C as a programming language on the iPhone and iPad, NetRexx on IOS works fine. This is what one should do to make it work:

1. Jailbreak¹¹ the device. This is necessary until a more sensible setup is used. I used Spirit; it synchs the phone with the hack and then Cydia is installed, an application that does package management the Debian way
2. Choose the "developer profile" on Cydia when asked. This applies a filter to the packages shown (or rather it doesn't) - but you need to do it in order to see the

¹¹Note that jailbreaking an iPhone is against your eula (well - Apple's eula) and might be illegal in some jurisdictions.

prerequisites

3. OpenTerminal will help you to do command line operations on the phone itself
4. The prerequisites are a Java VM (JamVM installs a VM and ClassPath, the open Java implementation) and Jikes, the Java compiler written in C and compiled to the native instruction set of the phone, which is ARM - most processors implementing this have *Jazelle*, a special instruction set to accelerate Java bytecode. However, this feature is seldom used.

The phone can also be logged on to using ssh from your desktop. Do not forget to change the password for the 'root' user and the 'mobile' user, as instructed in the Cydia package.

When this is done, NetRexxC.jar can be copied to the phone. I did this using 'scp NetRexxC.jar mobile@10.0.0.76:' (use the password you just set for this userid) (and because my router assigned 10.0.0.76 to the phone today). I crafted a small 'nrc' script that does a translate and then a Java compile using jikes (and I actually wrote this on the phone using an application called 'iEdit' - nano, vim and other editors are also available but I found the keyboard scheme to type in ctrl-characters a bit tedious - you type a 'ball' character and then the desired ctrl char, while shifting the virtual keyboard through different modes):

nrc:

```
java -cp ~/NetRexxC.jar COM.ibm.netrexx.process.NetRexxC $*
```

Now we can do a compile of the customary hello.nrx with './nrc -keep -nocompile hello' (notice that this is all in the home directory of the 'mobile' user, just like the jar that I just copied. The resulting hello.java.keep can then be mv'ed to hello.java and compiled with 'jikes hello.java'. This produces a class that can be run with 'java -cp NetRexxC.jar hello'

27.2 IBM Mainframe: Using NetRexx programs in z/OS batch

Traditionally the mainframe was a batch oriented environment, and much of the workload that counts still executes in this way. To be able to use NetRexx with Job Control Language (JCL) in batch address spaces, accessing traditional datasets and interacting with the console when needed, we need to know a bit more. This will be explained in these paragraphs.

A standard component of z/OS since version 1.8 or so is `jzos`, which acts as glue between the unix-like abstractions the JVM works with and the time tested way of working on z/OS, with its SAM and VSAM datasets, its Partitioned Data Set (PDS) file organization, the ICF Catalogs and console address space; all of which in existence long before Java reared its head in our IT environments.

The manuals will teach you that there are several ways to interact with HFS/OMVS resources in JCL, but the alternatives to `jzos` have so many drawbacks that it is the only sensible way to run NetRexx programs in the batch environment.

Translator inner workings

28.1 Method resolution

Until version 3.01 of the NetRexx translator a slightly different way of method resolution was used. The chances that this will ever impact your program are very small, but for the sake of history preservation (and to clarify the process that is used) the way in which the translator looks up and decides to find methods in the inheritance tree are documented here.

Index

Class, 51, 53, 68, 69
Options, 56
Properties, 41, 42
Rexx, 33, 45
SELECT, 53
arg, 15-17, 34, 35, 45, 51, 52
binary, 9, 43, 44, 55, 56, 58, 59, 67
catch, 17-19, 33, 42-45, 51-53, 56-59, 62-64
class, 9, 18-20, 34, 41-45, 55-59, 62, 63
constant, 16
dependent, 59
digits, 9, 33, 34
do, 17-19, 33, 34, 42, 45, 46, 51-53, 55-59, 62-64
else, 3, 17, 34, 35, 44, 46, 51-53
end, 7-9, 17-19, 33-35, 41-46, 51-53, 56, 57, 59, 62-64
exit, 17, 35, 51, 53, 56
extends, 19, 42-45, 58, 59
for, 8, 17, 41, 42, 51-53, 67, 68
forever, 8, 56
form, 9
if, 3, 8, 17, 33-35, 43-46, 51-53
implements, 20, 41, 43-45, 57, 62, 63
import, 16, 19, 20, 51, 52, 55-58, 62, 63, 67, 70, 71
indirect, 18, 19
inheritable, 41, 42, 58
interface, 20
interpret, 67, 70
leave, 8
loop, 7, 8, 17-19, 33-35, 41-45, 51-53, 56, 58, 59, 68
method, 9, 15, 16, 18-20, 33-35, 41-45, 55-59, 62-64, 68-70
nop, 46
numeric, 33, 34
options, 9, 33, 34, 43, 44, 59, 67
otherwise, 33, 56
package, 69
parent, 59
parse, 17, 34, 35, 45, 51, 52, 68, 69
private, 18, 33, 35, 55, 56, 59
properties, 18, 19, 55-59, 63
protect, 58, 59
public, 9, 33
queue, 55-57
return, 9, 17, 18, 20, 33-35, 43, 44, 46, 57, 58
returns, 15, 16, 18, 33, 44, 45, 57, 58, 68-70
say, iii, 3, 6-9, 15, 17-20, 33-35, 41, 42, 45, 51-53, 55-59, 62-64, 71
select, 33, 46, 53
signal, 33
signals, 33, 59, 63
sourceline, 9
static, 9, 15, 16, 33-35, 45, 55, 56, 58, 63, 64, 68
super, 9, 20, 42, 59
then, 3, 8, 17, 33-35, 43-46, 51-53
this, 9, 18-20, 43-45, 57-59, 62, 64
to, 7, 9, 33-35, 44, 51-53, 63, 64, 67
trace, 9
until, 63
upper, 33
volatile, 58
when, 33, 46
while, 17-19, 43-45, 51-53, 58, 59

applets for the Web, writing, 43
application programming interface, for
 interpreting, 67
ArchText example, 43

binary arithmetic, used for Web applets, 43

capturing translator output, 15
compiling, from another program, 15
completion codes, from translator, 15
constructor, in NetRexxA API, 68

exiting method, in NetRexxA API, 69

getClassObject method, in NetRexxA API, 69

HTTP server setup, 46

interpreting, API, 67
interpreting, using the NetRexxA API, 67
interpreting/API example, 67

NervousText example, 43

- NetRexxA, API, 67
- NetRexxA, class, 67
- NetRexxA/constructor, 68
- parse method, in NetRexxA API, 68
- PrintWriter stream for capturing translator output, 15
- ref /API/application programming interface, 67
- return codes, from translator, 15
- runtime/web server setup, 46
- Web applets, writing, 43
- Web server setup, 46
- WordClock example, 44

ISBN 978-90-819090-0-6

